

Capital Allocation: An OPINIONATED Survey

GARY G VENTER, GUY CARPENTER INSTRAT

A number of methods of allocating capital to business unit, e.g., line of business, profit center, etc., are discussed. Goals of capital allocation include testing the profitability of business units and determining which units could best be grown to add value to the firm. Methods of approaching these questions without allocating capital are included in the discussion.

CAPITAL ALLOCATION: AN OPINIONATED SURVEY

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. Pricing for risk is another.

Return-on-capital thinking would look to allocate capital to each business unit, then divide the units' profits by that capital. Of course if profit were negative, you would not need to divide by anything to know it is not sufficient. But this approach would hope to be able to distinguish the profitable-but-not-enough-so units from the real value-adders.

The same issue can be approached without allocating capital, using a theory of market risk pricing. The actual pricing achieved by each business unit can be compared to the risk price needed. This would require having a good theory of risk pricing, where the previous approach would depend on having a good theory of capital allocation. Since both are addressing the same decisions, both will be included in this survey. For those who like to phrase the issue as one of return on capital, the pricing method can be put into allocation terminology after the fact by allocating capital to equalize the ratio of target return to capital across business units.

Rating business units by adequacy of return is not necessarily the final purpose of the exercise. The rating could be used in further decisions, such as compensation and strategies for growth. For strategic decisions another question is important – 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. In some cases a company might not need too much more than it already has for the target growth, in which case not much additional profit would be needed to make the growth worthwhile.

This is the marginal pricing approach, and is a basic tenet of financial analysis. It differs from capital allocation in that for marginal-cost pricing not all capital has to be allocated to reach a decision. Only the cost of the capital needed to support the strategy has to be determined, to see if it is less than the profit anticipated. Methods of quantifying the cost of marginal capital will be reviewed here as well, as again this is aiming at answering the same strategic questions.

Finally, another way to determine which business units are adding most to the profitability of the firm is to compare the insurer to a leveraged investment fund. Sometimes this is called the cost-of-float approach. The overall return of the insurer can be evaluated by finding the borrowing rate that would equalize its risk and return after tax to a leveraged investment fund. If the fund would have to be able borrow significant funds at a particularly low rate of interest to match the insurer's risk and return, then the insurance business is clearly adding value. The business units can be ranked based on their impacts on this borrowing rate.

Thus while the general topic is capital allocation, this survey is looking at methods for answering questions that capital allocation is addressing. To summarize, four basic approaches will be reviewed:

1. Selecting a risk measure and an allocation method and using them to allocate capital
2. Comparing actual vs. model pricing by business unit
3. Computing the cost of the marginal capital needed for or released by target strategies
4. Evaluating profitability in comparison to a leveraged mutual fund

The time period for evaluation is an issue for all of these methods, and this is addressed in Appendix 1.

APPROACH 1 – ALLOCATING VIA A RISK MEASURE

Table 1 lists a number of risk measures that could be used in capital allocation. To summarize briefly, VaR, or value at risk, is a selected percentile of the distribution of outcomes. For instance, the value at risk for a company might be the losses it would experience in the worst year in 10,000.

Table 1: Risk Measures

<input type="checkbox"/> VaR
<input type="checkbox"/> EPD
<input type="checkbox"/> Tail VaR
<input type="checkbox"/> X TVaR
<input type="checkbox"/> Standard Deviation
<input type="checkbox"/> Variance
<input type="checkbox"/> Semi-Variance
<input type="checkbox"/> Cost of Default Option
<input type="checkbox"/> Mean of Transformed Loss

EPD is expected policyholder deficit, i.e., the expected value of default amounts. It can also be generalized to include the expected deficit beyond some level, rather than beyond default. If b is the target amount, the EPD beyond b is: $\Pr(X > b)E[(X - b) | X > b]$.

Tail value at risk is the expected losses in the event that losses exceed the value-at-risk target. If the target loss level is b , this is $E(X | X > b)$.

