

When Do Stop-Loss Rules Stop Losses?*

Kathryn M. Kaminski[†] and Andrew W. Lo[‡]

This Draft: January 3, 2007

Abstract

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

Keywords: Investments; Portfolio Management; Risk Management; Performance Attribution; Behavioral Finance.

*The views and opinions expressed in this article are those of the authors only, and do not necessarily represent the views and opinions of AlphaSimplex Group, MIT, or any of their affiliates and employees. The authors make no representations or warranty, either expressed or implied, as to the accuracy or completeness of the information contained in this article, nor are they recommending that this article serve as the basis for any investment decision—this article is for information purposes only. We thank Dimitris Bertsimas, Margret Bjarnadottir, Michael Brennan, Tom Brennan, Florian Ederer, Mike Epstein, Jasmina Hasanhodzic, Dirk Jenter, Carola Frydman, Leonid Kogan, Gustavo Manso, Stewart Myers, Bernhard Nietert, Jun Pan, Sebastian Pokutta, Michael Stutzer, Svetlana Sussman, and participants at the EURO2006 Conference, the INFORMS 2006 Annual Meeting, the MIT Sloan Finance Lunch, the MIT Sloan Finance Seminar, and SSgA for helpful comments and discussion. Research support from the MIT Laboratory for Financial Engineering is gratefully acknowledged.

[†]MIT Operations Research Center, 77 Massachusetts Ave, E40-149, Cambridge, MA 02139, (617) 253-7412 (voice), katykam@mit.edu (email).

[‡]MIT Sloan School of Management, and AlphaSimplex Group, LLC. Corresponding author: Andrew Lo, MIT Sloan School, 50 Memorial Drive, E52-432, Cambridge, MA 02142-1347, (617) 253-0920 (voice), (617) 258-5727 (fax), a.lo@mit.edu (email).

Contents

1	Introduction	1
2	Literature Review	3
3	A Framework for Analyzing Stop-Loss Rules	5
3.1	Assumptions and Definitions	5
3.2	Generalizations and Qualifications	9
4	Analytical Results	10
4.1	The Random Walk Hypothesis	11
4.2	Mean Reversion and Momentum	12
4.3	Regime-Switching Models	14
5	Empirical Analysis	17
5.1	Basic Results	18
5.2	Performance Attribution	21
5.3	A Comparison of Empirical and Analytical Results	22
5.4	A Behavioral Regime-Switching Model	28
6	Conclusion	31
A	Appendix	34
A.1	Proof of Proposition 1	34
A.2	Proof of Proposition 2	34
A.3	Behavioral Regime-Switching Likelihood Function	35
A.4	Additional Empirical Results	37
	References	42

1 Introduction

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

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

In this paper, we investigate the efficacy of such behavior in the narrow context of stop-loss rules, i.e., rules for exiting an investment after some threshold of loss is reached and re-entered after some level of gains is achieved. We wish to identify the economic motivation for stop-loss policies so as to distinguish between rational and behavioral explanations for these rules. While certain market conditions may encourage irrational investor behavior—for example, large rapid market declines—stop-loss policies are sufficiently ubiquitous that their use cannot always be irrational.

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

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

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

expected return of the portfolio is higher with the stop-loss rule than without it, we conclude that the stop-loss rule does, indeed, stop losses.

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

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

Of course, focusing on expected returns does not account for risk in any way. It may be the case that a stop-loss rule increases the expected return but also increases the risk of the underlying portfolio, yielding ambiguous implications for the risk-adjusted return of a portfolio with a stop-loss rule. To address this issue, we compare the variance of the portfolio with and without the stop-loss rule and find that, in cases where the stop-loss rule involves switching to a lower-volatility asset when the stop-loss threshold is reached, the unconditional variance of the portfolio return is reduced by the stop-loss rule. A decrease in the variance coupled with the possibility of a positive stopping premium implies that, within the traditional mean-variance framework, stop-loss rules may play an important role under certain market conditions.

To illustrate the empirical relevance of our analysis, we apply a simple stop-loss rule to the standard asset-allocation problem of stocks vs. bonds using monthly U.S. equity and bond returns from 1950 to 2004. We find that stop-loss rules exhibit significant positive stopping premiums and substantial reductions in variance over large ranges of threshold values—a remarkable feat for a buy-high/sell-low strategy. For example, in one calibration, the stopping premium is approximately 1.0% per annum, with a corresponding reduction

in overall volatility of 0.8% per annum, and an average duration of the stopping period of less than 1 month per year. Moreover, we observe conditional-momentum effects following periods of sustained losses in equities that seem to produce scenarios where long-term bonds strongly dominate equities for months at a time. These results suggest that the random walk model is a particularly poor approximation to monthly U.S. equity returns, as Lo and MacKinlay (1999) and others have concluded using other methods.

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

2 Literature Review

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

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

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

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

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

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

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

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

⁴For example, lack of age-dependence in allocation, lower wealth and lower education with greater non-participation and under-diversification, and greater sophistication in higher wealth investors have all been considered (see Ameriks and Zeldes, 2004).

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

3 A Framework for Analyzing Stop-Loss Rules

In this section, we outline a framework for measuring the impact of stop-loss policies on investment performance. In Section 3.1, we begin by specifying a simple stop-loss policy and deriving some basic statistics for its effect on an existing portfolio strategy. We describe several generalizations and qualifications of our framework in Section 3.2, and then apply our framework in Section 4 to various return-generating processes including the Random Walk Hypothesis, momentum and mean-reversion models, and regime-switching models.

3.1 Assumptions and Definitions

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

- (A1) The returns $\{r_t\}$ for the portfolio strategy P are stationary with finite mean μ and variance σ^2 .
- (A2) The expected return μ of P is greater than the riskfree rate r_f , and let $\pi \equiv \mu - r_f$ denote the risk premium of P .

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

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

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

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

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

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

$$s_t \equiv \begin{cases} 0 & \text{if } R_{t-1}(J) < -\gamma \text{ and } s_{t-1} = 1 \text{ (exit)} \\ 1 & \text{if } r_{t-1} \geq \delta \text{ and } s_{t-1} = 0 \text{ (re-enter)} \\ 1 & \text{if } R_{t-1}(J) \geq -\gamma \text{ and } s_{t-1} = 1 \text{ (stay in)} \\ 0 & \text{if } r_{t-1} < \delta \text{ and } s_{t-1} = 0 \text{ (stay out)} \end{cases} \quad (2)$$

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

$$r_{st} \equiv s_t r_t + (1 - s_t) r_f . \quad (3)$$

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

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

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

Definition 2 The stopping premium $\Delta_\mu(\mathcal{S})$ of a stop-loss policy \mathcal{S} is the expected return

difference between the stop-loss policy \mathcal{S} and the portfolio strategy P :

$$\Delta_\mu \equiv \mathbb{E}[r_{st}] - \mathbb{E}[r_t] = p_o (r_f - \mathbb{E}[r_t | s_t = 0]) \quad (4)$$

$$\text{where } p_o \equiv \text{Prob}(s_t = 0) \quad (5)$$

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

$$\frac{\Delta_\mu}{p_o} = r_f - \mathbb{E}[r_t | s_t = 0] . \quad (6)$$

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

The stopping premium (4) measures the expected-return difference per unit time between the stop-loss policy \mathcal{S} and the portfolio strategy P , but this metric may yield misleading comparisons between two stop-loss policies that have very different parameter values. For example, for a given portfolio strategy P , suppose \mathcal{S}_1 has a stopping premium of 1% and \mathcal{S}_2 has a stopping premium of 2%; this suggests that \mathcal{S}_2 is superior to \mathcal{S}_1 . But suppose the parameters of \mathcal{S}_2 implies that \mathcal{S}_2 is active only 10% of the time, i.e., 1 month out of every 10 on average, whereas the parameters of \mathcal{S}_1 implies that it is active 25% of the time. On a total-return basis, \mathcal{S}_1 is superior, even though it yields a lower expected-return difference per-unit-time. The stopping ratio Δ_μ/p_o given in (6) addresses this scale issue directly by dividing the stopping premium by the probability p_o . The reciprocal of p_o is the expected number of periods that $s_t=0$ or the expected duration of the stop-loss period. Multiplying the per-unit-time expected-return difference Δ_μ by this expected duration $1/p_o$ then yields the total expected-return difference Δ_μ/p_o between r_f and r_t .

