

Specifying and managing tail risk in portfolios a practical approach

Pranay Gupta

Sven Skallsjo

Current Version: 28 May, 2013

Abstract

Tail risk arises at multiple stages in the investment management process – from the high level asset allocation decision down to the individual portfolio manager’s process for selecting securities. We believe that conventional practices followed in these investment decision processes, largely ignore intra-horizon risk. We believe this leads to sub-optimal assessment of risk of assets, particularly in the context of potential tail risk, and leads to the construction of portfolios, which are not in sync with the risk aversion of the client.

In the present paper we propose a composite risk measure which simultaneously captures the risk of breaching a specified maximum intra-horizon drawdown threshold, as well as the risk that the performance is not met at the end of the investment horizon. We believe this captures the ‘true’ risk of a portfolio, much better than traditional end of horizon risk measures.

We find that intra-horizon risk can represent a substantial part of the total risk, and thus needs to be managed explicitly which constructing a portfolio of assets, strategies or asset classes. We propose that varying the investment horizon and implementing a customized stop loss for each asset can help construct a portfolio where portfolio risk is kept within bounds of tolerance, and can improve performance over time.

JEL Classification

G1, G3

SSRD ID : 2270914

Disclaimer

All views expressed in this article are solely of the authors, and do not in any reflect the views of Lombard Odier Darier Hentsch or any of its employees. The authors do not make any representation or warranty, express or implied, as to the accuracy and completeness of the information set forth herein, nor do the authors recommend that the attached information serve as the basis of any investment decision. This document is meant exclusively for the use of the person to whom it has been delivered to by the authors, and it is not to be redistributed or reproduced in any way.

1. Introduction

Generically, the full investment management process can be articulated in the following steps:

- An asset owner or allocator, based on his return requirement and risk tolerance, performs an asset allocation to allocate assets into various asset class silos.
- Various strategies or managers are selected to fulfill the allocation in each silo.
- The portfolio manager of each strategy invests the assets into market securities.

We examine these steps with a focus on the generation, specification and management of tail risk, and try to identify methods with which the tail risk can be managed within specified limits. Specifically, we focus on the asset allocation decision in step 1, and the security selection decision in step 3.

The asset allocation decision is of interest as this decision is done with low breadth, and the drawdown caused by asset class movements is the primary cause of asset liability funding gaps at asset owners. Gupta and Skallsjo (2013) detail a new framework for plan sponsor asset allocation decisions, which can mitigate tail risk in plan sponsor portfolios. The security selection decision is of interest as that is the ground level where tail risk management can perhaps be implemented.

The remainder of the paper is organized as follows. Section 2 discusses the basic practical setting of how assets are managed. With this as motivation we propose in Section 3 a new risk measure, which we believe captures the true risk in managing portfolios, incorporating both end-of-horizon and intra-horizon risk. Section 4 presents an extension of the standard model, where it is assumed that the mean return is not known with certainty. Although this is an uncertainty aspect with long and extensive history in the literature, its implications for risk in a dynamic context appears to be limited. Finally in Section 5 we study the related subject of stop losses and how they can be used to improve performance when model parameters are not known with certainty.

2a. Portfolio Management - the Practical Setting

Consider the case of an active equity portfolio manager investing in securities of, say, the S&P500 index, with the objective of beating the stock benchmark. In order to achieve this aim, he uses his judgment combined with various fundamental or quantitative techniques, to arrive at those stocks which he believes are likely to have a higher target price after a certain time period. He constructs a portfolio of these stocks, with the aim of achieving his target return objective for the fund, while seeking never to have a large enough loss at any time, where his client (the asset owner) would lose faith in the investment process and withdraw the assets.

In essence, two external parameters are given to the portfolio manager – a return target and a maximum drawdown threshold that the client is prepared to tolerate, breaching which would impose a stop loss decision. The risk that the portfolio manager takes is that the stop loss is breached intra-horizon or that the return target is not met at the end of the horizon.

Parameters inherent in the setting are the stock universe and the long term Sharpe ratio of the manager.

The controls which are in the hands of the portfolio manager are the number of stocks he selects in the portfolio, their volatilities and correlations, and the investment horizon over which the manager makes a price forecast for each stock.

Conceptually, uncertainty in all of these controls as well as the inherent parameters represent sources of risk in the final portfolio. While this is a very standard situation in asset management, in practice portfolio managers seldom use the investment horizon as a parameter, nor do they specifically incorporate the management of intra-horizon risk in their portfolio management process.

2b. Asset Allocation - the Practical Setting

The setting for the asset allocation decision is the same as that for security selection, where the assets are asset classes rather than individual securities. Here again a return target is usually specified along with an intra-horizon drawdown threshold. However, when constructing the allocation these constraints are often not fully incorporated. Instead, a common approach is to make an assessment of what long run Sharpe ratio is attainable using the main asset classes. The risk exposure is then calibrated to align the portfolio's expected return with the specified target return. This is an approach which completely ignores intra-horizon risk. The portfolio might achieve its target return in expectation, but does so at a certain risk of breaching the intra-horizon drawdown threshold. In case this occurs, the portfolio is terminated with the immediate consequence that the return target is not achieved. Should intra-horizon risk be incorporated into the initial allocation decision, it could lead to a different allocation. In fact, a more appropriate allocation design could be to ask what expected return is feasible subject to a constraint that the probability of breach does not exceed some pre-specified value.

We therefore seek to develop a model which incorporates both end of horizon and intra-horizon risk, and where we are able to define the control parameters to meet the requirements of the asset owner. Specifically, given a target return and a maximal loss acceptance intra-horizon for both the asset allocation and stock selection decision, we seek to decide:

1. the investment horizon that would be permissible in each of the investment decisions,
2. the skill level required by the portfolio manager or asset allocator, which would enable him to meet the target return without breaching the intra-horizon threshold,
3. the stop loss policy which would be consistent with the investment process being deployed,
4. the relative skill in different segments of the portfolio (or portfolio managers), given that the stock universes are different.

Another factor which is a determinant of tail risk is the amount of leverage used. We assume however an unleveraged portfolio, and further assume that the manager responsible for investment decisions has positive skill, though this can vary over time. Finally, tail risk can also be a function of an exogenous shock, creating event risk. Gupta et al. (2006) discuss the management of this risk in portfolios, and we will omit this also from our discussion here.

As a starting point, we first define a risk measure, which incorporates both intra and end-of-horizon risk.

3. Creating a real risk measure : End-of-horizon vs intra-horizon risk

In either of our previously defined settings, a straightforward definition of risk is the probability that either condition is violated, which effectively comes to formulate risk as a function of two variables, end-of-horizon performance and intra-horizon loss.

Setting focus on these two sources of risk, we introduce the two variables X_t and Y_t defined as

$$X_t = \log(P_t/P_0), \quad Y_t = \min_{0 \leq s \leq t} \log(P_s/P_0), \quad (3.1)$$

where P_t is the value of the portfolio at time t . Here X_t is simply a normalization of the portfolio, where expressing it as the logarithm of the portfolio somewhat simplifies the exposition. The variable Y_t represents our notion of intra-horizon loss – the minimal value that X_t assumes intra-horizon. With regards to intra-horizon risk we thus focus on loss since inception. While the literature on intra-horizon risk pays significant attention to peak to trough maximum drawdown and variants thereof, in our setting we believe loss since inception is more relevant.

With the two risk variables developed, Definition 3.1 gives a formal definition of the risk measure. Note that this makes no use of volatility. Volatility is a means by which risk can be controlled, and often the

investment manager is also subjected to a constraint that volatility does not exceed a certain level. However, the risk that is under the investment manager's responsibility is that of breaching the intra-horizon or end-of-horizon threshold.

