2nd AFIR Colloquium 1991, 3: 199-217

Risk Based Pricing of Life Insurance Products

Matthew S. Easley & Stephen A. J. Sedlak

Nationwide Life Insurance Company, One Nationwide Plaza, Columbus, Ohio 43216, U. S. A.

Summary

In today's rapidly changing and competitive environment, narrowing margins and shifting experience require a much better evaluation of product risks and more refinement in how they are recognized in the pricing process than has previously been the case. This paper describes some approaches which can be used to better reflect risk in the pricing process.

Résumé

Prix des Produits d'Assurance Vie Basé sur la Risque

Dans l'environnement concurrentiel d'aujourd'hui en évolution rapide, les marges de plus en plus étroites et l'expérience changeante exigent une évaluation beaucoup plus fine des risques de produits et une plus grande précision qu'auparavant dans la façon d'en tenir compte dans le processus de tarification. Cet article décrit quelques approches qui peuvent être utilisées pour mieux refléter le risque dans le processus de tarification.

I. Introduction

Until recent times, most insurance products were priced using methods which at best recognized their inherent risks using rules of thumb or estimated margins. While the tides of emerging experience were favorable and competition allowed fairly good margins, these techniques served fairly well. However, recent events such as the volatile interest rates of the 1980s, steeply increasing health care costs and the emerging HIV epidemic have served to change any beliefs we might have had about how stable our environment was. This in turn increased the need to recognize the various risks of our products in their pricing. In addition, competitive pressures have increased (at least in the U.S.) causing product margins, previously available to absorb risk, to shrink.

For many actuaries, it has become apparent that insurers must charge for the risks inherent in their products. The commonly used "best estimate" or "most likely" pricing methods are not adequate since they do not reveal the volatility and the resulting risk in the product being priced. Even when product volatility is known or estimated, such methods, by their very nature, do not systematically charge for it.

II. Risk Margins

One time honored means used to handle risks in pricing is to use risk margins or charges, such as increasing mortality 10% above expected levels. These charges are determined using methods which range from educated guesses to studies of actual experience and its volatility. Thus, the 10% referred to above could've been adopted because 10% looked like a good round number which management would accept. Alternatively, a study of mortality experience which revealed that actual mortality deviated more than 10% above expected levels only 25% of the time might be its basis.

Risk margins have the advantage of being simple, easily included in pricing projections and are fairly easy to explain to management. On the other hand, they suffer from a number of drawbacks:

1. While such margins appear to be conservative, this may not always be true. For example, mortality selection on inforce business increases during periods of high lapse.

2. Such margins aren't generally dynamic and there are many cases when they probably should be increased if changes in other assumptions occur. (The above example also serves here).

3. Without adjustment, they bias pricing conservatively due to the fact that they don't allow for more favorable results. On the other hand, they don't specifically allow for worse results either (e.g., expenses much higher than expected).

4. These margins recognize the volatility of pricing results and the resulting risks only indirectly.

Thus, while commonly used and certainly better than nothing, simply using risk margins in pricing can be improved upon. This is probably true even if such margins are determined using a considerable amount of sophistication.

III. Risk Surplus

One technique for handling risk is to recognize that insurance products use the surplus of the company to support their risks as well as to fund acquisition costs. This method is generally used in conjunction with pricing methods which are based on the book value of gains or losses, such as Anderson's method, which are commonly used in the U.S. Under this method, as described by Sondergeld (1), the amounts of risk surplus (also known as "target," "required" or "benchmark" surplus) are treated as if they were product reserves. Thus, investment income on such amounts is included in book gains and changes in risk surplus increase or decrease book gains.

This approach to recognizing pricing risk is attractive for several reasons, one of which is its intuitive appeal. Since surplus is obviously needed to support even the most profitable in force business in a viable company, it only seems proper to recognize this dependence up front in the pricing process. This approach also meshes quite well with the concept of the cost of capital and requiring an adequate return on its use. It does this by treating the risk surplus as being unavailable to the company for any purpose other than that of supporting the associated business over its lifetime (or at least its pricing horizon).

