Risk Attribution

[Nematrian website page: RiskAttributionTheory, © Nematrian 2015]

Abstract

The aim of the attached pages is to summarise some of the theory behind the attribution of portfolio risk. See also more general introductions to the topic such as Kemp (2005) or Kemp (2009).

Contents

1. Introduction
2. Implied Alphas
3. Beta-adjusted attribution

References

1. Introduction

[RiskAttributionTheory1]

1.1 Traditionally, risk attribution (if the risk model is characterised by a covariance matrix) proceeds as follows. We assume that there are *n* different instruments in the universe in question. We assume that the portfolio and benchmark weights can be represented by vectors $\mathbf{p} = (p_1, \dots, p_n)^T$ and $\mathbf{b} = (b_1, \dots, b_n)^T$ respectively. The active positions are then $\mathbf{a} = \mathbf{p} - \mathbf{b}$. If the risk model is characterised in the parsimonious manner involving a factor covariance matrix, $\widehat{\mathbf{V}}$, and a sparce idiosyncratic matrix, $\mathbf{Y}$, e.g. as described in Kemp (2009), then $\sigma^2(\mathbf{a}) = \sigma^2(\mathbf{p} - \mathbf{b}) = \big(\mathbf{F}(\mathbf{p} - \mathbf{b})\big)^T \widehat{\mathbf{V}}\big(\mathbf{F}(\mathbf{p} - \mathbf{b})\big) + (\mathbf{p} - \mathbf{b})^T \mathbf{Y}(\mathbf{p} - \mathbf{b})$. The matrix describing the covariance structure between factors, i.e. $\widehat{\mathbf{V}}$ corresponds to a *projection* of an *n* dimensional space onto a smaller *m* dimensional space.

1.2 Factors might be further grouped into one of, say, z different factor types, using what we might call a *factor classification*, $\mathbf{T}$, i.e. a $z \times m$ projection matrix that has the property that each underlying factor is apportioned across one or more 'super' factor types. By apportioned we mean that if $T_{k,j}$ corresponds to the exposure that the *j*'th factor has to the *k*'th factor type then the sum of these exposures for any given factor is unity, i.e. $\sum_k T_{k,j} = 1 \, \forall j$.

1.3 Usually such a factor classification (at least in equity-land) would involve unit disjoint elements, i.e. each factor would be associated with a single 'super' factor type. For example, equity sector classification structures are usually hierarchical, so each industry subgroup is part of a (single) overall market sector. More generally, factors might be apportioned across more than one factor type. The aggregate (relative) exposure to the different factor types is then, in matrix algebra terms, equal to $\mathbf{TFa}$.

1.4 To decompose (or 'attribute') the tracking error into its main contributors it is usual to decompose the tracking error, $\sigma(\mathbf{a})$, in the manner described in Kemp (2005), Kemp (2009) or Heywood, Marsland and Morrison (2003), i.e. in line with partial differentials (scaled if necessary by a uniform factor so that the total adds up to the total tracking error). For example, if the aim is to identify the risk contribution coming from each individual security then we might calculate the *marginal* *contribution to tracking error*, m_i, and the *contribution to tracking error*, c_i, from the *i*'th instrument as follows:

$$m_i \equiv \frac{\partial\sqrt{\sigma^2(\mathbf{a})}}{\partial a_i} = \frac{\partial\sqrt{\mathbf{a}^T\mathbf{F}^T\widehat{\mathbf{V}}\mathbf{F}\mathbf{a} + \mathbf{a}^T\mathbf{Y}\mathbf{a}}}{\partial a_i} = \frac{1}{\sigma(\mathbf{a})}\left(\mathbf{F}^T\widehat{\mathbf{V}}\mathbf{F}\mathbf{a}\right)_i$$
$$c_i = a_i m_i$$

1.5 This has $\sum_i c_i = \sigma(\mathbf{a})$, so the sum of the individual contributions assigned to each instrument is the total tracking error of the portfolio. Simetimes writers instead focus on decomposing the variance rather than the standard deviation. However the answers are the same up to a scaling factor. This is to be expected since for any two functions, $f(x)$ and $g(x)$ with first differentials f' and g' we have $\partial f\left(g(\mathbf{x})\right)/\partial x_i = f'\left(g(\mathbf{x})\right)\partial g(\mathbf{x})/\partial x_i$, i.e. the vector of partial differentials is the same, up to a scaling factor for all functions of same underlying risk measure. Variance and standard deviation in this context relate to the 'same' underlying risk measure, since variance is the square of standard deviation.