Definition 3.1

For a portfolio P_t and an investment horizon $0 \leq t \leq T$, let X_t and Y_t be as in (3.1). Also, let x be the end-of-horizon target return and let y be the intra-horizon drawdown threshold, both expressed with reference to the logarithm of P_t/P_0 . The *tail risk* of P_t is denoted ψ and is given by

$$\psi(x, y) = \mathbb{P}(X_T \leq x \text{ or } Y_T \leq y). \quad \square$$

A convenient representation of ψ is obtained by breaking it up in two components corresponding to intra-horizon and end-of-horizon risk. Intra-horizon risk is the probability of breaching y at some point before T regardless of end-of-horizon performance, and it is given by $\psi(y, y)$. End-of-horizon risk is the risk of being below x at the end of the horizon while not breaching y at any point prior to that. This results in a decomposition of the risk measure according to

$$\psi(x, y) = \psi_{IH}(y) + \psi_{EH}(x, y), \quad (3.2)$$

where

$$\psi_{IH}(y) = \psi(y, y), \quad \psi_{EH}(x, y) = \psi(x, y) - \psi(y, y).$$

In the special case that only intra-horizon risk matters, (3.2) can be applied with $x = y$ since this ensures $\psi_{EH} = 0$. In the opposite extreme case that risk is only a function of end-of horizon performance, (3.2) can be applied letting $y \rightarrow -\infty$. The risk decomposition is illustrated in Figure 3.1, where the straight dotted line represents the portfolio's target performance. The solid line portfolio path exemplifies an intra-horizon breach, while the circled line is an example of end of horizon shortfall.

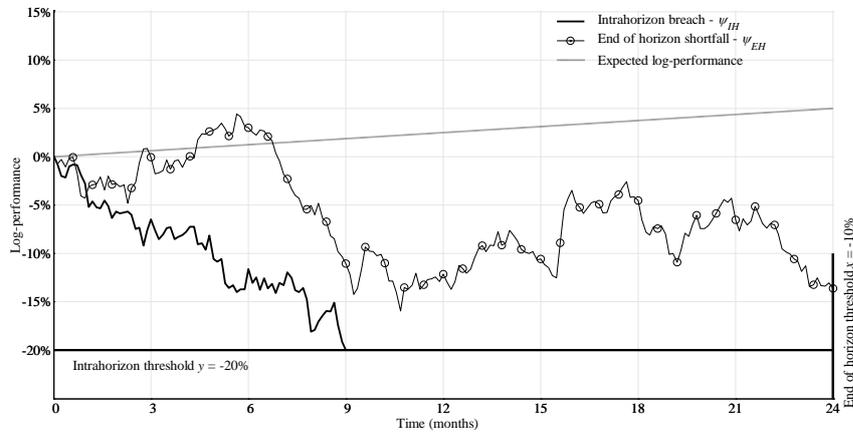


Fig 3.1: Illustration of the risk decomposition in intra-horizon and end of horizon risk. Settings: End-of-horizon threshold, $x = -10\%$, intra-horizon threshold $y = -20\%$, volatility $\sigma = 10\%$ and mean return $\mu = 3\%$.

To examine how ψ behaves under standard assumptions about asset prices we derive explicitly a formula in the case that the portfolio follows a geometric Brownian motion with drift. Suppose therefore that the portfolio evolves according to

$$dP_t = \mu P_t dt + \sigma P_t dW_t, \quad (3.3)$$

with the initial portfolio value $P_0 = 1$. Here W_t is a standard Brownian motion starting at zero, the drift term μ is the annualized expected return over the investment horizon and σ the annualized volatility.

Theorem 3.1

Suppose that the portfolio P_t evolves according to (3.3). Fix an investment horizon $0 \leq t \leq T$ and let X_t and Y_t be as in (3.1). Then, for any x, y with $y \leq 0$ and $x \geq y$ the tail risk ψ is given by

$$\psi(x, y) = \mathcal{N}(d_1) + \exp\left(\frac{2y\mu_{\log}}{\sigma^2}\right) \mathcal{N}(d_2),$$

where

$$d_1 = \frac{x - \mu_{\log}T}{\sqrt{\sigma^2 T}}, \quad d_2 = \frac{2y}{\sqrt{\sigma^2 T}} - d_1, \quad \mu_{\log} = \mu - \sigma^2/2$$

and \mathcal{N} denotes the standard normal distribution function. □

A proof is presented in Appendix A.

Figure 3.2 gives an example of ψ as a function of the time horizon using the decomposition (3.2). Parameters are selected to exemplify a pension fund with a 60/40 equity-fixed income allocation, which typically results in a total volatility around 10%. Assuming a Sharpe ratio of 0.3 and a risk free rate of zero this implies a μ of 3%. We investigate an end-of-horizon threshold of -10% and an intra-horizon drawdown tolerance of 20%, thus setting $y = -20\%$.

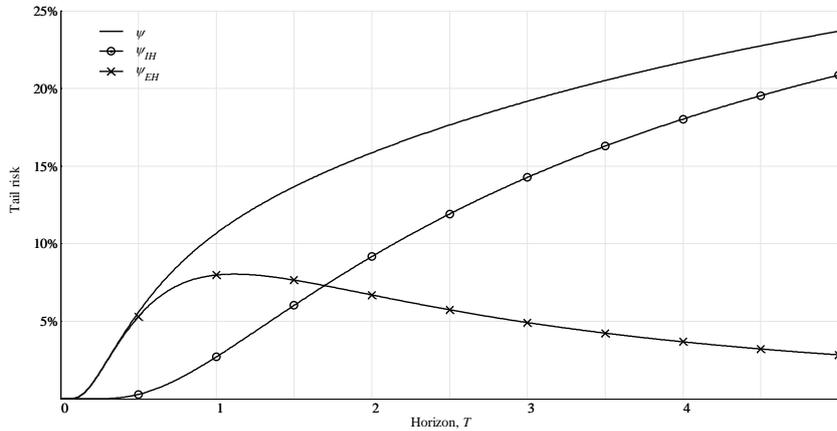


Fig 3.2: Decomposition of the risk function into intra- and end-of-horizon risk. Settings: End-of-horizon threshold, $x = -10\%$, intra-horizon threshold $y = -20\%$, volatility $\sigma = 10\%$ and mean return $\mu = 3\%$

The figure demonstrates a hump shaped form for end-of-horizon risk ψ_{EH} . Initially the risk goes up simply as a result of volatility in the asset price process. However, due to a positive mean return the probability of ending below a given threshold diminishes over longer horizons. In contrast, intra-horizon risk increases monotonically with the time. At the five year horizon, although the fund shows a mere 3% risk of being below -10%, the risk of having breached -20% in the meantime is 21%.¹ As the time horizon increases further, intra-horizon risk continues to go up and end-of-horizon risk continues to go down. In the current example total risk ψ levels out at around 37%, all resulting from intra-horizon risk.

The importance of intra-horizon risk addresses the first point listed in section 2b. In the presence of a tangible maximum intra-horizon drawdown threshold, it is not in the interest of the asset owner to have a long term investment horizon, as such a stance would inordinately increase the probability intra-horizon of breaching a tolerable threshold, thus leading to closure or demise. This is in sharp contrast to the currently accepted practice followed at most pension plans and sovereign wealth funds, where a long term investment horizon is standard procedure.

¹ Note that the 3% probability of being below -10% is while not having touched the intra-horizon threshold in the meantime. At the five year horizon the probability of being below -10% regardless of intra-horizon drawdowns is 16%, thus making the difference less striking but still significant. Note also that the difference accentuates for higher T .

