The Hidden Cost of the Stoploss

Robert Macrae, Arcus Investment May 2005

Risk control is a complex subject, and an area where simple and robust guides are particularly valuable. Few concepts are simpler and more intuitive than the stoploss, and it has taken a vice-like grip on some investors' approach to the subject. During investor meetings, I have even been asked: "Let us talk about risk control... what kind of stoploss do you use?"

Unfortunately, despite its appeal the stoploss is no panacea. In fact, it is not even a good system for the vast majority of portfolio investors, as a simple simulation will illustrate. The stoploss rule I will consider is that any position is automatically closed when it loses more than a certain fixed percentage from the initial price at which the trade was made.

Before looking at why this may not work in portfolios, it is worth considering a situation in which it can work very well. Imagine a commodity trading adviser (CTA), trading concentrated positions in one or a handful of liquid assets, and making money following trends in the market. If the CTA loses money on a position then it has misjudged the existence of a trend and it should get out, bide its time, and again jump in the next time it thinks a trend has started.

It can easily be shown that for positively-autocorrelated (or trending) series, stoplosses can help because when the stop is hit your expectation is that further losses would have followed. This suggests that stoplosses can make money as well as reducing downside, but how realistic a model is this imaginary CTA of what most alternative investors do?

Most investors construct a diversified portfolio of assets. At any time a portfolio might have between 10 and 100 positions contributing significantly to the overall risk profile. In this situation the worst losses for the portfolio are likely to be caused by common movements in related groups of positions rather than by events that affect only a single position, and stoploss is quite blind to this portfolio aspect of risk.

Some alternative investors follow trades that are mean reverting, so that returns are negatively autocorrelated. Examples might include value investors betting that disliked companies will move back into the mainstream, or arbitrage and relative value traders betting that assets will move towards some economic relationship. On mean-reverting trades, stoplosses systematically take you out of your best positions.

Many alternative investors trade in relatively illiquid assets, in which case slippage - the difference between the level at which you place the stop and your average trade price - will cost you a significant amount every time your stop is hit.

Those are three serious problems, but obviously they do not apply to all investors all the time. However, there are two problems that are much more fundamental because they do almost always apply. Provided that hitting the stop does not change our expectation of future return or future volatility, it is easy to show that stoplosses will increase trading cost and portfolio risk, and this should concern any investor who cares about their Sharpe ratio.

Figure 1

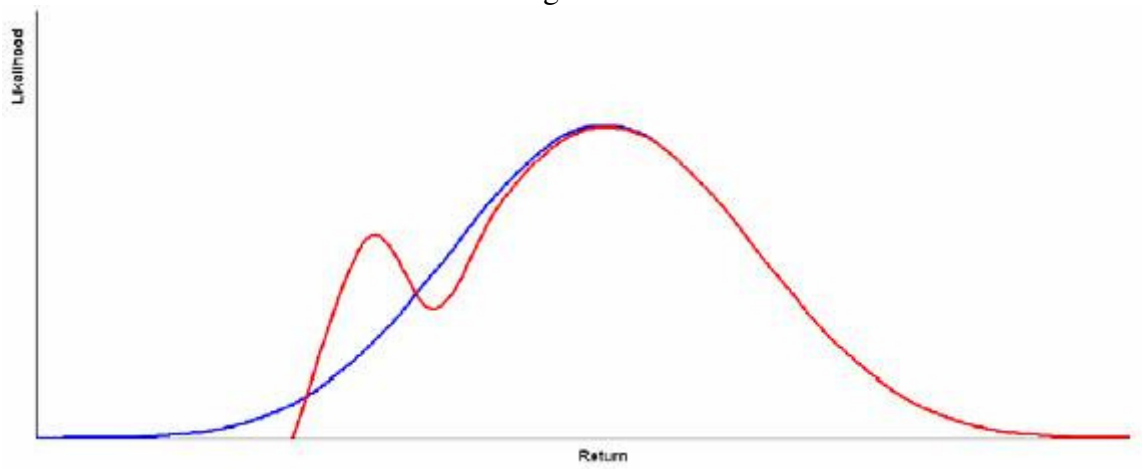

The blue bell curve in fig 1 represents a trade, with return on the X axis and likelihood on the Y. We may hope that the mean return is somewhere positive, but what we know for certain is that as soon as we put money on the table there is uncertainty, or risk. Now let's simulate a stoploss trading rule, shown in the red curve on fig 1. We can choose how tight to make the stop; the tighter it is, the more we restrict the worst loss we can take, but also the more likely that we are "stopped out" and exit the trade, and fig 2 shows a family of stoplosses that result from setting the stop at different levels. In all these curves, I have ignored frictional costs so that the average return of the trade is not changed by the addition of the stoploss.