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

easily by imposing a return penalty in (3):

$$r_{st} \equiv s_t r_t + (1 - s_t) r_f - \kappa |s_t - s_{t-1}| \quad (7)$$

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

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

Of course, this simple interpretation of a stop-loss policy's efficacy is based purely on expected return, and ignores risk. Risk matters because it is conceivable that a stop-loss policy with a positive stopping premium generates so much additional risk that the risk-adjusted expected return is less attractive with the policy in place than without it. This may seem unlikely because by construction, a stop-loss policy involves switching out of P into a riskfree asset, implying that P spends more time in higher-risk assets than the combination of P and \mathcal{S} . However, it is important to acknowledge that P and \mathcal{S} are dynamic strategies and static measures of risk such as standard deviation are not sufficient statistics for the intertemporal risk/reward trade-offs that characterize a dynamic rational expectations equilibrium.⁷ Nevertheless, it is still useful to gauge the impact of a stop-loss policy on volatility of a portfolio strategy P , as only one of possibly many risk characteristics of the combined strategy. To that end, we have:

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

$$\Delta_{\sigma^2} \equiv \text{Var}[r_{st}] - \text{Var}[r_t] \quad (8)$$

$$= \text{E}[\text{Var}[r_{st}|s_t]] + \text{Var}[\text{E}[r_{st}|s_t]] - \text{E}[\text{Var}[r_t|s_t]] - \text{Var}[\text{E}[r_t|s_t]] \quad (9)$$

$$= -p_o \text{Var}[r_t|s_t = 0] + p_o(1 - p_o) \left[(r_f - \text{E}[r_t|s_t = 0])^2 - \left(\frac{\mu - \text{E}[r_t|s_t = 0]}{1 - p_o} \right)^2 \right] \quad (10)$$

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

⁷See Merton (1973) and Lucas (1978), for example.

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

$$\Delta_{\text{SR}} \equiv \frac{\text{E}[r_{st}] - r_f}{\sigma_s} - \frac{\mu - r_f}{\sigma}. \quad (11)$$

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

3.2 Generalizations and Qualifications

The basic framework outlined in Section 3.1 can be generalized in many ways. For example, instead of switching out of P and into a completely riskfree asset, we can allow F to be a lower-risk asset but with some non-negligible volatility. More generally, instead of focusing on binary asset-allocation policies, we can consider a continuous function $\omega(\cdot) \in [0, 1]$ of cumulative returns that declines with losses and rises with gains. Also, instead of a single “safe” asset, we might consider switching into multiple assets when losses are realized, or incorporate the stop-loss policy directly into the portfolio strategy P itself so that the original strategy is affected in some systematic way by cumulative losses and gains. Finally, there is nothing to suggest that stop-loss policies must be applied at the portfolio level—such rules can be implemented security-by-security or asset-class by asset-class.

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

However, an important qualification regarding our approach is the fact that we do not derive the simple stop-loss policy (2) from any optimization problem—it is only a heuristic, albeit a fairly popular one among many institutional and retail investors. This is a distinct departure from much of the asset-pricing literature in which investment behavior is modelled as the outcome of an optimizing individual seeking to maximize his expected lifetime utility

⁸See Sharpe (1994), Spurgin (2001), and Lo (2002) for details.

by investing in a finite set of securities subject to a budget constraint, e.g., Merton (1971). While such a formal approach is certainly preferable if the consumption/investment problem is well posed—for example, if preferences are given and the investment opportunity set is completely specified—the simple stop-loss policy can still be studied in the absence of such structure.

Moreover, from a purely behavioral perspective, it is useful to consider the impact of a stop-loss heuristic even if it is not derived from optimizing behavior, precisely because we seek to understand the basis of such behavior. Of course, we can ask the more challenging question of whether the stop-loss heuristic (2) can be derived as the optimal portfolio rule for a specific set of preferences, but such inverse-optimal problems become intractable very quickly (see, for example, Chang, 1988). Instead, we have a narrower set of objectives in this paper: to investigate the basic properties of simple stop-loss heuristics without reference to any optimization problem, and with as few restrictions as possible on the portfolio strategy P to which the stop-loss policy is applied. The benefits of our narrower focus are the explicit analytical results described in Section 4, and the intuition that they provide for how stop-loss mechanisms add or subtract value from an existing portfolio strategy.

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

4 Analytical Results

Having defined the basic framework in Section 3 for evaluating the performance of simple stop-loss rules, we now apply them to several specific return-generating processes for $\{r_t\}$, including the Random Walk Hypothesis in Section 4.1, mean-reversion and momentum processes in Section 4.2, and a statistical regime-switching model in Section 4.3. The simplicity of our stop-loss heuristic (2) will allow us to derive explicit conditions under which stop-loss policies can stop losses in each of these cases.

4.1 The Random Walk Hypothesis

Since the Random Walk Hypothesis is one of the most widely used return-generating processes in the finance literature, any analysis of stop-loss policies must consider this leading case first. Given the framework proposed in Section 3, we are able to derive a surprisingly strong conclusion about the efficacy of stop-loss rules:

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

$$r_t = \mu + \epsilon_t, \quad \epsilon_t \stackrel{\text{IID}}{\sim} \text{White Noise}(0, \sigma_\epsilon^2) \quad (12)$$

then the stop-loss policy (2) has the following properties:

$$\Delta_\mu = p_o(r_f - \mu) = -p_o\pi \quad (13a)$$

$$\frac{\Delta_\mu}{p_o} = -\pi \quad (13b)$$

$$\Delta_{\sigma^2} = -p_o\sigma^2 + p_o(1 - p_o)\pi^2 \quad (13c)$$

$$\Delta_{\text{SR}} = -\frac{\pi}{\sigma} + \frac{\Delta_\mu + \pi}{\sqrt{\Delta_{\sigma^2} + \sigma^2}} \quad (13d)$$

Proof: See Appendix A.1. ■

Proposition 1 shows that, for any portfolio strategy with an expected return greater than the riskfree rate r_f , the Random Walk Hypothesis implies that the stop-loss policy (2) will always reduce the portfolio's expected return since $\Delta_\mu \leq 0$. In the absence of any predictability in $\{r_t\}$, whether or not the stop-loss is activated has no information content for the portfolio's returns; hence, the only effect of a stop-loss policy is to replace the portfolio strategy P with the riskfree asset when the strategy is stopped out, thereby reducing the expected return by the risk premium of the original portfolio strategy P . If the stop-loss probability p_o is large enough and the risk premium is small enough, (13) shows that the stop-loss policy can also reduce the volatility of the portfolio.

The fact that there are no conditions under which the simple stop-loss policy (2) can add value to a portfolio with IID returns may explain why stop-loss rules have been given so little attention in the academic finance literature. The fact that the Random Walk Hypothesis was widely accepted in the 1960's and 1970's—and considered to be synonymous with market efficiency and rationality—eliminated the motivation for stop-loss rules altogether. In fact, our simple stop-loss policy may be viewed as a more sophisticated version of the “filter rule” that was tested extensively by Alexander (1961) and Fama and Blume (1966). Their

conclusion that such strategies did not produce any excess profits was typical of the outcomes of many similar studies during this period.

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

4.2 Mean Reversion and Momentum

In the 1980's and 1990's, several authors documented important departures from the Random Walk Hypothesis for U.S. equity returns,⁹ and, in such cases, the implications for the stop-loss policy (2) can be quite different than in Proposition 1. To see how, consider the simplest case of a non-random-walk return-generating process, the AR(1):

$$r_t = \mu + \rho(r_{t-1} - \mu) + \epsilon_t, \quad \epsilon_t \stackrel{\text{iid}}{\sim} \text{White Noise}(0, \sigma_\epsilon^2), \quad \rho \in (-1, 1) \quad (14)$$

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

This simple process captures a surprisingly broad range of behavior depending on the single parameter ρ , including the Random Walk Hypothesis ($\rho = 0$), mean reversion ($\rho \in (-1, 0)$), and momentum ($\rho \in (0, 1)$). However, the implications of this return-generating process for our stop-loss rule are not trivial to derive because the conditional distribution of r_t , conditioned on $R_{t-1}(J)$, is quite complex. For example, according to (4), the expression for the stopping premium Δ_μ is given by:

$$\Delta_\mu = p_o(r_f - E[r_t | s_t = 0]) \quad (15)$$

but the conditional expectation $E[r_t | s_t = 0]$ is not easy to evaluate in closed-form for an

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

AR(1). For $\rho \neq 0$, the conditional expectation is likely to differ from the unconditional mean μ since past returns do contain information about the future, but the exact expression is not easily computable. Fortunately, we are able to obtain a good first-order approximation under certain conditions, yielding the following result:

Proposition 2 *If $\{r_t\}$ satisfies an AR(1) (14), then the stop-loss policy (2) has the following properties:*

$$\frac{\Delta_\mu}{p_o} = -\pi + \rho\sigma + \eta(\gamma, \delta, J) \quad (16)$$

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

$$\frac{\Delta_\mu}{p_o} \geq -\pi + \rho\sigma \quad (17)$$

Proof: See Appendix A.2. ■

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

In fact, (17) implies that a sufficient condition for a stop-loss policy with reasonable parameters to add value for a momentum-AR(1) return-generating process is

$$\rho \geq \frac{\pi}{\sigma} \equiv \text{SR} \quad (18)$$

where SR is the usual Sharpe ratio of the portfolio strategy. In other words, if the return-generating process exhibits enough momentum, then the stop-loss rule will indeed stop losses.

