IIJ Institutional Investor Journals

The Voices of Influence | iijournals.com

The Journal of Portfolio Management

VOLUME 27 NUMBER 4 www.iijpm.com SUMMER 2001

Tracking Error Allocation

DAVID C. BLITZ AND JOUKE HOTTINGA

Tracking Error Allocation

Three steps to an optimum allocation framework.

David C. Blitz and Jouke Hottinga

Suppose you are an active investment manager of a global balanced portfolio looking back at the investment decisions you made during the past year, such as asset allocation, country allocation, and stock selection. You are sure that during this period you made more good decisions than decisions that turned out to be wrong, yet your portfolio still did not outperform its benchmark.

This means something went wrong in translating your investment views into the overweights and underweights you implemented relative to the benchmark. Intuitively, the more value one expects to be able to add with a particular investment decision, the more impact this decision should have on performance.

In this article we provide a transparent framework that formalizes this intuition. Using the tracking error concept for measuring relative risk, we show how an active manager should allocate a partial tracking error to each investment decision so as to maximize total value-added. If an active manager adopts this framework, one should not look back with surprise.

Grinold [1989] argues that maximizing value-added is equivalent to maximizing the information ratio of the portfolio, i.e., the ratio of active return to tracking error. Furthermore, he shows that both skill and breadth drive the information ratio. *Skill* is measured using the information coefficient, which is defined as the correlation of one's forecasts with actual outcomes. *Breadth* is the number of independent forecasts one can make.

DAVID C. BLITZ is a senior researcher in the Quantitative Research Department of the Robeco Group in the Netherlands.

JOUKE HOTTINGA is a senior researcher in the Quantitative Research Department of the Robeco Group in the Netherlands.

Grinold [1994] shows how forecasts should be translated into input parameters for a mean/variance optimization procedure. A constant information coefficient is assumed in this analysis, which is not unreasonable when we focus on one investment decision, such as stock selection. Our approach is more general in the sense that we look at different investment decisions, for which a manager may have structurally different forecasting skills.

Goodwin [1998] and Gupta, Prajogi, and Stubbs [1999] analyze realized information ratios of different classes of active managers. Apart from this ex post-oriented analysis, Gupta, Prajogi, and Stubbs also give an example of an ex ante allocation of assets over different classes of managers. Mean/variance optimization is used to maximize the expected information ratio, based on assumptions with respect to future risk and return characteristics. An important feature of this approach is that asset allocation and tracking error allocation are considered simultaneously.

In our view, strategic asset allocation and benchmark determination should be treated separately from issues relating to tracking error allocation. We concentrate our analysis on the latter.

We extend the literature by deriving a simple rule that active managers can use to maximize their information ratio and thus their value-added. Our argument is that one should strive for an investment structure, with independent investment decisions. Given such a structure, the optimum allocation is characterized by partial tracking errors that are proportional to the expected information ratios of the investment decisions. Further, we will demonstrate why a distinction needs to be made between tracking error limits and target tracking errors. This framework is illustrated with examples that show some interesting practical consequences of an optimum tracking error allocation.

TRACKING ERRORS AND INVESTMENT DECISIONS

Concepts such as tracking error, partial tracking error, and investment decision play a key role in the tracking error allocation framework, so we must first define what we mean by these concepts. We then provide a numerical example that will further clarify the necessity of a structured tracking error allocation.

The *tracking error* of a portfolio is defined as the annualized standard deviation of the return differences between that portfolio and its benchmark index. It measures the extent by which the portfolio returns deviate from those of the benchmark. Tracking errors can be measured both ex ante (based on present weightings and historical return data) and ex post (based on realized active returns).

In modern portfolio management, most investment managers are subject to a limit on ex ante tracking error. This limit quantifies the maximum potential permitted for active positions. A portfolio manager can use this potential for multiple investment decisions. We define an *investment decision* as an allocation over certain similar investment opportunities. For example, the investment opportunities for a global balanced portfolio can be divided into four investment decisions as follows:

1. Asset allocation (allocation between equities and bonds in its simplest form).
2. Country and/or sector allocation within each asset category.
3. Currency allocation (which can be separated from country allocation if currency derivatives are used).
4. Security selection within each country and/or sector within each asset category.

