# *Discussion of "Capital Allocation for Insurance Companies" by Stewart C. Myers and James R. Read Jr.*

Gary G. Venter, FCAS, MAAA

## DISCUSSION OF
## "CAPITAL ALLOCATION FOR INSURANCE COMPANIES"

Gary G Venter, Guy Carpenter Instrat

"Capital Allocation for Insurance Companies[1]" is a useful and insightful paper for casualty actuaries. However it does not provide the denominator for a return-on-capital ranking of business units that many actuaries have sought. It does provide the basis for an alternative framework for evaluating business unit profitability.

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

## Discussion of "Capital Allocation for Insurance Companies"

My first introduction to the Myers-Read method was at a CAS session where Richard Derrig of the Massachusetts auto insurance bureau proclaimed "The capital allocation problem has finally been solved." Naturally I was glad to hear that, but as the session continued I began to suspect that he was talking about a different capital allocation problem than many actuaries had been addressing.

In the Massachusetts ratemaking scheme, insurers are permitted to incorporate into their rates a charge for the frictional costs of carrying capital. Since capital supports all lines of business, it is problematic how much of this cost can be attributed to any particular contract or even line of business. The Myers-Read approach does appear to provide an excellent methodology for this issue.

What a number of other actuaries have been seeking is a capital allocation to business unit in order to calculate the return on capital for each unit. This in turn would govern decisions as to which units to grow, shrink, revamp, drop, reward managers of, etc. I'll call this the problem of ranking profitability.

It now seems that the ranking problem and the allocation of frictional costs are distinct problems probably with different solutions. For instance, even in Massachusetts other elements of profitability are allowed into the ratemaking formula than just recovery of the frictional costs of carrying capital: carrying risk is rewarded within a CAPM framework over and above the frictional costs. This might be regarded by many actuaries as not much reward, but it illustrates that profit for bearing risk is not treated the same as recompense for frictional capital costs. In fact, the return for bearing risk is not even proportional to the allocated capital, indicating that the allocation is not intended to be the basis of a return calculation.

Nonetheless, I will argue later on that the risk pricing framework that Myers-Read presents does give a useful direction for solving the problem of ranking business units by profitability.

The remainder of this discussion has three main sections: the capital allocation problem, the Myers-Read solution, and an evaluation of applications and limitations.

## THE CAPITAL ALLOCATION PROBLEM

Initial actuarial approaches to capital allocation tend to allocate using some risk measure. The chosen risk measure is used to quantify the risk of the overall firm and each business unit. Then these risk measurements are combined in an allocation method to spread the capital to business unit. A simple example would be allocating capital in proportion to the variance of each unit's operating results.

There are numerous risk measures and allocation methods that can be used in this schemata. For examples, see the papers presented at the CAS

DFA seminar of June 2001. Many of the allocation methods look at marginal impact – i.e., the increase in the risk measure of the overall firm due to a given business unit, either in total or from its last small increment of exposure. The idea is to charge each unit only for the increase in capital it generates for the firm as a whole. It usually turns out that the sum of these marginal capital contributions is less than the entire capital, so the rest has to be allocated somehow.

Often the solution presented is to allocate the remaining capital in proportion to the marginal capital. But this could lead to inappropriate conclusions about business unit profitability. This is analogous to the problem of fixed and marginal production costs for a manufacturer, as illustrated in the following example.

Suppose a company has invested a lot of money in making an assembly line to produce hand phones. This line can produce phones for $2 each, but to recoup the investment costs the company wants to charge wholesalers $8 each. But suppose that after a while there is an oversupply in the market, and it can only charge $5 for each phone. Since each one costs only $2 to make, it decides to keep using the assembly line and keep selling phones. But if it required the fixed costs to also be covered, it would shut down, giving up the $3 per phone profit.

A similar situation can arise in insurance. If a line is generating enough profit to cover its marginal costs, including the marginal costs of capital, but not enough to cover some allocated fixed charges, allocating fixed capital in proportion to marginal could shut it down when in fact it is contributing to the overall profitability of the firm. Of course if every line is in this situation, the firm is going to have to find some strategy to cover its fixed costs, such as growing like mad, or merging, etc. This is a different problem that should not be buried in the by-line profitability analysis.

