VU Research Portal

Risk allocation under shortfall constraints

Molenkamp, Jan Bertus

2003

document version

Publisher's PDF, also known as Version of record

[Link to publication in VU Research Portal](#)

citation for published version (APA)

Molenkamp, J. B. (2003). *Risk allocation under shortfall constraints*. (Research Memorandum; No. 2003-9). Faculty of Economics and Business Administration, Vrije Universiteit Amsterdam.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

ET

d

35025

2003
009

Risk allocation under shortfall constraints

Research Memorandum 2003-9

Faculteit der Economische Wetenschappen
en Bedrijfskunde (FEW&B)

vrije Universiteit

amsterdam

Risk allocation under shortfall constraints

Research Memorandum 2003-9

Jan Bertus Molenkamp

Risk allocation under shortfall constraints

Improved allocation by incorporating parameter uncertainty

Jan Bertus Molenkamp^{1,2}

Abstract

Risk budgeting interpreted as efficient portfolio allocation is often based on expected outperformance, alpha or information ratio. Once these crucial parameters have been estimated, they are being treated as fixed. In this paper we develop some sense, both theoretical and practical, on the magnitude of the uncertainty and present a method to cope with uncertainty of the expected active returns. We will use examples to demonstrate the impact of the uncertainty. The developed model is widely applicable and can be used to make an optimal risk allocation on different levels of investment strategy (asset allocation, styles, managers, etc.).

Keywords: risk budgeting, information ratio, tracking error, volatility

¹ Vrije Universiteit Amsterdam, jmolenkamp@feweb.vu.nl

² Mn Services Vermogensbeheer

Introduction

Risk budgeting is a widely used term but gives room for several interpretations. According to Chow and Kritzman [2001] the predominant definition is the efficient portfolio allocation. This can of course be analyzed at several levels. We will focus here mainly on the active risk allocation, that is, after the strategic benchmark has been set via e.g. an ALM-process and given a pre-specified maximum level of active risk.

Most of the literature presents optimization models based on expected outperformance, alpha or information ratio. Once these crucial parameters have been estimated, they are being treated as certain. In this paper we develop some sense on the magnitude of the uncertainty and present a method to cope with uncertainty of the expected active returns, by broadening the approach followed by Blitz and Hottinga [2001]. We will use their examples to demonstrate the impact of the model changes. The developed model is widely applicable and can be used to make an optimal risk allocation to the several investment processes (styles, managers, etc.).

Prior research on risk allocation

The active risk allocation problem can usually be stated as maximizing some utility function given some (risk) constraints. In Waring et al. [2000] a utility function is introduced in which the active return minus a risk aversion parameter times the variance of the active return (tracking error) is being maximized. An efficient frontier is being calculated for several levels of risks. It is assumed that the different investment processes have a correlation of zero, although the model allows for non-zero correlation. The underlying active returns per investment process are fixed parameters. The Blitz and Hottinga [2001] framework present a framework in which optimal partial tracking errors are determined given an overall target tracking error under the assumption of independent investment processes (correlation is zero) and a fixed, certain return/risk ratio. Their approach can be viewed as calculating a specific point on the efficient frontier in the model of Waring et al. [2000], expressed in partial tracking errors.

Lee and Lam [2001] take into account the possible correlation between investment processes and, more importantly, the uncertainty on the expected active returns, which they refer to as "information risk". They make an assumption on the distribution of

the hit ratio of investment processes, i.e. how often will the investment process, for which the information ratio is a given, outperform its benchmark. The problem there is how to arrive at an ex-ante estimate for the hit ratio.

Here we have chosen to extend the Blitz and Hottinga [2001] model to incorporate correlation between investment processes and uncertainty in the expected information ratio. First we will establish some sense for the magnitude of uncertainty in information ratios.

Information ratio volatility: theoretical & empirical evidence

Based on Lo [2002] who describes the statistics of the Sharpe Ratios under different return distribution assumptions we can determine the standard error for the information ratio estimator asymptotically under the assumption that returns are IID¹.

$$\sigma(i\hat{r}) = \sqrt{\frac{1 + \frac{1}{2}i\hat{r}^2}{T}}$$

Where T is the sample size and $i\hat{r}$ is the information ratio estimator. With this standard error estimate we can create confidence interval around the estimated information ratio e.g. a 95% confidence interval is

$$i\hat{r} \pm 1.96 \times \sqrt{\frac{1 + \frac{1}{2}i\hat{r}^2}{T}}$$

If e.g. T=12 and the estimated information ratio is 0.5 then the standard deviation of the information ratio estimator is 0.3 with a 95%-confidence interval of the information ratio of [-0.1;1.1].

