


Guy Carpenter Views

Reprinted from www.guycarp.com

August 2003

Allocating Capital By Risk Measures



A Systematic Survey

By Gary G. Venter and John A. Major

Analysts have proposed methods of better addressing the risk nature of allocating capital.

Almost all large corporations face decisions on capital allocations. By correctly allocating capital to each line of business, a company is greatly empowered to make the best strategic decisions possible.

Purpose Of Insurance Company Capital Allocation

In property/casualty insurance, the capital allocated to a line of business is designed to act as a cushion against unexpected losses. In the past, insurers used a fairly straightforward method of providing this cushion: The bigger the line of business in terms of net written premium, the greater the amount of capital allocated to that line. For example, if personal auto insurance constituted 70 percent of a book of business, as measured by net written premium, the company assigned 70 percent of its capital to that line and calculated metrics used for policy guidance based on that allocation.

Drawbacks Of Premium-Based Allocation Method

This method has its drawbacks. Although it recognizes that some lines are bigger than others in terms of net premium written, it does not take into account the fact that two lines of insurance writing the same amount of business may face very different exposures to unexpected losses.

For example, personal auto and medical malpractice may each comprise 50 percent of a company's written premium. However, personal auto is a stable line of insurance, and its loss profile can generally be gauged with reasonable accuracy. Medical malpractice, on the other hand, is a notoriously volatile line and often incurs losses that could not have been anticipated. Therefore, the capital cushion provided to the latter line of business should be significantly higher than that reserved for the former.

Risk-Based Capital Allocation

Given the drawbacks of this historical method of allocating capital, analysts have proposed methods of better addressing the risk nature of this process.

There are a variety of risk measures that can be used in capital allocation, each of which provides a unique way of characterizing the probability distribution of specific financial measures (returns, profit, losses, etc.) and representing them with a single number. Capital allocation is generally not an end in itself, but rather an intermediate step in a decision-making process. Trying to determine which business units are most profitable relative to the risk they bear is a typical example. (That assessment could be used in further decisions, such as compensation.) Pricing for risk is another example. Determining which units could best be grown to add value to the firm is a third.

Having allocated capital, computing return-on-capital (divide the unit's profits by that capital) is the natural second step. Of course if profit were negative, you would not need to divide by anything to know it is not sufficient. But return-on-capital thinking hopes to be able to distinguish the profitable-but-not-enough-so units from the real value-adders.

For growth decisions, the question is subtly different – not how much capital a business unit uses, but how much more is needed to support the target growth. In general it will be profitable to grow if the additional return exceeds the cost of the additional capital.

Risk Measures

Table 1 lists a number of risk measures that could be used in capital allocation. These measures will be discussed in turn. Each risk measure provides a different way of characterizing the probability distribution of outcomes of some financial measure (returns, profit, losses, etc.) and representing it with a single number. It should be remembered that the full information is contained in the probability distribution and that a risk measure collapses this information down to one number by taking a particular perspective on what is important about the distribution.

Table 1: Risk Measures

- Mean (μ)
- Variance (σ^2)
- Standard Deviation (σ)
- Semi-variance
- Risk of Ruin
- Value at Risk (VaR) or Probable Maximum Loss (PML)
- Expected Policyholder Deficit (EPD)
- Tail Value at Risk (TVaR)
- Excess Tail Value at Risk (XTVaR)
- Cost of Default Option
- Mean of Transformed Loss

In the following, some mathematical notation will be used to aid the precision of the discussion. Let X represent the insured

losses random variable (in another application, it could be another financial variable such as income, net profits, etc.) with $F(x)$ its cumulative distribution function. The notation $E[\textit{something}]$ is used for mathematical expectation with respect to the distribution F . The notation $E[\textit{something}|\textit{condition}]$ is used for conditional expectation, that is, mathematical expectation with respect to the related distribution that is restricted to those outcomes where the condition is true.

Mean (μ) The mean (average) outcome, $E[X]$, is the measure most familiar to us. Often, it is not considered a risk measure because that term is reserved for measures that go beyond the mean in specifically addressing adverse possibilities. The mean simply averages good and bad outcomes together. The mean is often used as a measure of reward.