4. Model uncertainty

In a practical scenario, while an asset manager makes an assessment of the assets' expected returns, and constructs an allocation or portfolio to maximize the portfolio expected return subject to certain constraints, the manager acknowledges the possibility that the assessments about the expected returns may be wrong. We therefore investigate an extension of the standard model incorporating this facet, which we find has important consequences for the risk measure that we have developed.

Parameter uncertainty is of course only one of the possible extensions of the model (3.3) which has implications for intra-horizon risk. Other directions which have been investigated in the literature include non-normal return distributions, event risk, time-varying volatility, autocorrelation in returns. These are all extensions with implications for the tail risk. Our choice to investigate uncertainty about the mean return is not to discount the relevance of other factors, but it is something which we find has some important consequences.

Keeping the assumption that the portfolio is governed by (3.3), we assume that the drift term μ is a draw from a distribution of drift terms which is normal according to

$$\mu \sim \mathcal{N}(\bar{\mu}, \nu^2). \quad (4.1)$$

Here $\bar{\mu}$ represents the manager's expectation of the portfolio return, while the presence of other possible return drivers contribute with a standard deviation ν about the mean.

The model is in fact well-established, dating back to the Black and Litterman (1992) framework for portfolio allocation. We believe however that its implications for intra-horizon risk has not been well investigated. The literature on asset returns and parameter uncertainty also extends well beyond Black and Litterman, and it appeared a subject of investigation already in Zellner and Chetty (1965). Following a large amount of related research Avramov and Zhou (2010) provided a literature overview of parameter uncertainty in the context of portfolio analysis.

Of particular interest to the asset allocation problem is Pastor and Stambaugh's (2009) investigation of the statistical properties of stock market returns over various horizons from the view of an investor having imperfect knowledge of the mean return. Their paper reverses some previous findings that stock market volatility tends to decrease with the investment horizon. Instead they find that uncertainty about the mean return causes the return distribution to broaden faster than the square root of time over longer investment horizons. The mechanism driving their finding is the same as in the model presented here, even though our model is more reduced.

An important implication of the above setup is that the standard deviation about the log-price at a future date t is given by

$$\sqrt{\sigma^2 t + \nu^2 t^2}.$$

This shows that while the risk contribution of regular price volatility σ grows with the square root of time, the impact of parameter uncertainty ν is linear in time. Thus, regular price volatility remains the primary contributor to risk in the short run, but for longer horizons, parameter uncertainty takes over and dominates.

The difference in time dependency points to an interesting difference when comparing with one period models, since this easily misses a relevant aspect of parameter uncertainty. If the investment horizon is fixed at say one year, then in the above model the expected change in the log-stock price is normal with mean $\bar{\mu}$ and variance $\sigma^2 + \nu^2$. This implies that when it comes to the assessment of risk it matters little whether the uncertainty is the result of regular price volatility or parameter uncertainty. However, in a multi-period setting the difference is important since regular price risk propagates with the square root of time while risk related to parameter uncertainty grows linearly with time.

The observation in turn has implications for optimal allocation. In the Black-Litterman model the portfolio optimization problem can be reformulated in a standard Markowitz form, only that the covariance matrix is replaced with one that is the sum of two matrices, one coming from regular price volatility and one coming from uncertainty about the mean. However, when risk is also a function of intra-horizon drawdowns this approach is no longer feasible. It matters where risk comes from.

Theorem 4.1 gives an analytical expression for the risk measure taking uncertainty about the mean return into account, thus providing a generalization of Theorem 3.1.

Theorem 4.1

Given assumptions as in Theorem 3.1, assume in addition that the mean return μ is a random variable as specified by (4.1). Then for any x, y with $y \leq 0$ and $x \geq y$ it holds that

$$\psi(x, y) = \mathcal{N}(d_1) + \exp\left(\frac{2y}{\sigma^2}(\bar{\mu}_{\log} + \kappa^2 y)\right) \mathcal{N}(d_2),$$

where

$$d_1 = \frac{x - \bar{\mu}_{\log} T}{\sqrt{\sigma^2 T}}, \quad d_2 = \frac{2y}{\sqrt{\sigma^2 T}}(1 + \kappa^2 T) - d_1, \quad \bar{\mu}_{\log} = \bar{\mu} - \sigma^2/2,$$

\mathcal{N} denotes the standard normal distribution function and $\kappa^2 = \nu^2/\sigma^2$. □

A proof is presented in Appendix A.

The dependency on parameter uncertainty is illustrated in Figures 4.1a and b where the volatility, mean return and lower threshold have all been fixed.

Figure 4.1a gives the intra-horizon risk ψ_{IH} as a function of the investment horizon, and Figure 4.1b gives the total risk ψ . The lines correspond to different uncertainty about the mean. The view is altered in Figures 4.2a and b, where ψ_{IH} and ψ are given as functions of the uncertainty about the mean while keeping the investment horizon fixed.

It is evident that apart from an increase in time horizon, with a greater level of uncertainty in the mean return there is also an increase in intra-horizon and total risk. For example, at an uncertainty level of 10%, the intra-horizon risk more than triples, as the horizon goes from 1 to 3 years. Similarly, for a horizon of 2 years the intra-horizon risk doubles, as the uncertainty level goes from 2% to 9%.

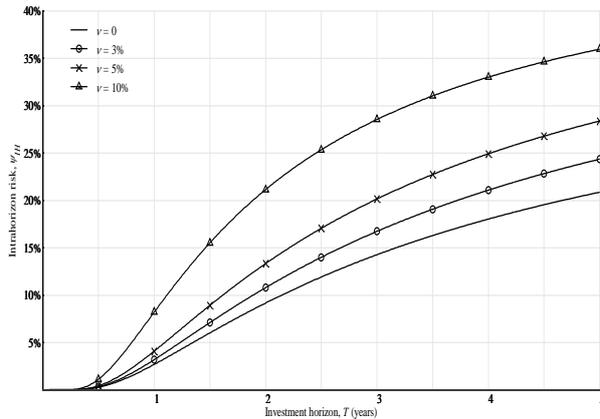


Figure 4.1a: Intra-horizon risk $\psi_{IH}(y)$,

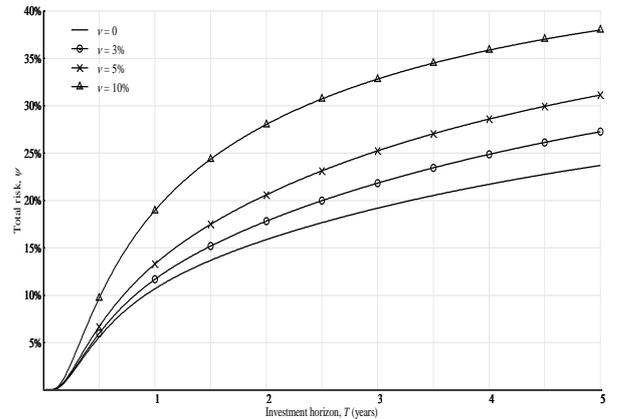


Figure 4.1b: Total risk $\psi(x, y)$

with $x = -10\%$ and $y = -20\%$ for an investment process with volatility $\sigma = 10\%$ and expected mean return $\bar{\mu} = 3\%$. The different lines correspond to different uncertainty about μ as indicated.

