Risk Measures and Capital Allocation

The tail risk measures VaR and TVaR are well known and are useful for benchmarking capital vs. tail risk. It is informative, for instance, to know that capital is x times $\text{VaR}_{99.6\%}$ or y times $\text{TVaR}_{99\%}$. However when capital allocation is used to compute return on allocated capital, the tail risk is not the only relevant risk, and risk below the tail should not be ignored. Capital is held to cover all losses, not just tail losses, and there is cost to bearing risk of any loss, not just the extreme losses. Thus for capital allocation purposes, tail risk alone is not a satisfactory basis.

The first and simplest alternative in this direction is to use TVaR at a lower percent, like 60%. The trouble with this is that TVaR is linear (in fact just the average) for losses above the threshold, which does not adequately capture the nature and value of risk. To do this, other risk measures are needed. A few are outlined below, then allocation methods are addressed.

Risk Measures For Capital Allocation

As the above comments suggest, what is needed are risk measures that give some recognition to all scenarios with subpar return, but that give higher weight to the more adverse scenarios. For illustration purposes it is assumed that there is a 55% chance that profit will meet target, so the 55^{th} percentile and above give problematic results. The user can change 55% to the relevant number for each company.

Multiple TVaR

An easy though ad hoc way to include all losses but weight the large losses more is to take an average, possibly weighted, of TVaR at a few levels. For example, TVaR at 55%, 80%, and 99% could be averaged. Since TVaR includes all losses above the threshold, the losses above the 80^{th} percentile would be counted double, and those above 99% triple. There are a number of options the user can select, with some standard defaults. This accomplishes the stated goals, but has artificial break points where the weight jumps, and treats all losses above the highest level equally, which are both disadvantages of this measure.

Semi-standard deviation

The standard deviation counts favorable and unfavorable deviations equally, but semi-standard deviation only looks at the unfavorable deviations. It is the square root of the average of the square of the amount each loss exceeds the mean, taken over all losses that are above the mean. It thus gives recognition to all losses above the mean, and gives increasing weight to increasingly larger losses. This gives continuously increasing weights to more adverse scenarios without the jumps and cap of multiple TVaR, but it has the weakness of being quadratic, where actual risk pricing seems to give more than quadratic weight in the tail.

Risk-adjusted TVaR (RTVaR)

TVaR is the average of losses above the threshold. RTVaR is TVaR plus a percentage of the standard deviation of losses above the threshold. If 30% is selected, RTVaR at 55% would be TVaR at 55% plus 30% of the standard deviation of losses above the 55^{th} percentile. Since these are larger losses already, the straight standard deviation of these tail losses is used, not the semi-standard deviation. RTVaR can be used down to a lower level like the 55^{th} percentile and still give increasing weight to the larger losses. 30% is a reasonable weight as it has been used in standard deviation pricing. Still this risk measure has the weakness of being quadratic in loss size, which may not give enough weight in the extreme tail.

Distorted probability measures

Risk-pricing theory has come to the conclusion that the cost of risk can best be quantified by taking the mean under distorted probabilities. Actually VaR and TVaR are already distorted probability measures. VaR can be considered to be the mean loss where 100% of the probability is put at the threshold. TVaR is the mean loss where zero probability is given to the losses below the threshold and the probabilities above the threshold are scaled uniformly to total 100%. However as mentioned above, these are not satisfactory risk pricing measures, as they ignore some risk that should not be given away for free. Two probability transforms that have been found to correspond to market risk pricing in some cases are the Wang transform and the Esscher transform.

Esscher transform

This is the mean under a certain set of transformed probabilities, and is often advocated on the grounds that these transformed probabilities are closest, among transforms with the same mean, to the original probabilities in a measure of information distance (entropy). To define the transform, assume that there are 100,000 simulations, each with probability 1/100,000. Those probabilities will be transformed to give more weight to the larger losses. It starts with a parameter C which is on the order of the 1:250 scenario, but which can be fine-tuned to meet specific criteria. If Y is the loss size (for now insurance losses, although possibly negative profit would work), first an intermediate value K is needed, which is the average of exp(Y/C). Then a scenario with loss Y is given the new probability exp(Y/C)/(100,000K). These will sum to 1.00 over the scenarios, so are legitimate probabilities. Then the Esscher transform is the mean of the losses under the new probabilities. The value C can be fine tuned to make the mean a desired amount, such as the original mean plus the expected profit for the company. This will be useful in capital allocation later to measure risk-adjusted return by business unit. The Esscher transform tends to give more than quadratic weight in the tail, so overcomes that problem of the semi-standard deviation and RTVaR.

Wang transform

Again this is the mean under transformed probabilities, but now the transform is defined on the distribution function, not the density. The cumulative probability p is transformed to a new cumulative probability q by a transform with two parameters ν and λ. Let T_ν denote the t-distribution with ν degrees of freedom, and Φ denote the standard normal distribution, which is also T_∞, which is an allowed degrees of freedom. Then the Wang transform is $q = T_\nu(\Phi^{-1}(p) - \lambda)$. Here $\Phi^{-1}(p)$ is the standard normal p-percentile, usually a number between -4 and 4 but sometimes a bit outside this range when p is very near 0 or 1. Then subtracting λ gives a lower percentile, corresponding to a lower probability. Having $q < p$ means that more probability is going to the higher values. Using lower ν pushes more probability into the extreme tails. A value of ν around 5.5 has been found to correspond to market prices, so this can be fixed and the transform calibrated to target prices using λ. Scenario probabilities can be computed as differences of the transformed cumulative probabilities. The t-distribution with fractional degrees of freedom can be calculated using the beta distribution. For instance in Excel this would be $T_\nu(x) = ½ + ½\ \text{sign}(x)\text{betadist}[x^2/(\nu+x^2), ½, \nu/2]$.

