

Capital Allocation for P&C Insurers: A Survey of Methods

GARY G. VENTER

Volume 1, pp. 215–223

In

Encyclopedia Of Actuarial Science
(ISBN 0-470-84676-3)

Edited by

Jozef L. Teugels and Bjørn Sundt

© John Wiley & Sons, Ltd, Chichester, 2004

Capital Allocation for P&C Insurers: A Survey of Methods

Capital allocation is never 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 **profit** by the capital for each. 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 depend on having a good theory of risk pricing, where the previous approach would depend on having a good theory of capital allocation. Since both address the same decisions, both will be included in this survey. For those who like to see returns on capital, the pricing method can be put into allocation terminology 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 future growth. For strategic decisions, another question is important – not how much capital a business unit uses, but how much more it would need to support the target growth. In general, it will be profitable to grow the business 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 it is a basic tenet of financial analysis. It differs from capital allocation in that, for marginal-cost pricing, not all capital needs 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 to be generated. Methods of quantifying

the cost of marginal capital will be reviewed here as well, as again this aims 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. The overall return of the insurer can be evaluated on the basis of the borrowing rate that would match its risk and the return on such a fund. If the fund would have to borrow at a particularly low rate of interest to match its risk, then the insurance business is clearly adding value. The business units can then be compared on the basis of their impacts on the borrowing rate.

Thus, while the general topic here is capital allocation, this survey looks at methods of answering the questions that capital allocation addresses. The following four basic approaches will be reviewed:

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

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 set at the losses it would experience in the worst year in 10 000. The expected policyholder deficit or EPD is the expected value of default amounts. It can also be generalized to include the expected deficit beyond some probability level, rather than beyond default. Tail value-at-risk is the expected loss in the event that losses exceed the value-at-risk target. X_{TvaR} , at a probability level, is the average excess of the losses over the overall mean for those cases over that level. Assuming a corporate form with limited liability, an insurer does not pay losses 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

2 Capital Allocation for P&C Insurers: A Survey of Methods

a risk measure. The other measures are the standard statistical quantities.

Typically, 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 such that the Tail VaR for the company at that level is the capital carried. Or 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. Once this has been done, an allocation method can be applied to get the capital split to the business unit level. Several possible allocation methods are given in Table 2.

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 credits each unit with a diversification benefit.

Table 1 Risk measures

☐	VaR
☐	EPD
☐	Tail VaR
☐	X TvaR
☐	Standard deviation
☐	Variance
☐	Semivariance
☐	Cost of default option
☐	Mean of transformed loss

Table 2 Allocation methods

☐	Proportional spread
☐	Marginal analysis
☒	By business unit
☒	Incremental by business unit
☐	Game theory
☐	Equalize relative risk
☐	Apply comeasure

Marginal analysis takes the risk measure of the company excluding a business unit. The savings in the implied total required 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 capitals of all the units. This usually allocates more than the marginal capital to each unit. The incremental marginal method is similar, but the capital savings is calculated for just the last increment of expected loss for the unit, say the last dollar. Whatever reduction in the risk measure that is produced 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. This is in accordance with marginal pricing theory.

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 coalition (set 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 Lloyd Shapley, the founder of game theory.

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 make the EPD for each business unit the same percentage of expected loss.

Comeasures can be thought of in terms of a scenario generator. Take the case in which 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, and the Tail VaR would be estimated by their average. The co-Tail VaR for each business unit would just be the average of its losses in those scenarios. This would be its contribution to the overall Tail VaR. This is 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 [6].