1.6 We can group the c_i in whatever manner we like, as long as each relative position is assigned to a unique grouping or if it is split across several groupings then in aggregate a unit contribution arises from it.

1.7 For example, suppose that we have a classification described by the $y \times n$ matrix $\mathbf{U}$ with elements $U_{q,i}$ being the contribution that the *i*'th instrument makes to the *q*'th classification. As it is a classification it needs to satisfy $\sum_q U_{q,i} = 1\ \forall q$ so $\sum_{q,i} U_{q,i}c_i = \sum_i (\mathbf{Uc})_q = \sigma(\mathbf{a})$, i.e. grouping in this manner is equivalent to calculating Uc where $\mathbf{c} = (c_1, \dots c_n)^T$.

1.8 A special case of such a classification would be to calculate the contribution to risk from a given issuer (rather than a given instrument), e.g. for a bond portfolio, where $U_{q,i}$ would be 1 if issue i is issued by issuer q and 0 otherwise.

1.9 The above approach calculates a *single* overall contribution to tracking error for each individual instrument (and then if necessary groups them). For bond risk analysis (and also in some instances for equity risk analysis) it may be more illuminating to subdivide these instrument specific contributions into several different sub-elements, each one relating to a given factor type. These factor types might be, say, currency, interest rate (duration), credit, sector, other factors and idiosyncratic.

1.10 To do this we need to subdivide $\mathbf{Fa}$ into several different elements, which cumulatively add up to $\mathbf{Fa}$, each relating to a different factor type, i.e. we define, say, $\bar{\mathbf{F}}_1, \dots, \bar{\mathbf{F}}_z$ each of which are $m \times n$ matrices, which have the property that the (j,i)'th element of $\bar{\mathbf{F}}_k$ is calculated as $(\bar{\mathbf{F}}_k)_{j,i} = (\mathbf{F})_{j,i}T_{k,j}$ for a given factor classification $\mathbf{T}$. The sum of the $\bar{\mathbf{F}}_k$ for all $k = 1, \dots z$ is then $\mathbf{F}$. We can then decompose the marginal contributions to tracking error into the following, where $m_{i,k} = \mathbf{F}^T\mathbf{V}\bar{\mathbf{F}}_k\mathbf{a}/\sigma(\mathbf{a})$ and $R_i = (\mathbf{Ya})_i/\sigma(\mathbf{a})$:

$$m_i = \frac{1}{\sigma(a)}\left(\sum_k \mathbf{F}^T\mathbf{V}\bar{\mathbf{F}}_k\mathbf{a} + \mathbf{Ya}\right)_i = \sum_k m_{i,k} + R_i$$

2. Implied Alphas

[RiskAttributionTheory2]

2.1 An individual security's marginal contribution to risk is closely allied to its *implied alpha*, α_i, i.e. the expected outperformance (or underperformance) you need to expect from the instrument if the portfolio is to be 'efficient' in the sense of optimally trading off risk against return. For the portfolio to be efficient we need to have, for some portfolio risk aversion parameter, λ, all of the following n equations simultaneously to be true (in a mean-variance world):

$$\frac{\partial}{\partial a_i}\left(\alpha_i a_i - \lambda(\mathbf{a}^T \mathbf{F}^T \mathbf{VFa})\right) = 0 \implies \alpha_i = C + \lambda \sigma(\mathbf{a}) m_i$$

2.2 Here C is an arbitrary constant that might be chosen so that the weighted average implied alpha of the benchmark is zero, since the implied alpha of a given portfolio or instrument is then more directly related to the expected excess alpha that such a portfolio might deliver versus the benchmark.

3. Beta-adjusted attribution

[RiskAttributionTheory3]

3.1 A problem that arises with the above approach for traditional long-only portfolios is that such portfolios are typically benchmarked against market indices and often have 'betas', i.e. exposures to the market, which are close to one. This can make traditional risk decompositions versus the benchmark of such portfolios very sensitive to small changes in the amount of cash held within the portfolio (because cash has a beta of zero, i.e. substantially different to the benchmark's beta of one), limiting the usefulness (or rather the stability) of the above decomposition for practical portfolio management. Better may be to decompose each instrument's contribution to risk into a beta component and a remainder, with only the latter then subject to further decomposition in the usual manner.

