CHAPTER FIVE
RATE OF RETURN

By Frank D. Pierson

It's easy to say that a company prices its products to achieve at least a minimum rate of return on equity. In reality, it is not that easy. The hardest part is to allocate capital to individual contracts or lines of business; capital is not really divisible by line of business. It is available, in its entirety, to any one contract or line of business, if that contract or line produces enough losses. That being said, it is useful to think about capital as divisible, at least for pricing purposes. Since the purpose of capital is to absorb adverse deviations from expected, any method to allocate capital is acceptable so long as it differentiates among contracts or lines of business based on the risk of adverse deviation. A method that is arbitrary and does not differentiate between risks is not a viable way to allocate capital. While it might mechanically allow someone to look at rates of return on equity, it will not provide any meaningful insight as to whether a particular contract or line of business provides an attractive return for the risk taken to achieve that return.

There are many different ways to create an allocation method that differentiates based on risk. This chapter outlines one such approach in use at the author's company, a reinsurer. This approach allocates capital based on the contract's or line of business's contribution to the overall risk of the portfolio of contracts. In other words, the approach looks at the marginal capital needed to support adding new exposures to those already on the books. The capital needed is based on a risk of ruin, i.e., the company wishes to maintain its capital such that there is a constant probability that it will become insolvent as it adds exposures.

This method follows on the work of Kreps. One key difference between the approach here and that suggested by Kreps is that there is no assumption that the shapes of the distributions before and after adding the new exposures are the same. (Kreps makes this assumption implicitly by assuming that "q" is unchanged.)

INTRODUCTION

Before we discuss capital allocation among contracts for a (re)insurer, it will be helpful to discuss how an investor, in general, and an insurer, in particular, looks at investing. Once we understand the dynamics of investing, then we can develop a framework and methodology for allocating capital that is consistent with how an insurer should look at investing and analyze risk.

Within market and economic constraints, an investor will always try to maximize his/her returns. Although it may be possible to achieve a required level of return without an analysis of the risk underlying his/her investments and allocating capital accordingly, it is not possible to optimize return versus risk without such analysis.

In addition, an investor cannot maximize his/her return if part of his/her capital is uninvested because the return on actual capital is diminished by the lack of return on the amount not invested. I doubt that anyone would argue that an insurance enterprise is an investor. It is obvious that an insurer is an investor when it invests the premium funds it receives in order to pay losses. It is less obvious that an insurer is also an investor when it underwrites a new policy. An insurance policy can be viewed as a "reverse" investment made by the insurer when it agrees to write the policy. One can see that this is true if one looks at an investment in rather generic terms as either i) an outflow from the investor on which he/she expects a return or ii) the amount that the investor could lose by making the investment. Using this view, it is clear that the insurer's investment, when it writes a policy, is not the premium, but the amount of loss payments it must make, or more specifically, the amount of the shortfall between the total loss payments and the premium. Under this view, premiums (and the interest on them) become the return for making the investment.

As respects an insurer, one might argue, then, that if the total capital is not yet allocated any policy with a positive return will increase the overall return and, therefore, the insurer should write the policy. This is not the case, however, for the following reasons:

1. Investing in marginal investments (i.e., policies) today might preclude the insurer from investing in more profitable investments later, inasmuch as once it is allocated, capital may remain allocated for a significantly long time.

2. At some point, as all of the capital is allocated, the overall return approaches the average return across all investments. At that point, prior marginal investments may prevent the insurer from achieving its minimum return hurdles.

Another way of seeing writing an insurance policy as an investment is to think of the policy as an option agreement issued by the insurer. The insurer sells an option to the insured to put losses specified in the policy to the insurer. The premium is, in reality, the fee the insurer charges for giving the insured a put option.

An insurer, just as any other investor would, wants to maximize its investment returns. As an investor of the premium funds, it is fairly easy to establish whether it is maximizing its return and it can determine how much of these funds is invested at any point in time. It is much more difficult, however, for an insurer to determine how much of its capital is invested through its underwriting. Current practice in the insurance industry assumes that capital is invested proportionally with the level of premiums written by the company. Unfortunately, the rationale for this approach does not have any real theoretical basis and is, therefore, inadequate to use as a means to judge returns.