It should also be noted that there are other additive approaches to allocation of capital that do not use marginal methods. A general class of such methods is outlined in Rodney Kreps' widely circulated working paper, *An Allocatable Generic Risk Load Formulation*, which shows how to create co-measures, analogous to covariance for the variance measure, that are totally additive across any partitions of an insurer's portfolio. A related procedure has been introduced by D. Tasche, *Risk contributions and performance measurement*, Zentrum Mathematik (SCA), U München, Feb. 2000, www-m4.mathematik.tu-muenchen.de/m4/pers/tasche/riskcon.pdf .

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

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

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

$$CoR(X_i) = E[(X_i - am_i)g(x) \mid \text{condition}]$$

Since expectations are additive, the sum of the CoR's of the n $X_i$'s is R(X).

For example, define the measure excess tail value at risk by:

$$XTVaR_q = E[X - m \mid X>x_q] \text{ where } F(x_q) = q. \text{ Then}$$
$$Co\text{-}XTVaR_q = E[X_i - m_i \mid X>x_q]$$

If capital is set by XTVaR, it would provide enough to cover losses above mean losses for the average of the years where losses exceeded the qth quantile. The capital allocated by Co-XTVaR to a line would be the line's average losses above its mean losses in those same adverse years. A constant loss would get no allocated capital in this procedure, for instance.

One issue this highlights is the arbitrary choice of risk measure. Does the company really know how much capital it needs, and how each business unit affects that? With arbitrary choices of risk measure and allocation method, unit managers are going to push for those that make them look better, and there will be no solid foundation to settle the matter.

This argues for some other approach to the ranking problem than allocating by risk measure. Some alternatives will be discussed in the evaluation section. But first the Myers-Read solution is addressed. It is at least able to avoid the problem of allocating fixed capital in proportion to marginal, simply because their marginal capital adds up to the total!

## THE MYERS-READ SOLUTION

Robert Butsic in *Capital Allocation for Property-Liability Insurers: A Catastrophe Reinsurance Application*, www.casact.org/pubs/forum/99spforum/99spf001.pdf, provides an extensive discussion and application of Myers-Read (MR). Butsic provides a slightly different derivation of the allocation formula than do Myers and Read themselves. 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 method seeks to allocate the frictional costs of holding capital. What does that mean? Essentially frictional costs accrue just by a company holding capital, even if it doesn't put the capital at risk. The return for bearing risk is not a frictional cost, but a separate input into insurance pricing. Examples of frictional costs include taxation, agency costs, liquidity costs, and reduced investment opportunities, as detailed below.

In some countries, insurer investment income is subject to taxation, so tax is a frictional cost in those jurisdictions. But even on small islands where insurer investment income is not taxed, there are frictional costs of holding capital. Unless the insurer has really vast amounts of money, it often has to invest more conservatively than the capital owners themselves would want to, due to the interests of policyholders, regulators, and rating agencies. Thus the reduced investment income due to an insurer's reduced scope of investment alternatives is a frictional cost. There is also a liquidity penalty from insurers holding of capital, in that investors do not have direct access to the assets purchased. Further, 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 fund. All of these costs accrue to the insurer whether or not it bears any risk.

MR uses capital allocation to allocate the frictional costs to policyholders. Every policyholder gets charged the same percentage of its allocated capital for frictional costs. Thus it is really the frictional costs that are being allocated, and capital allocation is a way to represent that cost allocation.

A key element of the MR development is the value of the default put option. Assuming it is an entity 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. The distributional assumptions will be discussed further in the evaluation section.

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

In other words, the default option value, as a percentage of expected losses, i.e., D/L, for the entire firm 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 overall 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 really specified, but it could be taken to be the amount of capital needed to get D/L to some target value. In practice, whatever D/L ratio the firm has can be taken to be the target. The allocation method then is based on the incremental marginal effect – the incremental dollar expected loss for a policy is charged with the amount of capital needed to keep the overall D/L ratio at its target. The typical problem of capital allocation by marginal methods – that fixed costs are allocated in proportion to marginal costs – is avoided because, unlike most marginal allocation approaches, the marginal capital amounts add up to the total capital of the firm with no proportional adjustment. This appears to be due to the additive nature of option prices.

