Attribution and risk-adjusted performance metrics

Andrew Colin
StatPro Australia
7th February 2007

1. Introduction

Consider two hypothetical fixed income fund managers. Last year, one of them made an impressive 5% return above benchmark by taking aggressive interest rate bets and making extensive use of highly leveraged derivative instruments. It's been a roller-coaster ride, but he's finally ahead of the game.

In contrast, his colleague from down the hallway only made a 2% return over the same benchmark. He carefully restricted his investments to AAA-rated gilt-edged stocks, never went more than six months long or short against benchmark, and regards derivatives as instruments of the Devil.

Who did the better job?

If you are not concerned about risks and potential future losses, then of course the first manager won out. But would you want him managing your money over the longer term? The prudent course might be to invest with the second manager. Although his returns are lower, he also takes lower *risks*, which means his returns will probably be more consistent over the longer term.

Unfortunately it's not always easy to get the true picture of the risks taken by money managers, particularly when trying to get a picture of what *type* of risks were taken. The reason is that running a portfolio requires to ability to juggle multiple risk factors, each of which can provide widely varying returns.

For instance, equity managers take stock selection and asset allocation decisions; fixed income managers look at duration, credit spreads and coupon returns; and more recently the industry is investing in instruments that transfer risk from the insurance industry, such as catastrophe bonds, weather and energy derivatives, mortality derivatives. It's even possible to buy bonds backed by future royalties on sales of popular music. Measuring the risks involved in exposing investor's money to each of these factors is not going to be an easy task.

The industry already makes wide use of portfolio-level measures of risk, such as tracking error or value at risk. The subject of this talk is how to move a step deeper and measure risk at the individual factor level. By the time we've reached the end, I hope you'll have an appreciation of how to measure, not only the *return* due to (say) stock selection and asset allocation bets, but also the *risk* involved in each decision – and whether that level of risk is appropriate for your institution. I'll concentrate mostly on fixed income, but you'll be able to apply the same approach to equity risks as well.

1.1. Defining risk

Numerous measures of portfolio risk against return exist. Some widely used metrics are

- Treynor measure (excess return above risk-free return, divided by beta for portfolio);
- Sharpe ratio (excess return above risk-free return, divided by standard deviation of current portfolio);
- Jensen alpha;
- Information ratio (excess return of portfolio over benchmark, divided by standard deviation of that excess return);

An alternative approach is to use RAROC measures. RAROC stands for 'Risk-Adjusted Return on Capital' and was initially developed in 1980s by Bankers Trust. A common definition is simply adjusted return on capital:

RAROC = (revenue – expenses – expected loss + income from capital) / capital

Alternative, absolute measures of risk are

- Value at Risk (or VaR);
- Expected tail loss (or ETL, or expected shortfall, or conditional VaR);
- Tracking error;
- Active duration

Some of these risk measures measure divergence about some mean (all those based on tracking-error type measurements) while others look at one-sided risk measurements (particularly those based on VaR). We refer to the former as *symmetric* risk measurements, and the latter as *asymmetric*. Active duration is a crude measure of sensitivity to parallel movements in the term structure, rather than being calculated from measured excess returns. It only refers to one specific type of risk, and is not considered further in this paper.

While allowing different portfolios to be compared, these measures all have the same drawback: they only provide an aggregated risk figure. In other words, they throw together risks from all sources without giving information about their relative sizes.

How can we break out, or decompose, these risks? The answer is to use *risk attribution*. Once you have a performance attribution capability in place, risk attribution is the logical next step, since it allows us to measure not just (for instance) the return made from credit decisions, but also how much of a risk the manager took to make those returns.

With the techniques I'll describe during this talk, you'll see how to use risk attribution not only to both measure the risk/reward ratio for different types of bets, but also where investment 'hot-spots' lie within individual credit or risk sectors, and whether some security holdings (or groups of security holdings) are riskier than others.