3.2 Beta is a measure of how much a portfolio (or of an individual security) might be expected to rise or fall as the market (i.e. benchmark) rises or falls. A beta of 1 means that, all other things being equal, a 1 basis point rise or fall in the market leads to a corresponding 1 basis point rise or fall in the portfolio value. A beta of more than one means that all other things being equal the portfolio should rise or fall more than the corresponding rise or fall in the market, a beta of less than one means it should rise or fall less than the corresponding rise or fall in the market. Betas can, of course, be negative (e.g. a put option would typically have a negative beta, since it rises in value as the underlying falls). Long only portfolios typically have betas that are not too far from one, this being by definition the average beta of the relevant index being used as the benchmark for the long only portfolio.

3.3 Beta is benchmark specific, i.e. a stock with a given beta against one market index may have a different beta against a different market index. The terminology 'beta' arises because in effect we are ascribing a security's (or an entire portfolio's) return in a manner akin to a regression analysis in which $y_{i,t} = \alpha_i + \beta_i x_t + \varepsilon_{i,t}$ where $y_{i,t}$ is the return on the *i*'th security in the *t*'th time period, x_t is the return on the market index in the *t*'th time period. Conventionally the intercept of this regression is typically referred to as the 'alpha' and the slope of this regression as the 'beta'.

3.4 From an ex-ante risk perspective, the *portfolio beta* can be calculated as:

$$\beta = \frac{\mathbf{p}^T\left(\mathbf{F}^T \widehat{\mathbf{V}} \mathbf{F} + \mathbf{Y}\right)\mathbf{b}}{\mathbf{b}^T\left(\mathbf{F}^T \widehat{\mathbf{V}} \mathbf{F} + \mathbf{Y}\right)\mathbf{b}}$$

3.5 The *active portfolio beta* is then:

$$\beta_{act} = \beta - 1 = \frac{\mathbf{a}^T\left(\mathbf{F}^T\widehat{\mathbf{V}}\mathbf{F} + \mathbf{Y}\right)\mathbf{b}}{\mathbf{b}^T\left(\mathbf{F}^T\widehat{\mathbf{V}}\mathbf{F} + \mathbf{Y}\right)\mathbf{b}} = \sum_i a_i\beta_i$$

where we have decomposed the overall active beta into contributions from each individual (active) position, β_i, where

$$\beta_i = \left(\frac{(\mathbf{F}^T\mathbf{V}\mathbf{F} + \mathbf{Y})\mathbf{b}}{\mathbf{b}^T(\mathbf{F}^T\mathbf{V}\mathbf{F} + \mathbf{Y})\mathbf{b}}\right)_i$$

3.6 The marginal contribution to tracking error from a portfolio's beta is then

$$M_\beta = \frac{\partial\sqrt{\left(\sigma(\mathbf{p} - \mathbf{b} + \varphi\mathbf{b})\right)^2}}{\partial\varphi} = \frac{\mathbf{b}^T V(\mathbf{p} - \mathbf{b})}{\sigma(\mathbf{a})} = \frac{(\beta - 1)\mathbf{b}^T\widehat{\mathbf{V}}\mathbf{b}}{\sigma(\mathbf{a})}$$

3.7 The portfolio's overall active contribution to tracking error from its beta is therefore:

$$C_\beta = (\beta - 1)M_\beta = \frac{(\beta - 1)^2\mathbf{b}^T\widehat{\mathbf{V}}\mathbf{b}}{\sigma(\mathbf{a})}$$

3.8 We can apportion the marginal contribution to tracking error from the portfolio beta across individual securities using, say, $m_{\beta,i} = \beta_i M_\beta$ and thus identify the 'residual' (non-beta) element of each security's marginal contribution to tracking error as, say, $m_{r,i} = m_i - m_{\beta,i}$. The active 'residual' (non-beta) contribution to tracking error by security would then be $c_{r,i} = a_i m_{r,i}$.

References

[RiskAttributionTheoryRefs]

Heywood, G.C., Marsland, J.R. and Morrison, G.M. (2003). Practical risk management for equity portfolio managers. *British Actuarial Journal*, **8**, pp. 1061-1123

Kemp, M.H.D. (2005). Risk Management in a Fair Valuation World. *British Actuarial Journal*, **11**, No. 4, pp. 595-712

Kemp, M.H.D. (2009). *Market consistency: Model calibration in imperfect markets*. John Wiley & Sons [for further information on this book please see MarketConsistency]