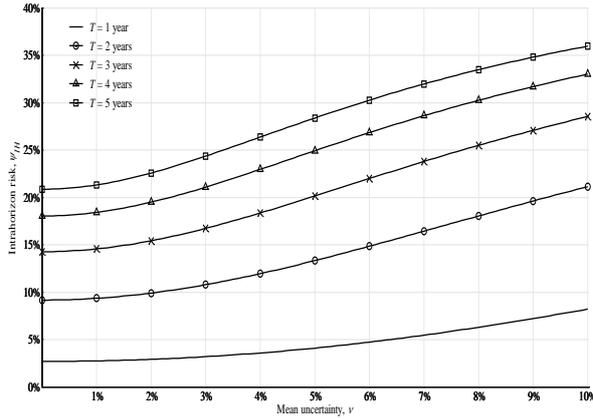


Figure 4.2a: Intra-horizon risk $\psi_{IH}(y)$

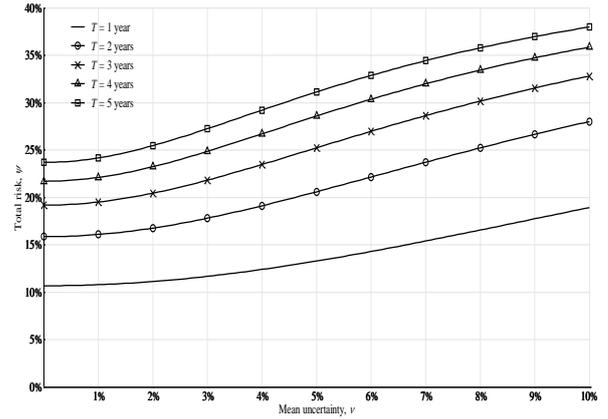


Figure 4.2b: Total risk $\psi(x, y)$

with $x = -10\%$ and $y = -20\%$ for an investment process with volatility $\sigma = 10\%$ and expected mean return $\bar{\mu} = 3\%$. The different lines correspond to different time horizons as indicated.

In Figures 4.3a and b the risk levels are kept fixed and the intra-horizon threshold y is given as a function of the investment horizon. For example, Figure 4.3a shows that at the three year horizon an intra-horizon risk threshold of -15% is breached with a 30% probability. In Figure 4.3b note that the line corresponding to $\psi = 10\%$ approaches $-\infty$ at an investment horizon approximately equal to nine months. This is where the probability of breaching the end-of-horizon threshold x exceeds 10%.

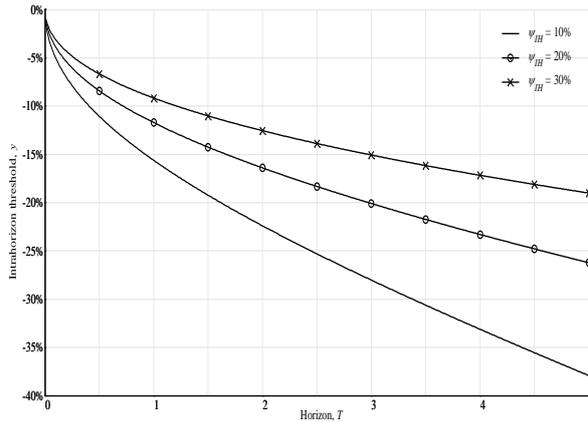


Figure 4.3a: Iso-intra-horizon risk ψ_{IH} lines, for various intra-horizon threshold y and horizon $(0, T)$.

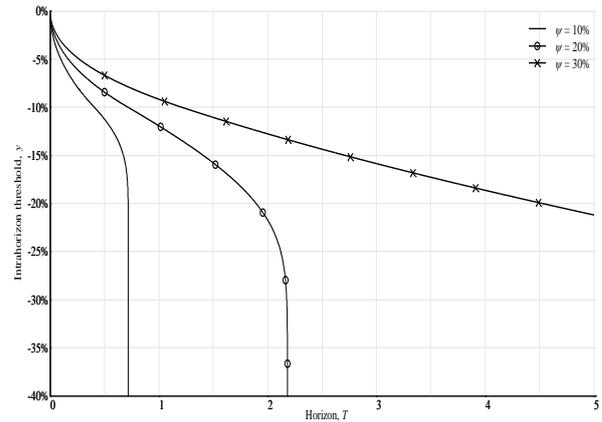


Figure 4.3b: Iso-total risk ψ lines for various intra-horizon threshold y and horizon $(0, T)$.

Settings: End-of-horizon threshold $x = -10\%$, volatility $\sigma = 10\%$, expected return $\mu = 3\%$, uncertainty about the mean $v = 5\%$.

We note that if an asset owner specifies that he does not want more than a 10% probability of breaching a maximum intra-horizon drawdown threshold of say 15%, then he cannot have an investment horizon of more than 11 months. Having an investment horizon longer than this would mean that he has a higher risk of breaching his intra-horizon risk level. This therefore means that the mere definition of an intra-horizon risk tolerance by an asset owner automatically implies a maximum investment horizon for his investment process. This facet is rarely seen either in asset allocation or portfolio management processes.

5. Stop-losses

When the risk is defined in terms of end of horizon and intra-horizon performance the question arises as to how risk exposure should optimally be managed over time. By applying a dynamic rebalancing rule it is often possible to reduce the probability of shortfall, as studied in e.g. Grossman and Zhou (1993), Carpenter (2000) or Gupta and Skallsjo (2007). Common for these approaches is that portfolio risk exposure is gradually reduced as performance approaches the lower threshold. In practice a common method is to apply a stop loss, which simply cuts risk exposure before the lower bound is breached.

When assets' mean returns are not known with certainty stop losses have the additional benefit that they can be used as a screening device for positive mean returns. If a stop loss is breached, it is more likely to have been generated by a process with a low mean return. By consistently exiting these investments – and replacing them with other with similar initial evaluation – the portfolio can be managed to maintain a higher mean return on average.

To study this mechanism we propose in this section a simplified model where a portfolio manager is invested in only one stock at the time. Each stock purchase is subjected to a maximum holding period and a stop-loss level, and together these two parameters determine how the portfolio is managed over time. The maximum holding period is primarily determined by the investment manager's style – i.e. the method which the portfolio manager uses to estimate expected returns and it is therefore less at the manager's control. The stop loss on the other hand can be set at the manager's discretion to optimize performance over time.

It should be noted that a stop loss need not be the optimal screening mechanism in the given context, and it might for instance be that a time-varying boundary actually proves more efficient for filtering out high returns. However, this is a topic abstracted from in the current paper. A stop loss set as a maximally accepted price fall is something which is frequently applied in practice and we stay with this convention.

We assume that each stock that the manager selects to include in the portfolio follows a process as in (3.3) where – from the view of the investment manager – the mean return is uncertain as in (4.1). When a stock reaches either its maximum holding period or its prescribed stop loss it is immediately replaced with a new stock with similar volatility and with its mean return picked from the same distribution for the mean, independent of previous stocks. Replacing the stock involves selling the old and buying the new, each of which is assumed to incur a relative transaction cost of $1 - \exp(-\delta)$ for a $\delta \geq 0$. For the new stock the investment horizon and stop loss are reset similarly to what was initially prescribed for the first stock. The process is repeated indefinitely. An illustration is given in Figure 5.1.

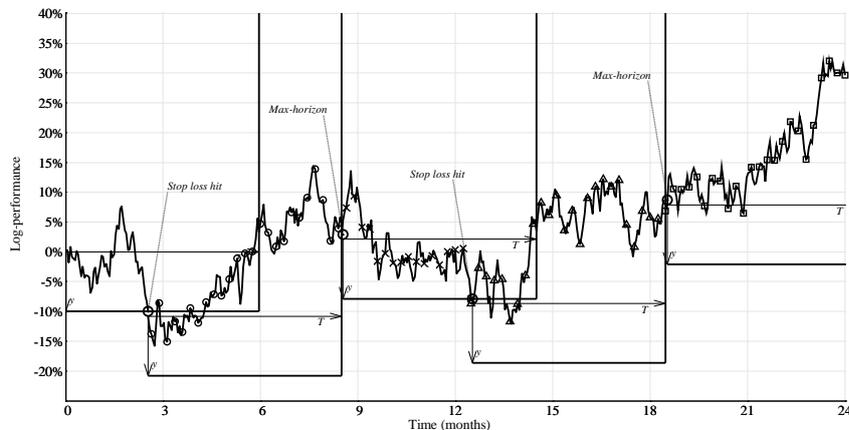


Fig 5.1: Model illustration. Settings: stop loss = -20% , maximum holding period $T = 6$ months, stock volatility $\sigma = 30\%$, average mean $\bar{\mu} = 3\%$, mean uncertainty $\nu = 5\%$ and transaction cost $\delta = 40$ basis points.