I like to claim that 'Performance tells us about the fund. Attribution tells us about the manager'. To this I now add: 'Risk attribution tells us even more about the manager.'

1.2. Risk attribution

Naturally, there are some preliminary steps to go through before we can actually arrive at a risk attribution figure. The first of these (and, in fact, the hardest) is to run a performance attribution analysis on which to base the risk calculation.

Attribution can be defined as *the decomposition of a portfolio's active return into returns generated by different sources of risk*, thus allowing us to measure the actual value added by, for instance, asset allocation, credit bets, or any of a number of other factors in which you may have an interest.[1] Note that all attribution does is to provide returns – it says nothing about the distribution of those returns.

Since the yield curve is the main driver of most fixed income portfolios, the fixed income attribution analyst typically breaks down the changes in the yield curve over the calculation interval into the changes due to each risk factor – such as parallel shift, twisting, curving, credit spreads, roll-down return, and so on. Using a pricing function that uses one or more yields, the security is then re-priced at the start of the interval, at the end, and with all intermediate yields. The difference between the prices then calculated can then be used to give the components of return due to each risk factor.

The two main problems in this process are

(i) deciding how to represent changes in the yield curve in terms that traders can follow. Shift and twist are widely understood concepts, but their exact meaning is not always clear when measured in tandem, and it is usually necessary to decide upon an arbitrary maturity about which a curve twists, and stick to it – which can lead to non-intuitive results. Using principal component analysis is mathematically elegant but does not give as much insight, especially if your traders make their decisions in other, unrelated terms.

(ii) translating changes in yield into changes in price. This is less a mathematical issue than a resources one. If all you're interested in is performance, all you have to do is download a price. Attribution, on the other hand, requires that each security in the portfolio and benchmark be re-priced accurately, several times; and then the end price still has to be checked against the market price, and any discrepancies corrected.

    The complexity and costs of this process lead many to move to alternate models, especially perturbational attribution, where the price/yield sensitivity is represented by one or more sensitivity measures, such as modified duration and convexity; however, these risk numbers still need to be downloaded or calculated. A more advanced approach to this problem that does away with the need for risk numbers entirely is presented in Colin (2007).

So what's a suitable measure of risk for attribution? It probably depends on the investment style of the manager.

For an indexed fund, it may be appropriate to ensure that the fund is hedged against all the different types of yield curve movement, and this case the curvature movement of the

---

[1] See Colin, 2005, for a comprehensive list of fixed income risk factors.

portfolio should exactly equal that of the benchmark. In this case, one of the symmetric risk measurements above is appropriate.

For a fund that intends to add excess return, or *alpha*, an asymmetric risk measurement may be more appropriate. In this context, attribution is intended to measure the causes of out-performance, not how closely we can stick to a predefined return.

As most of the audience is probably more concerned with the latter category, the rest of this talk will refer to value at risk (VaR). But don't overlook the other risk measurements; they may be right for your investment context.

2. A brief introduction to risk measures

2.1. Value at Risk

Firstly, some background. Value at Risk is a simple way to measure the magnitude of likely losses of a portfolio. It's defined as a one-sided confidence interval on potential portfolio losses over a specific horizon.