This may seem like a rather high hurdle, especially for hedge-fund strategies that have Sharpe ratios in excess of 1.00! However, note that (18) assumes that the Sharpe ratio is calibrated at the same sampling frequency as ρ . Therefore, if we are using monthly returns in (14), the Sharpe ratio in (18) must also be monthly. A portfolio strategy with an annual Sharpe ratio of 1.00—annualized in the standard way by multiplying the monthly Sharpe ratio by $\sqrt{12}$ —implies a monthly Sharpe ratio of 0.29, which is still a significant hurdle for ρ but not quite as imposing as 1.00.¹⁰

4.3 Regime-Switching Models

Statistical models of changes in regime, such as the Hamilton (1989) model, are parsimonious ways to capture apparent nonstationarities in the data such as sudden shifts in means and variances. Although such models are, in fact, stationary, they do exhibit time-varying conditional means and variances, conditioned on the particular state that prevails. Moreover, by assuming that transitions from one state to another follow a time-homogenous Markov process, regime-switching models exhibit rich time-series properties that are surprisingly difficult to replicate with traditional linear processes. Regime-switching models are particularly relevant for stop-loss policies because one of the most common reasons investors put forward for using a stop-loss rule is to deal with a significant change in market conditions such as October 1987 or August 1998. To the extent that this motivation is genuine and appropriate, we should see significant advantages to using stop-loss policies when the portfolio return $\{r_t\}$ follows a regime-switching process.

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

$$r_t = I_t r_{1t} + (1 - I_t) r_{2t} \quad , \quad r_{it} \stackrel{\text{iid}}{\sim} \mathcal{N}(\mu_i, \sigma_i^2) \quad , \quad i = 1, 2 \quad (19a)$$

$$A \equiv \begin{matrix} & \begin{matrix} I_{t+1}=1 & I_{t+1}=0 \end{matrix} \\ \begin{matrix} I_t=1 \\ I_t=0 \end{matrix} & \begin{pmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{pmatrix} \end{matrix} \quad (19b)$$

where I_t is an indicator function that takes on the value 1 when state 1 prevails and 0 when state 2 prevails, and A is the Markov transition probabilities matrix that governs the transitions between the two states. The parameters of (19) are the means and variances of the two states, $(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2)$, and the transition probabilities (p_{11}, p_{22}) . Without any loss

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

in generality, we adopt the convention that state 1 is the higher-mean state so that $\mu_1 > \mu_2$. Given assumption (A2), this convention implies that $\mu_1 > r_f$, which is an inequality we will make use of below. The six parameters of (19) may be estimated numerically via maximum likelihood (see, for example, Hamilton, 1994).

Despite the many studies in the economics and finance literatures that have implemented the regime-switching model (19), the implications of regime-switching returns for the investment process has only recently been considered (see Ang and Bekaert, 2004). This is due, in part, to the analytical intractability of (19)—while the specification may seem simple, it poses significant challenges for even the simplest portfolio optimization process. However, numerical results can easily be obtained via Monte Carlo simulation, and we provide such results in Sections 5.

In this section, we investigate the performance of our simple stop-loss policy (2) for this return-generating process. Because of the relatively simple time-series structure of returns within each regime, we are able to characterize the stopping premium explicitly:

Proposition 3 *If $\{r_t\}$ satisfies the two-state Markov regime-switching process (19), then the stop-loss policy (2) has the following properties:*

$$\Delta_\mu = p_{o,1}(r_f - \mu_1) + p_{o,2}(r_f - \mu_2) \quad (20)$$

$$\frac{\Delta_\mu}{p_o} = (1 - \tilde{p}_{o,2})(r_f - \mu_1) + \tilde{p}_{o,2}(r_f - \mu_2) \quad (21)$$

where

$$p_{o,1} \equiv \text{Prob}(s_t=0, I_t=1) \quad (22a)$$

$$p_{o,2} \equiv \text{Prob}(s_t=0, I_t=0) \quad (22b)$$

$$\tilde{p}_{o,2} \equiv \frac{p_{o,2}}{p_o} = \text{Prob}(I_t=0 | s_t=0) . \quad (22c)$$

If the riskfree rate r_f follows the same two-state Markov regime-switching process (19), with expected returns r_{f1} and r_{f2} in states 1 and 2, respectively, then the stop-loss policy (2) has the following properties:

$$\Delta_\mu = p_{o,1}(r_{f1} - \mu_1) + p_{o,2}(r_{f2} - \mu_2) \quad (23)$$

$$\frac{\Delta_\mu}{p_o} = (1 - \tilde{p}_{o,2})(r_{f1} - \mu_1) + \tilde{p}_{o,2}(r_{f2} - \mu_2) . \quad (24)$$

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

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

$$\tilde{p}_{o,2} \geq \frac{\mu_1 - r_f}{\mu_1 - \mu_2} \quad (25)$$

the corresponding stopping premium Δ_μ will be non-negative. By convention, $\mu_1 > \mu_2$, and by assumption (A2), $\mu_1 > r_f$, therefore the sign of the right side of (25) is positive. If r_f is less than μ_2 , then the right side of (25) is greater than 1, and no value of $\tilde{p}_{o,2}$ can satisfy (25). If the expected return of equities in both regimes dominates the riskfree asset, then the simple stop-loss policy will always decrease the portfolio's expected return, regardless of how accurate it is. To see why, recall that returns are independently and identically distributed within each regime, and we know from Section 4.1 that our stop-loss policy never adds value under the Random Walk Hypothesis. Therefore, the only source of potential value-added for the stop-loss policy (2) under a regime-switching process is if the equity investment in the low-mean regime has a lower expected return than the riskfree rate, i.e., $\mu_2 < r_f$. In this case, the right side of (25) is positive and less than 1, implying that sufficiently accurate stop-loss policies will yield positive stopping premia.

Note that the threshold for positive stopping premia in (25) is decreasing in the spread $\mu_1 - \mu_2$. As the difference between expected equity returns in the high-mean and low-mean states widens, less accuracy is needed to ensure that the stop-loss policy adds value. This may be an important psychological justification for the ubiquity of stop-loss rules in practice. If an investor possesses a particularly pessimistic view of the low-mean state—implying a large spread between μ_1 and μ_2 —then our simple stop-loss policy may appeal to him even if its accuracy is not very high.

5 Empirical Analysis

To illustrate the potential relevance of our framework for analyzing stop-loss rules, we consider the performance of (2) when applied to the standard household asset-allocation problem involving just two asset classes: stocks and long-term bonds. Using monthly stock- and bond-index data from 1950 to 2004, we find that stop-loss policies produce surprising conditional properties in portfolio returns, stopping losses over a wide range of parameter specifications. Our simple stop-loss rule seems to be able to pick out periods in which long-term bonds substantially out-perform equities, which is especially counterintuitive when we consider the unconditional properties of these two asset classes historically.

For our empirical analysis, we use the monthly CRSP value-weighted returns index to proxy for equities and monthly long-term government bond returns from Ibbotson and Associate to proxy for bonds. We also consider Ibbotson’s short-term government bond returns for purposes of comparison. Our sample runs from January 1950 to December 2004, the same time span used by Ang and Berkart’s (2004) study of regime-switching models and asset allocation. In Section 5.4, we consider the longer time span from January 1926 to December 2004 to reduce estimation error for our behavioral regime-switching model estimates.