X TVaR is similar to Tail VaR, but rather than the mean of all cases over a level, it is the average for those cases of the excess of the losses over the overall mean, i.e., $E[X - m | X > b]$.

A company with limited liability does not pay once its capital is exhausted. So the insurer holds an option to put the default costs to the policyholders. The value of this option can be used as a risk measure. The other measures are standard statistical quantities.

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 the capital carried. The capital could also be expressed as a multiple of the risk measure. For instance, the company could have a goal that the average loss

■ Table 2: Allocation Methods

- Proportional Spread
- Marginal Analysis
- By business unit
- Incremental by business unit
- Game Theory
- Equalize Relative Risk
- Apply Co-Measure

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 99% X TVaR. 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.

Proportional spread 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.

Marginal analysis measures the risk of the company with and without a specified business unit. The difference in required total capital is then the marginal capital for the business unit. The total capital can then be allocated by the ratio of the business unit marginal capital to the sum of the marginal capital of all the units. This usually allocates more than the marginal capital to each unit. The incremental marginal method is similar, but the change in capital is calculated for just the last 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) and applied to the entire unit to get its implied incremental marginal capital to use in the allocation.

The game theory approach is another variant of the marginal approach, but the business units are allowed to form coalitions with each other. The marginal capital for a unit is calculated for every group of units it could be a part of, and these are averaged. This gets around one objection to marginal allocation – that it treats every unit as the last one in. This method is sometimes called the Shapley method after a founder of game theory.

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 put, 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 Appendix 2.

Equalizing relative risk 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 make the EPD for every unit the same percentage of expected loss.

Co-measures were introduced by Rodney Kreps as a way of allocating capital in an additive manner that is nonetheless consistent with the overall risk measure used to define total capital. Appendix 3 discusses these in greater detail. They can be most easily thought of in terms of a scenario generator. Take the case where the total capital requirement is set to be the tail value at risk at the 1-in-1000 probability level. Then in generating scenarios, about 1 in 1000 would be above that level. The co-Tail VaR for each

business unit would just be the average of its losses in those scenarios. This is its contribution to the overall Tail VaR.

Co-measures provide a totally additive allocation. Business units could be combined or subdivided in any way and the co-Tail VaR's 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 for peril or other business categories as well.

Commentary on Allocation by Risk Measures

VaR could be considered to be a shareholder viewpoint, as once capital is exhausted, the amount by which it has been exhausted is of no concern to them. EPD, default option cost, X TVaR, and Tail VaR relate more to the policyholder viewpoint, as they are sensitive to the degree of default. And indeed the shareholders might do well when they consider policyholder needs. All of these measures ignore risk below the critical probability selected. VaR also ignores risk above that level, while the tail measures evaluate that risk linearly, which many consider to be an underweighting.

Variance does not distinguish between upward and downward deviations, and so could provide a distorted view of risk when these directions are not symmetric – which is the usual case. Semi-variance looks only at adverse deviations, so accounts for this. Taking the mean of a transformed loss distribution is a risk measure aiming at quantifying the financial equivalent of a risky position, and it can get around the problems of the tail methods. More exploration of transformations could be useful.

Allocating by marginal methods is accepted in financial theory. However, allocating more than the pure marginal capital to a unit it 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.

Myers-Read was introduced as a method of allocating the frictional costs of holding capital. These are discussed more in Appendix 2, but as a definition I would propose that costs which arise from holding capital even if no risk is written are frictional costs. Corporate tax on investment income is an example. A more delicate issue is any lower investment income resulting from taking less investment risk in order to give policyholders greater security. I would hold that this is a frictional cost as well. Even though it results from the intent to sell insurance, this does not differentiate it from other frictional costs.