Unlike the Esscher transform, the Wang probabilities do not use the simulation values, just their ranks. With 100,000 simulations and $\nu = \infty$, $\lambda = 0.4$, the two largest simulations get probabilities about 0.00005 and 0.00006. With $\nu = 5.5$ these increase to 0.00036 and 0.00216.

Allocation Methods

Proportional allocation is done by computing the risk measure for each business unit and in total, and allocating proportionally. The Shapley method is a game theory method that looks at the risk of each unit proportionally, averaged across all possible coalitions of business units. Another method described below is marginal allocation. It calculates the contribution of each unit to the risk measure for the whole company, and the contributions add up to the whole risk measure. It only works this way for homogenous risk measures, which here basically means risk measures that can be quantified in monetary terms. This includes VaR, TVaR, and all the measures discussed above. It would not include variance, which is quantified in squared dollars.

For homogeneous measures, the marginal allocation agrees with a more general game-theory method called Aumann-Shapley that also works for non-homogenous measures and is used in cost accounting. It is also called the Euler method, as the additivity for homogenous measures follows from a theorem of Euler.

An advantage of marginal allocation is that it is consistent with the economic principle of pricing in accord with marginal cost. It has another advantage called suitability. If a business unit with lower than average return on capital is reduced a bit proportionally, perhaps by quota share reinsurance, the capital need of the company is also reduced a bit. Doing this should increase the return on capital for the company as a whole, and it will when the allocation is by marginal allocation, but may not otherwise. Thus marginal allocation is the only way to identify business units that are increasing or decreasing overall company return.

In mathematical terms, marginal allocation is the derivative of the risk measure with respect to the volume of each business unit. By Euler's theorem, these allocations add up to the overall risk measure for homogenous risk measures.

The derivative has been worked out for many risk measures. For distorted means, the allocation of the risk measure is just the distorted mean for each business unit, using the scenario probabilities for the whole company. When the random variable is incurred losses, or loss plus allocated loss adjustment expense, the Wang and Esscher transforms can be calibrated so that the risk measures they produce are mean loss (maybe plus ALAE) plus expected profit for the company. Using the resulting scenario probabilities for each business unit gives the mean plus risk-adjusted profit for the units, calibrated to the risk and profit of the company as a whole. Then capital can be allocated in proportion to the profit portion, which would make the return on capital for each unit a proper risk-adjusted return. Instead of calibrating to mean profit, an alternative might be to calibrate to target profit.

This works for TVaR as well – that is, the allocation of TVaR is the average of the business unit losses in the scenarios where the company loss is above the threshold. In theory it works for VaR but in practice VaR is usually the value at a single simulation, and the business unit values at that simulation are too random to use. An alternative is to allocate "blurred VaR" which is the average, possibly weighted by distance from the threshold, of a range of simulations around the VaR value. The allocation of capital is not always the allocation of the risk measure, however, as discussed in more detail below.

For standard deviation, the marginal allocation for a business unit with losses X, with Y being the company losses, is cov(X,Y)/stdev(Y). Those covariances add up across the business units to be the variance of Y, so the allocation adds up to the standard deviation.

For RTVaR the allocation is the allocated TVaR plus the percentage of the allocated standard deviation. I.e., all the probability is assigned to the scenarios above the company TVaR threshold, and the new mean is calculated for each business unit using the overall company excess probabilities to give the TVaR allocation. After this, the covariance of each unit with the company is calculated across the excess scenarios to allocate the standard deviation portion. Similarly, for multiple TVaR, each individual TVaR is allocated to business unit, then these are weighted together.

Allocation of capital vs. allocation of risk measure

Capital should not be allocated by rote in proportion to the allocation of the risk measure. In large part the difference arises from what random variable the risk measure is applied to. Often the risk measure is applied to the insurance claims incurred, and this risk measure is allocated to business unit. In that case, the expected claims are supposed to be covered by the premium, so capital is needed for the claims in excess of the mean. What would make sense then would be to subtract the mean of the company and each business unit from their allocated risk measures, and allocate capital in proportion to the results. In the marginal case, both the allocated risk measure and the mean would add up to those for the company, so the differences would add up too. This would work for all the risk measures except semi-standard deviation, which is not designed to be more than the mean and is in fact the same if you subtract mean claims.

Where it gets more complicated is if the random variable is underwriting profit or net profit including underwriting and investments. In those cases, usually the negative of profit is the random variable used, as most of the risk measures are designed to attribute more risk to higher values. For the tail measures, the risk measure itself can now be used as the basis of allocation. For the Esscher and Wang transforms, however, there is a question of how to calibrate the free parameter in this case. What the theory says is that the transformed mean for the whole company should be zero. Since the mean of negative profit should be negative, the transform gets to zero by increasing the weights on the unfavorable outcomes. The way to convert this to a risk measure is to subtract the actual mean from the transformed mean. Again the allocation is by the mean of the unit negative profit using the transformed probabilities from the company as a whole less the actual mean of the unit. This will give more risk and thus more capital to the units that have bad results when the whole company has bad results.

This can be done on underwriting profit, but that does not capture the difference in target underwriting profit due to investment income differences. Perhaps a better method is to first discount the losses of each business unit in each scenario, using either simulated or target investment returns.