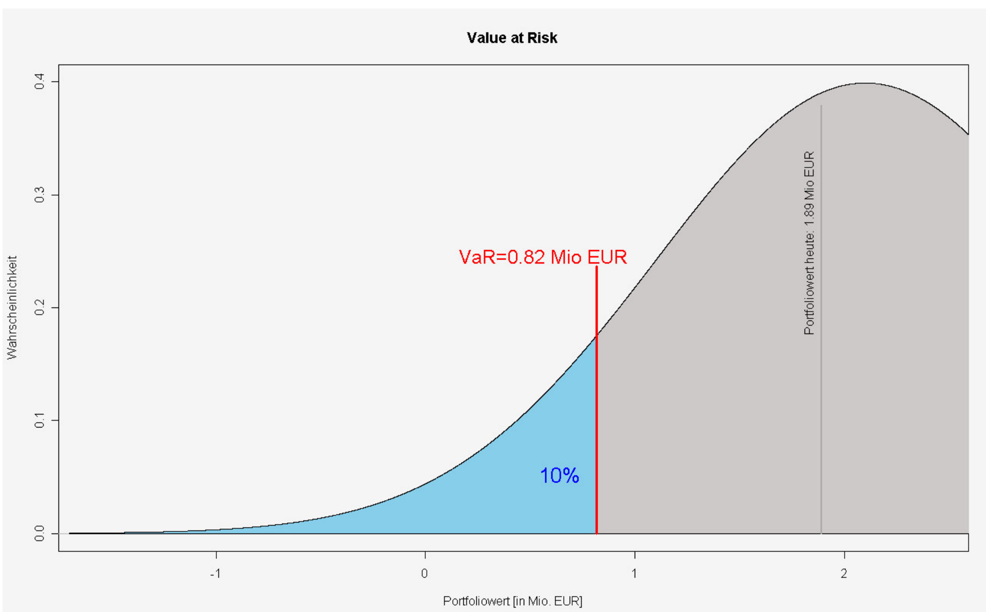

VaR is the most widely used risk measure due to European Capital Adequacy directives (CAD). However, it has a number of widely understood shortcomings, including the validity of its statistical approaches, and (in particular) that VaR is not a *coherent* risk measure. For a detailed discussion of these shortcomings, see Dowd (2006).

To take a classic example (Artzner et al, 1997):

- Trader A has sold an out-of-the-money put that is far out of the money, with one day to expiry.

- Trader B has sold an out-of-the-money call that is also far out of the money with one day to expiry.
- The probability of either option ending up in the money is 4%.

Taken individually, each trader has a holding that has a 96% of not losing money, so the individual 95% VaR amounts are zero.

but…

The combined portfolio has a 92% chance of not losing money, so the 95% VaR amount for the portfolio is non-zero!

So the risk of the combined portfolio is greater than the risks of the individual positions.

One of the best introductions to VaR is Jorion (2007). Various pungent criticisms of the technique, such as '*VaR is charlatanism, a dangerously misleading tool'* and other pungent remarks may be found at http://www.fooledbyrandomness.com/

2.2. Expected tail loss

ETL is defined as the average of the worst $100(1-\alpha)\%$ of losses. Unlike VaR, ETL *is* a coherent risk measure. ETL tells us what to expect in periods of loss and assigns a number to that loss, while VaR only tells us to expect a loss larger than the VaR value.

ETL can usually be calculated using similar techniques to VaR. The interested reader will again find a full discussion in Dowd (2006).

2.3. Tracking error

Tracking error is simply calculated as the standard deviation of the active return of portfolio over benchmark. Unlike VaR and ETL, which are usually calculated for ex-ante use, tracking error is usually calculated ex-post .

3. Calculating VaR for risk attribution

We already have a definition for a portfolio's VaR. Rather than go straight to a definition of how VaR is affected by different risk factors, let's look at how it is affected by changes in individual security positions.

First of all, let's define a quantity called *marginal VaR*, (or VaR_i^M). This is defined as the change in the overall portfolio VaR resulting from a change in position by a small amount (say, one dollar) of instrument i – which is the same as the partial derivative with respect to that exposure.

VaR is linearly homogenous, so we can apply Euler's theorem and write

$$VaR_P = \sum_i w_i \frac{\partial VaR_P}{\partial w_i} \tag{3.1}$$

The term under the summation sign is called the *component VaR*, defined as

$$VaR_i^C = w_i \frac{\partial VaR_P}{\partial w_i} \tag{3.2}$$

The term without the weighting factor is called the marginal VaR, defined as

$$VaR_i^M = \frac{\partial VaR_P}{\partial w_i} \tag{3.3}$$

From (3.1), we have

$$\sum_i VaR_i^C = VaR_P \tag{3.4}$$

(3.4) shows that component VaR has some extremely useful properties when it comes to decomposing risk. It allows contribution to overall VaR to be broken down by sector, by instrument and even – as we shall see – by risk factor.

