**RECORD OF SOCIETY OF ACTUARIES
1978 VOL. 4 NO. 1**

CLASSIFICATION MODELS

Moderator: CHARLES A. HACHEMEISTER. Panelists: JOSEPH FERRIERA**, ARNOLD F. SHAPIRO, OAKLEY E. VanSLYKE*

1. Models demonstrating the effect of a classification system superimposed on a merit rating system.

2. Models demonstrating the subsidization of poorer risks.

MR. JOSEPH FERRIERA: Proponents of unrestricted risk classification often point to statistics as the justification for their position. They argue that it is both efficient and equitable for each identifiable class to be "self-supporting" -- efficient since such an approach permits the private market to work smoothly and economically to satisfy the demand for insurance; and equitable since insureds with similar characteristics pool their risk by paying a price statistically estimated to match the average annual cost per insured in their class.

Such a rationale might be appropriate if these class averages were especially meaningful, so that each insured's risk was indeed close to the average for his class and distinctly different from those in other classes. But it is unrealistic to expect risk classes based on demographic and geographic characteristics to be that accurate. For inaccurate classes, the class mean has no such special meaning -- why not charge the median rate or the modal rate? The arguments for class mean pricing in such cases have more to do with efficiency considerations and private market mechanisms than with equity. Determining prices that are equitable to the insureds requires a closer look at how risk classification redistributes costs.

This afternoon I shall discuss my view of how these equity issues should be examined. The basic theme is that imperfect classes should be priced in a way that spreads the errors as evenly as possible among insureds. But such an approach can lead to equitable rates quite different from the class average rates that might arise in a perfectly competitive private market. I believe that this distinction is an important part of the current debate about classification and it is cause for some rethinking of the customary view of what is actuarially fair -- a view that, at least for auto insurance, now appears too simplistic.

Some discussion of the various risk spreading and public interest concerns that relate to risk classification will help place these notions of equity in perspective. I think that interesting differences in points of view about equity and some rationale for modification of classification have more

*Mr. Hachemeister, not a member of the Society, is a Fellow of the Casualty Actuarial Society.

**Mr. Ferriera, not a member of the Society, is an Associate Professor of Urban Studies and Operations Research at Massachusetts Institute of Technology.

to do with how these concerns are measured and put together than with a debate over their general relevance. What I shall do today is sketch four different perspectives and then relate them to my recent experiences and research working at the State Rating Bureau in Massachusetts during the past year.

The conventional perspective focuses on the predictability of classification factors and encompasses what one might call the "traditional rate making approach". Accordingly, the primary goal is to price insureds based on their expected costs. Underlying this goal is the notion that the price of say auto insurance ought to vary in direct relation to the risk inherent in the particular policyholder so that insurance spreads only the fortuitous risk associated with the random timing of claims. This focus has been part of the rationale for pricing each class subcell at its estimated average cost and for identifying and using new factors that are statistically correlated with the inherent risks of insureds. This perspective, carried to its fullest, says that predictability is the primary concern of classification systems and that there must be severe problems with the unfettered use and pricing of class factors before government should intervene. Hence, legal and social welfare issues are viewed as constraints and the overriding concern is to look for improved predictability with whatever information is available.

One alternative approach might be to view classification plans as useful only to the extent that they provide the right incentives. An element of this rationale might be that differentiating on the basis of convenient, economically viable factors, particularly factors that are uncontrollable, tends to be counterproductive. If one believes that feasible class plans do poorly in terms of individual accuracy, then the economic benefit obtained by using them may not be that large in terms of allocative efficiency gains. In fact, the use of a number of such imperfect and uncontrollable factors might hide the impact of the few factors that do provide the right incentives and are controllable. Rather, controllable factors such as merit rating and reliable mileage indicators ought to be the focus of attention.

A third approach is what I'd call, for want of a better word, a government responsibility approach. Here one is saying that there are externalities that affect the expected cost associated with one individual which are not his responsibility as a result of say traffic congestion or high crime neighborhoods. These externalities add another dimension to the insurance issues just discussed. To the extent that auto insurance is compulsory or a social necessity, these issues may bring to bear a broader government responsibility. Likewise, affordability issues raise questions about broader government responsibility.

Yet another perspective might be what I'd call a due process or casualty approach. Here the overriding consideration is the error resulting from the use of crude proxies for casually related variables. The argument is that, at best, such factors as age, sex, and territory are simply administratively convenient factors that happen to be correlated with a more fundamental set of reasons as to why insureds have different inherent risk. This approach says that, because of due process reasons, variables such as these may be economically convenient and statistically significant but are still suspect and may be used only by those who demonstrate the lack of alternatives and the certainty of unduly burdensome economic consequences.