The following example illustrates how premium is unrelated to the risk assumed. Suppose that the industry standard is to allocate $1 of capital for each $2 of premium written. If during one year, the company wrote $100 of premium, it would allocate $50 of capital. If rates doubled on January 1st of the next year and the company renewed all of its annual policies on that date at the new rate level, this method would indicate that the company needed to allocate $100 of capital based on the new premium level of $200. If the company did not increase its capital to $100, it would be considered under-capitalized compared to the previous year. If risk is measured by the potential for losses to exceed premiums plus interest, the company in the later year is better capitalized than in the earlier year. Obviously, this method produces conclusions that are exactly opposite to reality.

If one accepts the concept of underwriting as investing, then financial market tools used to evaluate investment should apply equally to insurance; in fact, there has been a great deal of interest in applying such financial tools as Option Pricing Theory ("OPT") and Capital Asset Pricing Model ("CAPM") to insurance. CAPM is already in use in at least one state (Massachusetts). In addition to financial market tools, the insurance industry is trying to develop its own tools, e.g., Ruin Theory ("RT"). A few comments about the use of these methods of allocating capital or determining risk are in order (a more detailed explanation is beyond the scope of this paper).

CAPITAL ASSET PRICING MODEL

CAPM attempts to set the premium at a level that will allow a company to achieve an appropriate return in the expected case. The appropriate return, in this case, equals the risk-free rate plus a risk adjustment. The risk adjustment is dependent on "Beta" which represents the covariance of returns between the insurer and the market. CAPM asserts that an investor should be rewarded for accepting systematic risk only and not for accepting diversifiable risk, i.e., the investor is not rewarded in proportion to the risk inherent in any single investment, but is rewarded for the risk he assumes for holding a well diversified portfolio of investments. A sophisticated investor will look at a new investment by analyzing how his overall portfolio will perform with and without that investment. Only if the overall performance of his portfolio improves should the investor add that investment.

If, as discussed above, an insurer's investment is truly represented by losses, not premiums, then premiums based on current CAPM methods are not correct because they try to generate appropriate returns on premium, not loss. The basis for the analysis is inappropriate since calculating returns based on premium is equivalent to calculating returns based on the return itself rather than on the investment.

OPTION PRICING THEORY

OPT is applicable to insurance if one views an insurance contract as an option contract, i.e., the insured pays an option premium to the insurer in order to call cash to pay its losses. Obviously, the insured does not have to recover losses from the insurer. In fact, with certain loss sensitive policies there are incentives not to report losses (if losses are loaded for expenses and/or profit) once the premiums are greater than minimum levels.

What many insureds are just realizing is that there are hidden options contained in their contracts, e.g., if the insurer runs out of funds it can put the losses back to the insured. OPT attempts to calculate the equilibrium price for all of the options embedded in an insurance contract.

RUIN THEORY

European actuaries having been exploring Ruin Theory ("RT") for some time. The basic goal of RT is to calculate the minimum amount of capital required to reduce the probability of insolvency below some selected level, e.g., 1/10 of one percent, over a fixed or indefinite time horizon. RT models usually simulate the operations of the insurer over the time horizon and the initial or minimum surplus is set so that the number of iterations that result in insolvency are less than the desired level.

One can use RT to determine the level of capital needed to support a given portfolio of contracts and determine the increase in the capital required by adding an additional contract. This marginal capital can be used as the basis for allocating capital to an individual policy. For example, if writing an additional policy increases the required surplus by $1 million, then the surplus allocated to that policy is $1 million.

RT usually produces a level of surplus needed to avoid insolvency to some specified degree of confidence. Unfortunately, an insurer would be out of business long before its surplus was depleted to the point of insolvency due to the lack of confidence a low amount of surplus would generate (in theory anyway). RT must be adjusted to accommodate a different threshold.

ALLOCATED RISK CAPITAL

The rest of this chapter presents an approach to allocate capital to individual contracts in order to determine the rate of return on equity that incorporates features of CAPM, OPT and RT.

There are two levels at which the issue of capital allocation must be addressed. The first level is the allocation of capital based on market constraints. Allocation at this level is usually based on simple Surplus:Premium (*S:P*) or Surplus:Reserve (*S:R*) ratios. The amount of capital so allocated can be referred to as "Market Perception" Capital ("*MPC*"). It has been shown that the ratios used to calculate *MPC* have little or no theoretical foundations. (See above for discussion of problems of *S:P* ratios). The ratios are based more on tradition than on any risk analysis.

