Journal of Financial Markets 34 (2017) 1–15

Contents lists available at ScienceDirect

Journal of Financial Markets

journal homepage: www.elsevier.com/locate/finmar

Stop-loss strategies with serial correlation, regime switching, and transaction costs ☆

Andrew W. Lo [a,b,*], Alexander Remorov [c]

[a] MIT Sloan School of Management and CSAIL, 100 Main Street, E62-618, Cambridge, MA 02142-1347, United States
[b] AlphaSimplex Group, LLC, United States
[c] MIT Operations Research Center, United States

ARTICLE INFO

Article history:
Received 20 January 2015
Received in revised form
1 February 2017
Accepted 16 February 2017
Available online 28 April 2017

JEL classification:
G11
G12

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

ABSTRACT

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

1. Introduction

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

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

☆ Research support from the MIT Laboratory for Financial Engineering and the Natural Sciences and Engineering Research Council of Canada is gratefully acknowledged. We thank Dimitris Bertsimas, Hui Chen, Jayna Cummings, Leonid Kogan, Glen Martin, and two anonymous referees for helpful comments and discussion. The views and opinions expressed in this article are those of the authors only and do not necessarily represent the views and opinions of any other organizations, any of their affiliates or employees, or any of the individuals acknowledged above.
* Corresponding author.
E-mail addresses: alo-admin@mit.edu (A.W. Lo), alexrem@mit.edu (A. Remorov).

http://dx.doi.org/10.1016/j.finmar.2017.02.003
1386-4181/

A natural question is whether these departures from randomness can be exploited using a dynamic investment strategy, including stop-loss policies.

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

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

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

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

From a broader perspective, the use of stop-loss strategies can correct for the tendency of investors to hold on to losers too long and to sell winners too early, a behavioral bias known as the "disposition effect" first documented by Shefrin and Statman (1985). The presence of behavioral biases such as this has been well documented in the finance literature.[2] While most of this research has focused on the empirical evidence for these biases and the theoretical models to explain them, few studies have proposed methods for investors to actively avoid or protect against such biases. Stop-loss strategies are an important first step in this direction.

2. Literature review

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

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

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

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