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

Table 1: Summary statistics for the CRSP Value-Weighted Total Market Index, and Ibbotson Associates Long-Term and Short-Term Government Bond Indexes, from January 1950 to December 2004.

In Table 1, we summarize the basic statistical properties of our dataset. To be consistent with practice, we implement our stop-loss policies using simple returns, but also provide means and standard deviations of log returns for equities and bonds in Table 2 to calibrate some of our simulations. The results in Table 1 are well known and require little commentary: stocks outperform bonds, long-term bonds outperform short-term bonds, and the corresponding annual volatilities are consistent with the rank-ordering of mean returns.

In Section 5.1, we present the performance statistics of our stop-loss policy applied to our stock and bond return series. We provide a more detailed performance attribution of

the stop-loss policy in Section 5.2. In Section 5.3, we compare our empirical findings to simulated results under the Random Walk Hypothesis, momentum and mean reversion, and regime switching. We conclude that stop-loss rules apparently exploit momentum effects in equities and long-term bonds following periods of sustained losses in equities.

5.1 Basic Results

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

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

Figure 1 also shows that Δ_μ seems to decrease with larger cumulative-return windows, especially for $J=6$ and 12 months. This finding is consistent with Δ_μ increasing in p_o when the riskfree rate r_f is higher than the conditional expected return of equities, conditioned on being stopped out (see equation (15)). For reference, we plot p_o in Figure 2.

For reference, we also plot p_o in Figure 2 and find that p_o is monotonically decreasing with γ as we would expect. In addition, p_o generally ranges between 5% and 10% implying that stop-loss rules stop-out rather infrequently, approximately once a year or once every two years. Nevertheless, these infrequent decisions seem to add considerable value to a buy-and-hold equity portfolio.

Figure 1 also contains plots of the stopping ratio Δ_μ/p_o and the figure shows that the stop-loss policy yields an incremental 0.5% to 1% increase in expected return on a monthly basis. The worst Δ_μ/p_o occurs for the 3-month cumulative-return window, and the best Δ_μ/p_o is obtained for large thresholds with an 18-month window size. For the shorter window lengths, smaller thresholds provide less value-added but the value remains positive. However, for the 18-month window, larger thresholds perform better. This connection be-

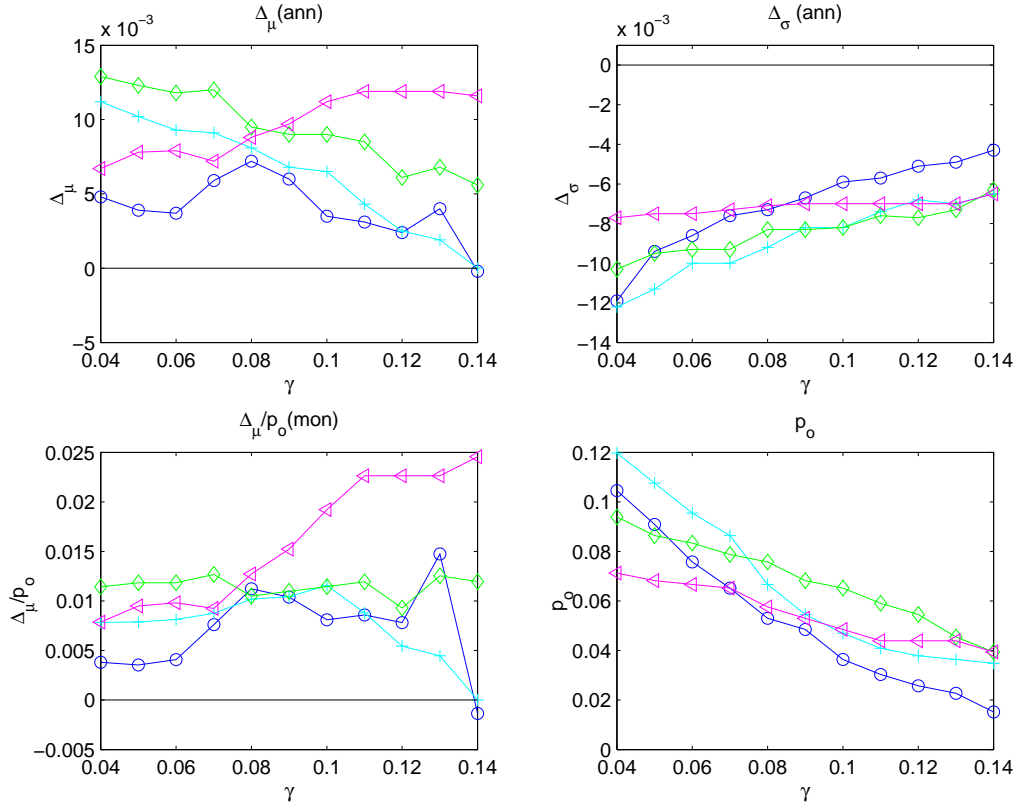


Figure 1: Stop-loss performance metrics Δ_μ , Δ_σ , $\frac{\Delta_\mu}{p_o}$, and p_o for the simple stop-loss policy over stopping thresholds $\gamma = 4\text{--}14\%$ with $\delta = 0\%$, $J = 3$ months (\circ), 6 months ($+$), 12 months (\diamond), and 18 months (\triangle).

tween stop-loss threshold and cumulative-return window size suggests that there is some fundamental relation—either theoretical or behavioral—between the duration of losses and their magnitude.




Figure 2: Stop-loss performance metrics for Δ_{SR} for the simple stop-loss policy over stopping thresholds $\gamma = 4\text{--}14\%$ with $\delta = 0\%$, $J = 3$ months (\circ), 6 months ($+$), 12 months (\diamond), and 18 months (\triangle).

In Tables A.2 and A.3 of Appendix A.4, we examine the performance of equities and bonds during stopped-out periods for stop-loss thresholds $\delta = 0\%$ and $\delta = 2\%$, and find that bonds have significantly better performance with the same level of volatility whereas stocks show reduced performance and increased volatility. We apply a Kolmogorov-Smirnov test to compare the returns before and after stop-loss policies are triggered, and find statistically significant p -values, indicating a difference between the marginal distribution of returns in and out of stop-out periods (see Table A.4).

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

Figure 1 also includes a plot of Δ_σ , which shows that volatility is always reduced by

the stop-loss policy, but the reduction is smaller for larger stopping thresholds γ . This is to be expected because larger thresholds imply that the stop-loss policy is activated less often. Nevertheless, the reduction in variance is remarkably pronounced for a strategy which so rarely switches out of equities (see Tables A.2 and A.3 for the relative frequency and duration of stop-outs). This reduction seems to be coming from two sources: switching to a lower-volatility asset, and avoiding subsequently higher-volatility periods in equities.

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

Based on our empirical analysis, we conclude that stop-loss policies could indeed have added value to the typical investor when applied to equities and long-term bonds from 1950 to 2004. In the next two sections, we provide a more detailed analysis of these results by conducting a performance attribution for the two assets, and by examining several models for asset returns to gauge how substantial these effects are.

5.2 Performance Attribution

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

In Figure 3, we plot the empirical cumulative distribution functions (CDFs) for equities, long-term bonds, and their difference for stopped-out and non-stopped-out returns, and in Figure 4, we plot the corresponding return histograms for equities and long-term bonds during stopped-out periods, non-stopped-out periods, and both. Figure 3 shows that for long-term bonds, returns during stopped-out periods seem to first-order stochastically dominate returns during non-stopped-out period, and that stopped-out returns exhibit a much larger positive skew. In contrast, equities have larger negative returns and a few larger positive returns, coupled with larger volatility.

When we examine the difference between long-term bonds and equities, we find that the




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

CDF of the stopped-out periods almost first-order stochastically dominates the CDF of the non-stopped-out periods, and the positive skew is due to both the increased positive skew in long-term bonds and the large negative returns in equities. The stopped-out difference does not stochastically dominate the non-stopped out periods due to the few large positive returns in equities during stopped-out periods. By examining these conditional CDFs, we conclude that performance during stopped-out periods is generally good because equities tend to have persistent negative performance and long-term bonds generate excess performance during the periods following negative equity returns. In addition, long-term bonds do not stochastically dominate equities because of the few large reversals in equity returns.