Figure 2

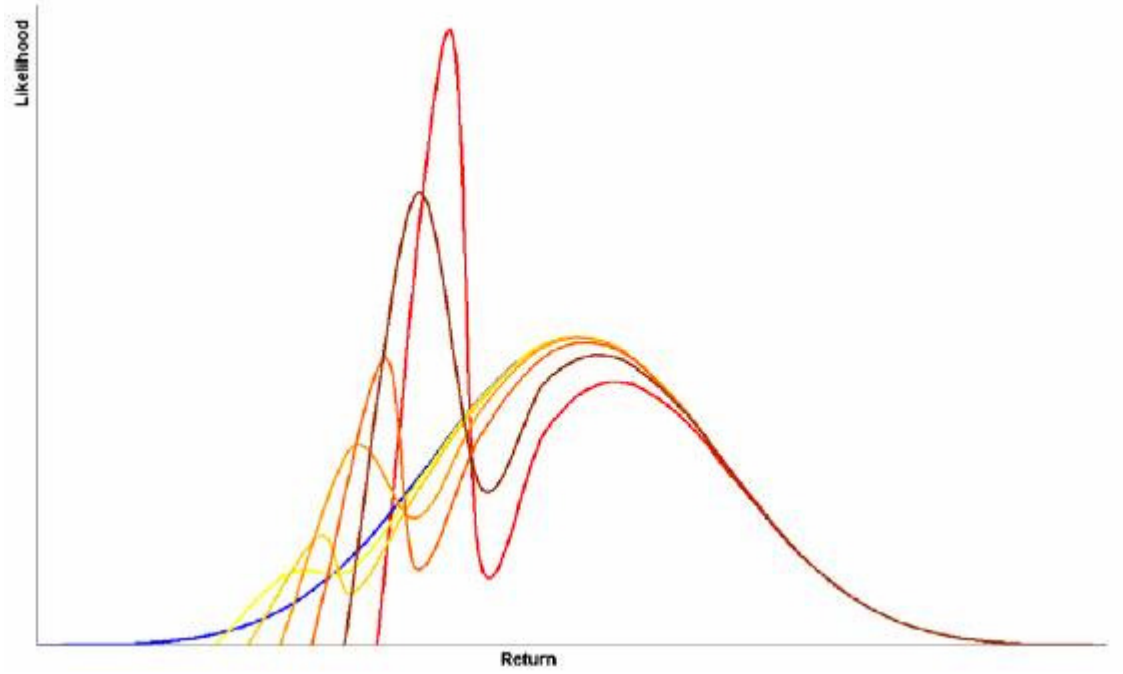

Table 1

Stop	Position Volatility	Leverage	Portfolio Volatility
0%	20.0%	100%	20.0%
4%	19.9%	104%	20.8%
6%	19.8%	107%	21.2%
12%	19.7%	113%	22.2%
17%	19.4%	120%	23.3%
26%	18.9%	136%	25.7%
36%	18.2%	156%	28.3%

Table 1 illustrates why, even in this favourable simulation, the stoploss hurts performance. The first two columns of the table show the likelihood the stop is hit and the standard deviation (volatility) of the trade that results. I've assumed the pure trade had a standard deviation of 20%, and you can see a modest reduction in volatility due to the stoploss, because the elimination of largest negative results tends to outweigh the matching reduction in results near zero. However, even for a stop that is hit 36% of the time, the reduction is only from 20% to 18%. This may seem somewhat disappointing, but it gets much worse when we consider the effect on the portfolio as a whole.

The third column looks at the impact on leverage that is required if we wish to retain the same average exposure to the trade. If there is a 36% chance we are stopped out, we need to make bets 1 / (1 - 0.36) = 1.56x as large just to maintain the same average exposure.

We pay the price for focussing on the position and ignoring the portfolio by being forced to use additional leverage. At the portfolio level this has two horrible effects:

1. It increases our volatility, in this case to 18.2% x 1.56 = 28.3%.
2. It increases our transaction costs, in this case by something like 1.56x and so reduces our expected return.

Both sides of Sharpe ratio have been damaged, a disastrous result for "risk control".

Though the table shows results for volatility, the same basic argument applies to all other risk measures. Stoploss takes us out of some of the available trades, so we are making inefficient use of diversification, so we are running more risk than we should.

For Table 2, I have picked a measure chosen to flatter stoploss, the worst simulated result for a single position. This is exactly what stoploss is designed to reduce. However, when adjusted for the required increase in average position size it is clear that stoplosses cease to improve this statistic if the stop is being hit more than a small percentage of the time. Even in terms of its own objectives, stoploss fails.

Table 2

Stop	Worst Position	Leverage	Worst Portfolio
0%	-51%	100%	-51%
4%	-30%	104%	-31%
6%	-27%	107%	-29%
12%	-25%	113%	-28%
17%	-22%	120%	-26%
26%	-19%	136%	-26%
36%	-17%	156%	-26%

For a portfolio investor in non-trending trades, stoploss is a bad approach. It looks in the wrong place, leads to unprofitable turnover, increases transaction costs, decreases expected returns and increases volatility. Just say no.

This article originally appeared in the April 2005 issue of the AIMA Journal and is republished with their permission.