The tracking error resulting from all investment decisions is the overall tracking error. We define the tracking error that results from one investment decision in isolation as a *partial tracking error*. In other words, the overall tracking error tells you to what extent the overall return of the portfolio and the benchmark can differ, while the partial tracking errors explain the composition of the overall tracking error in terms of the underlying investment decisions.

Suppose we have a benchmark consisting of both equities and bonds in two countries. We have forecasts for the decisions regarding equity country allocation, bond country allocation, and currency allocation. With respect to all other investment decisions, e.g., asset allocation, we are neutral. One way to implement our view is to use equal percentage bets for each investment decision.

The percentages we use for this implementation are summarized in Exhibit 1.

EXHIBIT 1
EQUAL PERCENTAGE ALLOCATION

	Benchmark Weights		Active (relative) Positions	
	Country A	Country B	Country A	Country B
Equities	25%	25%	–10%	+10%
Bonds	25%	25%	+10%	–10%
Currency	50%	50%	+10%	–10%

Now suppose country A equities outperform country B equities by 24%, country A bonds outperform country B bonds by 8%, and the currency of country A rises 12% against that of country B. These returns are comparable one standard deviation moves, loosely based on actual market data.

Although our positions are correct for two out of the three investment decisions, the overall active return is negative:

$$[-10\% \times 24\%] + [10\% \times 8\%] + [10\% \times 12\%] = -0.4\%$$

In this example, we have forgotten to take into account two important considerations. First, we should consider the tracking errors resulting from each decision. Because equity markets are much more volatile than bond markets and currencies, the partial tracking error of the 10% equity country allocation bet is much higher than that of the 10% bond country and currency allocation bets. In other words, we used equal percentage bets for each decision, although the risk associated with the bets is definitely not equal. The effect of the equity country allocation bet therefore outweighs the two other (successful) bets.

Second, we should take our confidence in each of the forecasts into account. Clearly, the more confidence we have in our capability to forecast a particular investment decision, the higher the tracking error allocated to that decision should be. The essence of the tracking error allocation framework is that it formalizes these intuitive insights.

There are three steps that need to be taken to establish an optimum tracking error allocation:

1. Identify independent investment decisions.
2. Rank forecasting capabilities.
3. Allocate partial tracking errors to each investment decision.

IDENTIFYING INDEPENDENT INVESTMENT DECISIONS

The basis for active management is a good investment structure. Ideally, all investment opportunities should be divided into separate investment decisions that are independent, meaning that the active returns resulting from each decision are uncorrelated. If this is the case, then each investment decision can be made independently from the others—possibly even by different people.

Because tracking error allocation is the goal of our analysis, we will not try to provide the best investment structure. Throughout, we assume an investment structure with independent investment decisions. The assumption of independence is essential for deriving a transparent tracking error allocation.

The four decisions in the country example are investment decisions that are (largely) independent. If global equities outperform global bonds (asset allocation), then in general this does not imply anything about the return difference between equities in country A and equities in country B (country allocation). One caveat might be that certain countries have a higher beta with respect to global developments than other countries, but imposing a beta-neutrality restriction on the country allocation effectively deals with this problem. By striving for independence in this way, we avoid making what is actually an asset allocation decision in our country allocation decision. Similar arguments can be made for security selection within countries and/or sectors.

Even if the returns that can be achieved with different investment decisions are correlated to some extent (which they often will be in practice), this does not automatically imply that the investment decisions are also correlated. Since the active returns resulting from an investment decision are determined by both the returns of the investment opportunities and the allocation over these opportunities, it follows that a second condition for independence is that allocations are independent for each investment decision.

In other words, if our bond country allocations are, over time, uncorrelated with our equity country allocations, then bond country allocation and equity country allocation are independent investment decisions, regardless of the fact that the returns of the underlying investment opportunities may be correlated.

Although there will never be complete independence between and among all investment decisions in practice, independence can be achieved to a large extent by design of the investment structure and/or by looking at different types of variables for different investment decisions. We would also note that although we assume independence to derive a transparent tracking error allocation, risk management should always take the full picture into account, i.e., the overall portfolio tracking error including all correlation effects.

RANKING FORECASTING CAPABILITIES

After the investment decisions have been identified, the next step is to rank our forecasting capabilities. In step 3, the actual tracking error allocation, a higher partial tracking error will be allocated to investment decisions for which we have more forecasting capability.