Each of these four perspectives was discussed in the course of the 1978 Massachusetts rate hearings for auto insurance. However, the analytical models that were used to obtain the specific pricing recommendations of the State Rating Bureau at that hearing were much less ambitious in scope and tended to be consistent with the conventional predictive accuracy perspective. Even so, these models indicated substantial departures from the more traditional ratemaking approach. Examining the reasons for these differences will indicate the theoretical basis for the equity concerns that I raised earlier. What I shall do is briefly indicate which aspects of the Massachusetts decision made use of theoretical arguments that fell under one of these various perspectives and then focus on one particular theoretical argument that I developed concerning possible inequities when each risk class is priced at one's best estimate of its break-even cost.

The major changes in the use of classification factors reflected in the decision are, I believe, these four: (1) the elimination of age, gender, and marital status factors as allowable factors, (2) the flattening of certain expenses, including the deficit of the Reinsurance Facility, (3) the use of an additive rather than a multiplicative approach toward combining class and territory factors and (4) the pricing of Massachusetts' 24 territories at amounts intentionally different from the estimated territorial average costs. The arguments concerning the elimination of age, gender, and marital status factors dealt with each of the four perspectives but relied heavily upon what I have called the due process/casualty point of view. The due process disadvantages associated with the use of these factors were considered to outweigh the statistical accuracy and predictive ability benefits.

Each of the other changes has, as a theoretical basis, a conventional predictive ability perspective. Nevertheless, several unconventional conclusions resulted from taking a close look at the implications of expected cost pricing. For example, if we are really interested in expected cost pricing, it behooves us to examine the distribution of expenses across risk classes as carefully as we focus on claim costs. Certainly, the proportional loading of all expenses appears to be extreme and even the question of how best to allocate any deficit from assigned risk or reinsurance pools is debatable. Similarly, the discussion of additive and multiplicative approaches toward combining class factors is consistent with expected cost pricing. The idea is to identify a more accurate method of estimating the expected cost of insureds with various combinations of class factors.

Even the adjustment of territorial relativities in the 1978 Massachusetts rates is consistent with an expected cost pricing goal. These adjustments involve a so-called tempering of the territorial relativities that brings all of them closer to unity than are the indicated territorial average pure premiums. For example, this was accomplished by taking the indicated pure premium relativity for the 24 Massachusetts territories and raising each of them to the 0.9 power (with balancing) before they were used to determine territorial rates.

The reasons for this tempering are both the most complicated aspect of the decision and yet, in many ways, the most interesting. I would like to spend the remainder of my time discussing these theories since they raise some fundamental concerns about the equity of pricing individual risk classes at their class means.

In the 1978 Massachusetts rate decision, Commissioner Stone cited three reasons for pricing territories at something other than their class means. One had to do with the overlap between merit rating and class plans. Here the argument is that territorial relativities should be developed in a way that reflects the simultaneous presence of Massachusetts' new merit rating plan. Because merit rating affects the premiums paid by individual drivers, it is in itself a classification plan. To the extent that it is useful in distributing the burden of premiums toward those who have claims, other classification variables should become less important. This implies that the range of relativities found for any classification variable should be narrower when a merit rating plan is in operation than when it is not. Since the data used to estimate indicated territorial relativities in Massachusetts did not take merit rating into account, some such narrowing is needed in order for the combined effect of merit rating and territorial relativities to yield unbiased estimates of the average claim costs for each subclass.

A second reason cited for the adjustment of territorial relativities had to do with a concern about congestion effects in the city whereby suburbanites driving in the city make the driving environment more dangerous for city residents. If such externalities are judged to place an inequitable share of the costs on city residents, then a reduction of territorial relativities for the cities is in order. Such adjustments follow from what I have called the government responsibility perspective rather than from the conventional expected cost pricing goal. However, the theoretical rationale that was actually used in the decision to determine the numerical adjustments in territorial relativities came from yet a third rationale that is in fact consistent with an expected cost pricing goal. This rationale has to do with errors that arise when rather imperfect class factors are used and with the way in which the choice of a ratemaking technique affects how the burden of these pricing errors is spread among the insureds. The technical development of this approach is explained in a paper that I prepared for last Fall's rate hearing. What I shall do here is sketch the central elements.

The notion that a ratemaking technique that is equitable to insureds is one that does not price each class at its class mean is both counterintuitive and unconventional. However, I believe that this conclusion necessarily follows if one accepts four quite plausible premises that are consistent with a conventional expected cost pricing goal.

The first premise of the argument is that a good, credible estimate of the average loss pure premium for a territory isn't very typical of each and every insured driving in the territory. That's because the classes aren't especially homogeneous. The second premise of the argument is that the classes, each of these territories, tend to be more or less equally accurate in percentage terms. That is, two-thirds of the people are within say 20 percent of the mean for each class. The third premise is that what matters from the point of view of equity in pricing classes is not the percentage accuracy but the dollar accuracy. That is, you're worried about dollar differences between the price an insured pays and the expected costs associated with his inherent risk. The fourth element of the argument is that, in comparing pricing errors of different dollar amounts, you are disproportionally concerned about large overcharges. Hence one $200 overcharge is of more concern than ten $20 overcharges.

