MICHAEL G. FOSTER SCHOOL OF BUSINESS | UNIVERSITY OF WASHINGTON

Re-examining the Hidden Costs of the Stop-Loss

Kira Detko
B.A. Finance

Wilson Ma†
B.A. Finance, B.S. Economics

Guy Morita‡
B.A. Finance

In this paper, we present general implications of the impact of stop-losses to future returns. The use of stop-losses change return distributions, but not in the way that one would typically expect. We find that while stop-losses can reduce position volatility, hidden costs offset perceived benefits in terms of altering future returns. Use of both stop-losses and profit-taking stops separately or in conjunction offer no statistically significant difference in expected return but have a meaningful impact in returns with drift, as the expected return converges to that of the underlying.

Contact Author Information:
†Wilson Ma | tireg@u.washington.edu | 206.354.8530
‡Guy Morita | guymo@u.washington.edu | 206.240.5846

UW BUSINESS SCHOOL

In the investment world, many rules of thumb are taken for granted. Buy low, sell high. The trend is your friend. Don't put all your eggs in one basket. Cut your losses short. Often offered without compelling support, these recommendations are often nonetheless followed by investors in faith. A simple analogy can be drawn to a child's belief in Santa Clause. Similar to how the impression of Santa changes a child's behavior, basic investment principles influence investors in their decisions. This notion can be thought of in terms of benefits and costs. The benefit for the child is that she receives presents in return for good behavior. Analogously, the investor receives improved returns or smaller losses in return for following the investment truisms. However, benefits do not come without costs. For the child, the presents she wishes for arrive exogenously, presumably paid for by Santa. Yet we know that the costs of the presents are actually borne by the parents. Likewise, for the investor who employs stop-losses to reap the benefits of risk management, there are costs that may not be evident. Just as costs exist for Christmas presents, they exist for stop-losses. We explore the hidden costs as well as the perceived benefits of stop-losses in this research essay.

We present four general implications of the impact of stop-losses to future returns: First, the use of stop-losses alter the shape of return distributions, but not in the way that one would typically expect. Expected return does not change but approaches that of the underlying asset. In other words, we find that there is no inherent edge to be had in using stop-losses or profit taking stops. Second, stop losses can add significant reduction to volatility. We found that as volatility of the underlying asset's returns increased, the impact of stop-losses increased as well. Third, the implication of the findings regarding stop-losses can be applied to profit-taking stops, or stops located on the positive return side of the distribution; this is intuited by examining short positions with stop-loss orders on the positive side of distribution. Fourth, the results hold true for combinations of stop-losses and profit taking stops. We conclude by discussing practical application of the implications described above.

The typical mention of the stop-loss in investing books, magazines, and articles refers to it as a risk management tool. The commission for a stop typically is the same as that of a normal market order; as a result stop-loss orders are widely used in investment circles because of their benefits. The stop-loss gives the investor the ability to keep losses small by exiting a trade at some predetermined threshold, avoiding potential further losses. It also allows for peace of mind in not having to actively monitor positions minute-by-minute with one's finger hovering over the 'sell' button. However, the common literature fails to address the true effect that stops have on altering expected future returns. Hence, the motivation of our research stems from the initial survey of a hidden cost of the stop-loss detailed in an article written by Robert Macrae of Arcus Investments.[1] The implications of Macrae's article are that stops impact expected return distribution and create higher portfolio volatility, the opposite of what is intended. The reasoning for this is that increased leverage must be used to maintain same exposure. Macrae's conclusion is that this increased volatility is the "hidden cost" of the stop-loss.[2]

Common misconceptions of the impact of the stop-loss to the expected return distribution are illustrated in Figure 1:

[1] Macrae, Robert. "The Hidden Cost of the StopLoss." Arcus Investments, AIMA Journal 2005.
[2] Further background to our response to this paper can be found in Appendix B.

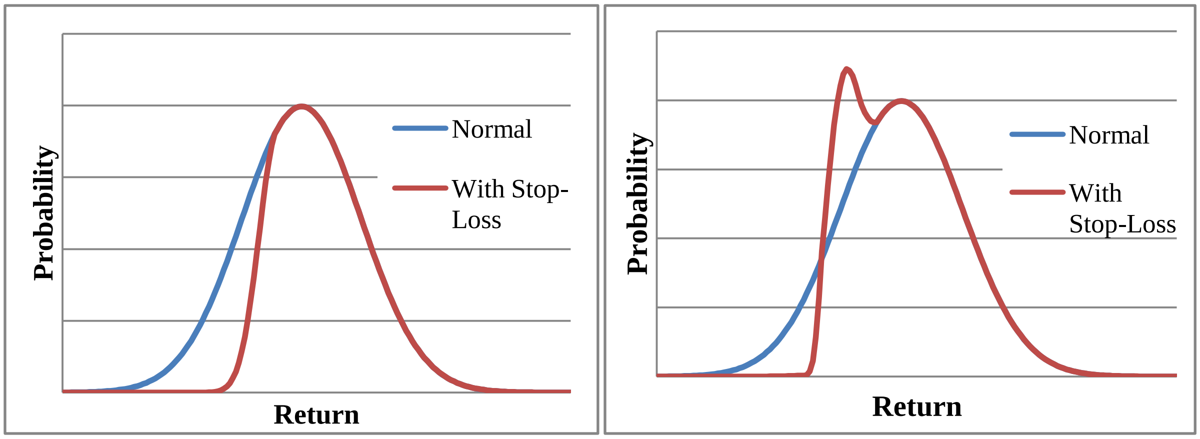

Figure 1: Common misconceptions of the impact of stop-losses

In these figures, the approximately normal distribution of returns realized by not utilizing a stop-loss is contrasted with the perceived expected return caused by having a stop. Therefore, the difference of the areas of the two distributions measures the impact of the stop-loss. The first diagram illustrates the fallacy of the belief that the stop will only act to prevent losses beyond some specified point. One problem with this line of reasoning is that it does not take into account the higher probability of a small loss occurring at the price where the stop is placed. In effect, this implies that only a portion of the negative side of the return distribution is removed so that the area of the stop-loss distribution is effectively smaller. The second diagram corrects for this so that the total areas of both distributions are the same. It indicates that a stop loss at some location will increase probability of hitting the stop in the vicinity. For example, entering a long position on Google at $470 with a stop-loss set at $465.30 (1% below the entry price), almost guarantees getting stopped out due to daily price movement so the probability of being stopped out must increase at the stop price. At the same time the probability of achieving larger negative returns, such as -10%, is lessened because the stop mitigates the probability of greater losses. Note that slippage and gaps are indicated in the diagrams by a sharply decreasing (but not immediately zero) probability of returns to the left of the stop-loss. The second diagram is an improvement, but is still inaccurate as it misses a subtle result.