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 ith 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 proportionality 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 change in the firm's overall default cost due to a 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 one unit increases capital by $c_i$ units, and the $c_i$ is found that will keep D/L constant. Thus the formal requirement that $\partial D/\partial L_i = D/L$ means that $c_i$ is determined so that the change in $c_iL_i$, the policy's capital, due to a small change in $L_i$ has to be the amount that keeps D/L constant.

The question then is, can allocation factors $c_i$ be found to satisfy both conditions $\sum c_iL_i = cL$ 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. Without going into the details of their derivation, the following reasoning shows why it is feasible.

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

Using this formula to expand the condition that $\partial D/\partial L_i = D/L$ requires the calculation of the partial derivative of D, and thus eventually c, w.r.t. $L_i$. Plugging in $\sum c_i L_i = cL$, the c derivative turns out to be $(c_i - c)/L$. This leads to an equation for each $c_i$ in terms of c. Thus the two conditions required combine to give equations for c and all the $c_i$'s. The derivation then consists of finding a convenient solution.

To show the resulting allocation formula, denote the coefficient of variation (CV) of total losses as $k_L$ and the CV of losses for the ith policy or business unit by $k_i$. Also define the policy beta as $\beta_i = \rho_{iL} k_i/k_L$, where $\rho_{iL}$ 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:

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

Note that Z does not depend on i, so $c_i$ is a linear function of $\beta_i$. Butsic provides a simple example of this formula. A company with three lines is assumed, with expect 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.

| | 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:** | 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 | |

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.

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. However 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.

## EVALUATION: LIMITATIONS AND APPLICATIONS

The cost of the default option per dollar of expected loss seems to be a reasonable quantity to keep constant. If a policyholder increases this ratio by a change in exposure, that would reduce the value of the other policies, and so would be unfair to the other policyholders. Also the allocation principle that each dollar of expected loss be charged the frictional costs of the capital needed to maintain the target ratio also appears reasonable. And the fact that the marginal capital allocations add up to the total eliminates the problem of some other allocation methods that fixed costs are allocated using marginal costs. Thus all in all MR seems to be a good method of capital allocation. However, there are several issues that need to be addressed.

### *Lognormal Assumption*

First of all, aggregate losses are assumed to be lognormally distributed. This is required only for the total company, not for individual lines or policies. This may or may not be reasonable depending on the company being analyzed. It is an assumption many actuaries are comfortable making, but should be evaluated for specific applications. It would be possible to extend the MR derivation to other distributions, but that would require an analogue of the Black-Scholes formula. That in turn would need a probability transform to a risk-neutral measure. That is not necessarily difficult to achieve. In many cases the problem is not of finding a transform, but of choosing among a number of possible candidates. Pricing papers for such situations often pick a transform with little justification for the choice. It would be interesting to see how the experts would handle this problem in the insurance pricing case.

### *Return on Allocated Capital*

Second, it is clear from the Massachusetts auto context that the MR allocation was not intended to be the basis of a return-on-capital calculation, since other profit elements are added that are not proportional to the allocated capital. But would it be wrong to use the allocation for this? MR appears to be as good as any of the risk-measure allocations for coming up with a value for the capital required to support a line of business. But there is no theory to suggest that equalizing the return on this capital – or that from any other risk-measure's allocation – would produce appropriate by-line pricing. Butsic tested this for MR with a risk loading method, but didn't like the results. This could be a problem with the entire enterprise of allocating capital by a logical but arbitrary measure then pricing to equalize return on that capital.

### *Using Pricing Measures for Ranking by-Line Profit*

MR is aimed at capital allocation for pricing. The pricing that results, including the costs of risk-bearing as well as the frictional costs, can be used for ranking by comparing it to the actual profitability realized. This could be put into a return-on-allocated-capital mode by reallocating capital by the combined risk-friction profit load in the model pricing. Shaun Wang suggested using pricing methods like this in *A Universal Framework For Pricing Financial And Insurance Risks*, ASTIN Bulletin, 2002, Volume 32, No. 2.