Now, what I suggest is that the expected cost pricing goal plus an acceptance of these four premises implies that the preferred rates are not going to be class mean rates. The preferred rates will be those which minimize the weighted sum of pricing errors within all the classes subject to the constraint that the total amount of premiums collected from all classes be adequate. These weights associated with pricing errors of different dollar magnitudes are set to reflect the disproportionate concern for large overcharges. Since the largest dollar overcharges will tend to arise in those classes with the highest class means, the preferred rates will charge amounts less than the class mean in the higher-rated territories and more than the class mean in the lower-rated territories.

Such an approach makes the premium for one class dependent upon the relative accuracy of other classes. If the classes provided accurate measure of each individual's risk and bore some identifiable relationship to inherent risk, then such dependence might be questioned. But the classes are not very accurate and the class factors tend to be administratively convenient factors such as age, sex, and place of residence that are economical to obtain. It is impractical and too costly for auto insurance ratemaking to utilize risk assessment procedures that include the more detailed techniques discussed in the highway safety literature. In such cases, the relatively low risks who fall in low-rated classes can be so identified only at the expense of other low-rated risks who, because of the inaccuracy of the class plan, are stuck in the high-rated classes.

Examining the question of equity along the lines that I have sketched above does not resolve the issue of how to sustain a preferred set of rates in the marketplace. However, I think that concern for the viability of a pricing scheme that deviated from class average prices or that ruled out certain class factors shouldn't prevent one from trying to explore in detail what the equity implications are. I do think there are ways to address some of these other concerns. We have already taken steps in this direction in Massachusetts. But such issues are best left for the question period.

In these opening remarks, I have tried to outline how some of the equity implications of various risk classification plans ought to be examined. I have sketched four different perspectives toward risk classification and suggested that even the most conventional of these -- the predictive ability approach -- does not necessarily lead to class mean pricing.

Let me conclude my remarks with a question that I find useful in trying to relate these equity issues to the interests of individual insureds. Once the choice of a class plan has been made, each insured acting in his self-interest would want the premium charged for his class reduced. But that is not the test that a regulator should consider in judging whether classification errors are distributed equitably. Instead, suppose a new auto insurance class plan were constructed that had about the same accuracy as a typical ISO class plan but was based on a new set of factors. Suppose you, as an individual insured, were quite confident that you were in fact a low risk but didn't know if you had erroneously been classified as a high risk under the new plan. In such circumstances would you prefer some amount of tempering in the rate structure?

MR. ARNOLD F. SHAPIRO: The more elaborate pension plan proposals and valuations typically append a pension cost projection to their report. Invariably, however, such projections are based on deterministic models and seldom, if ever, is there a quantified statement of the actuary's confidence in that projection. One suggested approach for overcoming this problem is to develop an array of projections based upon a spectrum of assumptions which encompass varying degrees of conservativeness in pension cost determinants, so as to provide some indication of the costs under diverse situations. While this approach has much to commend it, it suffers at least two serious shortcomings. First, it provides no mechanism for introducing credibility into pension cost projections. Ideally, attached to any estimate of projected pension costs should be a statement of the actuary's confidence in that estimate. Second, this approach provides no mechanism for incorporating the actuary's "feelings" regarding his confidence in underlying assumptions. Perhaps the most important attribute of an experienced pension actuary is his intuitive notion of what should be true and, ideally, there should be some vehicle for injecting this intuition into pension cost projections.

It has been argued that the refinement of pension cost projections is not an area of high priority. Proponents of this view reason that pension costs are funded sequentially over a number of years and that periodic actuarial valuations will uncover underfunding problems before they can materially affect the solvency of a plan. As a consequence, it is suggested that ex ante pension cost projections should be viewed strictly as rough (albeit best) estimates of ultimate pension plan costs. It is further suggested that the fact that such projections may not convey an accurate picture of ultimate cost has only marginal significance.

This view, however, presumes that the plan sponsor will be financially able to fund any deficiencies which arise. Furthermore, this proposition disregards the question of whether a particular plan or plan liberalization would have been introduced initially had the plan sponsor understood that the actual cost might be considerably in excess of the projected cost. These considerations have become increasingly important in light of the liability ERISA imposes on plan sponsors. Thus, while ex post reconciliation of pension cost estimates remains an important facet of pension cost funding, there are compelling arguments for developing techniques to measure the variability of ex ante pension cost projections.