We find that the realized impact of the stop-loss as a strategy manifests itself in hidden costs and perceived benefits, illustrated in Figure 2:

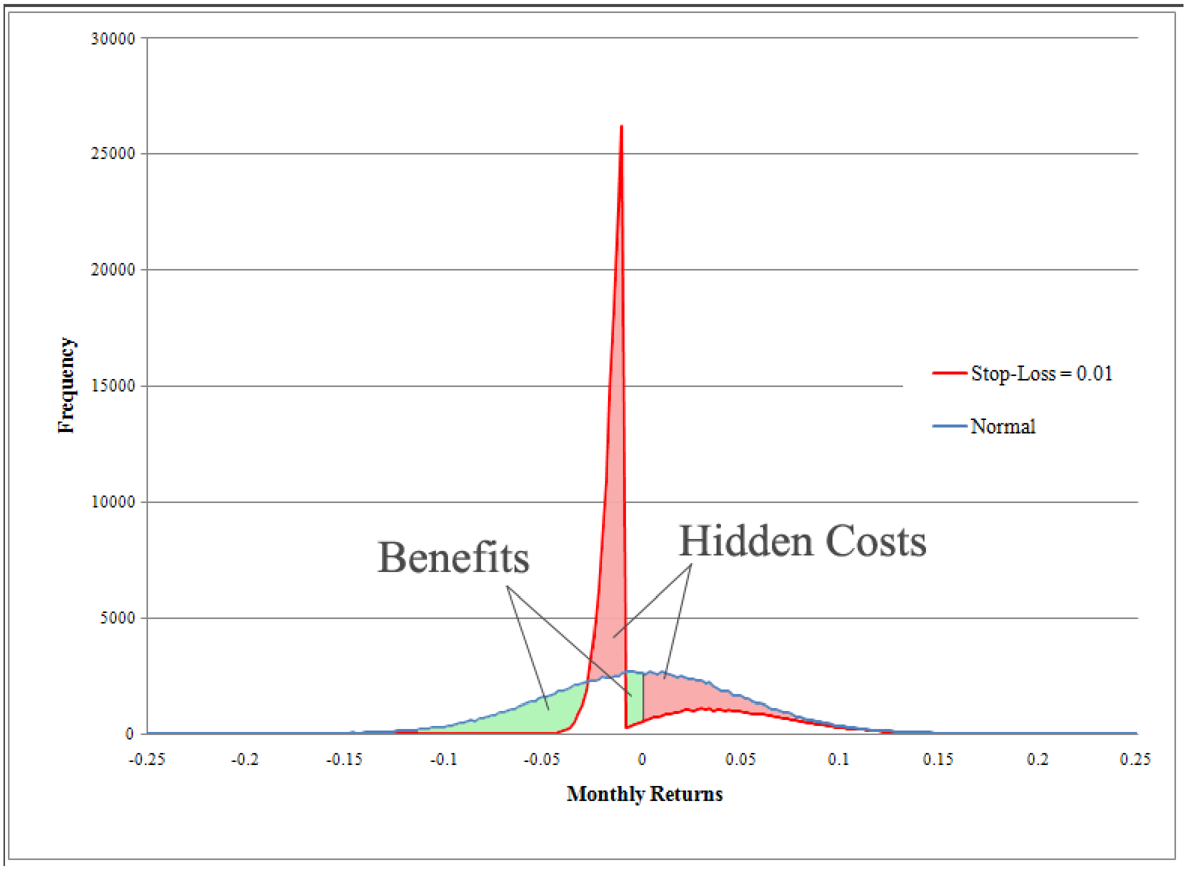

Figure 2: Actual impact of the stop-loss on future returns

Here, when the expected return of the asset is zero, the expected return of using a stop-loss is the same as that of not using a stop-loss. Graphically, this is demonstrated by the hidden costs, or red shaded area, having a total effective size as the benefits, represented by the shaded green area. This is due to the fact that the hidden costs cancel out the perceived benefits. Intuitively, this may not be a huge shock. If the costs outweighed the benefits, then no one would use stops and they would be deemed ineffective. If the benefits outweighed the costs, it follows that the average return for the asset would be higher when using a stop-loss. Managers and investors would be able to generate positive expected returns by just using stop-losses. Clearly this is not the case.

To derive general implications of stop-losses, we simulated asset returns by using stationary and nonstationary return models[3] to generate daily log returns for a random walk process:

Constant Expected Return (CER) model (stationary random walk):

$$r_t = \mu + \varepsilon_t \qquad t = 1, 2, \ldots, T$$

$$\varepsilon_t \sim iid\ N(0, \sigma_\varepsilon^2);\ Cov(\varepsilon_t, \varepsilon_s) = 0,\ t \neq s$$

Nonstationary random walk with drift:

$$r_t = r_{t-1} + \mu + \varepsilon_t \qquad t = 1, 2, \ldots, T$$

$$\varepsilon_t \sim iid\ N(0, \sigma_\varepsilon^2);\ Cov(\varepsilon_t, \varepsilon_s) = 0,\ t \neq s$$

Autocorrelated process with AR(1) model:

$$r_t = (1-\Phi)\mu + \Phi r_{t-1} + \varepsilon_t \qquad t = 1, 2, \ldots, T;\ -1 < \Phi < 1$$

$$\varepsilon_t \sim iid\ N(0, \sigma_\varepsilon^2);\ Cov(\varepsilon_t, \varepsilon_s) = 0,\ t \neq s$$

Returns were generated with varying monthly volatility at 1% and from 5% to 25% in five percent increments to simulate assets with low volatility to very high volatility. We simulated the random daily returns of 500 assets for each volatility level and measured the impact of stops from 1% below the entry price to 99% below the entry price for long positions (vice versa for short positions).

---

[3] Zivot, Eric. *Lecture notes for Economics 424: Computational Finance.* University of Washington: 2007.

The random walk models were ideal to generalize the results as they incorporate unrestrictive but simplifying assumptions and control for outside variables so that the pure mechanics of the stop-losses could be analyzed. The assumptions allow for changes in the parameters to be incorporated into a new model of returns but do not alter the implications.[4] The trading strategy employed to generate return figures is as follows: A new position was initiated at the start of the month. At this point, two stops were initiated: a time stop and the stop-loss. Daily returns were compared, and the time stop meant that the trade would be exited on the last day of the month if the stop-loss was not hit. Otherwise, if the stop-loss threshold was hit, the position would be closed and would not be re-entered for the remainder of that month. The process would then repeat at the start of the next month. This ensured that the sample size of returns was consistent throughout the simulations; every month we would record one return data point, equivalent to monthly return. For a passive long portfolio, this allows for monthly returns to be computed as well. An assumption made is that differences between price on the last day of the month and the start of the next month are negligible.

[4] See Appendix D for detailed explanation of assumptions.

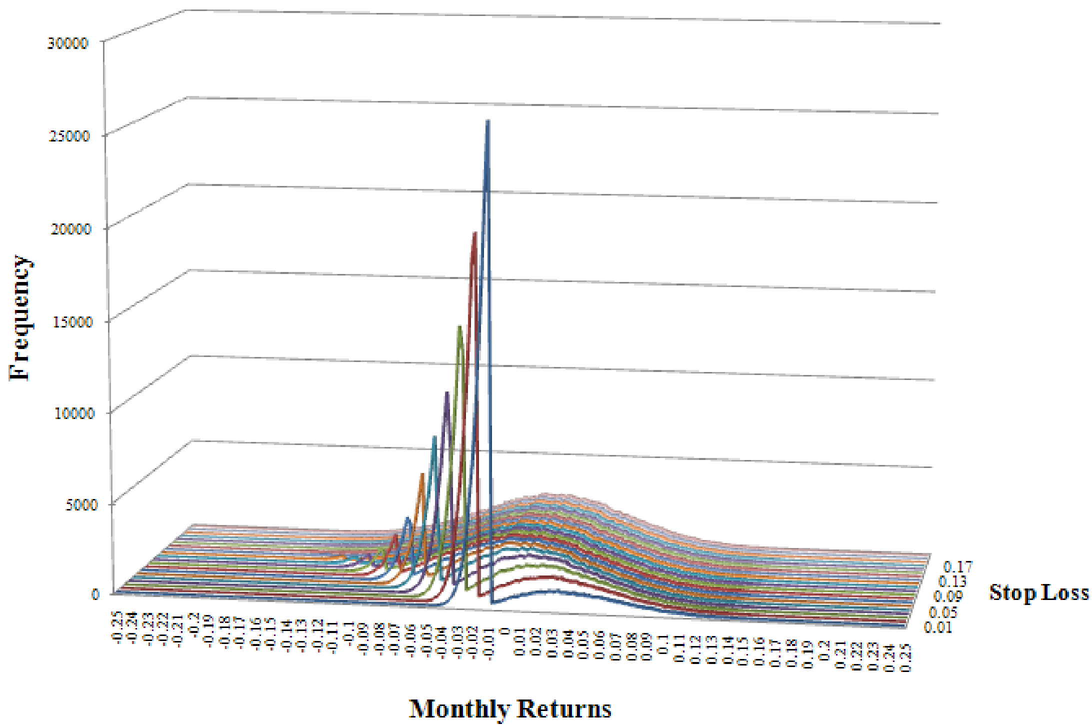

Figure 3: Expected return distribution surface with varying stop-losses on long strategy (E(μ)=0, E(σ) = 0.05)

We display the results of the simulation with expected return of zero and monthly volatility of 5% in Figure 3. Notice that for varying levels of stop-loss, the impact decreases until a point at which the stop is so far from the price that its impact is effectively the same as not having a stop at all. The impact of the stop-losses are illustrated by the peaks. A single slice of the surface is presented in Figure 4.