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

5.3 A Comparison of Empirical and Analytical Results

To develop further intuition for the empirical results of Section 5.1, we conduct several simulation experiments in this section for the return-generating processes of Section 4. These simulations will serve as useful benchmarks to gauge the economic significance of our empirical results, and can also provide insights into the specific sources of value-added of our stop-loss policy.

We simulate three return-generating processes: the Random Walk Hypothesis, an AR(1) with positive ρ (momentum), and the regime-switching model (19). For each process, we simulate 10,000 histories of artificial equity and bond return series, each series containing




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




Figure 5: Empirical CDFs of Ibbotson Associates Long-Term Government Bond returns (r_b) vs. CRSP Value-Weighted Total Market returns (r_e), for returns during stopped-out periods (50 data points, dotted line) and non-stopped out periods (610 data points, solid line) with stop-loss parameters $J=12$, $\gamma=8\%$, and $\delta=0\%$, from January 1950 to December 2004.

660 normally distributed monthly returns (the same sample size as our data), and calibrated to match the means and standard deviations of our data. The parameter estimates used for the IID and AR(1) cases are given in Table 2, and the regime-switching parameter estimates, estimated by maximum likelihood, are given in Table 3.

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

Given our analysis of the Random Walk Hypothesis in Section 4.1, it is clear that IID returns will yield a negative stopping premium. According to Proposition 1, we know the value of the stopping premia Δ_μ depends on our choice of stopping threshold only through p_o , and the value of $\frac{\Delta_\mu}{p_o} = r_f - \mu$ is constant. Figure 6 confirms these implications, and also illustrates the gap between the Random Walk simulations and the empirical results which are plotted using the symbol “o”. The t -statistics associated with tests that the empirical performance metrics Δ_μ , Δ_σ , and Δ_{SR} are different from their simulated counterparts are all highly significant at the usual levels, implying resounding rejections of the Random Walk Hypothesis. Alternatively, for our simulations to be consistent with our empirical findings, long-term bonds would have to earn a premium over equities of approximately 1% per month,

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

Table 2: Parameter estimates for monthly log returns under both IID and AR(1) return-generating processes for the CRSP Value-Weighted Total Market Index, and IID return-generating process for and Ibbotson Associates Long-Term and Short-Term Government Bond Indexes, from January 1950 to December 2004.

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

Table 3: Maximum likelihood estimates for a regime-switching model with constant transition probabilities for the CRSP Value-Weighted Total Market return, and Ibbotson Associates Long-Term and Short-Term Government Bond returns, from January 1950 to December 2004.




Figure 6: Empirical and simulated performance metrics $\Delta_\mu, \Delta_\sigma, \frac{\Delta_\mu}{p_o}$, and p_o for the simple stop-loss policy with stopping thresholds $\gamma = 4\text{--}14\%$, $\delta = 0\%$, $J = 12$ months. The empirical results (o) are based on monthly returns of the CRSP Value-Weighted Total Market Index and Ibbotson Associates Long-Term Bond Index from January 1950 to December 2004. The simulated performance metrics are averages across 10,000 replications of 660 monthly normally distributed returns for each of three return-generating processes: IID (+), an AR(1) (Δ), and a regime-switching model (*).

and equities would have to have much higher volatility than their historical returns have exhibited.

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

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

which is comparable to the stopping ratio under the Random Walk Hypothesis, -0.0045 .

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

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

Given empirical values for Δ_μ/p_o , we can back out the implied value of ρ under an AR(1); these implied values are given in Table 4. Clearly, these implied autocorrelations are unrealistically high for monthly equity returns, suggesting that simple AR(1) momentum cannot explain the empirical success of our stop-loss policy.

The third set of simulations is based on the regime-switching model (19) where long-term bonds are also assumed to vary across regimes, and the parameter estimates in Table 3 show some promise of capturing certain features of the data that neither IID nor AR(1) processes can generate. The conditional asymmetry of the two regimes is characterized by one regime with higher returns in equities and lower returns in bonds, and the other with lower returns in equities and higher returns in bonds. Using Proposition 3 (the case with a regime-switching riskfree asset), we can gauge the level of accuracy required of our regime-switching model to obtain a positive stopping premium. Recall from (24) that

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

Using this simple result, we see that the stop-loss strategy must correctly switch into bonds with 69.9% accuracy to yield a positive stopping premium. Given the level of volatility in asset returns, it is unrealistic to expect any stopping rule to be able to distinguish regimes with such accuracy. To confirm this intuition, we simulate the regime-switching model using the parameter estimates in Table 3 and plot the implied accuracy $\tilde{p}_{o,2}$ over a large range of stop-loss rules in Figure 7. The 3-month stopping window outperforms the other stopping windows, especially for large stopping thresholds γ , but none of the implied accuracies comes close to the required accuracy of 69.9% to yield a positive stopping premium. Despite the




Figure 7: The probability of correctly being out of equities during the low-mean regime for equities, or $\tilde{p}_{o,2}$, under Hamilton’s (1989) regime-switching model for Basic Stop Loss Rules over various stopping thresholds $\gamma = 4\text{--}14\%$ with $\delta = 0\%$, $J = 3$ (o), 6 (+), 9 (\diamond), and 18 (\triangle) months, threshold for positive Δ_μ .

intuitive appeal of the regime-switching model, it cannot easily account for the empirical success of our simple stop-loss policy.

5.4 A Behavioral Regime-Switching Model

Given the lack of success in the regime-switching model (19) to explain the empirical performance of the simple stop-loss policy, we propose an alternative based on the flight-to-safety phenomenon. The motivation for such an alternative is the mounting empirical and experimental evidence that investors have two modes of behavior: a normal state, and a distressed or panic state.¹¹ An implication of this behavior is that investors are asymmetrically impacted by losses, resulting in a flight to safety. The “distress state” is characterized by a lower mean in equities, as well as a higher mean in bonds, and one possible trigger is a sufficiently large cumulative decline in an investor’s wealth, e.g., a 401(k) account (Agnew, 2003)

This phenomenon can be captured parsimoniously by extending the regime-switching model (19) to allow the regime-switching probabilities to be time-varying and dependent on

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

a cumulative sum of past asset returns:

$$\text{Prob}(I_t=0 | I_{t-1}=1) = \frac{\exp(a_1 + b_1 R_{t-1}(J))}{1 + \exp(a_1 + b_1 R_{t-1}(J))} \quad (26a)$$

$$\text{Prob}(I_t=1 | I_{t-1}=0) = \frac{\exp(a_1 + b_1 R_{t-1}(J))}{1 + \exp(a_2 + b_2 R_{t-1}(J))}. \quad (26b)$$

The motivation for such a specification is to capture the flight-to-safety effect where the probability of switching to the distress state increases as cumulative losses mount, which implies a negative b_1 coefficient if we continue to adopt the convention that state 1 is the higher-mean state.¹² This behavioral regime-switching model can be estimated via maximum likelihood estimation following an approach similar to Ang and Bekaert (2004) (see Appendix A.3 for details), and the parameter estimates for our monthly equity and long-term bond return series are given in Table 5. With the exception of the case where $J = 18$, the b_1 coefficient estimates are indeed negative, consistent with the flight-to-safety phenomenon. Moreover, the coefficient estimates b_2 are positive and much larger in absolute value than the b_1 estimates, implying a stronger tendency to return to the high-mean state from the low-mean state given a cumulative gain of the same absolute magnitude. The fact that both b_1 and b_2 estimates are the largest in absolute value for the shortest horizon $J = 3$ is also consistent with the behavioral evidence that losses and gains concentrated in time have more salience than those over longer time periods.

Using the maximum likelihood estimates in Table 5, we can compute the implied accuracy $\tilde{p}_{o,2}$ required to achieve a positive stopping premium, and these thresholds are given in Table 6. These more plausible thresholds—for example, 58.9% for 3-month returns—show that a regime-switching model, modified to include time-varying transition probabilities based on cumulative returns, is capable of explaining the empirical results of Section 5. Moreover, a simulation experiment similar to those of Section 5.3, summarized in Table 7, also yields levels of implied accuracy levels required to yield positive stopping premia.

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

¹²According to (26a), a negative value for b_1 implies that cumulative losses would increase the probability of transitioning from state 1 to state 2.

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

Table 5: Maximum likelihood estimates of the behavioral regime-switching model for monthly and annual log-returns for the CRSP Value-Weighted Total Market Index and Ibbotson Associates Long-Term Government Bond Index, from January 1950 to December 2004, and for cumulative-return windows $J=3, 6, 12,$ and 18 months.