The foregoing observations suggest the need for a stochastic model for projecting pension costs. A straightforward procedure would be to base such a model on direct or deductive probabilities. One could assume, for example, that the number of participants who succumb to a particular decrement is binomially distributed and based upon a probability of decrement which is constant or is given by a degenerate distribution. This assumption of an underlying degenerate distribution, however, is questionable in actual practice. Probabilities of decrement, for example, are obtained either from intercompany experience, which at best may only approximate the actual experience of a particular firm, or else it is derived from the firm's own experience, which for the majority of firms is not very credible. Thus, what is needed is a model in which underlying parameters also may take on probability distributions.

These additional considerations lead naturally to a Bayesian approach to stochastic pension cost projections. Under this approach, not only are pension cost determinants, such as the number of decrements due to a given case and the fund accumulation factor, assumed to be stochastically distributed, but the parameters upon which these determinants depend are themselves assumed to be stochastically distributed.

This paper discusses the use of a Bayesian approach to persistency in the exploration of retirement costs projection variability. Both the projection of pension populations and confidence intervals for retirement cost projections are considered.

The Probability of a Given Number of Participants at Each Age

The number of participants at a given age in a pension plan may be regarded as a random variable that depends on the number of participants at the previous age, which is also a random variable except at the entry age. Thus, in order to determine the probability of a given number of participants at a given age, vectors of feasible numbers of participants at each previous age must be constructed. These vectors might be called feasible arrays. If, for example, the number of entrants at age 20 were equal to 100, then the feasible arrays consistent with 98 participants at age.22 would be (100,98,98), (100,99,98) and (100,100,98).

The probability that the number of participants at a given age is equal to a given value, then, is simply the sum of the probabilities of each feasible array associated with the number of participants.

A Conditional Probability Distribution Function for the Number of Participants at a Given Age

In order to calculate the probability of a given feasible array of participants it is necessary to specify the probabilty distribution function with which we are dealing. Under the assumption that valuations are based only on curtate ages, the number of participants who persist through a given age may be thought of as being generated by a Bernoulli process under which participants either persist as active members or leave the active group. A conditional distribution for this process is the binomial distribution.

In Figure 1 the binomial distribution issued to project the distribution of the number of plan participants at each age through age 65, assuming there are 100 entrants at age 20. The probabilities of persisting are based on mortality rates from the 1971 Group Annuity Mortality Table; disability rates used in the 1970 Civil Service Pension valuation; and Turnover Table 111 given by McGinn. This data base, which is used for illustrative purposes, will subsequently be referred to as "the decrement data". The curve to the far right represents the distribution of participants at age 21. The curve to the far left represents the distribution of the number of participants who will retire at age 65. The intermediate curves are associated with participants at intermediate ages.

It is apparent from these curves that, even under conditions of perfect information, the actual number of participants at a given age may vary considerably from the best estimate of the number of participants. While this is not surprising, what is interesting is the considerable disparity which is likely to occur. In the graph the locus of the modes of the distribution of participants is convex. The age at which the locus attains a minimum value represents the age at which the distribution of participants is most nearly symmetrical. Below this age the distribution of participants is negatively skewed and above this age the distribution of participants is positively skewed.

It is important to recognize that the binomial mass function is appropriate only under the assumption that the exact probabilities of persisting are known. This assumption, however, is generally not valid. Although it is true that estimates of the probability of persisting are often available, these estimates may or may not be valid for the particilar pension plan under