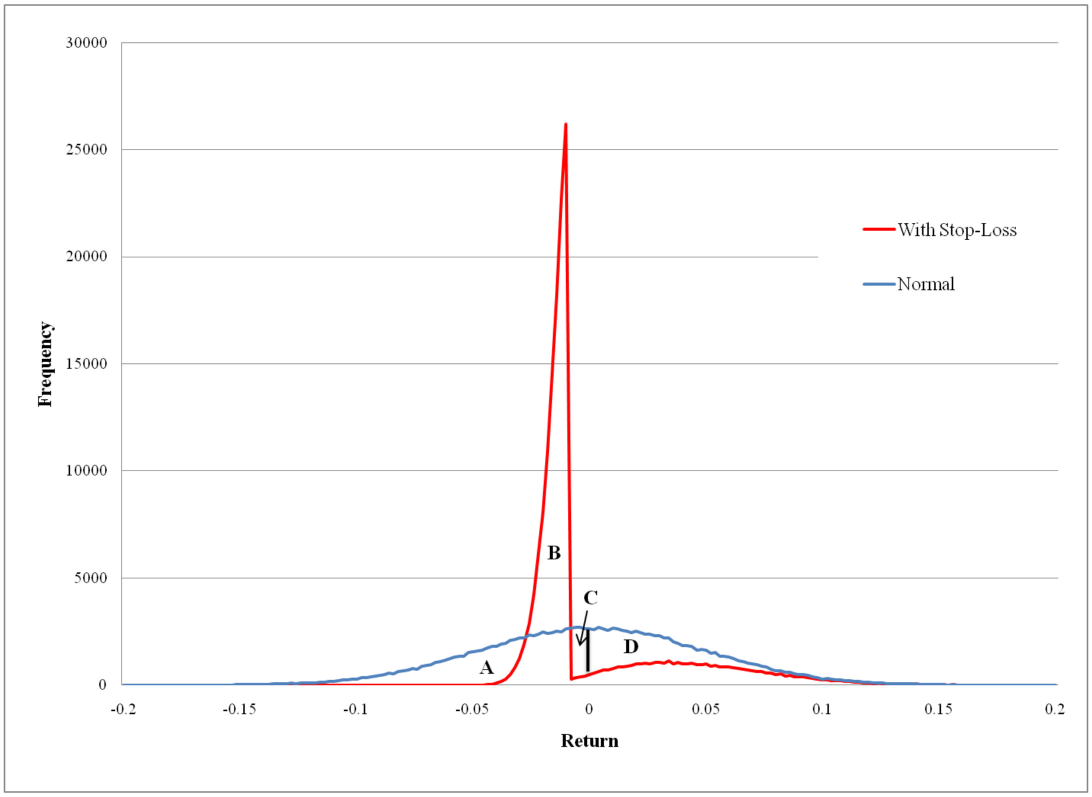

Figure 4: Return distribution generated by the stop loss at 1% below entry price

Results of the 1% stop-loss are presented to emphasize the difference between employing a stop-loss and passively holding the position. In Figure 4, the areas A and C denote the benefits realized by the stop-loss, while B and D label the costs. Area A represents the benefit of reducing the possibility of further losses, both in magnitude and in frequency. Area C also represents this; a reduction in slightly negative returns was found. Area B represents the increased probability of being stopped out; this is a hidden cost as the impact to the return distribution is negative. Note the relative steepness of the right side of B versus the left side. This is due to the fact that if the asset returns -0.99% the stop will not be breached and so a 'wall' forms to the side greater than the stop, whereas on the left the frequency declines less dramatically due to slippage. Area D represents an opportunity cost - the potential return given up if the price crosses the stop and then returns to the positive range. We dub B and D the true "hidden costs" of the stop-loss as they are not immediately obvious when considering the impact of stops. We found that the expected return of the distribution in Figure 4 generated by having a stop was 0.000129, while the expected return of not having a stop was 0.0000747, or effectively zero.[5] A paired t-test for means confirmed that the difference was insignificant with a t-score of 0.37070 (p = 0.71086). The implication derived here is that use of a stop-loss does not change the expected return of a trade when the expected return of the asset is zero.

Next we explored the implications of the stop in instances where the expected return was not zero.

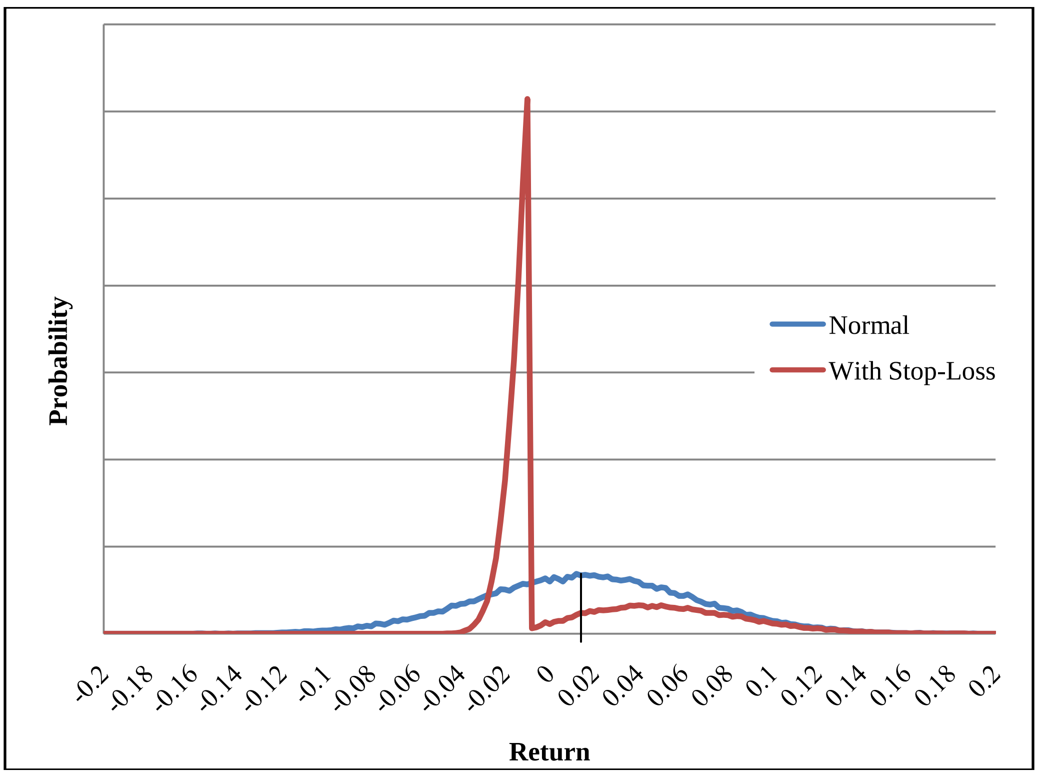

Figure 5: Return distributions with positive drift, 1% stop

Figure 5 shows the results of the simulation with positive drift and expected return of 10.18% annually (0.84% monthly) and monthly volatility of 5%, consistent with the

[5] Details of the statistical tests for means and variances of the distributions are provided in Appendices A.1-A.4.

historical return and volatility of the Dow Jones Industrial Average. The interpretation for the stop-loss in this situation is that it is equivalent to a "trailing-stop," or a stop-loss order where the stop-loss trigger price increases as the asset price increases so that it maintains some relative distance to the asset price. The t-statistic for the expected value of the distribution with the stop versus not having a stop was -34.421, indicating a very significant difference of the average returns. As a result, the stop had a meaningful impact on returns in this positively trending market.