Risk surplus, by various names, is used by rating agencies and state regulators to some extent currently. For example, as part of evaluating a life insurance company to assign a claims payment rating, Moody's Investors Service computes a ratio of capital to "required" surplus (2). The states of Utah and Wisconsin have "compulsory" surplus formulas (3) and may restrict operations of companies whose surplus falls below these amounts. Also, this concept is included as part of the Actuarial Standards Board's exposure draft on company appraisals (4) (which are a kind of pricing, albeit for companies or entire blocks of business).

Another appealing feature of this method is that it also relates to valuation approaches which recognize risk and attempt to set reserves accordingly (e.g., N.Y. Regulation 126). Consider a stochastic evaluation of a product line or company which determines an asset level such that additional assets are not needed to support future obligations at some level of confidence. To the extent that this differs from the statutory reserves held, the excess (or sufficiency) provides a very good and theoretically appealing definition of risk surplus. Of course, in jurisdictions where this is the normal method of reserve determination, risk surplus will be zero except to the extent that the confidence level used here exceeds that used in setting the statement reserves.

Attractive as the required surplus methodology is, it still has some serious drawbacks. In the first place, it may not always be possible to do the stochastic calculation referred to above. When it can be done, it is a fairly complex process to do the needed calculations at any point in time. To use risk surplus computed this way in a book value based pricing projection, it would appear that such a calculation would be needed at the end of each future period in order to establish the risk surplus at that time. Even if we have the raw computer power necessary to do this, if many pricing projections are needed (e.g., sensitivity testing, or the random walk process of reaching final agreement on a product's design), we will rapidly go far past the point of diminishing returns.

One practical solution to this problem is to use an approximate formula for risk surplus. This could be done by fitting some function (e.g., a linear combination of various product values such as reserves, mortality exposure and premium) to the results of the above stochastic calculation on similar current business. This would also have the advantage of being more understandable to other members of management and may lead to a more disciplined environment. Perhaps a more practical approach is to recognize that in one way or another external measures of risk surplus may impose limits on company operations. Accordingly, we can simply use a formula based on the jurisdictions and/or rating agencies which are either most important to the company or most conservative (if this can be lived with). In this regard, it might be a good idea to compare the resulting levels of risk surplus against the more theoretically "correct" computations to see how much extra conservatism, if any, is produced compared to what the company would hold if not concerned with external requirements. This relationship is likely to vary as the mix of business changes and new products are introduced. Even slight variations in product design can have a disproportionate impact on the amount of risk.

Another potential problem lies in the fact that risk surplus, while based on a distribution of possible results, contains only limited information as to what these results are for a given product. Even if our risk surplus was tailored to the product being priced, by itself it tells us little about how potential product results are distributed. While a given level of risk surplus protects against all but the worst losses, it doesn't tell us whether these are moderate or extreme. If the latter is true, one could argue that an extra charge should be made. However, the risk surplus mechanism, by itself, will not do this.

Another fundamental problem exists for this approach. At least for pricing methods which charge for capital use (where its use is most attractive), risk surplus can be shown to be a special form of risk charge as follows:

Let RS_t be the risk surplus at the end of period t
($RS_0=0$, at least when pricing new business)

i_t be the investment rate earned on risk surplus in period t

i^D be the discount rate used in discounting book gains

$PV_H(RS)$ be the present value of the risk surplus portion of book gains over the pricing horizon H

Assuming that the earnings on risk surplus arise only from amounts held at the beginning of each period and that each period's gain is discounted from the end of the period (5), we have:

$$PV_H(RS) = \sum_{t=1}^{H}\left[RS_{t-1}(1+i_t) - RS_t\right]/\left(1+i^d\right)^t$$

$$= \sum_{t=2}^{H} RS_{t-1}(1+i_t)/(1+i^d)^t + RS_o(1+i_1)/(1+i^d) - \sum_{t=1}^{H-1} RS_t/(1+i^d)^t - RS_H/(1+i^d)^H$$

$$= \sum_{t=2}^{H} RS_{t-1}(t+1_t)/(1+i^d)^t - \sum_{t=2}^{H} RS_{t-1}/(1+i^d)^{t-1} + RS_o(1+i_1)/(1+i^d) - RS_H/(1+i^d)^H$$

$$= \sum_{t=2}^{H} RS_{t-1}\left(i_t - i^d\right)/\left(1+i^d\right)^t + RS_o\left(i_1 - i^d\right)/\left(1+i^d\right) + RS_o - RS_H/\left(1+i^d\right)^H$$

$$= \sum_{t=1}^{H} RS_{t-1}\left(i_t - i^d\right)/\left(1+i^d\right)^t + RS_o - RS_H/\left(1+i^d\right)^H$$