FIGURE 1

DISTRIBUTION OF NUMBER OF PARTICIPANTS AT AGES 21 THROUGH 65, GIVEN THAT THEY ARE BINOMIALLY DISTRIBUTED

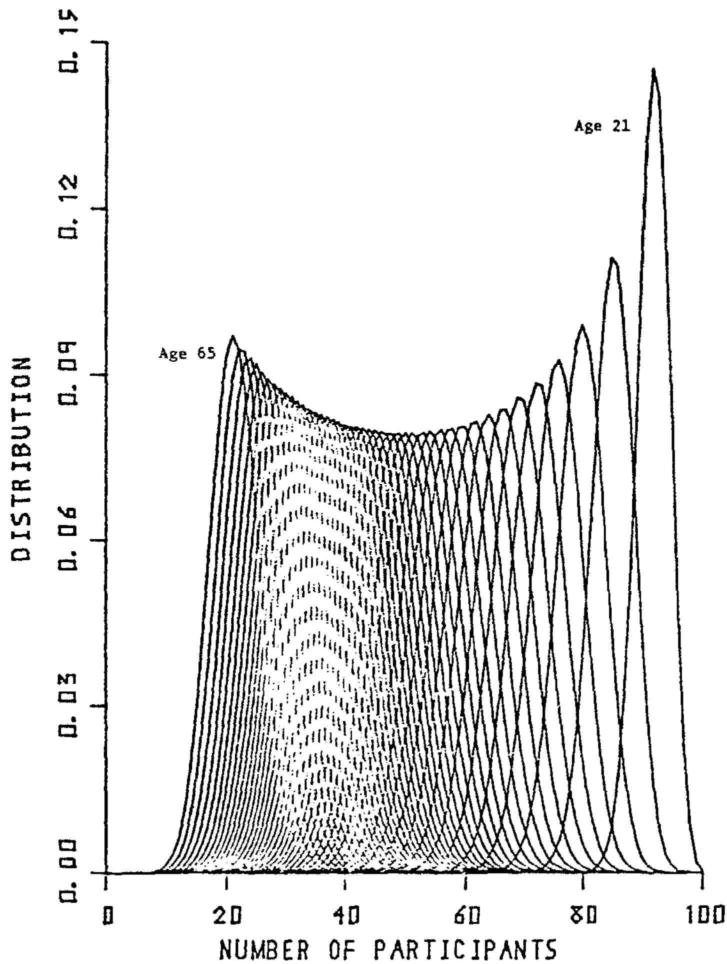

Data base: 1971 Group Annuity Table, rates of disability used in the 1970 Civil Service Retirement System Valuation, and Turnover Table III from Daniel F. McGinn, "Indices to the Cost of Vested Pension Benefits", TSA, Vol. XVIII (1966), pp. 235-236.

consideration. Furthermore, the binomial mass function provide no mechanism for the actuary to indicate the intensity with which he views the credibility of the estimated values. These criticisms suggest the need for a more general probability distribution function. What is needed is a distribution which is not conditional upon a degenerate probability, that is, an unconditional distribution.

An Unconditional Probability Distribution for the Probability of Persisting

Bayes' theorem offers a simple method for transforming the conditional probability of a given number of participants persisting to an unconditional probability. If it is assumed that the probability of persisting at a given age has a beta distribution, the unconditional distribution of the number of participants who persist through that age is beta-binomially distributed. Thus, a convenient choice for the probability of persisting is the beta distribution. The question, then, is whether this choice is reasonable.

Certain properties of the probability of persisting seem evident. First, it lies between zero and one, inclusive. Second, it may take any value on the continuum zero to one, so that it has a continuous distribution. Finally, for any given age, the probability of persisting may be concentrated at most one value, so that it has a single mode. For the purpose of this study it is assumed that any probability density function that is chosen to represent the probability of persisting must exhibit these properties. In addition to the empirical properties mentioned above, another desirable property stems from the fact that it may be impossible to specify the distribution of the probability of persisting exactly, due to a scarcity of relevant data. The distribution that is used to characterize the probability of persisting should lend itself to updating as more sample information becomes available.

The beta distribution satisfies all the empirical requirements mentioned above except that it is not necessarily unimodal. This condition is met, however, if its parameters are properly constrained.

The remainder of this paper assumes the beta distribution adequately describes the distribution of the probability of persisting, and that the beta-binomial distribution appropriately describes the distribution of the number of participants at a given age.

Estimating the Parameters of the Beta-Binomial Distribution

If it is assumed that the number of participants in a pension population is beta-binomially distributed, then the probability that the number of participants at age x will be equal to l_x is

$$\begin{bmatrix} l_{x-1} \\ l \end{bmatrix} \cdot \left[\frac{B(l_x + r,\ l_{x-1} - l_x + n - r)}{B(r,\ n-r)} \right], \quad \begin{array}{l} l_x = 0, 1, \ldots, l_{x-1} \\ n > r > 0 \\ \max\{r,\ n-r\} > 1 \end{array} \tag{1}$$

where

$$B(r,n-r) = \Gamma(r) \cdot \Gamma(n-r) / \Gamma(n)$$

The parameters of the distribution are n, which I shall call the precision parameter, and r, which I shall call the shape parameter. In practice, of course, it is unlikely that the exact values of the parameters are known, so it is necessary to estimate them. The method of moments in conjunction with subjective judgment provides a simple procedure for doing this.

It can be shown that if one assumes that the tabular probability of persisting (TPP) is approximately equal to the expected probability of persisting, where the latter has a beta distribution, then

$$r \doteq n \cdot TPP \qquad (2)$$

and the variance of the distribution will be approximately

$$\frac{TPP \cdot (1-TPP)}{n+1} \qquad (3)$$

It is clear that the estimated variance of the prior distribution of the probability of persisting will be inversely proportional to the size of the precision parameter. The greater the confidence in the TPP, the greater the precision parameter that should be chosen, that is, the smaller should be the estimated variance. Once an appropriate precision parameter is chosen, the shape parameter is determined by solving equation 2.

The Implementation of the Foregoing Procedure

Table 1 and Figure 2 exemplify the mechanics of the foregoing procedure by showing how one might determine a subjective prior distribution for the probability of persisting at age 20. Given the decrement data, the TPP is 0.918247. Table 1, which was developed by substituting this value into equation 3, shows the trend of the estimated variance for various choices of the precision parameter.