In this manner the portfolio evolves similarly to the standard model (3.3) except that every time a boundary is breached the process makes a jump downward as a result of the transaction cost after which its mean return changes. An analytical representation of the distribution function for the value of the portfolio at a given point in time does not appear tractable, and for the analysis we instead resort to a statistical moment $M(\mathcal{P})$ defined as

$$M(\mathcal{P}) = \rho^2 E \left[\int_0^\infty e^{-\rho t} \log P_t dt \right],$$

where we use \mathcal{P} and P_t to refer to the portfolio. We similarly use \mathcal{S} and S_t to refer to the stock first purchased into the portfolio and define the moment $M(\mathcal{S})$ accordingly.

For the derivation of the portfolio moment we initially assume that the mean return is known with certainty, i.e. that $\nu = 0$. The simplifying assumption also allows us to derive a representation of the portfolio moments from which it is possible to calculate moments for the more general case $\nu > 0$, at least numerically. The main result is given in Theorem 5.1.

Theorem 5.1

Suppose that $\mu = \bar{\mu}$ and $\nu = 0$. Then $M(\mathcal{P})$ satisfies the relation

$$(1 - \mathcal{R})(M(\mathcal{P}) - M(\mathcal{S})) = -2\delta\rho\mathcal{R},$$

where

$$\mathcal{R} = E \left[\int_0^{\tau \wedge T} e^{-\rho t} dt \right],$$

and $\tau = \inf\{t \geq 0: S_t = \exp(y)\}$. The above applies the conventions $\tau \wedge T = \min(\tau, T)$ and $\inf\{\emptyset\} = \infty$.

Applying the theorem requires calculating $M(\mathcal{S})$ and \mathcal{R} , something which is deferred to Lemma 5.1. Although the exposition is notation heavy, calculations involve at most the normal distribution function, which allows for a straightforward implementation.

Lemma 5.1

Suppose that $\mu = \bar{\mu}$ and $\nu = 0$.

- (i) The moment for the stock is given by

$$M(\mathcal{S}) = \mu_{\log}.$$

- (ii) The quantity \mathcal{R} can be written as

$$\mathcal{R} = \mathcal{Y} + e^{-\rho T} \mathcal{X},$$

where

$$\mathcal{Y} = E[e^{-\rho\tau}; \tau \leq T], \quad \mathcal{X} = \mathbb{P}(T > \tau).$$

- (iii) The quantities \mathcal{Y} and \mathcal{X} are given by

$$\mathcal{Y} = \exp(ab) \left(\exp(-ac) \mathcal{N}(a - c) + \exp(ac) \mathcal{N}(a + c) \right),$$

$$\mathcal{X} = \mathcal{N}(b - a) - \exp(2ab) \mathcal{N}(b + a),$$

where

$$a = \frac{y}{\sqrt{\sigma^2 T}}, \quad b = \frac{\mu_{\log} T}{\sqrt{\sigma^2 T}}, \quad c = \frac{\mu_{\log, \rho} T}{\sqrt{\sigma^2 T}},$$

$$\mu_{\log} = \mu - \frac{1}{2} \sigma^2, \quad \mu_{\log, \rho} = \sqrt{\mu_{\log}^2 + 2\sigma^2 \rho}.$$

In Theorem 5.1, for the special case that $\delta = 0$ it is readily verified that setting $M(\mathcal{P}) = M(\mathcal{S})$ satisfies the relation. This affirms the intuitive reasoning that when the mean return is known with certainty stop losses are inconsequential and do not improve portfolio performance. Moreover, if transaction costs are positive and \mathcal{R} is strictly between zero and one, then it must hold that $M(\mathcal{P}) < M(\mathcal{S})$, which says that performance deteriorates as a result of transaction costs.

The idea of Theorem 5.1 is that when the mean return is uncertain, all the relations hold in expectation, where the expectation is taken over the possible values of μ . We then have

$$E_{\mu}[(1 - \mathcal{R})(M(\mathcal{P}) - M(\mathcal{S}))] = -2\delta\rho E_{\mu}[\mathcal{R}],$$

where subscript ' μ ' denotes that the expectation is taken over the possible values of μ . This gives the portfolio moment as

$$M(\mathcal{P}) = \frac{E_{\mu}[(1 - \mathcal{R})M(\mathcal{S})] - 2\delta\rho E_{\mu}[\mathcal{R}]}{E_{\mu}[1 - \mathcal{R}]}.$$

Calculating the expectations involves some non-trivial integrals and in the application we apply a numerical integration routine, described in Appendix B.

A numerical example is presented in Figure 5.2, which gives the expected log-return as a function of the stop loss level. The different lines correspond to different transaction costs. It is clear from the figure that when the stop loss is set too tight the expected performance deteriorates rapidly. However, appropriately managed stop losses always improve expected performance. For higher transaction costs, the optimal stop loss is located further away from zero but it is still beneficial in terms of expected performance.

Figure 5.3 gives the expected performance as a function of the uncertainty about the mean. As uncertainty about the mean increases the stop loss becomes more efficient for screening processes with a higher mean return, which explains in the positive relationship seen in the figure. In Figure 5.4 expected performance is given as a function of the investment horizon. A longer investment horizon implies less frequent rebalancing and consequently less transaction costs, again resulting in a positive relationship.

Lastly, Figure 5.5 gives the optimal stop loss for different levels of uncertainty about the mean. Note the accelerating improvement in expected performance as uncertainty about the mean goes up. When uncertainty about the mean is higher it is better to close out strategies earlier, i.e. apply a tighter stop loss.

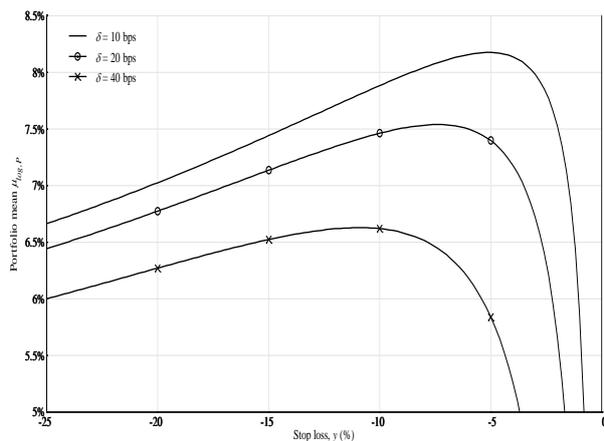


Fig 5.2: Expected log-performance $M(\mathcal{P})$ as function of the stop loss level y for different transaction costs. Avg mean $\bar{\mu} = 10\%$, volatility $\sigma = 30\%$, mean uncertainty $v = 15\%$, investment horizon $T = 1$ year, discount factor $\rho = 1/3$.