Based on empirical research Gupta et al. [1999] show that for a median manager there is not much consistency in performance over almost all categories. This would implicate a large uncertainty with respect to the information ratio². To give some insight into the empirical size of this uncertainty, we consider some actual data.

Based on the Wilshire Compass database we have researched the (historical) information ratio of managers in several investment categories over the period Q4 1996 up and until Q3 2001 (5 year manager performance). We have measured the volatility of the information ratio in two ways: the volatility of the 12 quarter rolling information ratios and the volatility of the 4 quarter rolling information ratios (or annual rolling information ratios)³⁴. In exhibit 3 the medians for managers in several asset classes are displayed.

Exhibit 3

	Median volatility 12 quarter rolling IR	Median volatility 4 quarter rolling IR
US large cap core S&P 500	0.5	1.4
US small cap core Wilshire 5000	0.4	1.5
US real estate Wilshire REIT index	0.4	1.4

The figures obtained indicate that the volatility of the information ratio can be substantial and can deviate for the several investment processes for which we would like to optimize the risk allocation. For the shorter horizon the volatility estimates are naturally higher than for the longer horizon. The level of volatility of the information ratio within an asset class can be quite different from manager to manager. If this was not the case there would be no information in the uncertainty of information ratios. Given the 4 quarter rolling figures, the consistency to achieve an constant level of performance can be considered weak if not totally absent. This is in agreement with the observations of Gupta et al. [2001].

The observed volatility in the information ratio's is not consistent with the above derived theoretical standard deviation. E.g. to obtain a standard deviation of 0.5 for US Large Cap with T=12, we would need an information ratio estimate of above 1. This is clearly not the median information ratio for this asset category⁵. Still the major conclusion that should be drawn from the presented material in this section is that uncertainty in information ratios can be quite substantial from both a theoretical expectation point of view and the practical observations.

Model sensitivity

How does the uncertainty of the information ratio translate into model outcomes? We will look at the original Blitz-Hottinga model to develop some feeling for the impact of uncertainty in the information ratio.

The objective is to allocate the partial tracking errors in such a way that the expected return, the product of the information ratio and the tracking error, is maximal given the overall target tracking error. A partial tracking error is the contribution to the overall tracking error of one activity in isolation. The underlying assumption is that in a separate process the target tracking error (risk budget) has been set.

$$\begin{aligned} \text{Max}_{TE_1, \dots, TE_n} \quad & \sum_{i=1}^n IR_i TE_i \\ \text{s.t.} \quad & \sum_{i=1}^n TE_i^2 = TE_{\text{target}}^2 \end{aligned} \quad (\text{BH})^6$$

The information ratio IR_i of investment process i is a forecast and is treated as a given constant. The partial tracking error of investment process i is denoted by TE_i . The investment processes are assumed to be independent. The optimal solution to this problem is

$$TE_i = TE_{\text{target}} \frac{IR_i}{\sqrt{\sum_{i=1}^n IR_i^2}}$$

The optimal partial tracking error allocation is dependent on the target tracking error and on the expected information ratios of the underlying processes, which are explicit forecasts⁷. The BH model is sensitive to changes in the expected information ratio. This can be seen looking at the derivatives from the partial tracking error to the information ratio.

$$\frac{\partial TE_i}{\partial IR_i} = TE_{i \text{ arg et }} \frac{\sum_{i=1}^n IR_i^2 - IR_i^2}{\left\{ \sum_{i=1}^n IR_i^2 \right\}^{\frac{3}{2}}}$$

The derivative of TE_i towards IR_i is mainly determined by the first two factors of the optimal solution. It is easy to see that the derivative is always greater or equal to zero⁸ as should be expected. The derivative will increase with lower information ratios and decrease with higher information ratios.

$$\frac{\partial TE_j}{\partial IR_i} = -TE_{i \text{ arg et }} \frac{IR_j IR_i}{\left\{ \sum_{i=1}^n IR_i^2 \right\}^{\frac{3}{2}}}$$

The derivative of TE_j towards IR_i is negative as can be expected; an increase in the information ratio of investment process i will cost the other investment processes a loss in partial tracking error as they become relatively less attractive. If we perturb the individual expected information ratios by e.g. 0.1, the partial tracking error allocation changes dramatically. In exhibit 1 this is demonstrated by a perturbation of the information ratio in the asset allocation process.

