

The Power of Scenarios

Risk measurement and the ranking of investment strategies

ANDREW AZIZ, ALGORITHMICS

Risk measurement and the legacy of Markowitz

Risk measurement best practice has followed closely the traditions of classical financial theory. A portfolio distribution of returns can be determined by 'mapping' portfolio holdings into some assumed distributional structure of risk factors. Risk is measured by a descriptor of the portfolio distribution that captures the extent of negative outcomes.

In the pioneering work of Markowitz, equities comprised the entire risk factor universe under an assumption of joint-normality – meaning that individual equity variances, and their correlations with each other, were all that was needed to fully define the distributional structure. The elegance of this model is that any portfolio of equities must have normally distributed returns that can be defined fully by their mean and variance. Risk becomes a function of a single parameter, variance, which captures all relevant information contained in the distribution.

The advent of value-at-risk (VaR) as the primary reporting measure for financial institutions borrowed liberally from this tradition, with the key advantage being that a parametric approach of this nature was easy to implement – the task was reduced to data acquisition and matrix multiplication. However, by extending the all-equity world of Markowitz to a multi-asset class world, two major issues arose – firstly, the risk factor universe typically extended beyond equities to include risk factors that fit less well with the assumption of joint-normality – and secondly, in many cases these additional risk factors were not observable securities like equities but abstract securities derived from observable prices.

The important implication of this second issue was that an additional mapping exercise became necessary to map actual securities into the derived risk factors – prior to the mapping of portfolio holdings into the actual securities. When the intermediate mapping could be accomplished in a linear manner – and when the risk factors were jointly normal – all the conclusions based on the Markowitz model still held and VaR, as a function of variance, became a good measure of risk. However, in too many cases, the relationship between actual securities and risk factors was non-linear and, as a consequence, the parametric approach often broke down.

The power of scenarios

When the composition of a portfolio is such that a parametric approach proves inappropriate, scenario-based simulation emerged as the methodology of choice. Under this approach, a set of scenarios is generated with each individual scenario providing a joint realisation of risk factors such that, across the full set of scenarios, some desired distributional

structure is modeled – unconstrained by the assumption of joint normality. Each security is then simulated across the full set of scenarios utilising some chosen valuation model that is a function of the input risk factors – either closed-form or stochastic – and unconstrained by the assumption of linearity. Portfolio holdings can now be mapped linearly into the individual security valuations to produce a set of portfolio values or returns across scenarios. With the relaxation of the assumptions underlying the parametric approach, the resulting empirical distribution of portfolio returns is now unconstrained. A non-parametric approach of this nature is effectively agnostic to risk factor type, asset class and investment strategy.

Until recently, the area where best practice had not kept up was in the adoption of an appropriate risk measure for non-normal distributions. On the one hand, a non-parametric version of VaR can be calculated for any empirical distribution – capturing the skewness and kurtosis of asymmetric portfolio returns. On the other hand, as this risk measure is no longer a function of a single parameter, important information contained in the individual scenarios is lost – the phenomenon of 'fat-tails' being a good example of this.

Increasingly popular risk measures such as 'expected shortfall' are designed to address this issue by capturing the expected loss in the tail of a portfolio distribution instead of focusing on a single confidence interval metric. These types of measures have an additional advantage over VaR of being coherent – meaning that the sum of the risk measures of child portfolios can never exceed that of the parent portfolio.

Another risk measure, 'downside regret', goes even further. Within each scenario, the return of a given portfolio deviates positively or negatively with respect to a benchmark – defined either as some fixed threshold or as another portfolio itself valued across the set of scenarios. Regret is defined as the expected deviation across scenarios – with deviation in each scenario expressed in absolute value terms. Downside regret is then simply the expectation across only the negative deviation scenarios.

Portfolio construction

Portfolio replication – how to create a portfolio to mimic the behaviour of some other portfolio or index – has recently become a hot topic. Once again, scenario-based simulation provides a powerful framework for addressing this task. Most portfolio optimisation techniques focus on an objective of minimising risk – traditionally variance but, increasingly, measures like expected shortfall – subject to a given level of return. Neither of these approaches, however, uses all the information

available in the full set of scenarios. In contrast, by identifying the portfolio to be replicated as the benchmark and focusing on an objective of minimising 'regret' – subject to an initial investment constraint – a best fit portfolio can be found that mimics the benchmark's behaviour across each and every scenario. Additional constraints can then be incorporated to place bounds on trades (long, short, gross or net), transactions costs and other portfolio measures. A multitude of applications are possible within this general model – replicating a given fund by removing certain undesirable components – replicating a benchmark where the underlying components are unknown – creating a synthetic exotic with a portfolio of vanillas – and so on.

Being even more adventurous, the replication problem can also be extended to a multi-time step world, where the objective becomes the minimisation of a weighted sum of the regret for each time step. Given a benchmark portfolio that may be rebalanced over time, this optimisation model has the unique ability to find the best 'static' portfolio to replicate a dynamic investment strategy.

Ranking investment strategies

Traditional performance measures such as the Sharpe ratio arose from the parametric legacy and are useful to the extent that a portfolio distribution can be fully defined by its mean and variance. When distributions are not normal, however, important information contained in the scenarios is once again lost – in the upside now as well as in the downside.

Again, regret provides an ideal framework for ranking portfolios on a risk-adjusted basis – long only or long/short – fully independent of investment strategy. A new performance measure, defined by the ratio of 'upside regret' to 'downside regret', captures all information contained in the scenarios and applies to all portfolio distributions. And as the benchmark can be a fixed threshold as well as some other portfolio, the increasingly popular omega ratio becomes a special case of this very general approach.

The bottom line is straightforward – if it is necessary to use scenario-based simulation to measure risk for asymmetric portfolio distributions, then it only makes sense to use all the available information contained in those scenarios for risk measurement, portfolio optimisation and performance measurement purposes. **THFJ**

Dr. Andrew Aziz is Executive Vice President of Risk & Capital Analytics at Algorithmics. He is also responsible for the newly launched, Algo Risk Service, a managed service offering targeted towards Hedge Funds, Asset Managers and Prime Brokers.