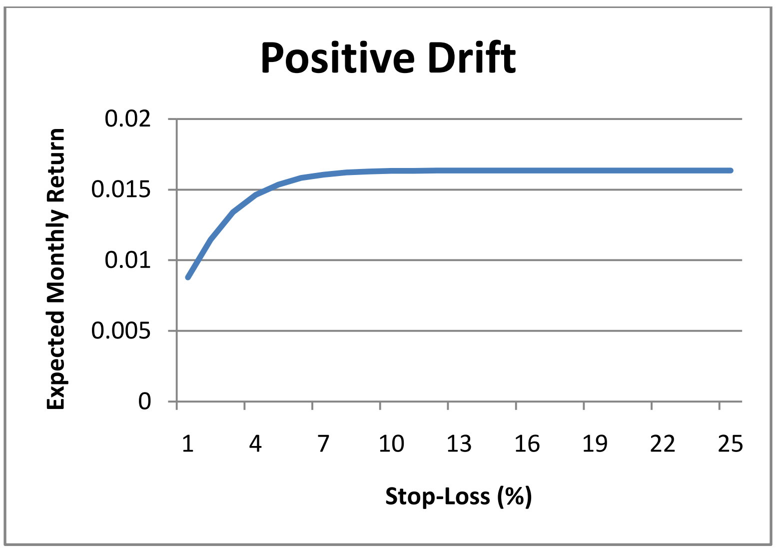

Figure 6: For assets with positive expected returns, stops actually reduce expected return.

We plotted the expected return versus the various stop-loss levels for this simulation and found the results shown in Figure 6. Surprisingly, here the monthly expected return of using the stop-loss versus not having a stop is actually lower. We see that as the stop is moved farther away from the price, the expected return rises until it is equivalent to that of not having a stop. The implication of this is that trailing stops actually reduce expected return when asset prices have positive drift or positive expected return. In other words, setting a 1% stop-loss actually decreases the probability of getting a positive 8% return. This stems from an opportunity cost of missing out on market recoveries where the stop price is breached and the asset price subsequently recovers. In a bull market, for example, the likelihood of this occurring is increased as positive drift causes the distribution to shift to the right, showing expected positive return from passive holding.

We found the opposite was true in markets where the drift was negative:

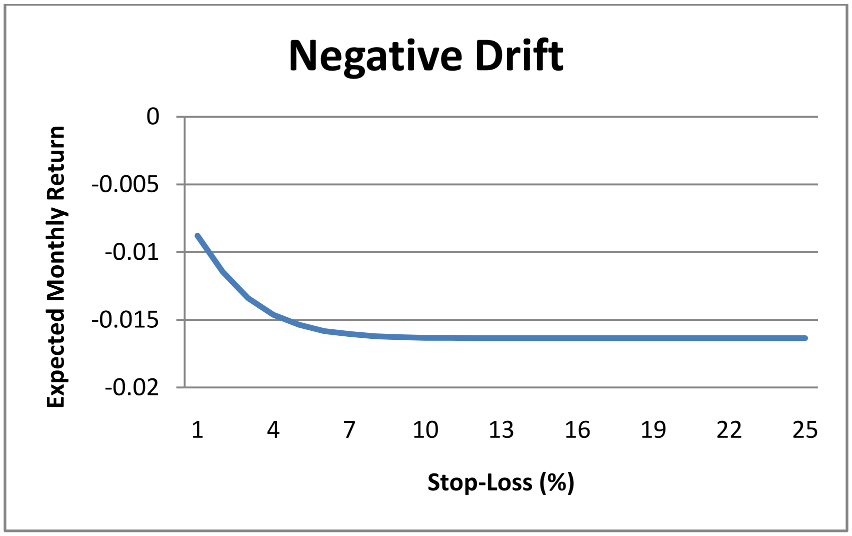

Figure 7: Inverse effect of stop losses in bear markets (long position)

Here the simulation with negative drift can be interpreted as a bear market condition. We see that the stops improved average return by limiting participation in the market fall. As the stop-loss increased, participation in the bear market increased and further losses were realized. These results are consistent with Macrae's assertion that for trending series, stop-losses can prevent further losses. We found that drift is a necessary requirement for stop-losses to have impact. Notice the similarity between Figures 4 and 5; essentially, the distributions are shifted but the total 'cost' and 'benefit' areas are the same form.

We also tested if the variance of the returns was different. We observed that volatility reduction occured, and F-tests for differences in variance confirmed statistically significant differences between using stops and not using stop-losses. Notice that F-values decrease as volatility of the asset increased, indicating increasing probability of reduced variance by using stops in volatile assets.

	Variance		Values	
Sigma	***Stop-Loss=0.01***	***No Stop-Loss***	**F-Value**	**P-Value**
5.0%	0.001065373	0.002374836	0.44860912	100.00%
10.0%	0.00329697	0.009557455	0.344963142	100.00%
15.0%	0.006750322	0.021506808	0.313869081	100.00%
20.0%	0.011397643	0.038017919	0.299796602	100.00%
25.0%	0.01707314	0.059221057	0.288295089	100.00%

Table 1: f-Values across assets with varying monthly volatilities

Given the nature of the distributions, we postulated that the same results would materialize using profit-taking stops instead of stop-losses. A profit-taking stop operates in the same way as the stop-loss, but instead of preventing further losses by exiting the position, the position is closed to lock in realized gain. This was intuited by the fact that short positions have stops on the positive side of the distribution. However, one difference in short positions is that the extent of gains and losses is converse to that of long positions; there is a maximum to the potential profits (when the asset becomes worthless) but the losses are potentially unlimited as asset prices such as stocks theoretically have no upper limit. This is one reason margin requirements exist for shorting. In this way stops play a unique role in short positions in creating a limit to maximum loss; the 'short stock with stop-loss' portfolio has the same payoff at expiration as a long put position.[6]

As we see in Figure 8, the results of having a profit-taking stop are symmetrical to those found in having stop-losses.

[6] Hull, John C. Options, Futures and Other Derivatives. Prentice Hall, 6th Ed. 2005.

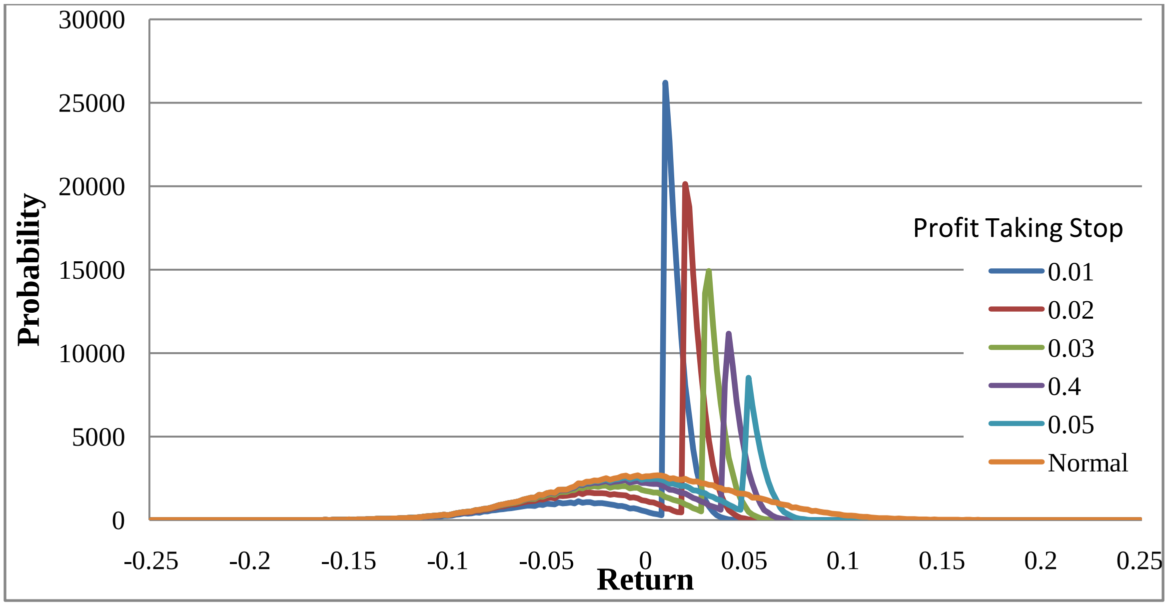

Figure 8: Distributions of profit-taking stops (E(μ)=0, E(σ) = 0.05)