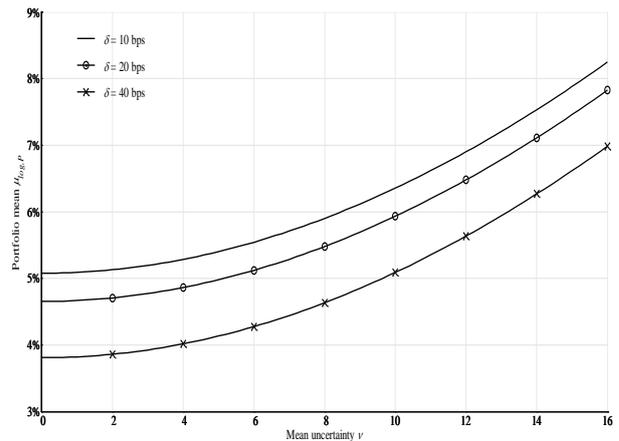


Fig 5.3: Expected log-performance $M(\mathcal{P})$ as function of mean uncertainty for different transaction costs. Avg mean $\bar{\mu} = 10\%$, volatility $\sigma = 30\%$, investment horizon $T = 1$ year, discount factor $\rho = 1/3$, stop loss $y = -10\%$.

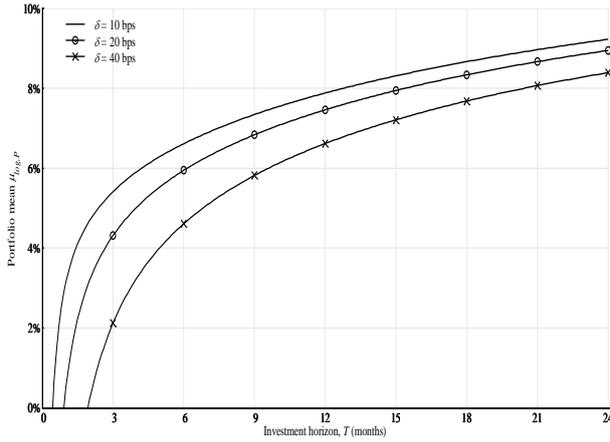


Fig 5.4: Expected log-performance $M(\mathcal{P})$ as function of investment horizon for different transaction costs. Avg mean $\bar{\mu} = 10\%$, volatility $\sigma = 30\%$, mean uncertainty $\nu = 15\%$, discount factor $\rho = 1/3$, stop loss $y = -10\%$.

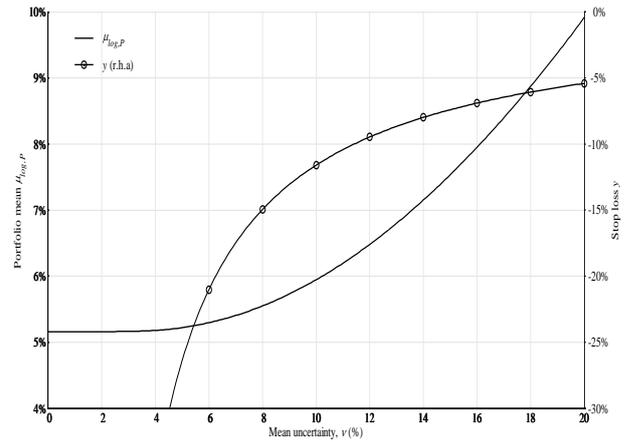


Fig 5.5: Optimal stop loss and associated expected log-performance $M(\mathcal{P})$ as functions of mean uncertainty. Avg mean $\bar{\mu} = 10\%$, volatility $\sigma = 30\%$, investment horizon $T = 1$ year, discount factor $\rho = 1/3$, transaction cost $\delta = 20$ bp.

We close this section with an example how stop losses might be applied in a portfolio management setting. Table 5.1 lists six hypothetical trading strategies corresponding to a long, medium and short investment horizon, and with high and low uncertainties about the mean return.² The optimal stop loss corresponding to different transaction costs is given to the right in the table. For the long term investment strategies, note how higher mean uncertainty implies a tighter stop loss, as was also illustrated in Figure 5.5. Comparing the long and medium term uncertain strategies we also see that the long term strategy applies a tighter stop loss. This results as the potential benefit of throwing out a bad strategy is more limited for the medium term strategy, since the position will be exited sooner.

	T	$\bar{\mu}$	ν	Transaction cost		
				10bps y^*	20bps y^*	40bps y^*
Long term certain	18 months	20%	15%	-3.9%	-5.7%	-8.3%
Long term uncertain	18 months	20%	30%	-1.9%	-2.7%	-3.9%
Medium term certain	6 months	10%	10%	-9.8%	-14.7%	-23.0%
Medium term uncertain	6 months	20%	30%	-2.7%	-3.9%	-5.7%
Short term certain	1 month	5%	3%	none	none	none
Short term uncertain	1 month	5%	30%	-5.4%	-8.3%	-13.4%

Table 5.1 Optimal stop losses for six investment strategies as characterized by the triple $(T, \bar{\mu}, \nu)$ under different transaction costs. Optimal stop loss y^* . Settings: Volatility $\sigma = 30\%$, discount factor $\delta = 1/3$.

² An extended version of the Table 5.1 including the expected mean returns is given in Appendix C.

6. Conclusions

In this paper we have reexamined the risk inherent in some typical settings within asset management. Focus has been on the asset allocation decision at one end of the investment management process and on individual stock selection strategies at the other. In both cases an investment mandate is issued together with an end of horizon return requirement and an intra-horizon drawdown tolerance. Naturally, the investment manager is expected to maximize expected performance subject to these requirements, but the risk inherent in the setting is that either of the constraints is breached.

Given this background we have proposed a new risk measure defined as the probability that an investment strategy falls short of either its end of horizon performance target or its maximally accepted drawdown intra-horizon. In the paper this is exemplified with a standard model for asset prices where it was found that as the investment horizon lengthens the intra-horizon threshold gains in importance, an aspect we argue is often ignored.

Extending the analysis we have investigated how uncertainty about assets' expected returns affect the risk over time. We found that especially over longer horizon uncertainty about the mean can have a significant impact on the distribution of returns.

The paper's last section examined how stop losses can be beneficial for managing a portfolio when stocks' mean returns are not known with certainty. By applying a stop loss the manager can to a certain extent screen stocks with a higher expected return, which on average improves portfolio performance.

Important tasks for future research involve how to construct a portfolio subject to a risk constraint as defined as in the current paper. Of particular interest is the case where the assets' expected returns are not known with certainty.