If beginning and ending risk surplus is zero, recognition of risk surplus as a capital need boils down to a charge of a percentage of risk surplus at the start of each period. Thus, this technique is subject to all of the drawbacks mentioned above for risk margins, no matter how sophisticated the actual risk surplus computation is. In this case the the method will become less conservative as assumed earned rates increase or if pricing discount rates are lowered for any reason. This can be especially insidious, since in either case the overall product risks may actually be increased. In particular, it should be noted that the charge actually becomes a subsidy if the earned rate exceeds the discount rate. Accordingly, this pricing methodology by itself should only be used with due regard for what it actually recognizes and charges for. This is primarily capital use; product risks are only indirectly addressed.

IV. Stochastic Pricing

One way to reflect actual product risks directly in pricing is to generate pricing projections under adverse conditions. A time honored way of doing this is to perform sensitivity testing of the various assumptions used in pricing. This can provide a good feel for the effects of deviations in pricing assumptions from our "most likely" or "best estimate" assumptions. However, reflecting this in the actual prices charged to customers can be a fairly subjective exercise.

A method which is becoming more popular in the U.S., especially with increased computer capabilities, is to project product results under a number of randomly chosen scenarios. Each scenario consists of a set of assumptions which affect the results of the product in question and differ from one another only in the one or more assumptions which are chosen stochastically. The results are then weighted by the probability of the occurrence of the scenario that generated them to arrive at weighted results from which the "price" or charges for the product are determined. Using this technique requires a good knowledge of the various assumptions used and their effect on the product in questions. Experience, judgement and perhaps some initial sensitivity testing are likely to be needed to obtain usable results.

Besides a fair amount of computer power, this technique implicitly requires that we have the ability to assign some reasonable probability distribution to the various assumptions used in projecting our results. Typically, this is done only for the one or two assumptions that are most volatile or are of most concern for other reasons. While there is no theoretical reason why several independent assumptions couldn't be modeled stochastically, the quality of the model's results may not be improved enough to justify the extra effort.

Another requirement for obtaining worthwhile results from stochastic models is to reasonably address the interactivity of the various assumptions used, such as interest rate levels and withdrawals of funds. This is not to say that such assumptions aren't used in deterministic pricing. To the contrary, they are implicitly made "up front" in choosing assumptions (e.g., mortality rates are based on anticipated persistency). However, such assumptions remain locked in rather than varying with other assumptions upon which they are logically dependent (e.g., if lapses increase, mortality tends to increase since impaired lives tend not to lapse). An important part of this process is to model the responses of the product, customer and company to various situations or changes in them. Thus, items such as the company's policy for crediting excess interest, experience rating and its investment strategy relative to the product in question are often specifically modelled.

An advantage of stochastic pricing is that preconceptions of what scenarios to use in pricing are minimized or avoided entirely. Of course, this is worthwhile only if the choice of scenarios is as unbiased as possible. If this is the case, the big advantage of stochastic pricing is that the distribution of possible results is much better revealed. This may be especially true for extreme results; it has been our experience that many times one or more stochastic scenarios will produce worse results than even contrived "worst case" assumptions generate. If a sufficient number of scenarios are generated, it is possible in many cases that the shape of the tails of the distribution are revealed. Since this is where the worst risk emerges, these tails are of major importance in risk based pricing.

V. Recognizing The Acceptability of Risks

One potential drawback of the above pricing methods is that they do not address the acceptability of various financial outcomes that can arise due to risk. For example, if we project that we could lose .1% of surplus on a product under one set of assumptions as to future experience, management might view this with a fair degree of equanimity. On the other hand, if another set of assumptions showed that we'd severely reduce our surplus (say by 20%), management might become quite concerned unless this was very unlikely to occur.

If we simply averaged these results, even if weighted by their probability of occurrence, we would not be taking these very different levels of discomfort into account. In economic terms, we would be assuming that insurers were indifferent to risk. However, this doesn't seem to be a very realistic view. For instance, if this view actually prevailed there would probably be a lot less reinsurance.