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

Table 6: Implied lower bound for the accuracy $\tilde{p}_{o,2}$ of the simple stop-loss policy to ensure a positive stopping premia, based on maximum likelihood estimates of the behavioral regime-switching model applied to monthly returns of the CRSP Value-Weighted Total Market Index and Ibbotson Associates Long-Term Government Bond Index, from January 1950 to December 2004.

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

Table 7: Simulated values for implied $\tilde{p}_{o,2}$, and thresholds for positive stopping premium based on maximum likelihood parameter estimates of the behavioral regime-switching model with behavioral cumulative-return windows of length n and stop-loss cumulative-return windows of length J .

that are more consistent with those in the data.

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

6 Conclusion

In this paper, we provide an answer to the question when do stop-loss rules stop losses? The answer depends, of course, on the return-generating process of the underlying investment for which the stop-loss policy is implemented, as well as the particular dynamics of the stop-loss policy itself. If “stopping losses” is interpreted as having a higher expected return with the stop-loss policy than without it, then for a specific binary stop-loss policy, we derive various conditions under which the expected-return difference—which we call the stopping premium—is positive. We show that under the most common return-generating process—

the Random Walk Hypothesis—the stopping premium is always negative. The widespread cultural affinity for the Random Walk Hypothesis, despite empirical evidence to the contrary, may explain the general indifference to stop-loss policies in the academic finance literature.

However, under more empirically plausible return-generating processes such as momentum or regime-switching models, we show that stop-loss policies can generate positive stopping premia. And when applied to the standard household asset-allocation decision between U.S. equities and long-term bonds from January 1950 to December 2004, we find a substantially positive stopping premium with a correspondingly large reduction in variance. These empirical results suggest important nonlinearities in aggregate stock and bond returns that have not been fully explored in the empirical finance literature. For example, our analysis suggests elevated levels of momentum associated with large negative returns, and asymmetries in asset returns following periods of cumulative losses.

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

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

Of course, our findings shed little light on the controversy between market efficiency and behavioral finance. The success of our simple stop-loss policy may be due to certain nonlinear aspects of stock and bond returns from which our strategy happens to benefit, e.g., avoiding momentum on the downside and exploiting asymmetries in asset returns following periods

of negative cumulative returns. And from the behavioral perspective, our stop-loss policy is just one mechanism for avoiding or anticipating the usual pitfalls of human judgment, e.g., the disposition effect, loss aversion, ambiguity aversion, and flight-to-safety.

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

A Appendix

In this appendix, we provide proofs of Propositions 1 and 2 in Sections A.1 and A.2, a derivation of the likelihood function of the behavioral regime-switching model (26) in Section A.3, and present some additional empirical results in Section A.4.

A.1 Proof of Proposition 1

The conclusion follows almost immediately from the observation that the conditional expectations in (4) and (6) are equal to the unconditional expectations because of the Random Walk Hypothesis (conditioning on past returns provides no incremental information), hence:

$$\Delta_\mu = -p_o\pi \leq 0 \quad (\text{A.1})$$

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

and the other relations follow in a similar manner. ■

A.2 Proof of Proposition 2

Let r_t be a stationary AR(1) process:

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

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

$$\text{E}[r_t | s_t = 0] = \text{E}[r_t | R_{t-1}(J) < -\gamma, R_{t-2}(J) \geq -\gamma] \quad (\text{A.4})$$

$$= \mu(1 - \rho) + \rho \text{E}[r_{t-1} + \epsilon_t | R_{t-1}(J) < -\gamma, R_{t-2}(J) \geq -\gamma] \quad (\text{A.5})$$

$$= \mu(1 - \rho) + \rho \text{E}[r_{t-1} | R_{t-1}(J) < -\gamma, R_{t-2}(J) \geq -\gamma] \quad (\text{A.6})$$

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

$$\mathbb{E}[r_t | s_t = 0] = \mu(1 - \rho) + \rho \mathbb{E}[r_{t-1} | R_{t-1}(J) < -\gamma, R_{t-2}(J) \geq -\gamma] \quad (\text{A.7})$$

$$= \mu(1 - \rho) + \rho \mathbb{E}_y[\mathbb{E}[r_{t-1} | r_{t-1} < -\gamma - y, r_{t-J-1} \geq -\gamma - y]] \quad (\text{A.8})$$

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

$$\mathbb{E}_y[\mathbb{E}[r_{t-1} | r_{t-1} < -\gamma - y]] \leq \mathbb{E}_{r_{t-J-1}}[\mathbb{E}[r_{t-1} | r_{t-1} < r_{t-J-1}]] \quad (\text{A.9})$$

$$\leq \mu - \sigma \quad (\text{A.10})$$

which implies:

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

$$\leq \mu - \rho\sigma \quad (\text{A.12})$$

■

A.3 Behavioral Regime-Switching Likelihood Function

The behavioral regime-switching model begins with the standard regime-switching model (19):

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

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

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

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

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

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

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

$$R_{t-1}(n) = r_{t-1} + \dots + r_{t-n} \quad (\text{A.15})$$

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

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

$$f(\mathbf{r}|\theta) = \prod_{t=1}^T \left(f(r_t|\mathcal{F}_{t-1}, I_t=1; \theta) \text{Prob}(I_t=1|\mathcal{F}_{t-1}; \theta) + f(r_t|\mathcal{F}_{t-1}, I_t=0; \theta) \text{Prob}(I_t=0|\mathcal{F}_{t-1}; \theta) \right) \quad (\text{A.16})$$

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

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

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

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

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

where

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

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

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

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

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

Then the likelihood function may be rewritten more compactly as:

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

$$p_{1t} = (1 - g_1(R_{t-1})) \frac{\phi_{1t-1}p_{1t-1}}{\phi_{1t-1}p_{1t-1} + \phi_{2t-1}p_{2t-1}} + g_2(R_{t-1}) \frac{\phi_{2t-1}p_{2t-1}}{\phi_{1t-1}p_{1t-1} + \phi_{2t-1}p_{2t-1}} \quad (\text{A.21b})$$

$$p_{2t} = g_1(R_{t-1}) \frac{\phi_{1t-1}p_{1t-1}}{\phi_{1t-1}p_{1t-1} + \phi_{2t-1}p_{2t-1}} + (1 - g_2(R_{t-1})) \frac{\phi_{2t-1}p_{2t-1}}{\phi_{1t-1}p_{1t-1} + \phi_{2t-1}p_{2t-1}} \quad (\text{A.21c})$$

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

A.4 Additional Empirical Results

In this section, we provide four additional tables to supplement the empirical results in the main text. In Table A.1, we present a more detailed set of summary statistics for the buy-and-

hold equities strategy of Section 5 with and without the stop-loss policy, including means, standard deviations, Sharpe ratios, and skewness and kurtosis coefficients for various stop-loss parameters (γ, δ, J) . In Tables A.2 and A.3, we present similar performance statistics, but only for returns from the stopped-out periods, assuming a re-entry threshold of 0% in Table A.2 and 2% in Table A.3. And in Table A.4, we report p -values of Kolmogorov-Smirnov test statistics designed to distinguish between the unconditional returns of our two asset classes and their conditional counterparts, conditioned on being stopped-out.

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

Table A.1: Performance statistics of a buy-and-hold strategy for the CRSP Value-Weighted Total Market return with and without a simple stop-loss-policy, where the stop-loss asset yields the Ibbotson Associates Long-Term Government Bond return, for stop-loss thresholds $\gamma = 4-14\%$, re-entry threshold $\delta = 0\%$, 2% , and window sizes $J = 3, 6, 12,$ and 18 months, from January 1950 to December 2004.

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

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

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

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

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

Table A.4: p -values of Kolmogorov-Smirnov tests for the equality of the empirical distributions of monthly returns unconditionally and after stop-loss triggers, for the CRSP Value-Weighted Total Market Index and Ibbotson Associates Long-Term Bond Index from January 1950 to December 2004.