Next we simulated returns utilizing both stops and returns; the interpretation of this is entering a long position with a stop loss below the price and an exit point above. We tested if the implications would hold true in this scenario. Figure 9 illustrates the distribution of returns from employing a stop loss and profit-taking stop at the same level above and below the price. Notice the trimodal distributions created by the employment of stop-losses and profit-taking stops. As the distance of the stops increase, the expected return distribution becomes more normal and approaches that of having no stop. At the extreme of 1% stop on each side, the distribution becomes bimodal as the probability of returns between -1% and 1% declines to zero. We conducted the same tests for mean and variance of the 1% stop distribution against the no-stop distribution, labeled "Normal" and found that the expected returns were the same and t-statistic and F-values were -0.2566453 and 0.6430071 respectively, indicating an insignificant difference in mean but significant difference in the variance of the return distributions. We tested this result with varying widths by using the function:

$\Pi = \theta s$

where Π is the profit-taking stop, *s* is the stop-loss, and unequal distances are achieved by varying θ. Simulations with non-equal stop-losses and profit-taking stops (such as 2% profit, 1% stop-loss) yielded the same expected return and decreased variance.

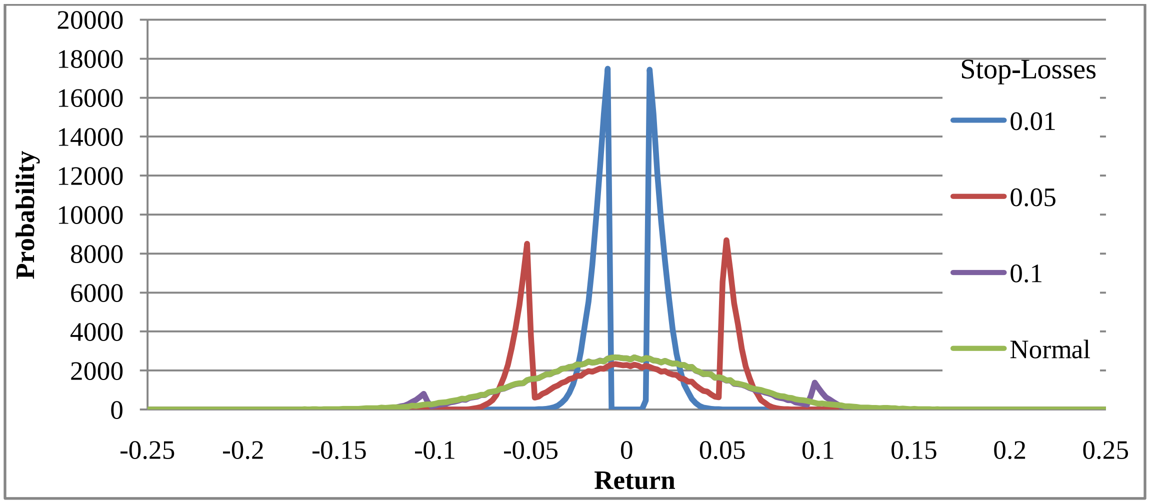

Figure 9: Distributions of returns using both stop-loss and profit-taking stop (E(μ)=0, E(σ) = 0.05)

The next step was to take derived above and see if they would hold true for historical returns of actual stocks and indices, confirming the validity of the results. We ran a backtest of the strategy for the past 27 years (Jan. 1980 to Dec. 2007) and display the resulting distributions at the 1% stop level and without stop. The detailed results of the statistical tests can be found in Appendix A.5 and A.6, but in general we found that as with our simulations, the means differed but were statistically insignificant, while the variances were statistically significant, indicating volatility reduction. Notice that the distributions here are much less defined than the distributions mentioned from the simulations. This is due to relatively small sample size (324 months in 27 years) that can be obtained in empirical data versus potentially unlimited samples generated by simulation. Hence an added benefit of simulation was that larger samples of data could be generated. By using a model with simple but flexible assumptions, the impact of variables can more easily be understood.