More formally, comeasures are defined when a risk measure R on risk X can be expressed, using a condition defined on X , a leverage function g and a mean ratio a , as the conditional expected value

$$R(X) = E[(X - aEX)g(X)|\text{condition}] \quad (1)$$

As an example, take $a = 1$ and $g(X) = X - EX$ with the condition $0X = 0$. Then $R(X)$ is the variance of X . Or for probability level q , take the condition to be $F(X) > q$, $a = 0$ and $g(x) = 1$. If $q = 99.9\%$, R is then Tail VaR at the 1-in-1000 level. When R can be so expressed, co- R is defined for unit X_j as:

$$\text{co-}R(X_j) = E[(X_j - aEX_j)g(X)|\text{condition}] \quad (2)$$

This defines the comeasure parallel to the risk measure itself, and is always additive. The g function is applied to the overall company, not the business unit. Thus in the variance case,

$$\text{co-}R(X_j) = E[(X_j - EX_j)(X - EX)],$$

which is the covariance of X_j with X .

In the Tail VaR case, co-Tail VaR(X_j) = $E[(X_j | F(X) > q)]$. This is the mean loss for the j th unit in the case where total losses are over the q th quantile, as described above. X TVaR is defined by taking $a = 1$ and $g(x) = 1$ with condition $F(X) > q$. Then

$$\text{co-}X \text{ TVaR}(X_j) = E[(X_j - EX_j) | F(X) > q] \quad (3)$$

This is the average excess of the business unit loss over its mean in the cases where total losses are over the q th quantile. Co-Tail VaR would allocate X_j to a constant X_j , while co- X TVaR would allocate zero.

The g function can be used to define the weighted version of Tail VaR and X TVaR. This would address the criticism of these measures that they weight very adverse losses linearly, where typically more than linear aversion is regarded as appropriate. Note also that the risk measure is not required to be subadditive for the comeasure to be totally additive. Thus, co-VaR could be used, for example.

Evaluation of Allocating by Risk Measures

VaR could be considered to be a shareholder viewpoint, as once the capital is exhausted, the amount by which it has been exhausted is of no concern to shareholders. EPD, default option cost, and Tail VaR relate more to the policyholder viewpoint, as they

are sensitive to the degree of default. X TVaR may correspond more to surplus in that enough premium is usually collected to cover the mean, and surplus covers the excess. All of these measures ignore risk below the critical probability selected. VaR also ignores risk above that level, while with $g = 1$ the tail measures evaluate that risk linearly, which many consider as underweighting. Variance does not distinguish between upward and downward deviations, which could provide a distorted view of risk in some situations in which these directions are not symmetrical. Semivariance looks only at adverse deviations, and adjusts 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.

Allocating by marginal methods appears to be more in accord with financial theory than is direct allocation. However, by allocating more than the pure marginal capital to a unit it could lead to pricing by a mixture of fixed and marginal capital costs, which violates the marginal pricing principle. The comeasure approach is consistent with the total risk measure and is completely additive. However, it too could violate marginal pricing. There is a degree of arbitrariness in any of these methods. Even if the capital standard ties to the total capital requirement of the firm, the choice of allocation method is subjective. If the owners of the firm are sensitive to correlation of results with the market, as financial theory postulates, then any allocation by measures of company-specific risk will be an inappropriate basis for return-on-capital.

Approach 2 – Compare Actual versus Model Pricing

One use of capital allocation could be to price business to equalize return-on-capital. However, there is no guarantee that such pricing would correspond to the market value of the risk transfer. The allocation would have to be done in accord with market pricing to get this result. In fact, if actual pricing were compared to market pricing, the profitability of business units could be evaluated without allocating capital at all. But for those who prefer to allocate capital, it could be allocated such that the return on market pricing is equalized across business units.

This method requires an evaluation of the market value of the risk transfer provided. Financial

4 Capital Allocation for P&C Insurers: A Survey of Methods

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 [9] has been used successfully in several markets to approximate risk pricing. The finance professionals appear to favor an adjusted CAPM approach that corrects many of the oversimplifications of the original formulation.

To use CAPM or similar methods, costs are first identified, and then a risk adjustment is added. Three elements of cost have been identified for this process: loss costs, expense costs, and the frictional costs of holding capital, such as taxation of investment income held by an insurer. The latter is not the same as the reward for bearing risk, which is separately incorporated in the risk adjustment. The frictional costs of capital have to be allocated in order to carry out this program, but because of the additional risk load the return on this capital varies among policies. Thus a reallocation of capital after the pricing is done would be needed, if a constant return-on-capital is to be sought across business units.