To recognize our sensitivity to risk, we can apply utility theory. To do this we need to establish a methodology for evaluating the various outcomes from the product being priced by their utility, defined as a rank ordering of the preference (or avoidance) we have for them (6). We must first define the outcomes we are ranking. Just about any pricing measure (e.g., present value of gains, asset share, rate of return, etc.) will work, but it should be based on the company's philosophy and objectives. If this is not done, it will be almost impossible to obtain meaningful and consistent statements of utility.

This may seem so obvious that it doesn't need to be stated. However, this has always been one of the most difficult parts of pricing because different parts of the company may have different views. For example, if the sales area is judged only on how much business they produce, they will tend to pick a different profit measure than other areas of the company. If left unresolved, this can lead to a condition of corporate schizophrenia where one side or the other "wins," but poor results usually follow to the detriment of all concerned.

The pricing measures mentioned above are based on aggregate results over the entire pricing horizon. However, for some products or lines of business, the primary concern may instead be with the period by period emergence of profit and loss. For example, consider a product with results that are very volatile from period to period. While results over the long term may appear generally favorable, large fluctuations can seriously affect the company's overall operating results or even reduce surplus. For example, certain accident and health coverages may have this pattern of earnings. One way this can be handled is to express each period's results in terms of its utility. This will have the effect of recognizing the effects of the extreme fluctuations rather than averaging them out. This will be likely to have a different pattern of utility than one based on the more usual aggregate results for all periods combined.

Given that management has defined the measure of product outcomes to use (and this is accepted by all areas of the company), we can proceed to establish their utility. Due to the large number of possible product outcomes and the fact that we don't necessarily know them in advance, it is useful to derive a utility function, U(O), which transforms any possible outcome, O, into a utility value, U(O), by fitting a curve to specific points, O, U(O). As long as our defined measure of outcomes is a positive attribute (e.g., profits, present value of gains or rate of return), theory tells us that the first derivative of this function, U'(O) (also known as the marginal utility), is positive. In other words, more is preferred to less (7). However, it is the second derivative of U(O) that has the most effect on pricing, for this reveals the degree of risk preference (or acceptability of possible outcomes). If U"(O) is positive, then utility increases faster as the value of O increases. This is known as risk seeking behavior (or sometimes a lottery mentality). On the other hand, if U"(O) is negative, utility drops more and more rapidly the worse the outcome is. This is called risk averse behavior (which most people usually display when faced with relatively large losses in relation to their net worth). Finally, if U"(O) equals zero there is no preference and behavior is said to be risk neutral (the typical response to situations where amounts to be gained or lost are relatively small).

As in the determination of the measure of outcomes, senior management should also be involved here since they make the ultimate decisions concerning corporate policy and bear the ultimate responsibility. In a very real sense, how top management feels about any given outcome determines how acceptable or how abhorrent it actually is to the company and the best way to find this out is to ask them. To do otherwise substitutes the judgement or prejudices of someone less informed and this is likely to introduce even more uncertainly into what is already a highly subjective process. Finally, senior management are probably more involved in the company to the point where they are less likely to consider its fortunes to be unrelated to their own. To the extent this is true, it tends to minimize the familiar tendency to be much less risk averse with other peoples' money.

A direct way to proceed is to ask management to rank several alternate outcomes in their order of preference by simply assigning a score to each. The quality of results here will depend on how well management can deal with assigning these quantitative values. The range of scores doesn't matter much, but duplicate responses should be avoided. Positive and negative values could be allowed to distinguish between attractive and repellent outcomes, but we should be cautious lest this introduce biases into the process.

Of much more importance is the range of outcomes chosen. This should be wide enough to encompass most of the possible outcomes, or a significant portion of risk may be ignored. The number of outcomes chosen to be ranked doesn't have to be too large, but they should not be clustered. As few as three or four points will give a good indication of our utility curve's shape and a few more points should refine things fairly well. Thus, a great deal of management's time shouldn't be needed in making actual responses, though gaining an understanding of the process and discussing results will also require some of their time. Similarly, this process need not be repeated for every product, since the company's risk aversion (or lack thereof) should be fairly similar over a broad range. However, if results are anticipated which are much more extreme than were previously evaluated or management changes significantly, it may be necessary to do another evaluation.