The second level of capital allocation is based on the risk of the contract being written. The amount of capital so allocated can be referred to as "Allocated Risk" Capital ("*ARC*"). In recognition of the shortcomings of using *S:P* or *R:S* ratios, the NAIC has adopted a new procedure that will calculate an insurer's *ARC* or, as the NAIC refers to it, "Risk Based" Capital ("RBC"), by applying industry ratios to various items on the balance sheet and income statement to arrive at the RBC needed as of a specific point in time. It does not include the impact of future business except to the extent that there is a risk charge for unearned premiums and written premiums (as a proxy for risk of running off the policies into the following year). One would expect that as the concept of RBC or *ARC* becomes accepted by the market place, *MPC* should approach *ARC* as old rules of thumb (P:S and R:S ratios) are no longer used. Unfortunately for our purposes, the NAIC's proposed calculation is based on the aggregate experience of the insurer rather than on the experience of an individual policy and, therefore, is not directly applicable to allocating capital to an individual policy.

The dilemma facing every investor is how best to invest all capital? Assuming an investor does not want to decrease the amount of capital he/she has to invest, one approach to investing the total capital is to allocate capital to each investment and to invest in only those investments that have returns on allocated capital greater than some minimum return. The investor maximizes his/her return by evaluating each investment opportunity individually and building a portfolio of investments, each meeting some pre-determined return. Depending on risk appetite, the investor might risk-adjust the returns before choosing the investments.

CONCEPTUAL FRAMEWORK

The distinction between *MPC* and *ARC* (and the concept of face capital developed below) gives us a framework to efficiently structure an insurer's capital, i.e., to determine an efficient mix between common equity, debt and/or preferred stock, and, at the same provide us with the means of determining the rate of return on capital at the individual policy level. Common equity usually bears the ultimate risk of any company and, in return, earns the highest yields. It is, therefore, closest in nature to the *ARC* in that the *ARC* is the amount of capital "at risk" for any given contract. Preferred stock or debt is usually used to augment yield on common equity and to supply "face" capital as needed. The excess of the *MPC* over the *ARC*, if any, could be viewed as "face" capital because it is needed only to calm outside observers and is excessive relative to the risk inherent in the book of business. Unfortunately, it is not prudent to ignore the *MPC*, at least in the long run. Market perception will dictate whether the company is viewed as strong or weak. If it is viewed as being under-capitalized, new business will not be written, or worse, only bad business will be offered to the company. In this manner, I believe that preferred stock or debt is closest in nature to the face capital.

If this characterization of capital is correct, then the proper base on which to measure return on equity ("*ROE*") is the *ARC*. Most methods in use today use *MPC* as the base on which to measure *ROE*. The cost of face capital should be included as an expense in the calculation of the *ROE* in the same way that payments on true debt or preferred stock would be included. Therefore, a charge to income would be included in the numerator rather than including the face capital in the denominator, i.e., using *MPC* as the base for the *ROE* calculation. In other words, the *ROE* should be calculated as:

$$ROE = \frac{p}{ARC} = \frac{P - L - E - r(MPC - ARC) - C}{ARC} \tag{1}$$

where p = present value profit,

P = present value of the net premium,

L = present value of the paid losses, if any,

E = present value of the expenses,

r = spread paid to borrow funds over what can earned on investing the same funds, and

C = present value of the profit commission, if any.

MPC is fairly easy to calculate. It can be based on the greater of the *S:P* ratio times the present value of the premium or the *S:R* ratio times the sum of the present value of the year-end reserves. It may be better to use year-end reserves since most analysts use the Annual Statement to evaluate financial strength. One would expect that this number should be replaced over time with the NAIC's statutory RBC calculation.

Unfortunately, *ARC* is not so easy to calculate. The proper level of *ARC* should reflect:

1. Variability in underwriting profit or loss for the contract in question, including underwriting and timing risk;

2. Interaction of that contract with all other contracts written by the company (i.e., the marginal *ARC* needed to write the additional contract);

3. Expense risk;

4. Credit risk stemming from underwriting (which may not be immaterial given the duration of our contracts);

5. Investment risk stemming from the mismatch of liabilities and assets;

6. Credit risk stemming from our investments; and

7. Off-balance sheet risk, regulatory changes, etc.

The *ARC* should also include recognition that risk exists over the entire life of a contract.