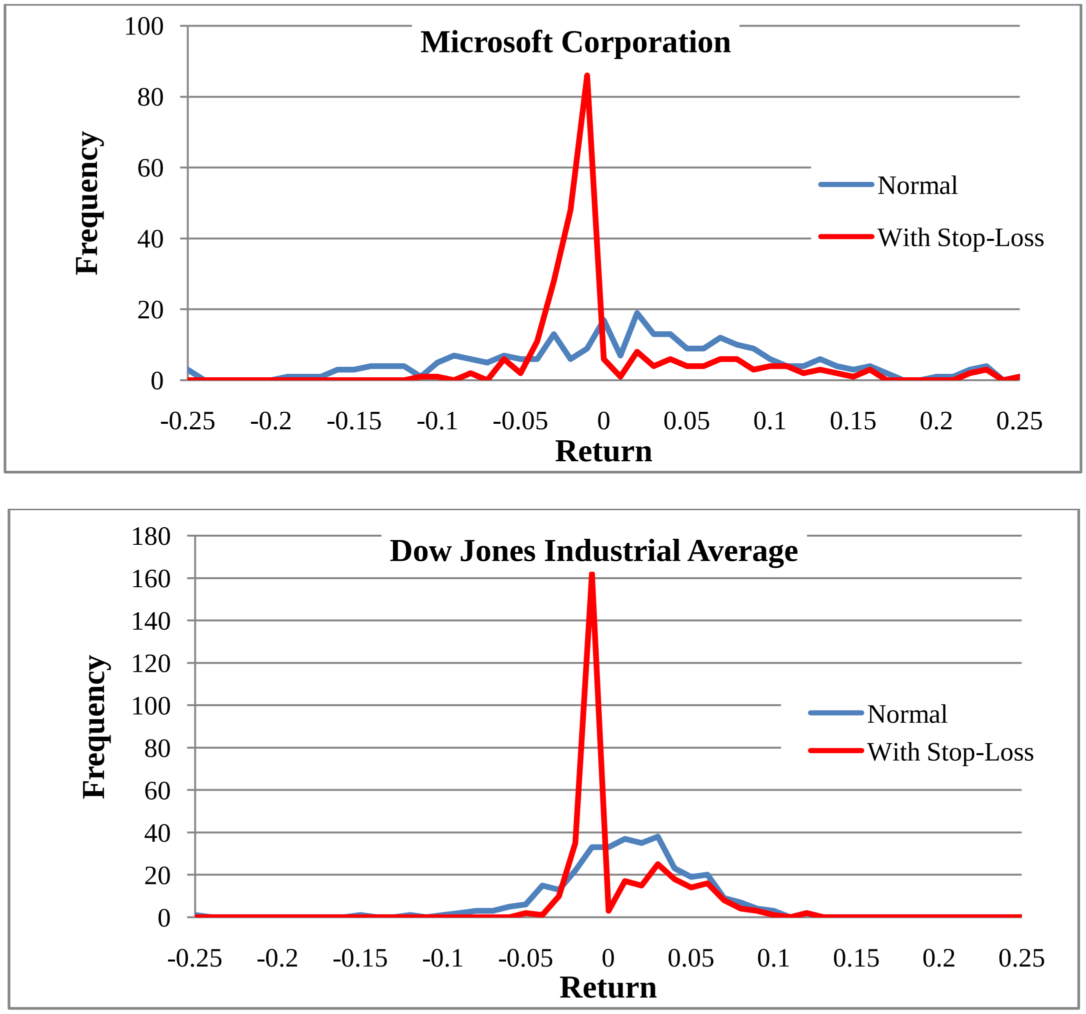

Figure 10: Stop-loss distributions on MSFT and DJIA

Practical Application

While the bulk of the paper so far has been in theory, practical application is the key to apply the results to the real world. So far, we've derived general implications of the stop-loss and then confirmed the results hold true for actual returns. From the implications of the impact of stop-losses in markets with drift, or "trending markets," we note that market conditions have an impact on what type of exit strategies to use. For example, in bull markets, use of a stop-loss actually reduces expected return by removing the possibility of the stock correcting and then recovering – setting a stop at -10% will reduce possibility of a +5% gain - while in bear markets stops increase expected return by preventing extended losses. Furthermore, we find that stops are potentially a way to reduce position volatility in a statistically significant manner, particularly in assets with higher volatility. Since use of stops as exits, profit-taking strategies, or both theoretically has negligible impact on expected return, the decision to use them and where to place them must be considered with other factors such as liquidity, frictional costs, and the context of the position with respect to strategy and portfolio structure as well as expected market conditions.

We summarize the implications in the following table:

	Volatility	Expected return		
	All Long	**Trailing Stop Long**	**Profit taking Long**	**Both Long**
Positive drift (trend up)	Decrease	Decrease	Decrease	Decrease
Flat (no trend)	Decrease	NC	NC	NC
Negative drift (trend down)	Decrease	Increase	Increase	Increase
	Volatility	**Expected return**		
	All Short	**Trailing Stop Short**	**Profit taking Short**	**Both Short**
Positive drift (trend up)	Decrease	Increase	Increase	Increase
Flat (no trend)	Decrease	NC	NC	NC
Negative drift (trend down)	Decrease	Decrease	Decrease	Decrease

Table 2: Implications of stop-loss strategies in asset price movement

Notice that regardless of the stop-loss strategy employed (trailing stop, profit-taking stop, or combinations of both), the impact on expected return is the same while volatility decreased throughout.

Peter Nilsson proposed a Money Management Matrix[7] which diagrams a way to apply entry and exit strategies with respect to market conditions for an investor always engaged in the market:

Stock/Market Trend	**Up**	**Sideways**	**Down**
Up	Trail/Pyramid	Trail	"2 for 1"
Sideways	Trail	"2 for 1"	P-target/Trail
Down	"2 for 1"	P-target/Trail	P-target

Table 3: Peter Nilsson's Money Management Matrix

The matrix implies that market trends and stock price movements have a positive relationship; stocks with significant betas would be ideal for such a strategy. By splitting conditions for stocks and market trends, the investor can take into account the impact of the market on stocks. Another assumption of the matrix is that the investor is able to reliably predict trends, or specifically, persistent drift in returns, the scope of which is outside this paper. In the matrix, "Trail" stands for trailing stop, while "2 for 1" indicates initiating trades only with risk to return profiles greater than two; the breakeven expectancy would require probability of 50% that the stock reaches its goal. "P-Target" stands for profit-target, equivalent to our profit-taking stop, while "Pyramid" indicates that one should add on to the position as it rises, so that each new position resembles the

[7] Nilsson, Peter. "Money Management Matrix." Technical Analysis of Stocks & Commodities. 2006.

layers of a pyramid in that they are decreasing. Using our implications, we make an update to the proposed strategy:

Stock/Market Trend	**Up**	**Sideways**	**Down**
Up	Pyramid/Trail	Trail	"2 for 1"
Sideways	Trail	"2 for 1"	P-target/Trail
Down	"2 for 1"	P-target/Trail	Trail/Pyramid

Table 4: Updated Money Management Matrix

The new results are consistent with trend-following or directional trading strategies.[8] Earlier, we mentioned that a short position with a stop loss was equivalent to a long put position. The same can be said about a long position with a stop-loss; its payoff is equivalent to a long call position. Since the return distributions were necessarily altered by having a stop-loss order out as shown in Figure 2, use of alternative risk management instruments such as options or offsetting hedges would allow one to continue participating in the upside while minimizing downside potential. Trailing stops are suggested for sideways markets so that reduction in volatility can be realized. However, since stops at +/- 1% yielded the most volatility reduction, taken at the limit, this would confirm intuition of simply not participating in the market at all in a sideways market with expected return equal to zero since we are not compensated even for the time value of money.

[8] Owen Katz, Jeffrey, and Donna L. McCormick. The Encyclopedia of Trading Strategies. 1st Ed. 2000.

Closing thoughts

Risk management is one of the most challenging elements of investing, even for something as fundamental as the stop-loss. Often the reason for this is because the understanding of risk involves in-depth analysis and understanding of dynamics that can be exceptionally ambiguous. However, it is important to have a firm grasp on all the tools used in investment management in order to more accurately assess how risk characteristics of the portfolio are altered.

We found that the perceived benefits of the stop-loss were largely balanced out by the hidden costs. As a result there is no inherent edge to be found in stop-losses or profit-taking stops. Furthermore, there is no clear optimal location to place a stop. As the stop moves further away from the asset price, the expected distribution of the stop-loss approaches that of the asset return distribution without a stop. Yet using a stop-loss price that is too close can lead to higher transaction costs without explicit or reliable gains. However, in certain situations the stop-loss may effectively reduce losses and volatility. In the end, it is up to the individual investor to consider the benefits and costs of the stop-loss in the context of the position, as illustrated by the Money Management Matrix. Given that one of the implications of the stop-loss was that it had a meaningful effect when prices exhibited persistent drift, further research may be dedicated to this area.

APPENDIX A.1: t-Test results for Means and F-Test for Variances (μ =0, μ = 10.18% annually)

t-Test: Two-Sample Assuming Unequal Variances

σ = 0.05	**μ =0**		**μ >0 (μ = 10.18% Annually, Bull Market)**	
	Stop-Loss=0.01	***No Stop-Loss***	***Stop-Loss=0.01***	***No Stop-Loss***
Mean	0.0001288	0.0000748	0.0078008	0.0153545
Variance	0.0010654	0.0023748	0.0015266	0.0023741
Observations	162000	162000	80999	80999
Hypothesized Mean Difference	0		0	
df	282997		154693	
t Stat	0.3706961		-34.4214072	
P(T<=t) two-tail	0.7108642		0.0000000	
t Critical two-tail	2.2414146		2.2414245	

F-Test Two-Sample for Variances

	Stop-Loss=0.01	***No Stop-Loss***	***Stop-Loss=0.012***	***No Stop-Loss3***
Mean	0.0001288	0.0000748	0.0078008	0.0153545
Variance	0.0010654	0.0023748	0.0015266	0.0023741
Observations	162000	162000	80999	80999
df	161999	161999	80998	80998
F	0.4486091		0.6430071	
P(F<=f) one-tail	0.0000000		0.0000000	
F Critical one-tail	0.9903081		0.9903081	

APPENDIX A.2: t-Test results for Means and F-Test for Variances (μ = -10.18% annually, μ =0)

t-Test: Two-Sample Assuming Unequal Variances

σ = 0.05	**μ<0 (μ = -10.18% Annually, Bear Market)**		**Stop loss and profit-taking stop, μ=0, σ = 0.05**	
	Stop-Loss = 0.01	***No Stop-Loss***	***Stop-Losses = 0.01***	***No Stop-Losses***
Mean	-0.0058067	-0.0149690	0.0000418	0.0000748
Variance	0.0007016	0.0023859	0.0003026	0.0023748
Observations	81000	81000	162000	162000
Hypothesized Mean Difference	0		0	
df	124842		202628	
t Stat	46.9295405		-0.2566453	
P(T<=t) two-tail	0.0000000		0.7974528	
t Critical two-tail	2.2414297		2.2414193	