To calculate marginal and component VaR, start from the definition of portfolio variance:

$$\sigma_P^2 = \sum_{i=1}^{N} w_i^2 \sigma_i^2 + 2\sum_{i=1}^{N} \sum_{j<i}^{N} w_i w_j \sigma_{ij} \tag{3.5}$$

where σ_i^2 is the variance of security i, σ_{ij} is the covariance of securities i and j, and w_i is the weight of security i. If we take the derivative of this expression and skip a couple of steps (see Jorion, 2007), it turns out that

$$\frac{\partial \sigma_P^2}{\partial w_i} = 2\sigma_P \frac{\partial \sigma_P}{\partial w_i} = 2\operatorname{cov}\left(r_i, r_P\right) \tag{3.6}$$

implying

$$\frac{\partial \sigma_P}{\partial w_i} = \frac{\operatorname{cov}\left(r_i, r_P\right)}{\sigma_P} \tag{3.7}$$

If we now define a quantity β_i given by

$$\beta_i = \frac{\operatorname{cov}\left(r_i, r_P\right)}{\sigma_P^2} \tag{3.8}$$

then equation (3.7) simplifies to the very simple form

$$\frac{\partial \sigma_P}{\partial w_i} = \beta_i \cdot \sigma_P \qquad (3.9)$$

So marginal and component VaRs are given by, respectively,

$$VaR_i^M = \beta_i \cdot VAR_P \qquad (3.10)$$

and

$$VaR_i^C = w_i \beta_i \cdot VAR_P \qquad (3.11)$$

We now have expressions for component VaR and marginal VaR.[2]

Note that marginal VaRs can be negative. This means that increasing the size of the corresponding instrument will decrease risk – it acts as a natural hedge.

Putting the result in an even simpler form, we can say that the contribution to VaR of component i is given by

$$\frac{VaR_i^C}{VaR^P} = w_i \beta_i \qquad (3.12)$$

giving us the risk contribution from each security. This shows which instruments are acting as natural hedges by decreasing risk, where the portfolio can be readjusted to reduce risk overall, and where the riskiest exposures lie.

Marginal VaR gives us the contribution to VaR from each security. In the next section, we'll show how to calculate the contribution to VaR from individual risk factors.

4. Calculating value at risk

4.1. Mean-variance VaR

If we assume the distribution of any asset's returns is normal, then one parameter describes the risk – the variance. This then allows use of the standard variance/covariance model, in which we use a matrix of covariances between all asset returns. This model is implemented in RiskMetrics from JP Morgan.

Unfortunately, real price distributions are not as simple as this, especially where fat-tailed distributions are concerned. A paper from by Litterman (1996) quotes, apparently quite without irony,

> *Given the non-normality of daily returns that we find in most financial markets, we use as a rule of thumb the assumption that four-standard deviation events in financial markets happen approximately once per year.*

[2] Without going into details, this relationship also holds for non-normal distributions.

A four standard deviation event has a one in 31574 chance in occurring, or approximately once every 143 years – a difference of two orders of magnitude between theory and practice!

4.2. Historical simulation VaR

Historical simulation can be a better way to go. This technique assumes asset returns in future will have same distribution that they had in the past (or at least over the sample supplied). Useful features of this approach are that it makes no assumptions about expected distribution of returns, and is very simple to implement. However, it requires large amounts of historical data.

4.3. Monte Carlo VaR

Monte Carlo simulation requires generation of a large number of market movements using some market model. If there are n samples, then the Var 95% loss is at position n * (1-0.95).