How our outcomes are expressed can be especially important when they are defined as dollar amounts. For example, if these are small percentages of total company surplus in all cases, very little risk preference is likely to be displayed and management will tend to indicate that they're risk neutral. This would also tend to happen if dollar amount outcomes are presented in the per unit mode typically used in pricing. However, this can happen even if the outcomes from the entire volume of the product estimated to be produced are used if the results are relatively small. The result will be that smaller lines and low volume products relative to total company operations will not reflect much of management's preference in their pricing. Thus, it may be better to scale our outcomes in terms of indices that relate well to product risk. For example, if our chosen outcome measure was in terms of accumulated surplus at the end of 20 years, we might express this in terms of a ratio to initial invested surplus or our target for accumulated surplus.

On the other hand, if management is willing to limit its writings of the product in question or to increase its price if volume is significant, low risk aversion may be a correct assumption. The problem with this approach is that if volume exceeds expectations, a process which might be called dis-economy of scale can occur where you can't really afford to write more business. The danger here lies in becoming dependent on the business to the point where either increasing price or limiting sales is unacceptable. Another potential pitfall can occur when limiting production if this kind of strategy is used many times for similar risks since a lot of small risks can add up and become a major concern.

In the ranking process, the outcome alone should be considered; unless part of the defined outcome, items such as sales levels, administrative costs and investment needs should not be taken into account. Care should also be taken to clearly describe outcomes and possibly their consequences as well to make choices more concrete. For example, if our defined outcomes are after tax returns on invested surplus (ROI), it might be worth explaining their calculation. In addition, items such as (i) company growth in excess of ROI can't be supported on a long term basis, (ii) the company can earn more at lower risk without selling the business if an ROI at or below the rate on high quality tax exempts is earned and (iii) anything below a zero ROI irrevocably consumes capital, might help to put the whole process in perspective.

Bias must be minimized as much as possible in this process. Besides obvious slips such as stating that we'll never sell anything above a certain profit outcome, more subtle forms of bias can occur. For example, if our outcomes are defined as the present value of gains at some desired rate of discount, exactly attaining this rate would produce a zero value and less favorable results would show values that were negative. If this negative number is interpreted as an actual loss by management, their ratings of such outcomes might be materially lower than otherwise. Generally, better results can usually be obtained if the outcomes being presented are arranged to avoid perceived discontinuities or prejudices.

Getting responses from several members of top management might be advantageous in several respects. It provides more data and can serve to smooth results. Also, if responses don't correlate well, it may indicate a lack of understanding, poor explanation of the process, bias or a some combination of all three. In addition, the inherent differences in risk aversion between various members of management will be indicated. In any event, the responses should be analyzed and discussed with all concerned parties. This not only lets them see results and possibly gain consensus, but may also provide an opportunity for a delphi-like process.

An alternate way to determine our utility function is to use a "reference lottery" (8) consisting of "good" and "bad" outcomes, O(G) and O(B) which we have chosen to rank. These needn't be the most extreme values we can think of, but they need to be at each end of the range of outcomes to which we want to assign utilities. We can then assign some arbitrary "scores," U(O(G)) and U(O(B)), to these. Next, a probability, P, of "winning" (getting the "good" outcome) is assigned by our management respondents. They are to assume that if they don't "win" then they achieve the "bad" outcome. This probability is picked for each of the intermediate outcomes, O(I), so that the respondent feels that the resulting reference lottery is just as attractive as being certain to obtain it. Then our utility for the intermediate value is equal to the expected value of the reference lottery:

$$U(O(I)) = p\,U(O(G)) + (1-p)U(O(B))$$

For example, assume that the "good" outcome chosen is +60 and the "bad" one is -100 and we assign them arbitrary scores of 6 and -10. Suppose that an intermediate outcome of zero is deemed by our respondent to be the same as a bet on the reference lottery with a 75% chance of a "win." The utility of our intermediate point is then 2=6x.75-10x.25.

At first glance, this positive utility value for an outcome of zero may appear quite strange. However, it should be remembered that our scores were arbitrary. More importantly, our objective here is to obtain the shape of the utility curve to use in weighting our various projected outcomes and the absolute utility values aren't important. Finally, it may be worth something simply to avoid a loss. This is indicative of conditions which foster reinsurance.