The return for actually putting the capital at risk is a different matter. This relates to the amount of risk taken, not the amount of capital allocated. In financial models beta is almost always a component of the return for bearing risk, but it is not generally a part of the frictional cost. Some actuarial pricing approaches have assumed that pricing to recoup frictional costs is sufficient, and this is encouraged by assertions that beta is zero for underwriting anyway. More recent theory, discussed below, shows that risk pricing is more than beta. This suggests that even if allocating capital by risk measure is sufficient for allocating frictional costs, there are other elements of return that will not be proportional to the amount of capital held and so should be measured in some other way.

APPROACH 2 – COMPARE ACTUAL VS. MODEL PRICING

A traditional use of capital allocation is to price business to equalize return on capital. However even if the allocation method is intuitively satisfying, there is no guarantee that such pricing would correspond to the market value of the risk transfer. If instead actual pricing were compared to value pricing, the profitability of business units could be evaluated without allocating capital at all (except to the degree this is necessary in the pricing to compute the frictional costs of holding capital). But for those who still prefer a single target return on capital, capital could be allocated after the pricing by equalizing the return on capital from the value prices.

This method requires an evaluation of the market value of the risk transfer provided. Financial methods for valuing risk transfer typically use transformations of the loss probabilities to risk-adjusted probabilities, with covariance loadings like CAPM being one special case. This is a fairly technical calculation and to date there is no universal agreement on how to do it. Some transforms do appear to give fairly good approximations to actual market prices, however. The Wang transform has been used successfully in several markets to approximate risk pricing. Finance professionals now appear to favor an adjusted CAPM approach that corrects many of the over-simplifications of the original formulation. For instance, a correlation with the insurer's own results may be as important as correlation with the market in determining the cost of risk transfer.

To use CAPM or similar methods, costs are first identified, then a risk adjustment added. Three elements of cost have been identified for this process: loss costs, expense costs, and the frictional costs of holding capital. The latter is not the same as the reward for bearing risk, which is separately incorporated in the risk adjustment.

The CAS Committee on the Theory of Risk is sponsoring the Risk Premium Project to look into how to do risk pricing right. Starting from CAPM, they are looking at are several considerations needed to get a realistic market value of risk transfer. Some issues in this area are:

- Company-specific risk needs to be incorporated, both for differential costs of retaining vs. raising capital¹ and for meeting customer security requirements.
- The estimation of beta itself is not an easy matter²
- Other factors besides beta are needed to account for actual risk pricing³
- To account for the heavy tail of P&C losses, some method is needed to go beyond variance and covariance^{4,5}
- Jump risk needs to be considered. Sudden jumps seem to be more expensive risks than continuous variability, possibly because they are more difficult to hedge by replication. Large jumps are an element of insurance risk, so need to be recognized in the pricing.

Commentary on Target Pricing

Measures of the market value of risk transfer are improving, and even though there is no universally accepted unique method, comparing actual profits to market-risk-model profits can be a useful evaluation. This can then be reformulated as a capital allo-

¹ Froot, Kenneth A. and Stein, Jeremy C., *A New Approach to Capital Budgeting for Financial Institutions*, Journal of Applied Corporate Finance, Summer 1998, Volume 11, Number 2

² Kaplan, Paul D. and Peterson, James D., *Full-Information Industry Betas* Financial Management 27 2 Summer 1998

³ Fama, Eugene F. and French, Kenneth R. *Multifactor Explanations of Asset Pricing Anomalies* Journal of Finance 51 1 March

⁴ Wang, Shaun *A Universal Framework For Pricing Financial And Insurance Risks*, ASTIN Bulletin, 2002, Volume 32, No. 2

⁵ Kozik, Thomas J. and Larson, Aaron M. *The N-Moment Insurance CAPM*, Proceedings of the Casualty Actuarial Society LXXXVIII, 2001

cation if so desired. The pricing can also be particularized to the company, considering that company costs of risk transfer may differ from the industry's. However the requisite pricing models are still under development.