References

- Agnew, J., P. Balduzzi, and A. Sunden, 2003, "Portfolio Choice and Trading in a Large 401(K) Plan", *American Economic Review* 93, 193–215.
- Agnew, J., 2003, "An Analysis of How Individuals React to Market Returns in One 401(K) Plan", Working paper, The College of William and Mary.
- Ait-Sahalia, Y. and M. Brandt, 2001, "Variable Selection for Portfolio Choice", *Journal of Finance* 56, 1297–1351.
- Alexander, S., 1961, "Price Movements in Speculative Markets: Trends or Random Walks", *Industrial Management Review* 2, 7–26.
- Ameriks, John, and Stephen Zeldes, 2004, "How Do Household Portfolio Shares Vary with Age?", working paper, Columbia University.
- Ang, A. and G. Bekaert, 2004, "How Regimes Affect Asset Allocation", *Financial Analysts Journal* 60, 86–99.
- Ang, A., Bekaert, G., and J. Liu, 2005, "Why Stocks May Disappoint", *Journal of Financial Economics* 76, 471–508.
- Balduzzi, P., and A.W. Lynch, 1999, "Transaction Costs and Predictability: Some Utility Cost Calculations", *Journal of Financial Economics* 52, 47–78.
- Barber, B. and T. Odean, 2000, "Trading is Hazardous to Your Wealth: The Common Stock Investment Performance of Individual Investors", *Journal of Finance* 55, 773–806.
- Barber, B. and T. Odean, 2005, "All that Glitters: The Effect of Attention and News on the Buying Behavior of Individual and Institutional Investors", unpublished working paper.
- Barberis, N., 2000, "Investing in the Long Run when Returns are Predictable", *Journal of Finance* 55, 225–264.
- Barberis, N., and Thaler, R., 2003, "A Survey of Behavioral Finance", in G. Constantinides, M. Harris, and R. Stulz (eds.), *Handbook of the Economics and Finance*. New York: Elsevier.
- Barrett, G. and S. Donald, 2003, "Consistent Tests for Stochastic Dominance", *Econometrica* 71, 71–104.
- Bell, D., 1982a, "Regret in Decision Making under Uncertainty", *Operations Research* 30, 961–981.
- Bell, D., 1982b, "Risk Premiums for Decision Regret", *Management Science* 29, 1156–1166.
- Bell, D., 1985, "Disappointment in Decision Making under Uncertainty", *Operations Research* 33, 1–27.
- Benartzi, S. and R. Thaler, 2001, "Myopic Loss Aversion and the Equity Premium Puzzle", *Quarterly Journal of Economics* 110, 73–92.

- Berkelaar, A. and R. Kouwenberg, 2000, “Dynamic Asset Allocation and Downside-Risk Aversion”, unpublished working paper.
- Berkelaar, A., Kouwenberg, R., and Post, T., 2003, “Optimal Portfolio Choice under Loss Aversion”, unpublished working paper.
- Biais, B., P. Hillion, and C. Spatt, 1995, “An Empirical Analysis of the Limit Order Book and the Order Flow in the Paris Bourse”, *Journal of Finance* 50, 1655–1689.
- Bollerslev, T., 1986, “Generalized Autoregressive Conditional Heteroskedasticity.” *Journal of Econometrics* 31, 307–327.
- Brandt, M., 2004, “Portfolio Choice Problems”, in Y. Ait-Sahalia and L.P. Hansen (eds.), *Handbook of Financial Econometrics*, forthcoming. New York: Elsevier.
- Brandt, M., Goyal, A., Santa-Clara, P., and J. Stroud, 2005, “A Simulation Approach to Dynamic Portfolio Choice with an Application to Learning About Return Predictability”, *Review of Financial Studies* 18, 831–873.
- Brennan, M. and Y. Xia, 2001, “Persistence, Predictability, and Portfolio Planning”, unpublished working paper.
- Brennan, M., Schwartz, E., and R. Lagnado, 1997, “Strategic Asset Allocation”, *Journal of Economic Dynamics and Control* 21, 1377–1403.
- Calvet, L., Campbell, J., and P. Sodini, 2006, “Down or Out: Assessing the Welfare Costs of Household Investment Mistakes”, Working Paper No. 12030, NBER.
- Campbell, J., Chan, Y., and L. Viceira, 2003, “A Multivariate Model of Strategic Asset Allocation”, *Journal of Financial Economics* 67, 41–80.
- Campbell, J., Lo, A., and C. MacKinlay, 1997, *The Econometrics of Financial Markets*. Princeton, NJ: Princeton University Press.
- Campbell, J., and L. Viceira, 1999, “Consumption and Portfolio Decisions when Expected Returns are Time-Varying”, *Quarterly Journal of Economics* 114, 433–495.
- Campbell, J., and L. Viceira, 2002, *Strategic Asset Allocation: Portfolio Choice for Long Term Investors*, Oxford University Press: New York, NY.
- Carr, P., and R. Jarrow, 1990, “The Stop-Loss Start-Gain Paradox and Option Valuation: A New Decomposition into Intrinsic and Time Value”, *Review of Financial Studies* 3, 469–492.
- Chakravarty, S., and C. Holden, 1995, “An Integrated Model of Market and Limit Orders”, *Journal of Financial Intermediation* 4, 213–241.
- Chang, F., 1988, “The Inverse Optimal Problem: A Dynamic Programming Approach”, *Econometrica* 56, 147–172.
- Clarke, R., Krase, S., and M. Statman, 1994, “Tracking Errors, Regret, and Tactical Asset Allocation”, *Journal of Portfolio Management* 20, 16–24.
- Constantinides, G. 1976, “Stochastic Cash Management with Fixed and Proportional Transaction Costs”, *Management Science* 22, 1320–1331.

- Constantinides, G., 1986. "Capital Market Equilibrium with Transaction Costs", *Journal of Political Economy* 94, 842–862.
- Constantinides, G., and M., Magill, 1976, "Portfolio Selection with Transaction Costs", *Journal of Economic Theory* 13, 264–271.
- Cox, J. and C. Huang, 1989, "Optimum Consumption and Portfolio Policies when Asset Prices Follow a Diffusion Process", *Journal of Economic Theory* 49, 33–83.
- Das, S., and R. Sundaram, 2000, "A Numerical Algorithm for Optimal Consumption-Investment Problems", unpublished working paper, Harvard University.
- Davidson, R., and J. Duclous, 2000, "Statistical Interference for Stochastic Dominance and for the Measurement Of Poverty And Inequality", *Econometrica* 68, 1435–1464.
- DeBondt, W. and R. Thaler, 1986, "Does the Stock Market Overreact?", *Journal of Finance* 40, 793–807.
- Easley, D., and M. O'Hara, 1991, "Order Form and Information in Securities Markets", *Journal of Finance* 46, 905–927.
- Engle, R.F., 1982, "Autoregressive Conditional Heteroskedasticity with Estimates of the Variance In United Kingdom Inflation", *Econometrica* 50, 987–1007.
- Fama, E., 1965, "The Behavior of Stock Market Prices", *Journal of Business* 38, 34–105.
- Fama, E. and M. Blume, 1966, "Filter Rules and Stock Market Trading Profits", *Journal of Business* 39, 226–241.
- Fama, E. and K. French, 1988, "Permanent And Temporary Components of Stock Prices", *Journal of Political Economy* 96, 246–273.
- Fischhoff, B. and P. Slovic, 1980, "A Little Learning...: Confidence in Multicue Judgment Tasks", in R. Nickerson (ed.), *Attention and Performance, VIII*. Hillsdale, NJ: Erlbaum.
- Fong, W., Wong, W., and H.H. Lean, 2003, "International Momentum Strategies: A Stochastic Dominance Approach", unpublished working paper, National University of Singapore.
- Genesove, D., and Mayer, C., 2001, "Loss Aversion and Seller Behavior: Evidence from the Housing Market", NBER Working Paper No. W8143.
- Gervais, S. and T. Odean, 2001, "Learning to Be Overconfident", *Review of Financial Studies* 14, 1–27.
- Gilboa, I., and D. Schmeidler, 1989, "Maximum Expected Utility with Non-Unique Prior", *Journal of Mathematical Economics* 18, 141–153.
- Goetzmann, W. and M. Massa, 2003, "Index Funds and Stock Market Growth", *Journal of Business* 76, pp. 1–27.
- Goetzmann, W., Massa, M. and G. Rouwenhurst, 1999, "Behavioral Factors in Mutual Fund Flows", Yale International Center for Finance, Working Paper No. 00–14, Yale University.