Table 1

Impact of Prior n On Variance

Prior Value of n	Variance
1+	.0359350*
2	.0250231
100	.0007507
500	.0001501
10000	.0000075
∞	.0000000

*Actual Value of n is 1.089+, which is the smallest value of n which is consistent with a unimodal beta distribution

An actuary who feels confident in the TPP might choose a precision parameter of 100 or more. This would result in a prior distribution which has a variance of .0007507 or less. For large precision parameters the distribution would, for all intent and purposes, be degenerate, and a binomial distribution might be used. On the other hand, an actuary may be satisfied that the TPP represents a good estimate of the mean of the distribution, but may, at the

FIGURE 2

EFFECT OF THE CHOICE OF THE PRECISION PARAMETER (n) ON THE BETA PROBABILITY DENSITY FUNCTION GIVEN A TPP OF 0.918

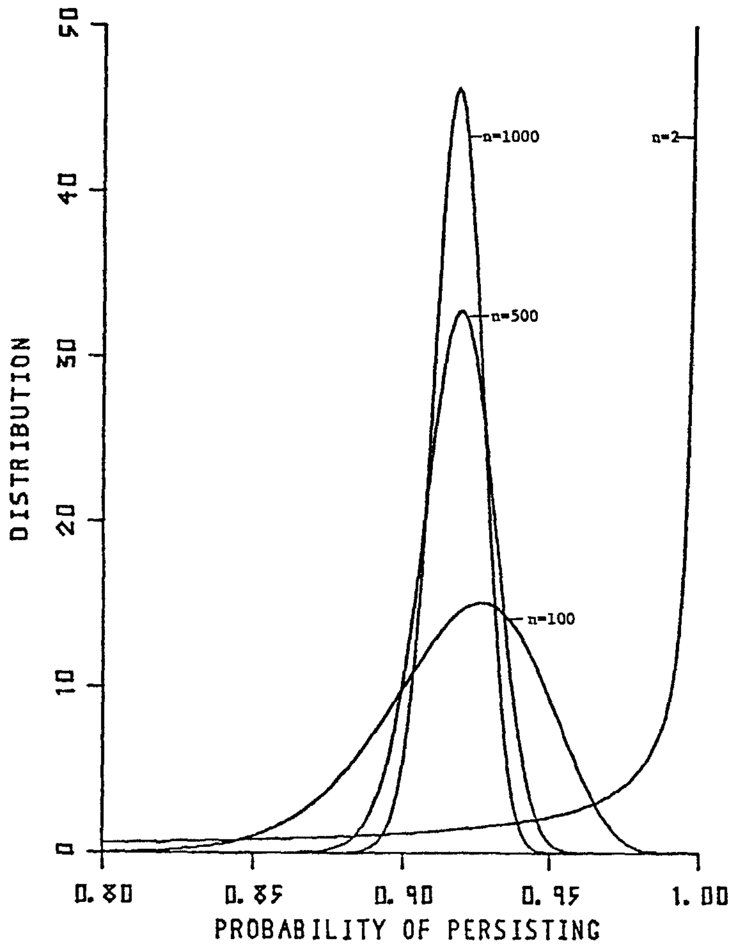

same time, feel that there is a considerable possibility of some other value. In an extreme situation of this kind, the actuary may be quite uncertain of the outcome and might choose to introduce considerable variability. This could be done by choosing a precision parameter equal to 2: a distribution with a mean of .918247 and a variance of approximately .025 would result. Any distribution between these two extremes would also be available.

It should be noted that the variance of the probability of persisting need not be the same for each age. The variance may, for example, be somewhat larger for the ages in the vicinity of the initial or full vesting age, where an actuary might be unsure of his best estimate. For other ages, where the impact of vesting might be slight, an actuary may have considerable confidence in his estimate and might choose a somewhat smaller variance.

Figure 2 shows the impact of various choices of precision parameter on the prior beta distribution of the probability of persisting at age 20. A precision parameter of 1000 causes the distribution to degenerate towards its mean. At the other extreme, a precision parameter equal to 2 results in a distribution which is highly skewed towards the origin and has a maximum value at unity.

The Projected Number of Retirees

Consider now the application of the beta-binomial mass function to the problem of projecting the distribution of retirees. Figure 3 shows the distribution of participants at age 65 resulting from 100 entrants at age 20, given the decrement data and various precision parameters. The projected number of retirees is 21.38. It is immediately apparent that as the probability of persisting at each age tends to degeneracy, the distribution of retirees tends to its limiting distribution. It should also be noted that the less credible the prior distribution of the probability of persisting, the greater the probability that the projected number of retirees will exceed the actual number of retirees.