References

- Avramov, Doron and Guofu Zhou, 2010, *Bayesian Portfolio Analysis*, Annual Review of Financial Economics 2010:2, pg. 25-47
- Black, Fischer and Robert Litterman, 1992, *Global Portfolio Optimization*, Financial Analysts Journal, Sep/Oct 1992, 48:5, pg 28-43
- Carpenter, Jennifer, 2000, *Does option compensation increase managerial appetite?*, Journal of Finance 55, 2311-2332
- Gupta, Pranay and Jan Straatman, 2006, *Skill Based Investment Management*, Journal of Investment Management, Vol 4, No. 1, 1-18.
- Gupta, Pranay, Jan Straatman, Sven Skallsjo and Alice Krol, 2005, *Risk Management in a Multi-Strategy Framework* Available at <http://ssrn.com/abstract=835244>
- Gupta, Pranay and Sven Skallsjo, 2013, *Rethinking the Asset Allocation Approach for Plan Sponsors*, Working paper, Available at <http://ssrn.com/abstract=2210695>.
- Gupta, Pranay and Sven Skallsjo, 2007, *Impact of Manager Compensation on Portfolio Risk and Return*, Working paper, Available at <http://ssrn.com/abstract=1021175>.
- Grossman, Sanford and Zhongquan Zhou, 1993, *Optimal Investment Strategies for Controlling Drawdowns*, Mathematical Finance, Vol 3:3, 241-276
- IImanen, Antti, 2012, *Do Financial Markets Reward Buying or Selling Insurance and Lottery Tickets?*, Financial Analysts Journal, Vol 68:5, 26-36
- Kaminski, Kathryn M and Andrew W. Lo, 2007, *When Do Stop-Loss Rules Stop Losses?*, EFA 2007 Ljubljana Meetings Paper, available at SSRN: <http://ssrn.com/abstract=968338>
- Magdon-Ismail, Malik and Amir Atiya, 2004, *Maximum Drawdown*, Risk October 2004
- Magdon-Ismail, Malik, Amir F Atiya., Amrit Pratap, and Yaser S. Abu-Mostafa, 2004, *On the Maximum Drawdown of a Brownian Motion*, Journal of Applied Probability 41, pg. 147-161
- Pástor, Ľuboš and Robert F. Stambaugh, 2012, *Are Stocks Really Less Volatile in the Long Run?*, The Journal of Finance, Vol 67:2, pg. 431-478, April 2012
- Pástor, Ľuboš and Pietro Veronesi, 2003, *Stock Valuation and Learning about Profitability*, The Journal of Finance, Vol 58:5, pg. 1749-1789, October 2003
- Wise, Mark B., Yonathan Schwarzkopf and Vineer Bhansali, 2011, *Fat Tails and Stop-Losses in Portable Alpha*, Journal of Investment Management, Third Quarter 2011
- Zellner, Arnold and V. Karuppan Chetty, V. K., 1965, *Prediction and decision problems in regression models from the Bayesian point of view*, Journal of the American Statistical Association, 60, pg. 608-616

Appendix A. Proofs

Proof of Theorem 3.1

The joint density for x and y is well known, see e.g. Karatzas and Shreve (1987) pages 196-197. It holds that

$$p(x, y) = \frac{2}{\sigma^2 t} \frac{2y - x}{\sqrt{\sigma^2 t}} \exp\left(\frac{\mu x - \mu^2 t/2}{\sigma^2}\right) \varphi\left(\frac{2y - x}{\sqrt{\sigma^2 t}}\right),$$

where φ denotes the standard normal density. Integrating over x and y gives the result in the main text.

Proof of Theorem 4.1

By the law of iterated expectations

$$\psi(x, y) = E[\mathbb{P}(x_t \leq x \text{ or } y_t \leq y | \mu)],$$

or

$$\psi(x, y) = \int_{-\infty}^{\infty} \psi(x, y; \mu) \frac{1}{v} \varphi\left(\frac{\mu - \bar{\mu}}{v}\right) d\mu,$$

where ' $'$; μ' indicates that the measure is applied with μ as drift term. The integral is split into two terms according to $\psi(x, y) = \mathcal{J}_1 + \mathcal{J}_2$ with

$$\begin{aligned} \mathcal{J}_1 &= \int_{-\infty}^{\infty} \mathcal{N}\left(\frac{x - \mu t}{\sqrt{\sigma^2 t}}\right) \frac{1}{v} \varphi\left(\frac{\mu - \bar{\mu}}{v}\right) d\mu, \\ \mathcal{J}_2 &= \int_{-\infty}^{\infty} \exp\left(\frac{2y\mu}{\sigma^2}\right) \mathcal{N}\left(\frac{2y - x + \mu t}{\sqrt{\sigma^2 t}}\right) \frac{1}{v} \varphi\left(\frac{\mu - \bar{\mu}}{v}\right) d\mu. \end{aligned}$$

Both \mathcal{J}_1 and \mathcal{J}_2 can be rewritten as integrals on the generic form

$$\mathcal{J} = \int_{-\infty}^{\infty} \mathcal{N}(a + bu) \varphi(u) du = \int_{-\infty}^{\infty} \int_{-\infty}^{a+bu} \varphi(v) \varphi(u) dv du.$$

Here the change of variables

$$\begin{pmatrix} w \\ z \end{pmatrix} = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} u \\ v \end{pmatrix}$$

applied with $\theta = \arctan(b)$ yields

$$\mathcal{J} = \int_{-\infty}^{\infty} \int_{-\infty}^{a \cos \theta} \varphi(w) \varphi(z) dw dz = \int_{-\infty}^{a \cos \theta} \varphi(w) dw = \mathcal{N}(a \cos \theta) = \mathcal{N}\left(\frac{a}{\sqrt{1 + b^2}}\right).$$

Applying this to \mathcal{J}_1 and \mathcal{J}_2 gives the result.

Proof of Theorem 5.1

The integral in the expectation is split into two parts, corresponding to the time before and after the selling of the first stock. This gives

$$M(\mathcal{P}) = m_{(0, \tau \wedge T)} + m_{(\tau \wedge T, \infty)}, \quad (\text{A1})$$

where we define

$$m_{(a,b)} = \rho^2 E \left[\int_a^b e^{-\rho t} \log P_t dt \right].$$

In the former expectation, factoring out $e^{-\rho(\tau \wedge T)}$ we can write

$$\begin{aligned} m_{(\tau \wedge T, \infty)} &= \rho^2 E \left[e^{-\rho(\tau \wedge T)} \int_{\tau \wedge T}^{\infty} e^{-\rho(t-\tau \wedge T)} \log P_t dt \right] \\ &= \rho^2 E \left[e^{-\rho(\tau \wedge T)} \int_{\tau \wedge T}^{\infty} e^{-\rho(t-\tau \wedge T)} \log \left(\frac{P_t}{e^{-2\delta P_{\tau \wedge T}}} \right) dt \right] \\ &\quad + \rho^2 E \left[e^{-\rho(\tau \wedge T)} \int_{\tau \wedge T}^{\infty} e^{-\rho(t-\tau \wedge T)} \log(e^{-2\delta P_{\tau \wedge T}}) dt \right]. \end{aligned}$$

We here use $P_{\tau \wedge T}$ to denote the value of the portfolio immediately before the old stock is sold and the new one is acquired. Using iterated expectations conditioning on information up to $\tau \wedge T$ shows that

$$\begin{aligned} m_{(\tau \wedge T, \infty)} &= E[e^{-\rho(\tau \wedge T)}] M(\mathcal{P}) + \rho^2 E \left[e^{-\rho(\tau \wedge T)} \int_{\tau \wedge T}^{\infty} e^{-\rho(t-\tau \wedge T)} (\log P_{\tau \wedge T} - 2\delta) dt \right] \\ &= \mathcal{R} M(\mathcal{P}) + \rho E[e^{-\rho(\tau \wedge T)} \log S_{\tau \wedge T}] - 2\delta \rho \mathcal{R}, \end{aligned}$$

where $\mathcal{R} = E[e^{-\rho(\tau \wedge T)}]$. The latter equality makes use of the fact that for $t < \tau \wedge T$ it holds that $P_t = S_t$. Turning to the term $m_{(0, \tau \wedge T)}$ in (A1) we again make use of this fact, which allows us to write

$$m_{(0, \tau \wedge T)} = \rho^2 E \left[\int_0^{\infty} e^{-\rho t} \log S_t dt \right] - \rho^2 E \left[\int_{\tau \wedge T}^{\infty} e^{-\rho t} \log S_t dt \right].$$

For the latter term on the right hand side, a similar procedure as for $m_{(\tau \wedge T, \infty)}$ shows that

$$\rho^2 E \left[\int_{\tau \wedge T}^{\infty} e^{-\rho t} \log S_t dt \right] = \mathcal{R} M(\mathcal{S}) + \rho E[e^{-\rho(\tau \wedge T)} \log S_{\tau \wedge T}].$$

Adding the acquired representations for $m_{(0, \tau \wedge T)}$ and $m_{(\tau \wedge T, \infty)}$ now yields

$$M(\mathcal{P}) = (1 - \mathcal{R}) M(\mathcal{S}) + \mathcal{R} M(\mathcal{P}) - 2\delta \rho \mathcal{R}.$$

Rearranging recovers the representation given in the main text.