Variance (σ^2) The variance is the average of the squared deviations from the mean, $E[(X-\mu)^2]$. It measures how far the possible outcomes spread out from the mean.

Standard Deviation (σ) This is the square root of the variance. It is useful because it is denominated in the same units as the mean, viz., dollars (instead of “dollars squared” for the variance).

Semi-variance Variance does not distinguish between upward and downward deviations, and so could provide a distorted view of risk when “upside” and “downside” risks are not considered equally important – which is usually the case. Semi-variance looks only at adverse deviations. Technically, it is $E[(X-\mu)^2|X>\mu]$ if the unfavorable outcomes are represented by X having higher values (e.g., X =losses). The inequality is reversed otherwise (e.g., X =profit).

Risk of Ruin In most financial risk assessment settings, there is the possibility of catastrophically bad outcomes, that is, outcomes which would lead to the ruin (insolvency) of the firm or cause a sudden and adverse qualitative shift in the status of the business unit. This measure is the probability of such an outcome. If values $X>r$ represent ruinous outcomes, then the risk of ruin is $1-F(r)$.

Value at Risk (VaR) or Probable Maximum Loss (PML) VaR is very familiar to the banking industry because of its endorsement by the Basle Committee [1996]. It is entirely equivalent to PML which is perhaps more familiar to property-casualty insurers. Under either guise, it is the mirror image of risk of ruin. Rather than specify a threshold value and measure the probability as in risk of ruin, VaR specifies the probability and measures the corresponding threshold value. For instance, the value at risk for a company’s losses might be the losses it would experience in the worst year in 10,000. Technically, with a specified probability level q (e.g., 0.01%), the VaR is the value¹ x_0 satisfying $1-F(x_0)=q$.

Expected Policyholder Deficit (EPD) EPD is also closely related to risk of ruin and VaR. It is the expected value of default amounts. If C is the firm’s capital available to fund losses, then outcomes $X>C$ represent insolvency, and the EPD is the average value of the shortfall, $E[(X-C)^+]$. Here the notation $(value)^+$ means the value if it is positive, otherwise zero. Of course, EPD can

¹ Technical note for quibblers – it may be the case that the distribution is not smooth and there is no value satisfying $F(x)=1-q$, or there may be “flat spots” and multiple values satisfying the equation. In such cases, the VaR is the least upper bound of the values x satisfying $F(x)\leq 1-q$.

be generalized to consider other variables and the expected deficit beyond some specified level other than default.

Tail Value at Risk (TVaR) Tail value at risk is closely related to EPD. It is the expected loss in the event that losses breach the value-at-risk target. If the specified probability is q and the corresponding threshold is r , i.e. $1-F(r)=q$, then TVaR is the conditional expectation $E[X|F(X)>1-q]$ which happens to equal $E[X|X>r]$ and also $r+EPD/q$. Algebra aside, there is a subtle distinction between the two concepts. EPD is specified by a dollar threshold; TVaR by a probability level. The linkage via $1-F(r)=q$ is not fixed if other aspects of the problem are changed. For example, in portfolio management, considering how changes in the portfolio affect the EPD or the TVaR will generally cause the relationship between them to change as well.

Excess Tail Value at Risk (XTVaR) XTVaR is similar to TVaR, but rather than the average of straight outcomes, it is the average excess beyond the overall mean, i.e., $E[X-\mu|X>r]$. While this is algebraically the same as $TVaR-\mu$, a practical difference emerges when allocating the risk measure down to business units, each with its own distinct μ_i .

Cost of Default Option A firm with limited liability does not pay once its capital is exhausted. In effect, the insurer holds an option to put the default costs to the policyholders. The fair market value of this option can be used as a risk measure. Determining that fair market value is not a simple exercise in statistics like it is for EPD, however.

Mean of Transformed Loss This, too, is motivated from the financial theory of option pricing. There, one distinguishes the “objective” probability distribution of possible outcomes from the “subjective” or “implied” probability distribution that is consistent with observed prices. Technically, the procedure is to replace the probability distribution $F(x)$ with another distribution $G(x)$ and to take the risk measure as $E_G[X]$, the mean with respect to the new distribution. Shaun Wang [2002], for example, has shown how systematically distorting catastrophe loss distributions can lead to a good fit to observed reinsurance and cat bond pricing. Christofides [1998] discusses how the transformed loss approach can sidestep entirely the need for capital allocation.