We quantify forecasting capabilities by formulating expected information ratios for each investment decision. An expected information ratio is the ratio between expected outperformance and the corresponding tracking error.

One way to establish expected information ratios is to use the historical track record of individual investment decisions. If quantitative models play an important role in the investment process, one might look at long-term backtest results of these models. One should be very careful when extrapolating historical realizations into the future, however. See Beckers [1997], who shows that short-term information ratios entail a large element of chance, and Goodwin [1998], who stresses that realized information ratios of different types of asset managers need not say much about the future. The arguments presented by these authors also hold true for the determination of expected information ratios of different investment decisions.

What ultimately matters are the expectations one has with respect to one's strengths and weaknesses. In particular, one should try to quantify the ability to forecast a decision relative to the other decisions. Actually, the absolute magnitude of the expected information ratios is not relevant. Only the ratio of each expected information ratio with respect to the other decisions influences the tracking error allocation. Therefore, it is essential that the expected information ratios express one's relative forecasting ability for a decision.

Suppose, for example, that we expect to be significantly better at stock selection than country and/or sector allocation, while we do not expect to be able to add any value with asset allocation. We might then use an expected information ratio of 0.5 for stock selection, 0.25 for country and/or sector allocation, and 0 for asset allocation. For tracking error allocation purposes, however, we might just as well use expected information ratios of 1.0, 0.5, and 0, respectively, because the absolute scaling of the expected information ratios turns out to be irrelevant. Note that the minimum expected information ratio is zero, when we assume no forecasting power at all.

ALLOCATING TRACKING ERRORS

Analysis Using a Target Overall Tracking Error

Given an investment structure with independent decisions, and given expected information ratios, we can derive the tracking error allocation that maximizes the overall information ratio. Define IR_i as the expected information ratio of investment decision i ($i = 1, \ldots, n$) and TE_i as the partial tracking error corresponding to investment decision i. Let TE_{target} be the target overall tracking error.

If we assume information ratios greater than or equal to zero, the problem can be written as:

$$\max_{TE_1,\ldots,TE_n} \frac{\sum_{i=1}^{n} IR_i TE_i}{TE_{target}} \tag{1}$$

$$\text{subject to } \sum_{i=1}^{n} TE_i^2 = TE_{target}^2 \tag{2}$$

Note that IR_i times TE_i equals the expected active return resulting from investment decision i. Therefore the numerator in (1) represents the overall expected active return. Dividing by the overall tracking error then yields the overall expected information ratio that we want to maximize. Equation (2) shows how the partial tracking errors combine into the overall target tracking error, using the assumption of independence among the investment decisions.

It is straightforward to obtain the optimum solution to this problem:

$$TE_i^* = \frac{IR_i}{\sqrt{\sum_{j=1}^{n} IR_j^2}} TE_{target} \tag{3}$$

Equation (3) says that the overall tracking error should be distributed over the different investment decisions in such a way that the partial tracking errors are proportional to the corresponding expected information ratios. This is one of the key results of our allocation framework.

A consequence of this result is that every investment decision with a positive expected information ratio is allocated a positive tracking error. In other words, one should not restrict attention to the best decision or decisions alone, but consider all decisions, including those for which one expects relatively low information ratios.

Finally, by substituting (3) in (1), we find that the optimum overall expected information ratio is equal to:

$$\sqrt{\sum_{i=1}^{n} IR_i^2} \quad (4)$$

Note that the overall expected information ratio is higher than each of the expected information ratios for individual investment decisions.

We can illustrate the optimum tracking error allocation by continuing the example. Suppose we have equal confidence in our forecasts for the equity country allocation, the bond country allocation, and the currency allocation. Each decision should then have an equal tracking error.

An example of an optimum allocation is given in Exhibit 2. Note that we take the volatilities of the different investment decisions into account, as mentioned earlier.

Consider, for example, the country A bond overweight, which is three times the size of the country A equity underweight. The partial tracking error of these positions is the same because the volatility of the equity country allocation decision is assumed to be three times that of the bond country allocation decision (24% versus 8%).

With the allocation in Exhibit 2, the resulting returns, for the same one-standard deviation return realizations used earlier, are equal for each investment decision. This leads to a positive active return:

$$[-5\% \times 24\%] + [15\% \times 8\%] + [10\% \times 12\%] = +1.2\%$$

The equity bet no longer dominates the overall performance, because instead of using equal percentage overweights and underweights we have allocated equal tracking errors to each decision.