Having established utility values for intermediate outcomes, we can fit a utility curve which is hopefully representative of our company's attitude toward risk. At this point, we can get back to actually pricing the product using the outcomes of our stochastic pricing process. Instead of averaging these, we average their utility values. Thus, if p(O(S)) is the probability of outcome S, then the expected utility value, $\tilde{U}$, is used instead of the expected value of outcome. Thus,

$$\tilde{U} = \sum_{S} p(O(S))U(O(S))$$

The final step here is to find a price or charges which produce an outcome such that its utility equals the expected utility. This process will be shown in more detail in the following section.

In the above, the utility functions are applied to an outcome for each scenario. However, this doesn't have to be the case and for certain products an alternative procedure might be worthwhile. Thus, if volatility in each period was the major concern, the average utility for each period could be computed instead and then used in the rest of the pricing process. This process can be used in conjunction with a hurdle rate on required surplus. In this way, both risk and surplus use can be reflected in the pricing process.

One would generally expect the utility function which applies to show the company to be risk averse given the widely held image of insurance companies as being prudent, responsible and conservative. On the other hand, this also used to be the image of savings and loan companies in the U.S. With this in mind, if your company is risk seeking, it is probably a good thing to know up front. However, if the conventional wisdom holds, pricing will become more conservative than it would be absent a risk averse posture. Therefore, if the company doesn't actually avoid risks, it will tend to want to charge more for them.

Unfortunately, in an era of reduced margins and increased competition, this makes it far more difficult to develop a product which is competitive in the market place. To the extent this is the case, there are several possible courses of action:

1. Other than avoiding the more down side risks altogether, it may be possible to reduce them by analyzing the product and its operation to find features of the product or methodologies associated with it that can be changed. Examples of this are restrictions on cash withdrawals and matching assets to liabilities.

2. The risks can be reinsured. This requires a reinsurer that is less risk averse than your company and may be prohibitively costly, depending on the risks being transferred and the competitiveness of the reinsurance market.

3. Reduced profitability can be accepted. Since this involves waiving senior management's agreed to profit goals and/or charges for risk aversion, only they should make this decision (though it is sometimes made by default). At least knowing the company's utility function enables the actuary to better quantify the effects of such a decision and to focus on those products which come closest to the desired goals.

4. Offsetting risk factors in the rest of the company might be found. To the extent that these can be matched with the risks of the product in question and are not already used elsewhere, the risk used in pricing can be reduced. Of course, one inherent danger here is that the offsetting risk may change or disappear and this should be a consideration in any decision.

VI. Case Study

The following is an example of the type of pricing approach discussed in this paper. The product used for this example is a single premium deferred annuity (SPDA) contract. In the U.S., this type of individual annuity contract generally provides for cash values equal to the contract balance less any surrender charges. The particular product used for this illustration has a seven year surrender charge beginning at 7% of the initial premium amount and declining 1% per year to zero after the seventh policy year. A single premium is accepted when the contract is issued. The deposit and any subsequent available cash are invested in six year bonds with a call provision at the end of the fifth year. The product has first been priced assuming that current rates do not change to return a 15% after tax return on invested surplus. This produces a required asset margin which does not vary depending on the investment strategy or the renewal rate setting strategy.

Next, product results were projected under a set of 90 different stochastic interest rate scenarios. These were generated using a log-normal distribution of the six year bond rate. Two different interest crediting strategies were tested. One method is to set the renewal year interest guarantees by subtracting the margin determined above from the rate earned on the underlying assets subject to a minimum interest credit of 3.5%. This is sometimes called a "credit-to-margin" strategy. The other method is to set the renewal rate equal to the the rates offered to new contracts ("credit-to-market" strategy).

If we simply average our results (assuming that each one has the same probability since each scenario was randomly selected) the additional margins needed to retain our 15% return goal were found to be 5.6 basis points (bp) for the credit-to-margin strategy and 16.5 bp for credit-to-market. Thus, the credit-to-market strategy would be have an initial interest guarantee 11 bp (16.5 bp - 5.6 bp) less than the credit-to-margin strategy.