Discussion

VaR and risk of ruin could be considered to reflect a shareholder’s viewpoint, because once capital is exhausted, the amount by which it has been exhausted is of no concern. EPD, default option cost, Tail VaR, and XTVaR relate more to the policyholder viewpoint, as they are sensitive to the degree of default. (This is not to imply that shareholders do not consider policyholder needs, too!) All of these measures ignore risks less severe than the critical probability selected. VaR also ignores more severe risk, while the tail measures evaluate that risk linearly, which might be considered an underweighting. Taking the mean of a transformed loss distribution takes all possible outcomes into consideration without treating them equally. It can get around some of the problems of the tail methods.

Allocation Methods

Often when allocating capital with a risk measure, the total capital is expressed as the risk measure for the entire company. For instance, the probability level can be found so that the Tail VaR for the company at that probability level is equal to the capital

carried. The capital could also be expressed as some multiple of the risk measure. For instance, the company could have a goal that the average loss in the 1-in-100 year or worse not use up more than premium plus $\frac{1}{3}$ of capital. This would make the capital goal three times the 1% XTVaR. This is consistent with the idea that renewal business has a value, so the goal should be to have enough capital to continue operating even in the identified adverse situation. Also, some amount of capital might be set aside as not being risk capital – it could be for acquisitions perhaps – and the remainder used to calibrate the risk measure. In any case, once the total capital has been associated with a risk measure, an allocation method can be applied to get that capital split to the business unit level by allocating the risk measurement. Several possible allocation methods are given in Table 2. Not all of these work with all of the risk measures.

Table 2: Allocation Methods

Proportional Spread
 Equalize Relative Risk
 Incremental Risk
 Marginal Risk
 Game Theory
 Myers-Read
 Co-measure

Proportional Spread This is the most direct method – apply the risk measure to each business unit and then allocate the total capital by the ratio of business unit risk measure to the sum of all the units’ risk measures. Usually the sum of the individual risks will be greater than the total risk, so this method is crediting each unit with a diversification benefit.

Equalizing Relative Risk This involves allocating capital so that each unit, when viewed as a separate company, has the same risk relative to expected losses. Applying this to the EPD measures, for instance, would allocate enough capital to each business unit to make the EPD for every unit the same percentage of expected loss.

Incremental Risk This measures the risk of the company with and without a specified business unit. The difference in required total capital is the incremental capital for the business unit. The total capital can then be allocated by the ratio of the business unit incremental capital to the sum of the incremental capital of all the units. This usually allocates more than the incremental capital to each unit.

Marginal Risk The marginal method is similar to the incremental method, but the change in capital is calculated for just the last small increment of expected loss for the unit, say the last dollar. Whatever reduction that is produced in the risk measure by eliminating one dollar of expected loss from the business unit is expressed as a capital reduction ratio (capital saved per dollar of expected loss). This ratio is applied to the entire unit to get its implied marginal capital to use in the allocation.

Game Theory This approach is another variant of the incremental approach, but the business units are allowed to form coalitions with each other. The incremental capital for a unit is calculated with respect to every group of units it could be a part of (not just

the totality of all others), and the results are averaged. This gets around one objection to incremental allocation – that the sum of incremental capital requirements does not equal the total capital required. With the game theory method, it does. This method is sometimes called the Shapley method because the calculated value is known as the “Shapley value” in game theory.²

Myers-Read The Myers-Read method also uses marginal allocation. It sets the marginal capital needed to support an exposure increase equal to the additional capital it would take to make the cost of the default option, as a percentage of expected losses, the same before and after. It has the advantage over other marginal methods that the marginal increments add up to the total capital. This method is discussed in detail in Butsic [1999] and Myers and Read [2001].

Co-measures These were introduced by Rodney Kreps [2003] as a way of allocating capital in an additive manner that is nonetheless consistent with the overall risk measure used to define total capital. The procedure naturally generalizes to other risk measures the relationship that exists between covariance and variance.

It can be most easily thought of in terms of a scenario generator. Consider an example where the total capital requirement is set to be the tail value at risk at the 1-in-1000 probability level for losses. Then 0.1% of generated scenarios would have losses above that level. The co-Tail VaR for a business unit would just be the average of *its* losses in just those scenarios. That would be its contribution to the overall Tail VaR.