APPROACH 3 – CALCULATING MARGINAL CAPITAL COSTS

A third approach to evaluating business unit profitability is to look at the last increment of business written by the unit to see whether the cost of the additional capital required is less than the profit it generates. This is not necessarily an allocation of capital, in that the sum of the marginal increments may not add up to the total capital cost of the firm, leaving some fixed capital not allocated. It does correspond, however, to the financial principle of marginal pricing. In basic terms, if the profit from adding an increment of business in a unit exceeds its marginal capital cost, then the unit should be expanded.

Because of the unallocated fixed capital charges, an anomalous situation could arise where each business unit is profitable enough on the margin but the firm is not so as a whole. In such cases further strategic analysis would be needed to reach an overall satisfactory position for the firm. One possibility might be to grow all the business units enough to cover the fixed charges. Another might be to look a merger possibilities.

One way to do the marginal calculation would be to set a risk requirement for overall capital, and then see how much incremental capital is needed to continue to meet this requirement after the small expansion of the unit. This is the same approach used in the incremental marginal capital allocation by risk measure, but there is no allocation. The cost of capital would be applied to the incremental capital and compared directly to the incremental expected profits.

Another way to calculate marginal capital costs is the options-

based method introduced by Merton and Perold. A business unit of an insurer could be regarded as a separate business operating without capital, but with a financial guarantee provided by the parent company. If the premium and investment income generated by the unit is not enough to pay the losses, the firm guarantees payment, up to its full capital. In return, if there are any profits, the firm gets them.

Both the value of the financial guarantee and the value of the profits can be estimated using option pricing techniques. The financial guarantee in effect gives the unit's policyholders an option that allows them to put any losses above the unit's premium and investment income to the firm. But this is not unlimited, due to the firm's limited resources, so the value of this guarantee is the difference between two put options: the option with a strike at losses equal to the sum of premium plus investment income, less the value of the insolvency put. The firm's call on the profits is a call option with strike of zero. If that is worth more than the financial guarantee provided, the business unit is adding value. These options would take some work to evaluate, however, in that the lognormal assumption of Black-Scholes would often be not sufficiently heavy-tailed. The options pricing could also reflect the specific cost to the firm of providing the guarantee, which would take into account guarantees provided to correlated business units. The managers of the unit could also be treated as having a contingent claim on the profits through incentive compensation.

Commentary on Marginal Capital Costs

This method directly evaluates marginal costs of decisions, so it can correctly assess their financial impact. If a large jump in business – upwards or downwards – is contemplated, the marginal impact of that entire package should be evaluated instead of the incremental marginals. There is still a potential arbitrary step of the criteria chosen for the aggregate capital standard, however. This is avoided in the financial guarantee approach, but that is

more difficult to calculate, in that some method of pricing heavy-tailed options would be required.

APPROACH 4 – MUTUAL FUND COMPARISON

An insurer can be viewed as a tax-disadvantaged leveraged mutual investment fund. It is tax-disadvantaged since a mutual fund does not usually have to pay tax on its earnings. It is leveraged in that it usually has more assets to invest than just its capital. An equivalent mutual fund can be defined as one that has the same capital and the same after-tax probability distribution of returns as the insurer. It can be specified by its borrowing rate, the amount borrowed, and the investment portfolio. This should provide enough variables to be able to find such a mutual fund. If there are more than one such, they could all be considered as strategic alternatives and the easiest one to create would be the equivalent.

The insurer can be evaluated by the equivalent borrowing rate. If the investors can duplicate the risk and return by not writing insurance but by borrowing at a high rate of interest, there is not much value in writing the insurance, as they could readily borrow the money instead. However if they have to be able to borrow at a very low rate to get an equivalent return, the insurer is producing a result that is not so easily replicated by a leveraged mutual fund.

This is first of all a method for evaluating the overall value added of the insurer, but it can be done excluding or adding a business unit or part of a business unit to see if doing so improves the comparison. If a business unit lowers the equivalent borrowing rate on the margin, making a loan more difficult to get by the equivalent mutual fund, it is increasing the value of the firm.