Exhibit 1

	IR_i	Optimal TE_i BH	Perturbed IR_i	Optimal TE_i	% change in TE_i
Asset Allocation	0.2	1.0%	0.3	1.4%	+43%
Equity Country Allocation	0.2	1.0%	0.2	1.0%	-5%
Bond Country allocation	0.3	1.5%	0.3	1.4%	-5%
Currency allocation	0.0	0.0%	0.0	0.0%	-
Stock selection	0.4	2.0%	0.4	1.9%	-5%
Bond selection	0.4	2.0%	0.4	1.9%	-5%
Total	0.7	3.5%	0.7	3.5%	

An increase in the information ratio of the asset allocation process causes a 0.4 percentage point change in its partial tracking error. As a result, the other processes see a decline of circa 5% in their partial tracking errors. Of course as a result of the

0.1 perturbation the percentage change in the expected information ratio is large. If however we would observe in practice an information ratio realization that is 0.1 percentage point of f , the ex-ante estimate would still be viewed as a good estimate.

Given the impact of changes in the expected information ratio we are left with some basic questions. How much will in practice the realized information ratio deviate from the expected information ratio? If these deviations are substantial, is there a possibility to make the model outcomes more robust for the described deviations?

Shortfall constraint model

There are several possibilities to accommodate for uncertainty with respect to the projected information ratios⁹. It depends on the preferences you take into account. Let us assume that it is not desirable to have a return below a certain threshold, then adding a shortfall constraint yields the following model:

$$\begin{aligned} & \underset{TE_1, \dots, TE_n}{\text{Max}} \sum_{i=1}^n IR_i TE_i \\ & \text{s.t.} \sum_{i=1}^n \sum_{j=1, j \neq i}^n (TE_i^2 + \rho_{ij} TE_i TE_j) = TE_{\text{target}}^2 \quad (1) \\ & \sum_{i=1}^n IR_i TE_i - z_\alpha \sqrt{\sum_{i=1}^n \sum_{j=1, j \neq i}^n (\sigma_i^2 TE_i^2 + \sigma_{ij} TE_i TE_j)} \geq r_{\min} \quad (2) \end{aligned}$$

Equation (1) is stating that the partial tracking errors are to be set in such a way that the correlated combination is equal to a target level of risk. The correlation coefficient ρ_{ij} , is the correlation between the information ratio's of the investment processes i and j . Although a zero-correlation assumption could be defended by qualitative arguments (different selection processes etc.) and is being used in the BH-model, in practice correlation will deviate from 0. The TE_{target} is a constant resembling the target level of risk.

Equation (2) describes how uncertainty with regard to the expected information ratio is being translated into the model. The equation can be interpreted as a type of Value-at-Risk constraint on model or parameter risk. The $r_{\min}$ is the minimal required return.

The square root term denotes the standard deviation of the expected return. The variance of the error terms around the information ratio is denoted by σ_i^2 and the covariance of the error terms is denoted by σ_{ij} . The derivation of the inequality is straightforward from the assumption that the active returns are normally distributed¹⁰. The z_α term is the normal z-value at a (one-sided) confidence level of α . The effect of this equation is obvious. Once we are less confident in an ex-ante information ratio estimate for a specific investment process we would be inclined to spend less risk on that investment process.

The optimization problem can easily be solved numerically. In exhibit 2 we extend the example of exhibit 1. Added are the standard deviations of the information ratio. The model has been solved given a 1% shortfall probability (where $z_{1\%} = 2.32$) below a minimum return assumption of 0.

Exhibit 2

	Expected information ratio IR_i	Standard Deviation of IR_i (σ_i)	Optimal partial tracking error (BH model)	Optimal partial tracking error $\rho_{ij}=0.2, \forall i \neq j$ (no shortfall)	Optimal partial tracking error (shortfall 1%)
Asset Allocation	0.2	0.3	1.0%	0.3%	0.5%
Equity Country Allocation	0.2	0.1	1.0%	0.3%	1.2%
Bond Country allocation	0.3	0.2	1.5%	1.1%	1.6%
Currency allocation	0.0	0.1	0.0%	0.0%	0.0%
Stock selection	0.4	0.6	2.0%	1.9%	0.6%
Bond selection	0.4	0.3	2.0%	1.9%	1.7%
Total			3.5%	3.5%	3.5%

Given a slightly positive correlation the optimal partial tracking errors are all decreasing versus the BH model outcome. This follows from the simple observation that there are less diversification opportunities. The largest decrease is in the lower information ratio processes, as one still strives for a maximal expected information ratio or return. So, the shifts are exactly what should be expected given the objectives. E.g. the allocation to asset allocation decreases by 70% (from 1.0% to 0.3%). When a shortfall constraint is added the tracking error allocation shifts towards the investment processes with the more certain information ratios. E.g. the allocation to

the stock selection process with a high information ratio uncertainty decreases by more than 60% (from 1.9% to 0.6%).