The *Myers–Read method* [8] for allocating capital costs has been proposed for the first allocation, the one that is done to allocate the frictional cost of carrying capital. It adds marginal capital for marginal exposure in order to maintain the cost of the default expressed as a percentage of expected losses. This method is discussed in detail in the appendix. To calculate the market value of risk transfer, simple CAPM is now regarded as inadequate. Starting from CAPM, there are several considerations that are needed to get a realistic market value of risk transfer. Some issues in this area are as follows:

- Company-specific risk needs to be incorporated, both for differential costs of retaining versus raising capital, and for meeting customer security requirements – for example, see [3].
- The estimation of beta itself is not an easy matter as in [4].
- Other factors that are needed to account for actual risk pricing – see [2].

- To account for the heavy tail of P&C losses, some method is needed to go beyond variance and covariance, as in [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 they need to be recognized in the pricing.

Evaluation of 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 be reformulated as a capital allocation, if so desired, after the pricing is determined.

Approach 3 – Calculating Marginal Capital Costs

A typical goal of capital allocation is to determine whether or not a business unit is making enough profit to justify the risk it is adding. A third approach to this question is looking at the last increment of business written by the unit, and comparing the cost of the additional capital this increment requires to 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. It does correspond, however, with the financial principle of marginal pricing. In basic terms, if adding an increment of business in a unit adds to the total value of the firm, then the unit should be expanded. This could lead to an anomalous situation in which each business unit is profitable enough but the firm as a whole is not, because of the unallocated fixed capital charges. 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.

One way to do the marginal calculation would be to set a criterion for overall capital, and then see how much incremental capital would be needed for the small expansion of the business unit. This is the same approach that is used in the incremental marginal allocation, 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 [7]. 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 gives the policyholders a put option that allows them to put any losses above the business unit's premium and investment income to the firm. Since this is not unlimited owing to the firm's limited resources, the value of this option is the difference between two put options: the option with a strike at premium plus investment income less the value of the insolvency put held by the firm. The firm's call on the profits is a call option with a strike of zero. If that is worth more than the financial guarantee provided, then the business unit is adding value.

Evaluation of Marginal Capital Costs

This method directly evaluates the 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 discussed. There is still a potential arbitrary step of the criteria chosen for the aggregate capital standard. This is avoided in the financial guarantee approach, but for that the options must be evaluated correctly, and depending on the form of the loss distribution, standard options pricing methods may or may not work. If not, they would have to be extended to the distributions at hand.

Approach 4 – Mutual Fund Comparison

An insurer can be viewed as a tax-disadvantaged leveraged **mutual investment fund**. It is tax-disadvantaged because a mutual investment fund does not usually have to pay taxes on its earnings, while an

insurer does. 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 after-tax probability distribution of returns as the insurer. It can be specified by its borrowing rate, the amount borrowed, and its investment portfolio. This should provide enough degrees of freedom to be able to find such a mutual fund. If there are more than one such, consider the equivalent one to be the one with the highest interest rate for its borrowing.

The insurer can be evaluated by the equivalent borrowing rate. If you can duplicate the return characteristics by borrowing at a high rate of interest, there is not much value in running the insurance operation, as you could borrow the money instead. However, if you have to be able to borrow at a very low or negative rate to get an equivalent return, the insurer is producing a result that is not so easily replicated.

While this is a method for evaluating the overall value added of the insurer, it could be done excluding or adding a business unit or part of a business unit to see if doing so improves the comparison.

Evaluation of Mutual Fund Comparison

This 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 best equivalent mutual fund. It would seem to be a useful step for producing an evaluation of firm and business unit performance.

Conclusions

Allocating by a risk measure is straightforward, but arbitrary. It also could involve allocation of fixed costs, which can produce misleading indications of actual profitability prospects. Pricing comparison is as good as the pricing model used. This would require its own evaluation, which could be complicated. The marginal-cost method shows the impact of growing each business unit directly, but 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 insight into the value of the firm and its business units. All of these methods have a time-frame issue not addressed here in detail:

lines of business that pay losses over several years have several years of capital needed, which has to be recognized.