Like covariance, co-measures provide a totally additive allocation. Business units could be combined or subdivided in any way and the co-Tail VaRs would add up. For instance, all the lines of business could be allocated capital by co-Tail VaR, then each of these allocated down to state level, and those added up to get the state-by-state capital levels for all lines combined. This could be done by peril or other business categories as well.

Hallerbach [1999] uses the term “component VaR” for this concept in connection with VaR but does not discuss other risk measures as Kreps does.

Discussion

Allocating by marginal methods is accepted in financial theory. However, allocating more than the pure marginal capital to a unit could lead to pricing by a mixture of fixed and marginal capital costs, violating the marginal pricing principle. Even when the total capital is the sum of the marginal increments, as in Myers-Read, there is no tie-in between the capital allocated to a line and the value of its risk. Thus it would be a great coincidence if this allocated capital were right for a return-on-capital ranking.

The co-measure approach is consistent with the total risk measure and is completely additive. Thus if the risk measure gives the right capital need overall, the co-measure shows each line’s contribution to that. But it too could violate marginal pricing.

The allocation method in the end depends on why you are allocating capital. Allocating by a risk measure is straightforward but subjective. It appears to be appropriate for allocating frictional capital costs, which are proportional to capital, but not for return

² Lloyd S. Shapley (currently at UCLA) is one of the founders of multiplayer game theory.

on risk bearing, which might not be proportional. If it also allocates fixed costs, it could produce misleading indications of actual profitability prospects.

Strong candidates for risk-measure allocations are Myers-Read and co-XTVaR. Both start with reasonable stories of the overall capital need – enough to keep the default cost low for MR and enough to be able to continue writing after the very bad year for XTVaR. Then they both allocate all the capital in an additive manner that directly reflects the individual contributions to the overall capital need. The capital standard for MR sounds a little stronger in theory, but the computational aspects are harder than they might appear. The value of the put option involves calculations way out in the tail of a distribution whose tail is not known that precisely. XTVaR can use a capital standard for partial loss of surplus, which is more reliably modeled than default.

Note: An extended discussion of these and related issues written by Gary Venter will be published as part of the 2003 Bowles Symposium.

References

- Basle Committee on Banking Supervision (1996) “Amendment to the Capital Accord to Incorporate Market Risks,” January.
- Butsic, Robert J. (1999) “Capital Allocation for Property-Liability Insurers: A Catastrophe Reinsurance Application,” Forum, Casualty Actuarial Society, Spring.
- Christofides, Stavros (1998) “Pricing for Risk in Financial Transactions,” GISG/ASTIN Joint Meeting, Glasgow, October.
- Hallerbach, Winfried G. (1999) “Decomposing Portfolio Value-at-Risk: A General Analysis” Technical Report, Erasmus University Rotterdam and Tinbergen Institute Graduate School of Economics.
- Kreps, Rodney E. (2003) “Riskiness Leverage Models,” In strat working paper, to appear.
- Myers, Stewart C and Read, James A., (2001) “Capital Allocation for Insurance Companies,” Journal of Risk and Insurance, 68:4, 545-580.
- Wang, Shaun (2002) “A Universal Framework For Pricing Financial And Insurance Risks,” ASTIN Bulletin, Volume 32, No. 2.

Guy Carpenter & Company, Inc. (“Guy Carpenter”) provides this publication for general informational purposes only. The data contained herein is based on sources Guy Carpenter believes to be reliable, as of a specified date, but Guy Carpenter in no way guarantees its accuracy. Guy Carpenter makes no representations or warranties, express or implied, concerning the statements made herein or the financial condition or solvency of reinsurers. Past performance does not guarantee future outcome, and readers are cautioned against placing undue reliance upon any forward-looking statements. Guy Carpenter does not undertake any obligation to publicly update or revise the material herein, regardless of new data, future events or otherwise.

This document is not an offer to sell, or a solicitation of an offer to buy any financial instrument or reinsurance program. This report, its data and the methodology utilized are proprietary to Guy Carpenter. The dissemination, reproduction or use of this report without Guy Carpenter's express written permission is prohibited.