EXHIBIT 2
EQUAL TRACKING ERROR ALLOCATION

	Benchmark Weights		**Active (relative) Positions**	
	Country A	**Country B**	**Country A**	**Country B**
Equities	25%	25%	–5%	+5%
Bonds	25%	25%	+15%	–15%
Currency	50%	50%	+10%	–10%

Analysis Using Tracking Error Limits

A target tracking error can be seen as the average, long-term, tracking error. In practice, an overall tracking error limit is usually given, instead of an overall target tracking error. It is also desirable to formulate tracking error limits for individual investment decisions.

A tracking error limit is a restriction on the (ex ante) tracking error at any particular time and is relevant for risk management purposes. Therefore, a target tracking error is less than or equal to the corresponding tracking error limit.

Tracking error limits can be incorporated in the allocation framework by determining for each investment decision the ratio between the maximum tracking error and the target tracking error of that decision expected over time. The ratio has a minimum value of 1.0 if we maintain a constant tracking error for the investment decision. Stock selection may be an example of such a decision. The more the tracking error for an investment decision fluctuates, the higher the ratio will be. This happens when bets are large at some moments, but small (or even neutral) at others. This is the case for asset allocation, if we restrict ourselves to taking prominent bets only when multiple indicators point in the same direction.

Given an overall tracking error limit TE_{limit}, and denoting the ratios by f_i, a restriction can be added to the problem:

$$\sum_{i=1}^{n} \left(f_i TE_i\right)^2 = TE_{limit}^2 \quad (5)$$

Equation (5) shows how the maximum tracking errors for each investment decision i, f_i times TE_i, combine into the overall tracking error limit. Note that if a value for the tracking error limit TE_{limit} is specified, the target overall tracking error TE_{target} should no longer be seen as an input parameter but instead as a variable in the optimization problem. The optimum solution is still characterized by target partial tracking errors that are proportional to the expected information ratios. The only difference is that the scaling is now such that Equation (5) is satisfied:

$$TE_i^* = \frac{IR_i}{\sqrt{\sum_{j=1}^{n} \left(f_j IR_j\right)^2}} TE_{limit} \quad (6)$$

In practice, it may be desirable to impose additional restrictions on the tracking error allocation. A minimum partial tracking error may sometimes be necessary, such as when we want to limit the maximum number of securities in the portfolio. Other examples are restrictions on maximum attainable tracking errors because of restrictions on short-selling, or maximum exposures in certain investment opportunities.

One can cope with these kinds of restrictions by including them in the optimization problem and resolving for the optimum solution. The restricted optimum solution is again characterized by tracking errors that are proportional to the expected information ratios, but with some tracking errors cut off at their maximum or minimum values.

APPLICATION

Let us now apply the complete framework to a global balanced portfolio. Suppose we distinguish the investment decisions and corresponding expected information ratios shown in Exhibit 3. The information ratios imply that we are most optimistic about our stock and bond selection capabilities. The expected information ratios for the other decisions are lower, but still positive, except for currency allocation, for which we do not expect to be able to add any value in this example.

The other required input consists of an overall tracking error limit of 5% and assumptions with respect to the ratios between partial tracking error limits and partial target tracking errors (f_i). For stock selection and bond selection, we expect to maintain a constant tracking error, implying factors f that are equal to 1.0. For asset allocation, equity country allocation, and bond country allocation, the size of our bets fluctuates more over time, leading to factors f estimated to be equal to 2.0.

Equation (6) is used to calculate the optimum partial target tracking errors and partial tracking error limits. It can be seen that the optimum partial target tracking errors TE_i^* are proportional to the expected information ratios IR_i. When they are multiplied by the factors f_i, the optimum partial tracking error limits are obtained. With the optimum tracking error allocation, we are able to reach an overall expected information ratio of 0.7, which can be verified using (4). Note that this value is higher than the information ratio for each individual decision.