If we incorporate these added margins, our results are now distributed as follows, expressed as averages of each decile in the form of present values of book profits as a percentage of initial premium:

Decile	Credit-to-Margin	Credit-to-Market
1	1.82%	4.05%
2	0.45	2.89
3	0.24	1.81
4	0.01	1.28
5	0	0.46
6	-0.08	-0.27
7	-0.21	-0.92
8	-0.36	-1.59
9	-0.62	-2.80
10	-1.36	-4.92

If the utility function approach is to be used, a particular function must be chosen. Two different families of functions are illustrated below. The first is based on a function used in a presentation made by Shane Chalke. What we desire to determine is the charge (here in terms of extra margin) to obtain an outcome, O, such that:

$$\left(1 - e^{-rO}\right) = \sum_{S} p(O(S))\left(1 - e^{-rO(S)}\right)$$

The value of r is used to determine the degree of risk averseness. The level of r determines the degree of the increased charges for higher levels of loss. This formula puts a limit on the contribution which a positive value can make while negative values can become quite large. Similarly, as r gets large, the weighted value approaches the worst case result. The following table presents results for several values of r n terms of additional margins needed to charge for the volatility of the product.

Factors	Credit-to-Margin	Credit-to-Market
r=.01	.3 bp	1.5 bp
r=.05	.8 bp	7.2 bp
r=.1	1.5 bp	14.4 bp
r=.3	4.1 bp	43.7 bp
r=.5	6.6 bp	71.0 bp

As can be seen from this table, the results are quite sensitive to the selection of the value of r, especially for the credit-to-market strategy. This is due to its higher relative volatility.

Another family of utility functions is the following:

$$U(O(s)= \ln (O(s)+1) \quad \text{for } O(s) \geq 0$$

$$(1-(1-O(s))^k)/k \text{ for } O(s) \leq 0$$

Under this formula, relatively less weight is given to the more extreme negative outcomes, but it treats positive ones somewhat more favorably. This utility function and its derivative are both continuous where the definition changes. It also shows that it is possible to construct our utility function in two or more segments if desired.

Factors	Credit-to-Margin	Credit-to-Market
k=1	4.3 bp	22.8 bp
k=2	21.7 bp	65.4 bp

Comparing these results to those above shows that, the credit-to-margin results are not quite so good compared to credit-to-market. This should illustrate the obvious need for being as accurate as possible in the determination of our utility function.

FOOTNOTES:

1 Sondergeld, D. R. "Profitability as a Return On Total Capital," Transactions Of The Society Of Actuaries, XXXIV(1982), pp. 415-429.

2 Moody's current formula for "required" surplus has three components: (i) asset value (C-1) risk which varies from 1% for investment graded bonds to 25% for common stock (and 100% for affiliated common), (ii) insurance (C-2) risk of 25% of life tabular cost, 10% of group A&H premium and 25% of individual A&H premium and (iii) interest rate (C-3) risk of 1% of group annuity reserves, 2% of ordinary life and 3% of individual annuity reserves.

3 Wisconsin requires a "compulsory" surplus of 15% of individual non-annuity premium plus 10% of group non-annuity premium plus the greater of 7.5% of annuity premium or 2% of annuity reserves. Utah's compulsory surplus is 5% of life insurance in force (net of reserves) plus 3.5% and 1% of reserves on annuities with and without cash value plus 10% of A&H premiums plus 5% of common stocks (excluding subsidiaries) and real estate plus 2% of other assets (except policy loans) less the mandatory securities reserve plus the compulsory surplus of subsidiaries.

4 Actuarial Standards Board. "Actuarial Appraisals Of Insurance Companies, Segments Of Insurance Companies, and/or Blocks Of Insurance Companies," Exposure Draft Actuarial Standard Of Practice (April, 1990), pp. 5.

5 This differs from Anderson's method which discounts to the start of each year at the earned rate. Using this convention would merely change the result shown by the ratio of one plus the discount rate to one plus the earned rate.

6 Hirshleifer, J. Price Theory And Applications. Englewood Cliffs, N. J.: Prentice-Hall, 1984, pp. 61.

7 Ibid, pp. 60.

8 Narragon, E. A. A Study Manual For Operations Research. Chicago, IL.: Education And Examination Committees, SOA and CAS, 1980, pp. 4-64 to 4-68.