Proof of Lemma 5.1

- (i) The moment for the stock is given by

$$M(\mathcal{S}) = \rho^2 E \left[\int_0^{\infty} e^{-\rho t} \log S_t dt \right] = \rho^2 \int_0^{\infty} e^{-\rho t} \mu_{\log} t dt = \mu_{\log}.$$

- (ii) The quantity \mathcal{R} can be written as

$$\mathcal{R} = E[e^{-\rho(\tau \wedge T)}] = E[e^{-\rho\tau}; \tau \leq T] + e^{-\rho T} E[1; \tau > T],$$

as asserted in the main text.

- (iii) For \mathcal{X} the result is a direct application of Theorem 3.1. The calculation of \mathcal{Y} is presented below.

Calculation of \mathcal{Y}

For the calculation of $\mathcal{Y} = E[e^{-\rho\tau}; \tau \leq T]$ we apply a change of a measure. Let \mathcal{F}_T be the filtration generated by W_t up until time T and consider the measure change

$$L_t = \exp(\lambda W_t - \lambda^2 t/2), \quad \frac{d\mathbb{Q}_\lambda}{d\mathbb{P}} = L_T$$

on \mathcal{F}_T , where λ is a constant to be determined. Writing \mathcal{Y} as an expectation under \mathbb{Q}_λ we then have

$$\mathcal{Y} = E^{\mathbb{Q}_\lambda}[\exp(-\rho\tau) I\{\tau \leq T\} L_T^{-1}].$$

Because L_t^{-1} is a \mathbb{Q}_λ -martingale we can apply iterated expectations to obtain

$$\mathcal{Y} = E^{\mathbb{Q}_\lambda} \left[E^{\mathbb{Q}_\lambda}[\exp(-\rho\tau) I\{\tau \leq T\} L_T^{-1} | \mathcal{F}_\tau] \right] = E^{\mathbb{Q}_\lambda}[\exp(-\rho\tau - \lambda W_\tau + \lambda^2 \tau/2) I\{\tau \leq T\}].$$

Further, the fact that at the time of intercept $X_t = \mu t + \sigma W_t = y$ allows us to write

$$\mathcal{Y} = \exp\left(-\frac{y\lambda}{\sigma}\right) E^{\mathbb{Q}} \left[\exp\left\{\left(-\rho + \lambda \frac{\mu_{\log}}{\sigma} + \frac{\lambda^2}{2}\right)\tau\right\} I\{\tau \leq T\} \right].$$

Selecting λ as $(-\mu_{\log} + \mu_{\log,\rho})/\sigma$ with

$$\mu_{\log,\rho} = \sqrt{\mu_{\log}^2 + 2\sigma^2\rho}$$

the quadratic expression vanishes and

$$\mathcal{Y} = \exp\left(-\frac{y\lambda}{\sigma}\right) \mathbb{Q}_\lambda(\tau \leq T),$$

where the breaching probability $\mathbb{Q}_\lambda(\tau \leq T)$ is given by Theorem 3.1. Note that under \mathbb{Q}_λ , the drift of X_t is $\mu_{\log} + \sigma\lambda = \mu_{\log,\rho}$. Inserting results gives \mathcal{Y} as in the main text.

Appendix B. Numerical integration for the expected values

The expected values are on the form

$$E_\mu[f(\mu)] = \int_{-\infty}^{\infty} f(\mu) \frac{1}{\sqrt{2\pi v^2}} \exp\left(-\frac{1}{2}\left(\frac{\mu - \bar{\mu}}{v}\right)^2\right) d\mu,$$

which via a change of variables to $z = (\mu - \bar{\mu})/\sqrt{2v^2}$ can be written as

$$E_\mu[f(\mu)] = \frac{1}{\sqrt{\pi}} \int_{-\infty}^{\infty} f(\bar{\mu} + \sqrt{2v^2}z) \exp(-z^2) dz.$$

Using the exponential term as weight function leads to Gauss-Hermite quadrature, which we apply with 10 abscissas. Numerical evaluation and comparison with other integration methods indicate good convergence.

Appendix C. Optimal stop-losses for six investment strategies

	T	$\bar{\mu}$	ν	Transaction cost								
				10bps y^*	20bps y^*	40bps y^*	10bps μ_{\log}^*	20bps μ_{\log}^*	40bps μ_{\log}^*	10bps μ_{\log}	20bps μ_{\log}	40bps μ_{\log}
LTC	18 months	20%	15%	-3.9%	-5.7%	-8.3%	18.8%	18.3%	17.5%	15.4%	15.3%	15.1%
LTU	18 months	20%	30%	-1.9%	-2.7%	-3.9%	29.4%	28.4%	27.0%	15.6%	15.5%	15.3%
MTC	6 months	10%	10%	-9.8%	-14.7%	-23.0%	5.6%	5.1%	4.1%	5.1%	4.8%	4.0%
MTU	6 months	20%	30%	-2.7%	-3.9%	-5.7%	23.5%	22.0%	19.9%	15.1%	14.8%	14.0%
STC	1 month	5%	3%	<i>none</i>	<i>none</i>	<i>none</i>	-1.9%	-4.2%	-9.0%	-1.9%	-4.2%	-9.0%
STU	1 month	5%	30%	-5.4%	-8.3%	-13.4%	-0.5%	-3.6%	-8.8%	-1.9%	-4.2%	-9.0%

Table C1 Extended version of Table 5.1 including maximal expected log-return μ_{\log}^* and expected log-return without stop loss μ_{\log} .

Appendix D. Notation and variables

x, X_t	Logarithm of the stock price S_t at time t .
y, Y_t	Running minimum of X_t up until time t : $\inf\{X_s; 0 < s \leq t\}$. Also the stop loss applied by the portfolio manager.
δ	Logarithm of one-way transaction cost.
φ	The probability density function for a standard normal variable.
κ^2	Auxiliary variable defined as ν^2/σ^2 .
μ	Drift term for the stock price.
$\bar{\mu}$	Expected value of μ when μ is a random variable.
μ_{\log}	Auxiliary variable defined as $\mu - \sigma^2/2$.
$\mu_{\log, \rho}$	Auxiliary variable defined as $\sqrt{\mu_{\log}^2 + 2\sigma^2\rho}$.
ν	Standard deviation of μ when μ is a random variable.
ρ	Time discount rate for calculation of $M(\mathcal{P})$ and $M(\mathcal{S})$.
σ	Volatility term for the logarithm of the stock price.
τ	The stopping time $\inf\{t > 0; x_t = y\}$.
$\psi(x, y)$	Risk function for end-of-horizon threshold x and intra-horizon threshold y .
$M(\mathcal{P})$	The first moment of the log-value of the portfolio.
$M(\mathcal{S})$	The first moment of the log-value of the stock first purchased into the portfolio.
\mathcal{N}	The probability distribution function for a standard normal variable.
\mathbb{P}	Probability.
\mathcal{P}, P_t	The portfolio manager's portfolio and the value of the portfolio at time t .
\mathcal{R}	The expectation $E[\exp(-\rho\tau)]$.
\mathcal{S}, S_t	The stock first purchased into the portfolio and the value of a stock at time t .
T	Risk time horizon for ψ . Also the portfolio manager's maximum holding period for a stock.
W_t	Standard Brownian motion with initial condition $W_0 = 0$.
\mathcal{X}	The expectation $E[1; \tau > T] = \mathbb{P}(\tau > T)$.
\mathcal{Y}	The expectation $E[\exp(-\rho\tau); \tau \leq T]$.