References

- [1] Butsic, R. (1999). *Capital Allocation for Property-Liability Insurers: A Catastrophe Reinsurance Application*, CAS Forum, Spring 1999, www.casact.org/pubs/forum/99spforum/99spforum001.pdf.
- [2] Fama, E.F. & French, K.R. (1996). Multifactor explanations of asset pricing anomalies, *Journal of Finance* **51**.
- [3] Froot, K.A. & Stein, J.C. (1998). A new approach to capital budgeting for financial institutions, *Journal of Applied Corporate Finance* **11**(2).
- [4] Kaplan, P.D. & Peterson, J.D. (1998). Full-information industry betas, *Financial Management* **27**.
- [5] Kozik, T.J. & Larson, A.M. (2001). The N-moment insurance CAPM, *Proceedings of the Casualty Actuarial Society* **LXXXVIII**.
- [6] Kreps, R.E. (2004). *Riskiness Leverage Models*, Intrat working paper, to appear.
- [7] Merton, R.C. & Perold, A.F. (1993). Theory of risk capital in financial firms, *Journal of Applied Corporate Finance* **6**(3), 1632.
- [8] Myers, S.C. & Read, J.A. (2001). Capital allocation for insurance companies, *Journal of Risk and Insurance* **68**(4), 545–580.
- [9] Wang, S. (2002). A universal framework for pricing financial and insurance risks, *ASTIN Bulletin* **32**, 213–234.

Appendix: The Myers-Read Approach

Myers-Read (MR) capital allocation presents a challenge to the classification of methods, in that it allocates all capital by a **risk measure**, provides a marginal capital cost, and can be used in pricing.

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, that is, 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 profits.

MR works for 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 capital allocation can be used for 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 allocation is a way to represent that cost allocation. The pricing 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 development is the value of the default put option. Assuming a corporate form 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 log-normal or normal distribution for the insurer's entire loss portfolio, so it 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, that is, 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 could be taken to be the amount of capital needed to get D/L to some 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. Unlike most marginal allocation approaches, the marginal capital amounts add up to the total capital of the firm with

no proportional adjustment. This might be expected from the additive nature of option prices.

The total capital is the sum of the individual capital charges, that is, $\sum c_i L_i = cL$, where $c_i L_i$ is the capital for the i th policy with expected losses L_i , and cL is the 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 [1] 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 by increasing L_i by a dollar increases capital by c_i , which is set to be enough to keep D/L constant. Thus, the formal requirement that $\partial D/\partial L_i = D/L$ means that the change in capital produced by c_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

$$D = L[N(y + v) - (1 + c)N(y)] \quad (4)$$

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 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 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 (standard deviation/mean) 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

$$c_i = c + (b_i - 1)Z, \text{ where } Z = \frac{(1 + c)n(y)k_L^2}{N(y)v(1 + k_L^2)} \quad (5)$$

Butsic provides a simple example of this calculation. A company with three lines is assumed, with expect losses, CV's, and correlations as shown in Table 3. The total capital and its volatility are also inputs. The rest of the table is calculated from those assumptions.

Table 3

	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	2500	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	

Changing the byline 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 kept by adding an amount to capital equal to the c 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 the assets constant, even though the 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 its 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 as 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 keeps 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 as 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 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 the management wants to express this as a return-on-capital, the MR capital would not be appropriate. Rather, the total capital should be reallocated so that the ratio of modeled target profit to allocated capital is the same for each unit. Then comparing the returns on capital would give the same evaluation as comparing the profits to target profits. MR capital allocation would be the basis of allocating frictional capital costs, but not for calculating the return-on-capital.

(See also Audit; Financial Economics; Incomplete Markets; Insurability; Risk-based Capital Allocation; Optimal Risk Sharing)

GARY G. VENTER