Calculating the *ARC* for all the factors listed above is theoretically complex—too complex and time consuming to be used in pricing each individual deal. The *ARC* corresponding to (4) through (7) (collectively referred to as "All Other *ARC*" or "*AO ARC*") is beyond the scope of this paper. It may be possible to include the *AO ARC* in pricing by calculating the ratio of the *AO ARC* needed for an "average" contract to some base; likely candidates are the assets, reserves or, possibly, the total *ARC* for the "average" contract.

For pricing purposes, it may not be necessary to calculate the *AO ARC* for an individual contract unless one were trying to develop an absolutely correct *ARC* for each contract. If, instead, one were trying to develop an *ARC* that can be used to determine which policies have the highest relative *ROE* and we assume that the *AO ARC* is approximately proportional to the total *ARC* for each contract, then calculating the *ARC* for (1) through (3) only (collectively referred to as "Underwriting *ARC*" or "*UW ARC*") for each contract should be sufficient. For simplicity, I will use just "*ARC*" in the rest of this paper to denote *UW ARC*.

Given the complexity of calculating the *ARC*, we need to develop a simple measure of the *ARC*. One approach would be to calculate the *ARC* in a manner similar to that used to calculate the *MPC*, i.e., apply Reserve:Surplus (*R:S*) ratios that vary by the amount of assumed risk. To do this, we would need to define a risk/financing continuum with very risky policies at one end (e.g., conventional cat covers) and low/no risks policies at the other (e.g., Time & Distance policies) and then subjectively assign R:S ratios to each end of the continuum and create a formula for the ratios in between (e.g., linear or log). Once the continuum is defined, the underwriter could then place each policy on the continuum and use the corresponding R:S ratio to calculate the *ARC*. One drawback of this approach is that it requires the underwriter to add another layer of assumptions on top of those used to price the deal. There is no guarantee that the placement of the policy on the continuum would be consistent with the risk implied by the distribution of profit and loss underlying the pricing. In addition, two underwriters might look at the same profit and loss distribution and place the policy in different places on the continuum.

Another approach would be to set the *ARC* equal to an amount that would guarantee that all liabilities would be honored at some specific confidence level, e.g., 90%. This approach is similar to the concept of ruin theory in that the capital is set so that the probability of insolvency or ruin is very remote. Although I use a 90% confidence level in the following examples, I believe that the right level is closer to 99% or higher (this latter level is typically used in ruin theory). The *ARC* for a single contract would be the *ARC* for the portfolio of contracts including the contract in question less the *ARC* for the portfolio without the same contract. Rodney Kreps uses a similar method to calculate a risk load for reinsurers.

There is currently much work being done, particularly in Europe, on ruin theory. Ruin theory concentrates on variability in the loss process and, therefore, it ignores some of the other risks faced by an insurance company as mentioned above, e.g., investment, credit or expense risk. To overcome that limitation, many actuaries are now attempting to simulate the entire insurance operation to incorporate the risks ignored by traditional ruin theory. Although I believe that simulating the entire insurance operation may ultimately work to estimate the total *ARC*(*UW ARC*+*AO ARC*) for a company as a whole, I do not believe that it would work, in practice, for pricing an individual policy.

To see how this would work in practice, let's consider a few examples. Assume that for policy number 1, there is a 90% probability of a $3,333 profit and a 10% probability of a $10,000 loss. Using the criteria of a 90% confidence level, the *ARC* would be, therefore, equal to $10,000. Given that the expected profit is $2,000, the total return on *ARC* (excluding the cost of face capital, if any) would be 20%.

Assume that for policy number 2, there is a 80% probability of a $5,000 profit and a 20% probability of a $10,000 loss. Using the criteria of a 90% confidence level, the *ARC* would be, therefore, equal to $10,000 (the amount closest to a 90% confidence level). Given that the expected profit is $2,000, the return on *ARC* (excluding the cost of face capital, if any) would be 20%, the same as for policy 1.