Under a condition of considerable uncertainty, a precision parameter equal to 2 for all ages, the probability that the actual number of retirements will be exceeded by the projected number is 59.88 percent. This is due to the extreme skewed nature of the distribution of retirements under a condition of high uncertainty. Attributing a high uncertainty to the estimated value of the probability of persisting is tantamount to assuming that the probabilities of decrement may be higher than their best estimates indicate. Thus, there is considerable likelihood that the actual number of retirees will be exceeded by the estimated number of retirees. Under a condition of high certainty, on the other hand, a precision parameter approaching infinity for all ages, the probability that the actual number of retirements will be exceeded by the projected number is 52.09 percent.

Estimating the Adequacy of Projected Pension Costs

As a final example of the implementation of the beta-binomial distribution, consider the problem of calculating the probability that projected retirement costs will exceed actual costs. For the purpose of example, assume a pension plan where entry takes place quinquennially from ages 20 through 40, inclusive, with the proportion of entrants at each age being .28, .24, .18, .12, and .08, respectively. Assume, also, that the total number of entrants is chosen so that, if entry were to take place annually, an ultimate population

FIGURE 3

EFFECT OF THE CHOICE OF THE PRECISION PARAMETER (n) ON THE DISTRIBUTION OF RETIREES AT AGE 65

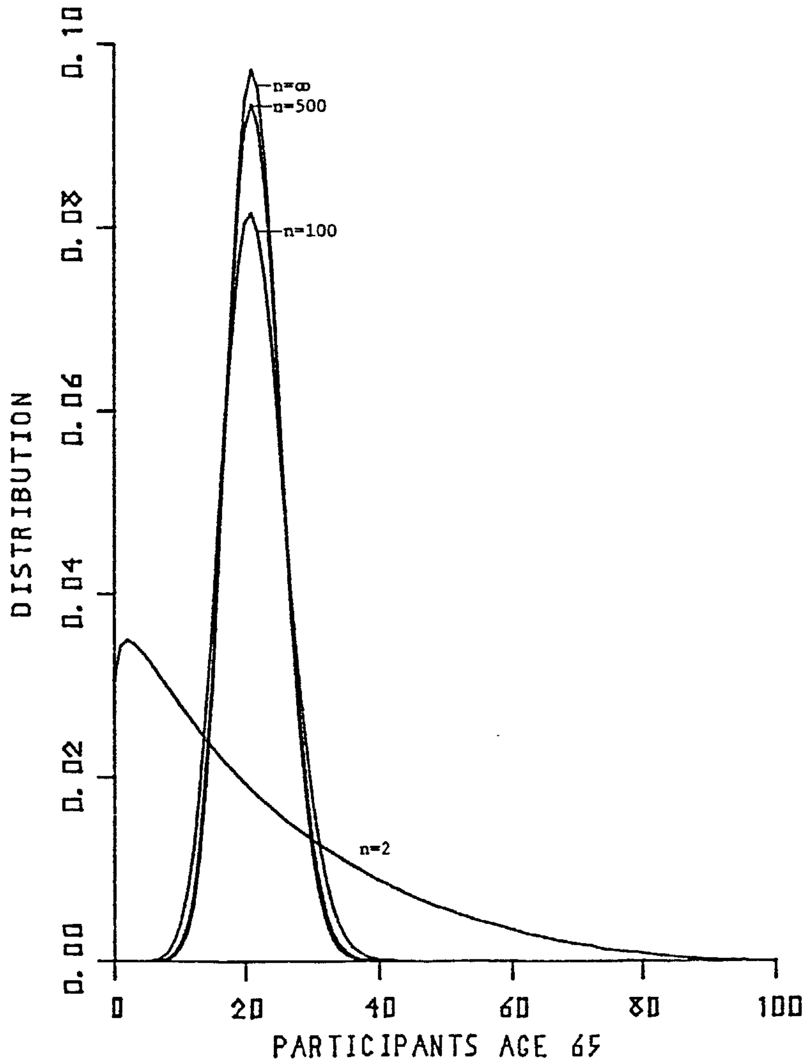

Data base: 1971 Group Annuity Table, rates of disability used in the 1970 Civil Service Retirement System Valuation, and McGinn's Turnover Table III.

of approximately 2,000 participants would result. In addition, the benefit function is based on 2 percent of final salary for each year of service, using Salary Scale 3 of the Actuary's Pension Handbook.

Figure 4 shows the results of this analysis. The probability that the projected cost will be adequate is greatest under a condition of high uncertainty regarding the probability of persisting. This is shown by the curve labelled "n=2". However, as a contingency charge is added, its impact is directly proportional to the confidence in the TPP. The greater the degeneracy of the prior distribution the smaller the contingency charge needed to obtain a given probability of adequate funds. This is shown by the curve labelled "$n=\infty$."