Monte Carlo VaR requires extensive number-crunching, and is subject to model risk – the risk that the function generating market movements may not be realistic.

4.4. Extending to ETL

Here we've talked about VaR; but it's very straightforward to extend this to ETL.

For VaR at the 99^{th} percentile from 1000 days: rank in order, VaR = P(990).
For ETL at the 99^{th} percentile from 1000 days: rank in order

$$ETL_{\alpha} = \frac{1}{1-\alpha} \sum_{i=N\alpha}^{N} P_i \tag{4.1}$$

In other words, expected tail loss is just the average of all returns at positions greater than the 95^{th} percentile.

4.5. Which is best for risk attributon?

For risk attribution, we recommend using historical simulation. The reasons are that

(i) we have to work out attributed return time series anyway, so the raw data will be available;

(ii) prices are trader-independent, but attribution returns are not (by definition). Attribution returns are affected by trading style, ability, and in fact everything we are trying to measure. So it is in principle not possible to download the distribution of attribution returns; it must be calculated internally as part of the risk attribution process.

(iii) Risk attribution is the easiest way to include changes in position from trading.

5. Risk and attribution

So far, we have only discussed VaR in the context of individual instruments. Is it instead possible to apply this approach to returns from sources of risk, rather than from individual instruments?

The answer is yes, *if* we have attribution returns available.

Consider the returns of a bond over a given interval. Ignoring residuals, an attribution analysis will decompose this return into components from coupon, duration and credit.

$$r = r_{Cpn} + r_{Duration} + r_{Credit} \tag{5.1}$$

To calculate risk attribution, we regard the portfolio as the union of several virtual instruments, each of which has the return shown by exactly one risk factor, and each with the same weighting. The return of these various virtual instruments will add up to the same return as the portfolio.

For instance, what is the risk from credit spreads? Calculate the attribution returns for all securities from credit, and aggregate them to get a per-day credit return, form a time series of credit returns. Calculate the covariance of this time series with the time series of portfolio returns. This will then give the contribution to VaR from credit.

Summarizing:

(i) Decide on your risk factors. Are you going to use shift and twist, principal components, or some other factor decomposition of the main risks affecting your portfolio?
(ii) Run an attribution analysis on the portfolio, and decompose security returns in terms of these risk factors.
(iii) Aggregate the factor returns over all securities.
(iv) Calculate the covariance of each factor return against that of the portfolio to give a factor beta.
(v) Given the portfolio's VaR, decompose this figure into contributions from risk sectors.
(vi) Calculate the attribution ratio – the ratio of each factor's return to its risk, defined as

$$R_{risk} = \frac{r_i}{VaR_i^C}$$

(vii) Compare and use the attribution ratios. Is one risk factor performing significantly worse than its peers? Are you taking too much (or too little) risk in one area, compared to the realized returns?

6. Summary

The key point of this paper is that the next step beyond providing a performance attribution capability is to provide a risk attribution capability. This allows us to measure not just the return made by different risk factors, but the risk taken by the manager in running exposures against these factors.

The key measure in this discussion is the attribution ratio $\mathbb{R}_i$

$$\mathbb{R}_i = \frac{r_i}{VaR_i} \tag{5.2}$$

where r_i and VaR_i are the return and value at risk due to risk factor i, respectively.
This simple ratio supplies the next level of insight beyond performance attribution. As I hope this paper has conveyed, it’s not actually that hard to calculate.

7. References

Artzer, P., Delbaen, F., Eber, J., Heath, D., *Coherent measures of risk*, Mathematical Finance 9 (July), 1999
Colin (2007), *Fixed income attribution with minimum raw material*, Journal of Performance Measurement, vol. 11 #1, Winter 2007
Dowd, K., *Measuring market risk*, John Wiley and Sons, 2005
Jorion, P., *Value at Risk*, McGraw-Hill, 2007
Litterman, P., *Hot Spots and Hedges*, Goldman Sachs, New York, 1996