The sensitivity analysis of the model can yield some interesting insights. If the probability of shortfall is increased then the optimal allocation shifts to the no shortfall solution. E.g. with a shortfall probability of 5% the optimal allocation to stock selection increases from 0.6% to 1.6%. At a 10% level the shortfall constraint is redundant.

If the tracking error budget is reduced by 50%, the risk allocations for the individual processes are also reduced by 50% leaving the relative differences in tact (linear scalability; see also next section).

The (boundary) values for the minimum return threshold, the shortfall probability and the tracking error budget for which the model is solvable are correlated. If the minimum required return increases beyond a certain level, then, to obtain a solution, the target tracking error together with the shortfall probability have also to be increased. To put it otherwise: one cannot boundlessly increase the minimum required return without having to take more risk.

A crucial assumption underlying the shortfall model is that we can discriminate with respect to uncertainty. In other words, the several investment processes have to have different levels of uncertainty with respect to their projected information ratios.

Depending on the choice of parameters the model is not always solvable. In that case a sensitivity analysis will yield a lot of insight. The shortfall constraint can also be redundant. This in fact tells you that the probability of achieving threshold return is higher than deemed necessary

Implementation issues

The optimal partial tracking errors can easily be translated into tracking error targets or Value-at-Risk targets for the individual investment processes⁷. An underlying assumption in the partial tracking error approach is that the risk of the investment processes is scalable and independent of information ratio. The ratio of the partial tracking error of an investment process to the target risk is therefore constant for each level of target risk. In Gupta et al. [1999] it is shown that the information ratio for active strategies in several asset classes is not constant under changes in the level of

active risk. In the proposed model this can be accommodated for by making the information ratio dependent on the partial tracking error i.e. describe the information ratio as a function of the partial tracking error. If these functions are convex or concave the problem is still numerically solvable.

One can correlate the historical monthly performance figures of the investment processes to get an ex-post estimate of the correlation parameters. Past correlation is not necessarily a good estimator of future correlation (e.g. different correlation in bull and bear periods and, in general, there is volatility in correlation). For an ex-ante estimate one could use a factor model applied to the actual holdings in the several investment processes (managers) to generate a correlation estimate.

The shortfall constraint model is about target risk. In practice, from a risk management point of view, there is also need for a risk budget. Blitz and Hottinga [2001] present a way of establishing a risk budget based upon the expected variation in the tracking error of the investment processes. The risk target or risk budget can be very powerful tools in terms of communication with the managers of the several investment processes. The discussion can be focussed on the usage of risk in addition to the specific bets of the underlying investment process.

If a shortfall constraint is being specified on the top level the optimization should be done on all managers/investment processes across the entire portfolio. Otherwise information will be lost and the optimality of the solution is questionable. There is no easy way to decompose the optimization problem due to the probability constraint. The model cannot explicitly be used for the decision between active and passive management. If a passive investment process is added to the problem with an expected information ratio of 0, assigning partial tracking error to this investment process will not add value to the goal function. However, the assumption that the tracking error is scalable implies that in the investment process a choice between active and passive is being made.

Once the risk targets/budgets have been set the actual situation can change the optimal allocation e.g. due to observed returns/volatility. One could formulate this as a

dynamic risk allocation problem, which seems to have a similarity with the problem of dynamic asset allocation (see e.g. Perold and Sharpe [1988]).

Concluding remarks

In each risk allocation model the main issue will be to arrive at a good estimate of the expected information ratio. Often the uncertainty of this estimate is being neglected. Theoretical and empirical evidence show that volatility of information ratios is there and is certainly not negligible. Once estimates are available, the framework presented here makes the risk allocation more robust in terms of the volatility of the outcomes. The resulting risk allocations can deviate substantially from the ones where parameter uncertainty is not taking into account. The framework is easy to implement, and can be enhanced to cope with issues as e.g. scalability. This can be based e.g. on the historically observed volatility of the information ratio.