A final example is based on our own practical experiences with implementing an optimum tracking error allocation. Consider an investment manager who is offering clients a 100% equity product and a 100% bond product. The manager expects to achieve the same information ratios for both products, but has set a higher target tracking error for the 100% equity product, because the clients for that product have a higher risk tolerance than those for the 100% bond product. The investment manager now wants to introduce a 50% equity/50% bond balanced product, with a target tracking error lying between that of the two products. A straightforward, practical way of managing this balanced product might be to invest in the

EXHIBIT 3
APPLICATION OF TRACKING ERROR ALLOCATION FRAMEWORK TO A GLOBAL BALANCED PORTFOLIO

	Expected Information Ratio IR_i	Optimum Target Partial Tracking error TE_i^*	Factor Tracking Error Limit / Target Level f_i	Optimum Tracking Error Limit ($f_i TE_i^*$)
Asset Allocation	*0.2*	**1.0%**	*2*	**2.0%**
Equity Country Allocation	*0.2*	**1.0%**	*2*	**2.0%**
Bond Country Allocation	*0.3*	**1.5%**	*2*	**3.0%**
Currency Allocation	*0.0*	**0.0%**	*not relevant*	**0.0%**
Stock Selection	*0.4*	**2.0%**	*1*	**2.0%**
Bond Selection	*0.4*	**2.0%**	*1*	**2.0%**
Total	**0.7**	**3.5%**		*5.0%*

Note: Input parameters represented in italics, calculated numbers in ***bold***.

existing 100% equity and 100% bond products and to implement asset allocation views by changing the weights in these existing products.

Unfortunately, this approach is suboptimal from a tracking error allocation point of view. Because the expected information ratios for the equity and bond parts of the balanced portfolio are equal, both parts should have the same target tracking errors. What goes wrong with the straightforward implementation is that the position in the 100% equity product has a much greater impact on the overall performance than the position in the 100% bond product, due to the different tracking error levels of these products.

We encountered this phenomenon when we implemented an optimum tracking error allocation for our mutual funds and mandates. An important consequence of the tracking error allocation framework for us has therefore been to shift tracking error from the equity part to the bond part of our balanced portfolios, so that the resulting allocation complies with the framework.

CONCLUSION

A systematic approach to structuring allocation decisions is essential for modern active portfolio management. If you lack forecasting capability, then a good allocation framework cannot help you. On the other hand, poor implementation can wreck the most excellent forecasting capability.

We have shown that investment managers need to complete three steps to obtain an optimum allocation framework:

1. Identify the different investment decisions, striving for independence as much as possible to avoid making the same decision at multiple stages in the investment process.
2. Rank one's forecasting capabilities by defining an expected information ratio for each investment decision.
3. Allocate partial target tracking errors proportionally to the expected information ratios of the decisions, to achieve a maximum expected portfolio information ratio. A simple consequence of this rule is that if we expect the same information ratio for two decisions, the percentage of overweights or underweights should be lower for the more risky (volatile) decision.

When implementing the framework, one should bear in mind that an optimum tracking error allocation is partly based on subjective estimates, in particular with respect to the expected information ratios. The assumption of independence between or among investment decisions will also not be satisfied exactly in practice. The aim, however, is not to derive a set of restrictive tracking error limits to be meticulously monitored, but rather to provide a framework for translating intuitive ideas about capabilities into practical and transparent allocation rules.

A useful side effect of adopting the tracking error allocation framework is that it helps to achieve consistent performance for the different portfolios under management. The investment decisions can be seen as the building blocks of active management applied to every portfolio.

An interesting topic for further research is the translation of our findings to a similar allocation problem: the allocation of tracking error over asset managers by a plan sponsor. The issue of independence is particularly interesting. Should a plan sponsor prefer managers with uncorrelated, positively correlated, or negatively correlated strategies? What are the advantages of specialized mandates compared to balanced mandates? And is there an optimum number of active managers?

ENDNOTE

The authors thank Lev Dynkin, Peter Ferket, and Angelien Kemna for their useful comments.

REFERENCES

Beckers, S. "Manager Skill and Investment Performance: How Strong is the Link?" *The Journal of Portfolio Management,* Summer 1997, pp. 9-23.

Goodwin, T.H. "The Information Ratio." *Financial Analysts Journal*, July/August 1998, pp. 34-43.

Grinold, R.C. "Alpha is Volatility Times IC Times Score." *The Journal of Portfolio Management*, Summer 1994, pp. 9-16.

——. "The Fundamental Law of Active Management." *The Journal of Portfolio Management*, Spring 1989, pp. 30-37.

Gupta, F., R. Prajogi, and E. Stubbs. "The Information Ratio and Performance." *The Journal of Portfolio Management,* Fall 1999, pp. 33-39.