F-Test Two-Sample for Variances

	Stop-Loss=0.01	***No Stop-Loss***	***Stop-Loss=0.012***	***No Stop-Loss***
Mean	0.0001288	0.0000748	0.0078008	0.0153545
Variance	0.0010654	0.0023748	0.0015266	0.0023741
Observations	162000	162000	80999	80999
df	161999	161999	80998	80998
F	0.4486091		0.6430071	
P(F<=f) one-tail	0.0000000		0.0000000	
F Critical one-tail	0.9903081		0.9903081	

APPENDIX A.3: t-Test results for Means and F-Test for Variances (σ = 0.10, σ = 0.15)

t-Test: Two-Sample Assuming Unequal Variances

Higher Variance	**μ =0, σ = 0.1**		**μ =0, σ = 0.15**	
	Stop-Losses = 0.01	***No Stop-Losses***	***Stop-Loss = 0.01***	***No Stop-Loss***
Mean	-0.0000182	0.0000998	0.0001229	0.0000330
Variance	0.0032970	0.0095575	0.0067503	0.0215068
Observations	162000	162000	162000	162000
Pearson Correlation	0		0	
Hypothesized Mean Difference	261881		254572	
df	-0.4191755		0.2150606	
t Stat	0.3375441		0.4148601	
P(T<=t) two-tail	2.2414156		2.2414159	

F-Test Two-Sample for Variances

	Stop-Losses = 0.01	***No Stop-Losses***	***Stop-Loss = 0.01***	***No Stop-Loss***
Mean	-0.0000182	0.0000998	0.0001229	0.0000330
Variance	0.0032970	0.0095575	0.0067503	0.0215068
Observations	162000	162000	162000	162000
df	161999	161999	161999	161999
F	0.3449631		0.3138691	
P(F<=f) one-tail	0.0000000		0.0000000	
F Critical one-tail	0.9903081		0.9903081	

APPENDIX A.4: t-Test results for Means and F-Test for Variances (σ = 0.20, σ = 0.25)

t-Test: Two-Sample Assuming Unequal Variances

Higher Variance	**μ =0, σ = 0.20**		**μ =0, σ = 0.25**	
	Stop-Loss = 0.01	***No Stop-Loss***	***Stop-Loss = 0.01***	***No Stop-Loss***
Mean	0.0004262	0.0002300	-0.0001511	0.0006898
Variance	0.0113976	0.0380179	0.0170731	0.0592211
Observations	162000	162000	162000	162000
Pearson Correlation	0		0	
Hypothesized Mean Difference	251122		248238	
df	0.3552889		-1.2253280	
t Stat	0.3611867		0.1102264	
P(T<=t) two-tail	2.2414161		2.2414163	

F-Test Two-Sample for Variances

	Stop-Loss = 0.01	***No Stop-Loss***	***Stop-Loss = 0.012***	***No Stop-Loss3***
Mean	0.0004262	0.0002300	-0.0001511	0.0006898
Variance	0.0113976	0.0380179	0.0170731	0.0592211
Observations	162000	162000	162000	162000
df	161999	161999	161999	161999
F	0.2997966		0.2882951	
P(F<=f) one-tail	0.0000000		0.0000000	
F Critical one-tail	0.9903081		0.9903081	

APPENDIX A.5: t-Test results for Means and F-Test for Variances (DJIA & MSFT)

t-Test: Two-Sample Assuming Unequal Variances

	Dow Jones Industrial Average		Microsoft	
	With Stop-Loss			
	Stop-Loss = 0.01	***No Stop-Loss***	***Stop-Loss = 0.01***	***No Stop-Loss***
Mean	0.0026893	0.0061097	0.0103114	0.0186031
Variance	0.0008946	0.0017787	0.0056309	0.0116224
Observations	336	336	264	264
Pearson Correlation	0		0	
Hypothesized Mean Difference	604		469	
df	-1.2126118		-1.0256715	
t Stat	0.1128762		0.3055752	
P(T<=t) two-tail	2.2470047		2.2486222	

F-Test Two-Sample for Variances

	Stop-Loss = 0.01	***No Stop-Loss***	***Stop-Loss = 0.012***	***No Stop-Loss2***
Mean	0.0026893	0.0061097	0.0103114	0.0145320
Variance	0.0008946	0.0017787	0.0056309	0.0093840
Observations	336	336	264	229
df	335	335	263	228
F	0.5029396		0.6000488	
P(F<=f) one-tail	0.0000000		0.0000327	
F Critical one-tail	0.8069178		0.7785942	

APPENDIX A.6: t-Test results for Means and F-Test for Variances (Ford & BRK-A)

t-Test: Two-Sample Assuming Unequal Variances

	Ford		Berkshire Hathaway	
	Stop-Loss = 0.01	**No Stop-Loss**	**Stop-Loss = 0.012**	**No Stop-Loss3**
Mean	0.0013649	-0.0003961	0.0072722	0.0108396
Variance	0.0029439	0.0079417	0.0019827	0.0035216
Observations	374.0000000	374.0000000	218.0000000	218.0000000
Hypothesized Mean Difference	0		0	
df	616		403	
t Stat	0.326408206		-0.709961497	
P(T<=t) two-tail	0.7442264		0.4781386	
t Critical two-tail	2.2468953		2.2498088	

F-Test Two-Sample for Variances

Column1	**Stop-Loss = 0.01**	**No Stop-Loss**	**Stop-Loss = 0.012**	**No Stop-Loss3**
Mean	0.0013649	-0.0003961	0.0072722	0.0108396
Variance	0.0029439	0.0079417	0.0019827	0.0035216
Observations	374.0000000	374.0000000	218.0000000	218.0000000
df	373	373	217	217
F	0.370687397		0.563009434	
P(F<=f) one-tail	0.0000000		0.0000135	
F Critical one-tail	0.8160524		0.7658244	

APPENDIX B: A response to the *Hidden Cost of the Stoploss*

The motivation of our research stems from the initial survey of a hidden cost of stoploss detailed in an article written by Robert Macrae of Arcus Investments, entitled "The Hidden Cost of the Stop-loss." In researching this area we found a dramatic lack of meaningful studies in the topic. The implications of the Arcus paper are that stops impact expected return distribution and create higher portfolio volatility, the opposite of what is intended. The reasoning for this is that increased leverage must be used to maintain same exposure. Macrae's conclusion is that this increased volatility is the "hidden cost" of the stop-loss.

In generating his results, Macrae uses the assumption that hitting the stop does not change the expectation of future return. Our research confirmed this assumption to arrive to similar conclusions, in that stop-losses are helpful in trending returns but harmful in mean-reverting returns. Macrae also finds that stop-losses actually cause increased volatility at the portfolio level, as the portfolio manager must lever up the initial position in order to maintain the same exposure. We note that in this situation one can choose either to minimize risk (losses) for the same level of expected return or maximize returns for the same level of risk as a passive position. Macrae's strategy employs the latter, resulting in larger volatility as increased leverage was used to maintain the same level of risk exposure. As a result he noted increased portfolio volatility, which is not surprising due to increased leverage. We illustrate the result with two scenarios:

Scenario 1: A long position is entered at a stock at $10. If you have no stop, then the maximum amount you can lose is $10 for investing in one share. Expected return for the position is the expected return for the stock.

Scenario 2: You enter the same position at $10, but you set a stop at $5. Theoretically if stop is hit, max loss = $5 (plus frictional costs); to maintain the same $10 exposure (max loss) as in Scenario 1, you would need to double the position to 2 shares - that way if stop is hit you lose

$10; expected return for position = 2x expected return of stock. Increased volatility ensues because of leverage as well as increased local probability of the stop being hit.