- Grinblatt, M, and Han, B., 2001, “The Disposition Effect and Momentum”, unpublished working paper, University of California, Los Angeles.
- Grinblatt, M, and Han, B., 2004, “Prospect Theory, Mental Accounting, and Momentum”, to appear in *Journal of Financial Economics*.
- Gul, F., 1991, “A Theory of Disappointment Aversion”, *Econometrica* 59, 667–686.
- Hamilton, J., 1989, “A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle”, *Econometrica* 57, 357–384.
- Hamilton, J., 1994, *Time Series Analysis*, Princeton University Press.
- Handa, P., and R. Schwartz, 1996, “Limit Order Trading”, *Journal of Finance* 51, 1835–1861.
- Harris, L. and J. Hasbrouck, 1996, “Market vs. Limit Orders: The SuperDot Evidence on Order Submission Strategy”, *Journal of Financial and Quantitative Analysis* 31, 213–231.
- Huberman, G. and T. Regev, 2001, “Contagious Speculation and a Cure for Cancer: A Nonevent That Made Stock Prices Soar”, *Journal of Finance* 56, 387–396.
- Ingersoll, J., 1987, *Theory of Financial Decision Making*. Totowa, NJ: Rowman & Littlefield.
- Jegadeesh, N., 1990, “Evidence of Predictable Behavior of Security Returns”, *Journal of Finance* 45, 881–898.
- Jegadeesh, N., and S. Titman, 1993, “Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency”, *Journal of Finance* 48, 65–91.
- Kahneman, D., and A. Tversky, 1979, “Prospect Theory: An Analysis of Decision Under Risk”, *Econometrica* 47, 263–291.
- Kahneman, D., and Tversky, A., 1992, “Advances in Prospect Theory: Cumulative Representation of Uncertainty”, *Journal of Risk and Uncertainty* 5, 297–323.
- Kaminski, K., 2006, “General Superposition Strategies”, unpublished manuscript.
- Kaur, B., Rao, S. and H. Singh, 1994, “Testing of Second-Order Stochastic Dominance of Two Distributions”, *Econometric Theory* 10, 849–866.
- Kim, T., and E. Omberg, 1996, “Dynamic Non-Myopic Portfolio Behavior”, *Review of Financial Studies* 9, 141–161.
- Kogan, L., and R. Uppal, 2000, “Risk Aversion and Optimal Portfolio Policies in Partial and General Equilibrium Economies”, unpublished working paper, University of Pennsylvania.
- Laibson, D., 1997, “Golden Eggs and Hyperbolic Discounting”, *Quarterly Journal of Economics* 62, 443–477.
- Lichtenstein, S., Fischhoff, B. and L. Phillips, 1982, “Calibration of Probabilities: The State of the Art to 1980”, in D. Kahneman, P. Slovic, and A. Tversky (eds.), *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge, UK: Cambridge University Press.

- Lintner, J., 1965, “The Valuation of Risk Assets and the Selection of Risky Investments In Stock Portfolios and Capital Budgets”, *Review of Economics and Statistics* 47, 13–37.
- Liu, J., 1999, “Portfolio Selection in Stochastic Environments”, unpublished working paper, University of California, Los Angeles.
- Liu, J. 2001, “Dynamic Portfolio Choice and Risk Aversion”, working paper, 2001.
- Lo, A., 1991, “Long-Term Memory in Stock Market Prices”, *Econometrica* 59, 1279–1313.
- Lo, A., 2002, “The Statistics of Sharpe Ratios”, *Financial Analysts Journal* 58, 36–50.
- Lo, A. and C. MacKinlay, 1988, “Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test”, *Review of Financial Studies* 1, 41–66.
- Lo, A. and C. MacKinlay, 1990, “An Econometric Analysis of Nonsynchronous Trading”, *Journal of Econometrics* 45, 181–212.
- Lo, A. and C. MacKinlay, 1999, *A Non-Random Walk Down Wall Street*. Princeton, NJ: Princeton University Press.
- Lo, A., MacKinlay, C. and J. Zhang, 2002, “Econometric Models of Limit-Order Executions”, *Journal of Financial Economics* 65, 31–71.
- Lo, A., Mamaysky, H., and J. Wang, 2004, “Asset Prices and Trading Volume under Fixed Transactions Costs”, *Journal of Political Economy* 112, 1054–1090.
- Lo, A. and J. Wang, 1995, “Implementing Option Pricing Models When Asset Returns are Predictable”, *Journal of Finance* 50, 87–129.
- Loomes, G., and Sugden, R., 1982, “Regret Theory: An Alternative Theory of Rational Choice under Uncertainty”, *The Economic Journal* 92, 805–824.
- Lucas, R., 1978, “Asset Prices in an Exchange Economy”, *Econometrica* 46, 1429–1446.
- Markowitz, H., 1952, “Portfolio Selection”, *Journal of Finance* 7, 77–91.
- Massa, M., and A. Simonov, 2004, “Hedging, Familiarity and Portfolio Choice”, CEPR Discussion Paper No. 4789.
- McFadden, D., 1989, “Testing for Stochastic Dominance, ” in T.B. Fomby and T.K. Seo (eds.), *Studies in the Economics of Uncertainty*. New York: Springer-Verlag.
- Merton, R. 1969, “Lifetime Portfolio Selection Under Uncertainty; The Continuous-Time Case”, *Review of Economics & Statistics* 51, 247–257.
- Merton, R., 1971, “Optimum Consumption and Portfolio Rules in a Continuous-Time Model”, *Journal of Economic Theory* 3, 373–413.
- Merton, R., 1973, “An Intertemporal Capital Asset Pricing Model”, *Econometrica* 41, 867–887.
- Odean, T., 1998, “Are Investors Reluctant to Realize Their Losses?”, *Journal of Finance* 53, 1775–1798.
- Odean, T., 1999, “Do Investors Trade Too Much”, *American Economic Review* 89, 1279–1298.

- Polkovnichenko, V., 2002, “Household Portfolio Diversification”, unpublished working paper, University of Minnesota.
- Poterba, J. and L. Summers, 1988, “Mean Reversion in Stock Returns: Evidence and Implications”, *Journal of Financial Economics* 22, 27–60.
- Samuelson, P. 1969, “Lifetime Portfolio Selection By Dynamic Stochastic Programming”, *Review of Economics and Statistics* 51, 239–246.
- Seppi, D., 1997, “Liquidity Provision with Limit Orders and a Strategic Specialist”, *Review of Financial Studies* 10, 103–150.
- Shapira, Z., and Venezia, I., 2001, “Patterns of Behavior of Professionally Managed and Independent Investors”, *Journal of Banking and Finance* 25, 1573–1587.
- Sharpe, W., 1964, “Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk”, *Journal of Finance* 19, 425–442.
- Sharpe, W., 1994, “The Sharpe Ratio”, *Journal of Portfolio Management* Fall, 49–58.
- Shefrin, H., 2005, *Behavioral Approaches to Asset Pricing*. New York: Elsevier.
- Shefrin, M. and M. Statman, 1985, “The Disposition to Sell Winners Too Early and Ride Losers Too Long: Theory and Evidence”, *Journal of Finance* 40, 777–790.
- Shiller, R., 2001, “Human Behavior and the Efficiency of the Financial System”, *Cowles Foundation for Research in Economics* Paper No. 1025, 1305–1340.
- Spurgin, R., 2001, “How to Game Your Sharpe Ratio”, *The Journal of Alternative Investments* 4, 38–46.
- Thaler, R., 1985, “Mental Accounting and Consumer Choice”, *Marketing Science* 4, 199–214.
- Thaler, R., 1994, “Psychology and Savings Policies”, *American Economic Review* 84, 186–192.
- Thaler, R. and H. Shefrin, 1981, “An Economic Theory of Self Control”, *The Journal of the Political Economy* 89,392–406.
- Tobin, J., 1958, “Liquidity Preference as Behavior Towards Risk”, *Review of Economic Studies* 25, 65–86.
- Tsay, R., *Analysis of Financial Time Series*, Wiley, 2002.
- Tschoegl, A., 1988, “The Source and Consequences of Stop Orders: A Conjecture”, *Managerial and Decision Economics* 9, 83–85.
- Tversky, A. and D. Kahneman, 1981, “The Framing of Decisions and the Psychology of Choice”, *Science* 211, 453–458.
- Wachter, J., 2002, “Portfolio and Consumption Decisions under Mean-Reverting Returns: An Exact Solution for Complete Markets”, *Journal of Financial and Quantitative Analysis* 37, 63–91.
- Xia, Y., 2001, “Learning about Predictability: The Effects of Parameter Uncertainty on Dynamic Asset Allocation”, *Journal of Finance* 56, 205–246.