Let's now consider the *ARC* for the two policies combined assuming that the two policies are independent. The following table will help determine the proper *ARC* given a 90% confidence level:

TABLE 1

Policy Number				
1	2	1+2	Probability	Cumulative
$ 3,333	$ 5,000	$ 8,333	72%	72%
$ (10,000)	$ 5,000	$ (5,000)	8%	80%
$ 3,333	$ (10,000)	$ (6,667)	18%	98%
$ (10,000)	$ (10,000)	$ (20,000)	2%	100%

Based on a 90% confidence level, the *ARC* would be $6,667. As you can see, this amount is much less than the sum of the *ARC*s for each policy. In fact, the *ARC* is less than the *ARC* for either policy written separately.

This is obviously an overly simplified example, but it highlights the fact that the *ARC* of a portfolio of mutually independent policies can be significantly less than the sum of the *ARC*s for each policy assuming that each policy is expected to be profitable. In fact, the law of large numbers is another manifestation of this underlying process.

If, on the other hand, the policies were 100% positively correlated (e.g., if one of the policies above had a profit or loss, the other policy would have a profit or loss, respectively) the *ARC* would be equal to the sum of the *ARC*s of each policy. Since not all insurance policies are truly independent from each other or perfectly correlated, the correct *ARC* is somewhere between these two extremes. The *ARC* calculation can be adjusted for mutual dependence to the extent that it can be estimated and modeled.

There are a number of issues that need to be addressed before we can use this approach to calculate the *ROE* of individual contracts. Among them:

There are a number of issues that need to be addressed before we can use this approach to calculate the *ROE* of individual contracts. Among them:

I. There is a small interpretational issue involved in this method. If the aggregate *ARC* for the company after writing policy 1 in the example above is $10,000 and is $6,667 after the company writes policy 2, does this method imply that the *ARC* for both policies is equal to $3,333 or is it $10,000 for policy 1 and ($3,333) for policy 2? The latter doesn't make much sense, yet it is true that the aggregate *ARC* goes down because an additional policy was written.

   In addition, the *ARC* for each new policy can be affected by the order in which policies are added to the portfolio.

II. We need to incorporate the change in the profit/loss profile of bound deals over time. In other words, the maturing of the existing portfolio will change the probability of ruin with or without writing any additional policies. We need to reflect the existing portfolio at its current stage of maturity. This will affect pricing indirectly because as the portfolio ages, the estimate of the incremental *ARC* will vary depending on when the new policy is added to the portfolio.

III. There are accounting issues that have to be considered when selecting the ratio to be used in calculating the *MPC*. For example:

   A. should reserves that are discounted be grossed up for the discount?

   B. should reserves be gross or net of subrogation and salvage and reinsurance?

   C. should the ratio be adjusted if contracts all have contractual limits?

   D. should reserves be gross or net of an explicit risk load?

IV. This methodology does not guarantee that the aggregate *ARC* equals the company's actual capital as some people believe it should. If the aggregate *ARC* as calculated by this method were less than the actual capital, it might tell management that it could safely write new business or that it should distribute some of the capital to its shareholders. On the other hand, if the *ARC* were greater than the actual capital, it might tell management that it should cut back on its writings or it should raise more capital.

V. The choice of the confidence level is arbitrary, but it must be fairly close to 100%. There are, among others, two quite different ways to set the confidence level. The first is the more straightforward approach, i.e., management subjectively selects the level and then the aggregate *ARC* is calculated. If the aggregate *ARC* is not equal to the actual capital, then some action as outlined in (4) should occur. The second approach would be to compare the actual capital to the probabilities of various loss amounts and set the confidence level so that the aggregate *ARC* equals the actual capital. The latter approach is not as robust as the former, nor does it provide a consistent benchmark for management to assess the adequacy or efficiency of its capital.

METHODOLOGY

The method described above would be very difficult or impossible to apply at the individual policy level because of the time involved in calculating the *ARC* with and without the policy in question. It may be possible, however, to approximate this method by making some simplifying assumptions about the profit/loss distribution of the current portfolio and the new contract. To develop this approximation, let's assume:

1. That U and S^2 equal the mean and the variance, respectively, of the existing portfolio (i.e., without the new policy) and u and s^2 equal the mean and variance, respectively, of the policy for which we wish to calculate the return on equity.