Commentary on Mutual Fund Comparison

This is a potentially useful analysis, but it would require modeling the distribution function of return for the entire firm, including all risk and return elements, and a potentially extensive search procedure for finding the equivalent mutual fund.

CONCLUSIONS

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 on risk bearing, which might not be. 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-X TVaR. 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 X TVaR. Then they both allocate all the capital in an additive manner which 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 involves calculations way out in the tail of a distribution whose tail is not known that precisely. X TVaR can use a capital standard for partial loss of surplus, which is more reliably modeled than default.

Pricing comparison is applicable to evaluating the actual realized pricing including frictional and risk transfer costs. However, it is only as good as the pricing model used, and that could be complicated.

The marginal cost method shows directly the impact of growing each business unit. It still requires a choice for the overall capital standard, unless the financial guarantee method is used, in which case it requires an appropriate option pricing formula.

The mutual fund comparison could be computationally intensive, but would provide qualitative insight into the value of the firm and its business units.

REFERENCES

Kreps, Rodney E, *Riskiness Leverage Models*, Instrat working paper, to appear.

Merton, Robert and Perold, Andre "Theory of Risk Capital in Financial Firms," *Journal of Applied Corporate Finance*, Fall 1993.

Myers, Stewart C and Read, James A. 2001, "Capital Allocation for Insurance Companies," *Journal of Risk and Insurance*, 68:4, 545-580.

APPENDIX 1: TIME FRAME FOR EVALUATION

Different business units will tend to pay their losses out over different time frames. This complicates the capital cost allocation issues. Generally speaking, capital will be needed to support reserves as they run off, and this should get into the allocation. More research would be useful to specify how to do this in each approach. An outline of some possibilities for this is below.

It is possible to quantify the remaining runoff risk for each year for each business unit. The years would be correlated, as issues in the claims environment could hit several years at once. Methods using risk measures could incorporate this runoff risk. To put the years together, a cost of capital could be applied to each year, and then discounted. Ongoing investment income on premiums not yet paid out could be discounted as well. This could be done historically on existing reserves or prospectively on the projected payout pattern.

Pricing transformations could use a similar approach. The adjusted probabilities for the cash flow stream could be transformed and discounted. One way Myers-Read could adapt to this is by considering a sequence of default put options – one at each year end as policies run off. These become increasingly more likely to be hit as the time frame expands. The prices of these options could be present-valued and summed up to get a total value of the default puts for current writings. Then for a small increase in writings in any business unit, the additional capital needed to keep this total put value constant, as a percent of expected losses, could be calculated and used as the basis of capital allocated to the unit. The marginal amounts seem likely to add up to total capital, as they would for each of the annual puts separately.

A similar method should work for pricing in the financial guarantee approach. The firm could be getting a sequence of call options

and providing a sequence of put options, whose total prices could be compared.

For the mutual fund comparison it would seem sufficient to look at the current annual risk to earnings including runoff risk for current liabilities. This would not be a totally prospective look at current strategies, but would still provide a valuable perspective on the financial status of the firm as it has been managed to date.

APPENDIX 2: THE MYERS-READ APPROACH

Myers-Read capital allocation presents a challenge to the classification of methods, in that it allocates all capital, it provides a marginal capital cost, and it can be used in pricing. But in the context of ranking returns, it is a risk-measure based method.

Butsic provides a slightly different derivation of the allocation formula than do Myers-Read themselves, and his approach is basically followed here, referred to as MR. You can get the same result from slightly different sets of assumptions, so this is not one of those situations where if you accept the assumptions you must accept the result. The results and assumptions can be evaluated from various viewpoints, and so the question is, does the whole approach work well?