Position-level volatility, on the other hand, is significantly reduced by the use of the stop-loss in both trending and non-trending returns. The table below shows returns and volatility for various stop-losses in a simulation with positive drift. Note that while volatility is reduced, expected return is reduced dramatically so that the risk-adjusted return metric (return / volatility, equivalent to Sharpe ratio with risk-free rate equal to zero) is harmed. This holds true even after position-sizing is re-adjusted to maintain volatility or return.

Stop (%)	E[Return]	Volatility	Return / Volatility	Volatility-adjusted weight	Levered return	Levered volatility	Return / Volatility	Return-adjusted weight	Levered return	Levered volatility	Return / Volatility
25.00	1.6353%	4.8725%	0.34	1.0000	1.6353%	4.8725%	0.34	1.0000	0.0164	4.8725%	0.34
20.00	1.6353%	4.8725%	0.34	1.0000	1.6353%	4.8725%	0.34	1.0000	0.0164	4.8725%	0.34
15.00	1.6352%	4.8731%	0.34	0.9999	1.6350%	4.8725%	0.34	1.0001	0.0164	4.8736%	0.34
10.00	1.6321%	4.8758%	0.33	0.9993	1.6310%	4.8725%	0.33	1.0020	0.0164	4.8855%	0.33
7.00	1.6051%	4.8818%	0.33	0.9981	1.6021%	4.8725%	0.33	1.0188	0.0164	4.9736%	0.33
5.00	1.5357%	4.8509%	0.32	1.0045	1.5426%	4.8725%	0.32	1.0649	0.0164	5.1656%	0.32
4.00	1.4624%	4.7950%	0.30	1.0162	1.4860%	4.8725%	0.30	1.1182	0.0164	5.3620%	0.30
3.00	1.3401%	4.6723%	0.29	1.0429	1.3976%	4.8725%	0.29	1.2203	0.0164	5.7015%	0.29
2.00	1.1449%	4.4133%	0.26	1.1041	1.2641%	4.8725%	0.26	1.4283	0.0164	6.3036%	0.26
1.00	0.8775%	3.9071%	0.22	1.2471	1.0944%	4.8725%	0.22	1.8635	0.0164	7.2811%	0.22
			(1)				(2)			(3)	

Stop-losses on underlying with positive drift.

1) Note that return / volatility ratio does not change with leverage but declines
2) In the volatility-adjusted weighting, volatility remains the same and expected return improves slightly.
3) In return-adjusted weightings, returns equal that of the underlying with nom stop but volatility increases.

We also note that in generating his proof of concept charts, Macrae used smooth functions, creating continuous bimodal distributions when in fact our findings suggest that bimodal peaks occur, but with a 'wall' - e.g. if stop is set at -10% below the entry price, if the price goes to -9.9%, the stop will never be hit and thus the frequency of -9.9% (or any area local to the stop but not surpassing it) should not be affected. Figure 4 illustrates this. In our investigation of the implications of Macrae’s article, we find that there are hidden costs to the stop-loss, but not due to the effect of increased volatility from leverage used to maintain effective exposures.

APPENDIX C: An illustration of hidden costs of a stop-loss using a coin flipping game

This illustration shows that the impact of the stop-loss is distribution-independent. Say you have a coin and flip it 3 times. What is the probability of getting at least one head [$P(H \geq 1)$] and how does that probability change if you stop flipping if you get one tail?

Eight outcomes are possible: HHH, HHT, HTH, HTT, THH, THT, TTH, TTT.

P(H >= 1) = 7/8 (the only time you'd not get at least one H is TTT). Now let's find out what happens when you STOP flipping after you get one T: HHH, HHT, HT, HT, T, T, T, T.

P(H >=1 | stop rule) = 1/2! This is the 'hidden' cost of the stop-loss and why the probability distribution has to be lower in areas greater than the stop-loss (The probability of getting a return of 0% is decreased!)

To prove expected value doesn't change, say that T = some 'bad' with a value of -1 and H = some 'good' with a value of +1, so that HT = 0 and HTH = 1 for example. The possible outcomes with [value in brackets]: HHH [3], HHT [1], HTH [1], HTT [-1], THH [1], THT [-1], TTH [-1], TTT [-3]. Adding up the values, we find the expected value is zero. Now with the 'stop' at T: HHH [3], HHT [1], HT [0], HT [0], T [-1], T [-1], T [-1], T [-1]. Here we also see that the expected value is 0. The two are equivalent in a game with a fair die.

Now call expected value μ, so that $\mu = 0$, and x is each individual run's value, so that we have two series, c and d, corresponding to having a stop versus not having a stop:

c = [3, 1, 1, -1, 1,-1, -1, -3]

d = [3, 1, 0, 0, -1, -1, -1, -1]

Now let's measure each variance:

var(c) = [(3-0)^2, 1^2... etc.] = 24

var(d) = [(3-0)^2, ... etc.] = 14

Lo and behold variance has decreased. Since standard deviation is the square root of variance, standard deviation of returns should also decrease. We found that for empirical data, assets with larger volatility experienced a more meaningful reduction in volatility from the stop-loss.

APPENDIX D: Details on the assumptions of the random walk model

i) Returns are normally distributed [$r_t \sim N(\mu, \sigma^2)$] and prices are lognormally distributed. This is approximately true for assets such as stocks, although much research has shown financial assets may tend to exhibit excess kurtosis, skew, and other non-normal artifacts. The implications of the model are such that the effects to using stops are indifferent to non-normal returns. An illustration of this is provided in Appendix C, where the results are replicated with a coin-flipping game (fair coins have a uniform distribution).

ii) Expected return is constant over time [$\mu_t = \mu_{t+1} = E(r_t)$]. When this assumption is relaxed the conclusion is the same as we demonstrate that expected return of the stop-loss strategy approaches that of the underlying asset. Results of a model with nonstationary returns with varying μ were simulated and we found the results held.

iii) Return volatility is constant over time [$\sigma^2 = var(r_t)$; $\sigma_{ij} = cov(r_{it}, r_{jt})$; $\rho_{ij} = cor(r_{it}, r_{jt})$]. Errors are *iid* and covariance stationary; homoskedasticity was assumed in the CER model. We generated returns of assets with varying σ to simulate expected returns for assets with low and high volatility. To relaxing this assumption, one would merely ‘shift’ from one volatility to another as we generated representative volatilities with σ = 0.01, 0.05, 0.10, 0.15, 0.20, and 0.25. One can picture the aggregate results as a set of level surfaces of varying σ. With nonstationary returns, $var(r_t) = \sigma^2 t$. We find that there is no difference in the expected return among these level surfaces.

iv) Expected value of error is zero [$E(\varepsilon) = 0$] and is normally distributed with variance σ^2. The interpretation of this is that 'good news' represented by a random positive error term and 'bad news' represented by a negative error term balance out in the long run, with a net result being 0 impact to expected asset return. In the linear representation these errors terms capture the residual impact of outside forces.

Works Cited

Hull, John C. Options, Futures and Other Derivatives. Prentice Hall, 6th Ed. 2005.

Owen Katz, Jeffrey, and Donna L. McCormick. The Encyclopedia of Trading Strategies. 1st Ed. 2000.

Macrae, Robert. "The Hidden Cost of the StopLoss." Arcus Investments, AIMA Journal. 2005.

Nilsson, Peter. "Money Management Matrix." Technical Analysis of Stocks & Commodities. 2006.

Rupert, David. Statistics and Finance. Springer. 2006.

Seykota, Ed. Risk Management. The Trading Tribe. 2003. http://www.seykota.com/tribe/risk/index.htm.

Woolridge, Jeffrey. Introductory Econometrics: A Modern Approach. South-Western College Pub:3rd edition, 2005.

Zivot, Eric. *Lecture notes for Economics 424: Computational Finance*. University of Washington: 2007.