2. That the *ARC* for the existing portfolio will be defined as an amount such that the probability that an aggregate loss is less than or equal to that amount is equal close to 100%. If we set the ARC_e so defined, equal to $q_1 S - U$, we can calculate q_1 as:

   $q_1 = (ARC_e + U)/S$, where ARC_c is the value at the selected confidence level.

3. That the mean of the new portfolio (i.e., including the new policy) would be equal to $U + u$ and that the standard deviation would be equal to $(S^2 + s^2 + 2cSs)^{\frac{1}{2}}$, where c is the correlation coefficient between the new policy and the existing portfolio. Let's denote $(S^2 + s^2 + 2cSs)^{\frac{1}{2}}$ by S_2.

Based on these assumptions, the ARC_n for the new portfolio is equal to:

$$
\begin{aligned}
ARC_n &= q_2(S^2 + s^2 + 2cSs)^{\frac{1}{2}} - (U + u) \\
&= q_2 S_2 - (U + u)
\end{aligned}
$$

where q_2 is calculated in a manner similar to q_1.

If we define, as suggested above, the ARC for the new policy as the difference between the ARC for the existing and new portfolios, then the ARC for the new policy would be equal to:

$$ARC_p = ARC_n - ARC_e = q_2 S_2 - q_1 S - u,$$

or

$$ARC_p = q_1 (S_2 - S) - u + (q_2 - q_1) S_2,$$

where

$$(S_2 - S) = (S^2 + s^2 + 2cSs)^{\frac{1}{2}} - S.$$

If we multiply the right hand of the formula for $(S_2 - S)$ by one in the form of (s/s), we get

$$\begin{aligned}(S_2 - S) &= (s/s)(S^2 + s^2 + 2cSc)^{\frac{1}{2}} - S(s/s) \\ &= s\left((S/s)^2 + (s/s)^2 + 2cSs/s^2\right)^{\frac{1}{2}} - S(s/s) \\ &= s\left((S/s)^2 + 1.0 + 2c(S/s)\right)^{\frac{1}{2}} - (S/s)s.\end{aligned}$$

The term (S/s) reflects the size of the variability in the existing portfolio relative to the volatility of the new policy. As this term gets larger, the value of $(S_2 - S)$ approaches the value of c. Therefore, if we assume that the term (S/s) is sufficiently "large," it can be shown that:

$$(S_2 - S) = cs, \, c \neq 0.$$

and the formula for ARC (from now on, the subscript "p" is dropped to keep the rotation cleaner) reduces to:

$$ARC = q_1 cs - u + (q_2 - q_1) S_2,$$

or

$$ARC = q_1 cs - u + K,$$

where $K = (q_2 - q_1) S_2$. K can be regarded as "portfolio adjustment factor."

This formula for ARC has a straightforward interpretation. For any new contract, the Allocated Risk Capital is equal to the maximum loss that the company will tolerate at a specified confidence level (represented by the term $q_1 cs$ and based on the volatility of the new contract) less the expected profit of an average contract (represented by the term u, assuming that all contracts are the same size) less an adjustment factor to reflect the fact that the contract will be written in the context of an existing portfolio (represented by the term K).

It should be noted that q so determined is at the portfolio level, not at the individual policy level. In other words, the value of $qs - u$ at the individual policy level may not correspond to a 99% confidence level. In fact, it may correspond to a level significantly less than that.

ANNUALIZING THE RETURN ON ARC

Now that we have a workable formula to calculate the *ARC* for an individual policy at inception of the policy, let's return to our formula for *ROE*, i.e., $ROE = p/ARC$. This *ROE* calculation does not reflect the fact that the *ARC* allocated to support the policy remains committed throughout the life of the policy as measured by the presence of risk. It would be great if the underwriter knew immediately after binding the policy whether it made or lost money. Unfortunately, the underwriter does not know the final outcome of a deal for many years after he/she writes it (an extreme example of this are all the insurers who wrote GL policies in the 40's who are now paying out massive amounts for asbestosis and pollution claims). Until we are certain about the outcome of a policy, the market will require a company to support it with equity or *ARC*. Obviously, the level of *ARC* does not stay constant over time because we learn incrementally about the results of a policy and risk is amortized away. *ARC* typically starts out at its maximum value (and it may stay at that level for some time) and then decreases over time, finally reaching zero when all uncertainty about the policy is extinguished. How can we determine how long *ARC* is committed to any policy?