End notes

The author likes to thank Andre Lucas and Cees Dert for their feedback.

References

- Blitz, D.C. and Hottinga, J.. "Tracking error allocation , steps to an optimum allocation framework", The Journal of portfolio management, Summer 2001, pp. 19-25.
- Chow, G. and Kritzman, M., "Risk budgets, Converting Mean-variance optimization into VaR assignments", The journal of portfolio management, Winter 2001, p56 – 60.
- Gupta, G., Prajogi R. and Stubbs E., "The information ratio and performance", The journal of portfolio management, Fall 1999, pp. 33-39.
- Lee, W. and Lam D.Y., "Implementing Optimal Risk Budgeting, managing uncertainty of alpha", The journal of portfolio management, Fall 2001, pp. 73 – 80.

Lo, A.W., "The Statistics of Sharpe Ratios", *Financial Analyst Journal*, July/August 2002, pp. 36-52.

Muralidhar, A.S., "Optimal risk-adjusted portfolios with multiple managers", *Journal of Portfolio Management*, Spring 2001, pp.97 – 104.

Perold, A.F. and Sharpe, W.F., "Dynamic strategies for asset allocation", *Financial Analyst Journal*, 1988, p16 – 27.

Waring, B., Whitney, D., Pirone, J. and Castille, C., "Optimizing manager structure and budgeting manager risk, a user's guide for active management", *The journal of portfolio management*, Spring 2000, p90 – 104.

¹ For the profound derivation of the formula the reader referred to the article of Lo [2002], which is very accessible for a technical paper. The basis formula is $v(IR)_{IID} = \left(\frac{\partial IR}{\partial \mu} \right)^2 \sigma^2 + \left(\frac{\partial IR}{\partial \sigma^2} \right)^2 2\sigma^4$. From the evaluation of the derivatives follows the formula in the text.

² For the top-quartile managers however there is some evidence in several asset categories that there is consistence in performance which would indicate a smaller uncertainty in the information ratio.

³ The measurement on basis of rolling periods is not ideal and can lead to wrong estimates of the volatility of the information ratio. Probably the main effect is a smoothing of the Information ratio series; this in turn implies that the stated figures are lower bounds to the "real" volatility of the information ratio. For the data used the average of non overlapping information ratio volatility was indeed slightly higher than the rolling period estimates. The number of datapoints per manager is however too low to draw significant conclusions.

⁴ One could also apply a bayesian approach in which each new information ratio data point leads to a better estimation.

⁵ The theoretical assumption of IID returns seems unlikely to hold. If we would have made a correction on the data as a result of survivorship bias that the database might have, the median volatility of the information ratios could even be higher. As a consequence the gap between theory and practice would increase.

⁶ Blitz & Hottinga divide the objective function by te_{target} to maximize the information ratio. As this parameter is a constant, it does not influence the results of the optimization. Exhibit 1 & 2 are based on the examples presented in the paper of Blitz & Hottinga.

⁷ To translate the partial tracking errors to weights and tracking errors the following procedure may be followed. Let A denote total assets, A_i denote allocation to investment process i and te_i the tracking error of process i. The partial tracking error can be translated into tracking error as follows: $te_i^p * A/A_i = te_i$, where A_i/A is the asset mix weight of investment process i. If there are limits on the tracking error of investment process i, then these limits can be (easily) translated into partial tracking error limits.

⁸ The numerator determines the sign of the derivative. The expression $\sum_i ir_i^2 - ir_i^2$ is clearly non-

negative and only zero when the information coefficients for the other investment processes are 0.

⁹ E.g. one could apply a similar approach as Waring et al. by creating a utility function. Then, e.g. the optimization problem could be stated as follows:

$$\begin{aligned}
& \underset{TE_1, \dots, TE_n}{Max} \sum_{i=1}^n IR_i TE_i - \lambda \sum_{i=1}^n \sum_{j=1; j \neq i}^n \left(\sigma_i^2 TE_i^2 + \sigma_{ij} TE_i TE_j \right) \\
& s.t. \sum_{i=1}^n \sum_{j=1; j \neq i}^n \left(TE_i^2 + \rho_{ij} TE_i TE_j \right) = TE_{target}^2
\end{aligned}$$

This problem formulation is a penalty function approach where as the shortfall probability formulation is more explicit on the minimum required return.

¹⁰ It could well be that the distribution of IR's is not normal and that it can be better approximated with a student's t distribution.