Carrying this out in practice would require a good theory of insurance pricing. Many actuaries are skeptical of CAPM because it does not take into account all sources of risk. However further financial research is refining the original CAPM assumptions and developing broader-based pricing formulas. For instance, company-specific risk needs to be added to CAPM pricing, as shown in 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, pp. 59-69. The estimation of beta itself is still an unresolved issue, with a new approach offered by Kaplan, Paul D. and Peterson, James D., *Full-Information Industry Betas* Financial Management 27 2 Summer 1998. Also other factors besides beta are needed to account for actual risk pricing, as discussed in Fama, Eugene F. and French, Kenneth R. *Multifactor Explanations of Asset Pricing Anomalies* Journal of Finance 51 1 March. Also, to have pricing that will account for the heavy tail of P&C losses, some method is needed to go beyond variance and covariance, such in as Wang's article above, or Kozik, Thomas J. and Larson, Aaron M. *The N-Moment Insurance CAPM*, Proceedings of the Casualty Actuarial Society LXXXVIII, 2001. Finally, the pricing of jump risk needs to be considered. Models for pricing the default risk of corporate bonds incorporate a risk element for the possibility of sudden jumps. The same degree of variability seems to be more expensive as a sudden jump than as a continuous movement, possibly because it is more difficult to hedge by replication. Large jumps are an element of some insurance risk, so need to be recognized in the pricing.

Some of the above elements of a risk pricing formula are being studied by the CAS Risk Premium Project, which is using MR for the frictional capital part of risk pricing. With a good understanding of the value of risk-bearing, insurers will be armed with better tools for comparing actual profitability to a risk-based target.

### Other Methods for Ranking by-Line Profit

Return on capital allocated by risk measure and comparison to risk-based pricing are not the only alternatives for the profit ranking problem. Another is using pure marginal costs of capital without allocating fixed capital. This could be done with a risk measure for overall target company capital, or it could quantify the marginal cost of capital by the value of the financial guarantee provided by the firm to the customers of the business unit. This is an approach supported by the paper Merton, R. and Perold, A., *Theory of Risk Capital in Financial Firms,* Journal of Applied Corporate Finance, Fall 1993. The value of the financial guarantee could be priced as a put option, where the customers put all losses in excess of the net premium and investment income of the business unit to the overall firm (up to the assets of the firm – so its really the difference between two puts). This is the default put for the business unit as a separate entity with no capital, so it is a different order of magnitude than the default put for the whole firm that MR considers.

Another alternative method for ranking profitability is to create a model of a leveraged mutual investment fund that borrows enough money at the right interest rate and invests in the right way to have the same probability distribution of after-tax returns as does the insurer. The borrowing rate would be a key measure of the financial viability of the insurer. Then the marginal impact of each business unit on the borrowing rate can be found and used to rank the units.

These approaches are discussed further in my paper *Capital Allocation: An Opinionated Survey*, CAS Forum, to appear.

### *Time Frame*

All the losses outstanding for an insurer would be affected by a default, so several accident or policy years share in the default risk. This complicates the capital allocation problem. The charge for frictional capital costs for a given policy year might consist of shares of a series of put options over several years, where the share could be based on the portion of policy reserves (loss plus unearned premium) represented. The more future years would have costlier options due to the time element in the options pricing formula. In fact, options in practice are priced by assuming even greater volatility for the longer-term options, using smile tables. This would further increase the prices of capital for the later year reserves, and so would tend to increase the proportion of capital allocated to the longer-tailed lines.

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 hypothetical equivalent mutual fund, it would seem sufficient to look at the current annual risk 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.

## TO WRAP UP

The Myers-Read methodology appears to accomplish its aim – to allocate to insurance policies the frictional costs of holding capital. My chief concern in that regard is the time frame for loss payments, with the lognormal assumption a potential issue.

Actuaries would like to have a method of allocating capital in order to rank business units by profitability. Myers-Read seems no better or worse than a number of equally arbitrary but reasonable methods for doing that. In combination with a risk pricing methodology it does lead to an alternative route to that goal: rank by comparing actual profit to the value of the risk transfer provided.