To understand how long the *ARC* is committed to a policy, it is necessary to understand the source of the risks underlying the policy. Some of sources of the risk underlying insurance policies (due to uncertainty about these items) are:

Subject premium volume

Subject exposure volume (which may be different from premium)

Ultimate loss levels

Timing of premium and loss payments

Catastrophe losses

Mix of business/classes

Mix by territory

Inuring reinsurance

Lack of actual data/immature of most recent historical data

Rate changes/Inflation

Risk amortizes away as the uncertainty about these items is reduced.

Uncertainty about the above items is eliminated over two different, overlapping time periods: the exposure period and the period over which the losses pay out. The first period generates Exposure Risk and the second one generates the Loss Risk. Exposure Risk covers most of the items listed above. I believe that this risk has a relatively very short life. For some lines of business, for example, property, this risk is near zero shortly after the policy expires. For other lines of business, for example, workers compensation, this risk is near zero by 12-24 months after the end of the policy period. If that is true, then this risk amortizes over the life of the policy plus 12-24 months, less in the first part of the policy period and more quickly later.

Loss Risk stems from uncertainty of the ultimate losses and their payment pattern. This risk is well known to actuaries as they try to set reserves each year. This risk amortizes over time as losses are paid and is zero when all losses have been settled. It is not clear whether this risk amortizes pro-rata or faster than the reserve run-off pattern. It is fairly easy to produce cases where this risk amortizes quicker than the reserve run-off pattern. I believe that it is conservative to assume is that this risk amortizes pro-rata as reserves run-off.

Unfortunately, the relative influence on the total risk of a policy is not constant for all classes of business. For example, much of the risk for a stop-loss on an auto liability book would expire within 12-24 months after the end of the policy period even if the ceded losses would not be paid for a number of years after that since the ultimate losses and the payment pattern should be well established at that time. For excess D&O or other long-tail lines, the bulk of the risk would remain outstanding for a much longer time.

We could factor in the influence of different classes of business by changing the weights given to the two patterns. Figure 1 shows how this might work. The Exposure Risk is assumed to run-off evenly over a three year period. The Loss Risk follows the loss reserve run-off (assumed to be zero for seven years and then 20% per year thereafter). The graph also shows weighted averages giving 80% and 50% to the Exposure Risk.