The context for the method is that there are frictional costs to holding capital. In some countries, insurer investment income is subject to taxation, so tax is a frictional cost in those jurisdictions. Unless the insurer has really vast amounts of capital, it often has to invest more conservatively than the owners themselves would want to, due to the interests of policyholders, regulators, and rating agencies. There is a liquidity penalty as investors cannot get their investments out directly, and there are agency costs associated with holding large pools of capital, i.e., an additional cost corresponding to the reluctance of investors to let someone else control their funds, especially if that agent can pay itself from the results.

MR assumes a pricing approach in which the policyholders are charged for these frictional costs. This requires that the costs be allocated to the policyholders in some fashion, and MR uses capital allocation to do that. Every policyholder gets charged the same percentage of its allocated capital for the frictional costs. Thus it is really the frictional costs that are being allocated, and capital allo-

ation is a way to represent that cost allocation. The formula can be adapted to include in the premium other risk charges that are not proportional to capital, so this capital allocation does not necessarily provide a basis for a return-on-capital calculation.

A key element of the MR approach is the value of the default put option. As a company with limited liability, an insurer does not pay losses once its capital is exhausted. So it can be said that the insurer holds an option to put the default costs to the policyholders. MR assumes a lognormal or normal distribution for the insurer's entire loss portfolio, so can use the Black-Scholes options pricing formula to compute D , the value of this put option.

Adding a little bit of exposure in any policy or business unit has the potential to slightly increase the value of the default option. But adding a little more capital can bring the value of this option back to its original value, when expressed as a percentage of total expected losses. The MR method essentially allocates this additional capital to the additional exposure that required it.

In other words, the default option value, as a percentage of expected losses, i.e., D/L , is held as a fixed target, and the last dollar of each policy is charged with the amount of extra capital needed to maintain that target option value. But any dollar could be considered the last, so the whole policy is charged at the per dollar cost of the last dollar of expected loss. The beauty of the method is that those marginal capital allocations add up to the entire capital of the firm.

In the MR development, the total capital requirement of the firm is never specified, but could be taken to be the amount of capital needed to get D/L to a target value. The allocation method is the incremental marginal effect method – the incremental dollar loss for the business unit or policy is charged with the amount of capital needed to keep D/L at its target.

The total capital is the sum of the individual capital charges, i.e., $\sum c_i L_i = cL$, where $c_i L_i$ is the capital for the i th policy with expected losses L_i , and cL is total capital. Thus each policy's (or business unit's) capital is proportional to its expected losses, and the capital allocation question becomes how to determine the allocation factors c_i .

Formally, MR requires that the derivative of D with respect to L_i be equal to the target ratio D/L for every policy. Butsic shows that this condition follows from some standard capital market pricing assumptions. This requirement means that the marginal change in the default cost due to a dollar (i.e., fixed, small) change in any policy's expected losses is D/L . Thus D/L does not change with an incremental change in the expected losses of any policy. How is this possible? Because increasing L_i by a dollar increases capital by c_i , which is set to be enough to keep D/L constant when L_i increases. Thus the formal requirement that $\partial D / \partial L_i = D/L$ means that the change in $c_i L_i$ due to a small change in L_i has to be enough to keep D/L constant.

The question then is, can allocation factors c_i be found to satisfy both $\sum c_i L_i = cL$ and $\partial D / \partial L_i = D/L$? That is, can by-policy capital-to-expected-loss ratios be found so that any marginal increase in any policy's expected losses keeps D/L constant, while the marginal capital charges sum to the overall capital? The MR derivation says yes.

In the MR setup, after expenses and frictional costs, assets are just expected losses plus capital, and so the Black-Scholes formula gives:

$$D = L[N(y+v) - (1+c)N(y)]$$

where v is the volatility of company results, $y = -\ln(1+c)/v - v/2$ and $N(y)$ denotes the cumulative standard normal distribution.

Using this to expand the condition that $\partial D/\partial L_i = D/L$ requires the calculation of the partial of c w.r.t. L_i . Plugging in $\sum c_i L_i = cL$, this partial derivative turns out to be $(c_i - c)/L$. This leads to an expression for c_i in terms of c and some other things, which is the basis of the allocation of capital. This is how the condition on $\partial D/\partial L_i$ leads to an expression for c_i .

To express the allocation formula, denote the CV of losses as k_L and the CV of losses for the i th policy or business unit by k_i . Also define the policy beta as $b_i = \rho_{iL} k_i/k_L$, where ρ_{iL} is the correlation coefficient between policy i and total losses. Myers-Read also considers correlation of assets and losses, but Butsic gives the following simplified version of the capital allocation formula, assuming that the loss-asset correlation is zero:

$$c_i = c + (b_i - 1)Z, \text{ where } Z = (1+c)n(y)k_L^2/[N(y)v(1 + k_L^2)]$$

Butsic provides a simple example of this calculation. A company with three lines is assumed, with expected losses, CV's, and correlations as shown below. The total capital and its volatility are also givens. The rest of the table is calculated from those assumptions.

Changing the by-line expected losses in this example allows you to verify that if you add a dollar of expected losses to any of the lines, the overall D/L ratio is maintained by adding an amount to capital equal to the c_i ratio for that line.

Some aspects of the approach can be illuminated by varying some of the input assumptions. The examples that follow keep the volatility of assets constant, even though assets vary, which seems reasonable.

	line 1	line 2	line 3	total	volatilities
EL	500	400	100	1000	
CV	0.2	0.3	0.5	0.2119	0.2096
corr 1	1	0.75	0		
corr 2	0.75	1	0		
corr 3	0	0	1		
variance	10,000	14,400	2,500	44,900	
beta	0.8463	1.3029	0.5568		
capital	197.872	282.20	19.93	500	0.2209
assets				1500	0.0699
c_i:	0.3957	0.7055	0.1993	0.5	
- y:	1.9457807	y+v:	-1.7249		
N(y):	0.0258405	N(y+v):	0.042277		
n(y):	0.0600865	1/n(y):	16.64267		
Z:	0.6784		D/L:	0.0035159	

First, consider what happens if the CV for line 3 is set to zero. In this case, the line becomes a supplier of capital, not a user, in that it cannot collect more than it's mean, but it can get less, in the event of default. Then the capital charge c_i for this line becomes -17% , and the negative sign appears appropriate, given that the only risk is on the downside. The size of the coefficient seems surprising, however, in that its default cost is only 0.3% (which is the same for the other lines as well), but it gets a 17% credit. Part of what is happening is that adding independent exposures to a company will increase the default cost, but will decrease the D/L ratio, as the company becomes more stable. Thus in this case, increasing line 3's expected losses by a dollar decreases the capital needed to maintain the company's overall D/L ratio by 17 cents. This is the incremental marginal impact, but if line 3 decides to go net entirely, leaving only lines 1 and 2, the company will actually need $\$19.50$ in additional capital to keep the same default loss ratio. This is the entire marginal impact of the line, which will vary from the incremental marginal.

Another illustrative case is setting line 3's CV to 0.335. In this case, its needed capital is zero. Adding a dollar more of expected loss maintains the overall D/L ratio with no additional capital. The additional stability from its independent exposures exactly offsets its variability. Again the marginal impact is less than the overall: eliminating the line in this case would require \$10.60 in additional capital for the other lines.

The risk measure of the cost of the default option per dollar of expected loss, and the allocation principle that each dollar of expected loss be charged the frictional costs of the capital needed to maintain the target ratio, both appear reasonable, and the marginal costs adding up to the total eliminates the problem that fixed costs are being allocated using marginal costs. However, this is only so for incremental marginal costs. The marginal impacts of adding or eliminating large chunks of business can have a different effect than the incremental marginals, and so such proposals should be evaluated based on their total impacts.