FIGURE 1

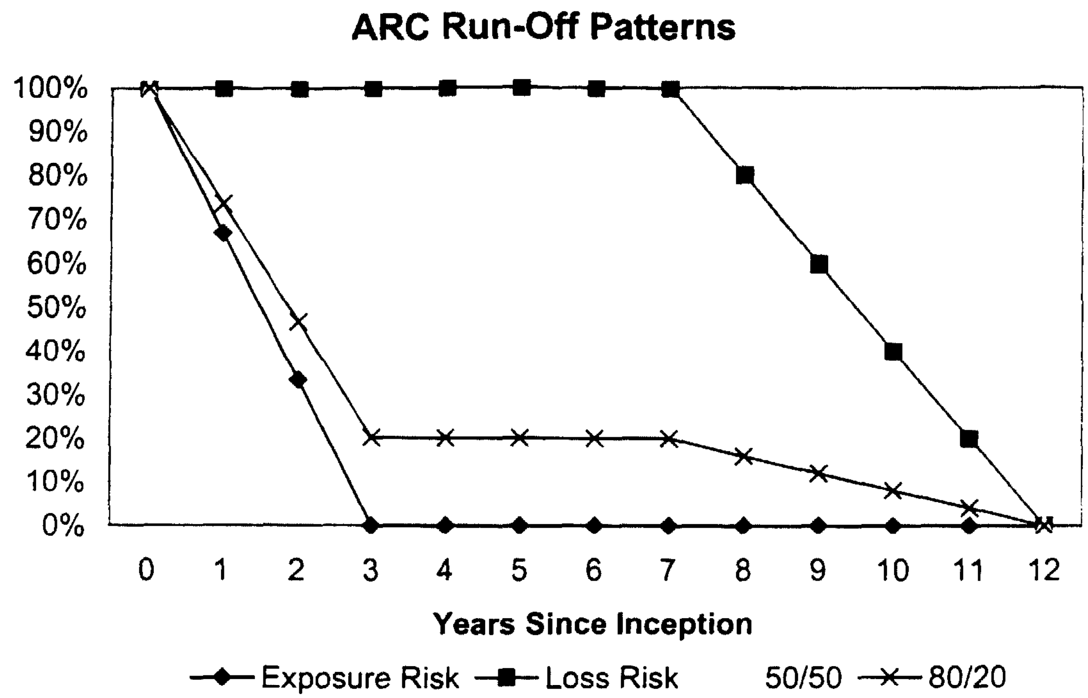

Now that we have a method to determine the *ARC* and the speed at which it amortizes, it is now possible to determine the return on the insurer's "investment" when it writes a policy. If we assume that the insurer "invests" an amount equal to the *ARC* in each new policy, we can then calculate the internal rate of return ("*IRR*") of the policy, i.e., the rate that equalizes the present value of the investment and the present value of the return that the investor receives for writing the deal. The return the "investor" receives is equal to the initial investment plus the profits which the investor receives over time. We assume that the initial investment is returned to the investor as the risk amortizes away. We assume that the investor receives the expected profits as they would be recognized in the insurer's financials.

An example may be helpful. Let's assume that the *ARC*, calculated as described above, is \$100 and that we expect the *ARC* needed at the end of years one through four is \$100, \$81, \$25 and \$0, respectively. Let's further assume that the profit of \$10 is recognized evenly in the first two years. The flows underlying this investment are, therefore:

TABLE 2

Year	ARC	ARC Flow	Profit Flow	Total Flow
0	$ 100	$ (100)	--	$ (100)
1	$ 100	$ 0	$ 5	$ 5
2	$ 81	$ 19	$ 5	$ 24
3	$ 25	$ 56	$ 0	$ 56
4	$ 0	$ 25	$ 0	$ 25

The *IRR* of the last column is approximately 5%.

REMAINING ISSUES

There are a number of issues that still need to be discussed or addressed:

1. There is nothing to prevent the *ARC* produced by the above formula from being negative. If the *ARC* is negative, the new policy is a net provider of capital rather than a consumer of capital. In this particular case, the company should always write these deals because as net equity providers these are an extremely cheap form of capital.

2. There is nothing to prevent the face capital from being negative. In other words, when face capital is negative, the policy will be a net provider of face capital. This only happens when $ARC > MPC$. These contracts help to reduce the need for face capital across the entire portfolio. In essence, other deals "borrow" this face capital and, therefore, a deal with negative face capital should be credited with the investment income it "earns" by loaning its face capital to other deals.

3. The above formula for *ARC* does not reflect parameter risk and unquantifiable extra contractual risks and a separate loading should be included in the *ARC* calculation. The question that naturally arises is how to add a load for parameter risk (if we could quantify it, it would no longer be parameter risk) or for extra-contractual risk (if we could quantify or identify it, we could eliminate it). There are a number of candidates that could serve as a basis for this loading, among them:

   Reinsurance Premium

   ARC

   MPC

   Limit

   Limit - Premium

And I am sure that there are others. My recommendation is the aggregate limit less the maximum premium. This value represents the maximum exposure (excluding the risk that the limit provision of the contract is not upheld). The above formula for *ARC* could be modified to include such a loading as follows:

$$ARC = qcs - u + w + K$$

where $w = 1\%(\text{limit} - \text{premium})$. The w value could be referred to as the "who knows factor."

4. At what level within a company should the value of q be determined? Should there be a single value of q for the entire company, by line of business or territory or some other market segment? I believe that the value of q should be set at the overall company level unless capital has been allocated down to line of business, etc. and the company wishes to allocate capital among policies at this lower level at different confidence levels.

5. The choice of q is dependent on the confidence level required by management and should reflect management's risk tolerance. But even after management has selected the confidence level, should the portfolio and the new policy be judged on an ongoing concern basis or in a run-off situation (the unearned portion of policies written are canceled)?

In any case, the value of q must be based on the existing portfolio including the policies that were written in the past and are in run-off. In other words, the value of q depends on more than just the current policy or accident year. As such, as the book matures it will affect the shape of the profit/loss distribution.