Butsic also considers adding a risk load beyond the capital charge to the pricing. The same derivation flows through, just with expected losses replaced by loaded expected losses, and the capital charge set to c_i times the loaded losses. This provides a pricing formula that incorporates both risk load and frictional capital charges.

Using this, business unit results can be evaluated by comparing the actual pricing to the target pricing. If management wants to express this as a return on capital, the MR capital would not be appropriate. Rather the total capital should be re-allocated so that the ratio of modeled target profit to allocated capital is the same for each unit. Then comparing returns on capital would give the same evaluation as comparing profits to target profits. MR capital allocation would be the basis of allocating frictional capital costs, but not for calculating return on capital.

APPENDIX 3: CO-MEASURES

Co-measures can be defined for any risk measure that can be expressed as a conditional expectation, which is most of them. Suppose a risk measure for risk X with mean m can be defined as:

$R(X) = E[(X - am)g(x) | \text{condition}]$ for some value a and function g .

Suppose further that X is the sum of n portfolios X_i each with mean m_i . Then the co-measure for X_i is:

$$\text{co-R}(X_i) = E[(X_i - am_i)g(x) | \text{condition}]$$

Here the condition is the same as in the definition of R , so it is a condition on X , not X_i . Since expectations are additive, the sum of the co- R 's of the n X_i 's is $R(X)$.

Variance

As an example, take $a=1$ and $g(X) = X - m$, with any condition that is always fulfilled, like $0X=0$. Then $R(X)$ is the variance of X . Thus,

$\text{co-R}(X_i) = E[(X_i - m_i)(X - m)]$, which is the covariance of X_i with X .

Value at Risk

Value at risk at probability level q can be defined as:

$$E(X | F(X)=q)$$

This is just the q th quantile of the distribution. Then the co-VaR is:

$$E(X_i | F(X)=q)$$

This would be the average value of portfolio i when total losses are at the qth quantile.

Tail Value at Risk

For probability level q, take $a=0$ and $g(x) = 1$, with condition $F(X)>q$. If $q=99.9\%$, R is TVaR at the 1-in-1000 level. Then:

$$\text{co-TVaR}(X_i) = E[(X_i | F(X)>q)]$$

This is the mean loss for the ith unit in the case where total losses are over the qth quantile.

Expected Policyholder Deficit

As another example, consider the expected policyholder deficit, or EPD. If X is all years' losses unpaid, b is total assets, and $S(b)=1 - F(b)$, then:

$$\text{EPD} = E[(X - b)S(b) | X>b]$$

This is the R(X) form with $a = 1$, $g(x) = S(b)(X - b)/(X - m)$ and condition $X>b$. With these, the co-measure is:

$$\begin{aligned} \text{Co-EPD}(X_i) &= E[(X_i - m_i)g(X) | X>b] \\ &= E[S(b)(X - b)(X_i - m_i)/(X - m) | X>b] \end{aligned}$$

Each gets a fraction of the overall expected deficit given by the ratio of its losses above mean to the total losses above mean when there is a deficit.

Excess Tail Value at Risk

Define the measure excess tail value at risk by:

$$\begin{aligned} X \text{ TVaR}_q &= E[X - m | F(X)>q], \text{ so} \\ \text{Co-X TVaR}_q &= E[X_i - m_i | F(X)>q] \end{aligned}$$

If capital is set by X TVaR, it would provide enough to cover losses above mean losses for the average of the years in which

losses exceeded the q th quantile. The capital allocated by Co-X TVaR to a line would be the line's average losses above its mean losses in those same adverse years. There should be some probability level q for which X TVaR or a multiple of it makes sense as a capital standard, as the mean loss should be already collected in premium. Using co-X TVaR for allocation would not charge capital to a unit for its mean losses. If by some chance the unit did not have losses above its mean in the average of the scenarios above the q^{th} quantile for the entire company, it would not be charged any capital. This makes sense if capital is indeed being held for the adverse outcomes.