Note that, on the basis of the decrement data, even with perfect information, a contingency charge of 40 percent of projected cost would be required to attain a 99 percent probability of adequate funds. Note, also, under a condition of high uncertainty a contingency charge of 100 percent of projected cost would be needed to raise the probability of adequate funds to 99 percent.

Comment

It is hoped that this paper will help to stimulate both theoretical and empirical research into the stochastic nature of pension costs. As regards the latter, although this paper did attempt to obtain certain specific results, these results were intended primarily as examples of the implementation of a stochastic retirement cost projection model, and, as such, were far from exhaustive. Future researchers should find the study of the stochastic nature of pension costs a fruitful area for exploration, particularly if they have at their disposal an accommodating computer facility.

MR. OAKLEY E. VAN SLYKE: By way of introduction, let me define my subject to be what Joe Ferriera defined as Class Mean Pricing. I'm not going to worry about the social aspects of the fact that the rates whose computation is described here are something other than what is socially desirable from the insurance buyer's point of view. I want to talk today about how to design class plans that are socially acceptable: plans that produce the expected loss cost for each individual policy and yet are the most socially acceptable class plans.

In the last session, Bill Gillam referred to the interrelationships between the various rating variables. In some lines such as pension or life insurance there aren't very many variables and there aren't many interrelationships. But the casualty lines, especially private passenger auto, have a great many variables. Exploring the interrelationships is the heart of the approach to designing socially acceptable class plans described here.

The basic mathematical framework used is known by a number of names. Taken all at once the name would be something like empirical Bayesian credibility theory. It is based on work begun by Albert Whitney (Proceedings of the Casualty Actuarial Society, Vol. IV), a trend of thought that was finally put in its clearest form by Hans Bühlmann in 1967. A number of people, including the moderator of this panel, have added significantly to the literature in this field.

Basically, this mathematical technique says that the experience average in a class is not the best estimate of the loss cost of individuals in that

FIGURE 4

PROBABILITY OF ADEQUATE FUNDS

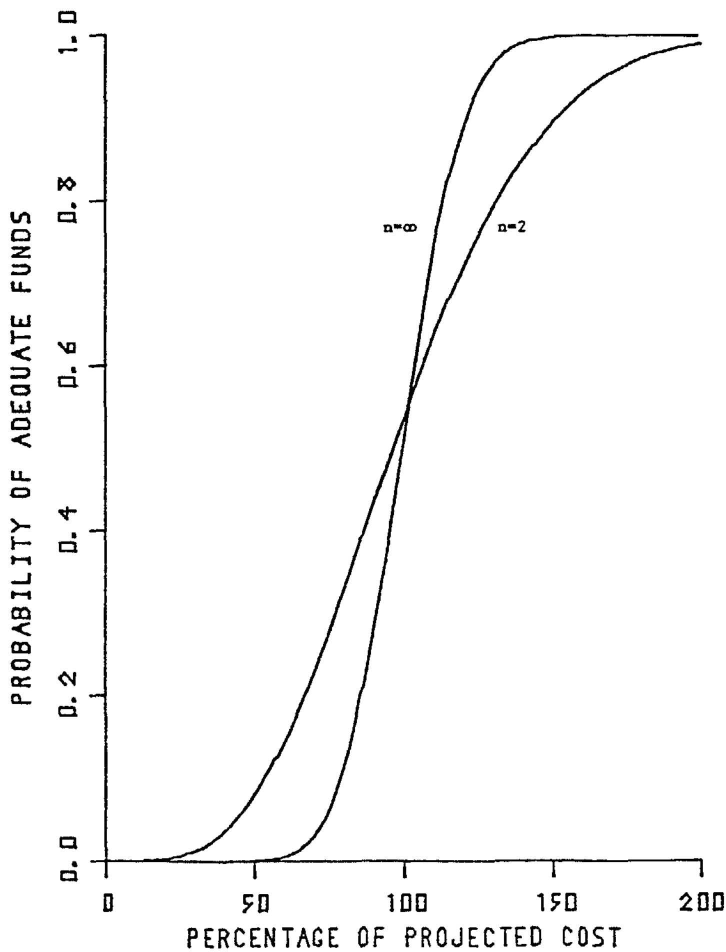

Data base: 1971 Group Annuity Table, rates of disability used in the 1970 Civil Service Retirement System Valuation, and McGinn's Turnover Table III, and Salary Scale 3 of the Actuary's Pension Handbook.

class, given that we have some information about what's going on in the other classes. This is not to say that the class average wouldn't be the best estimate if we knew only that class's experience. But once we start knowing something about other classes then we can bring these other bits of information into play in determining the class's expected loss cost.

Bayesian credibility theory does just this. In general, it produces least squares estimates of the class rates. In many cases, because the underlying probability distributions happen to be of certain types, it produces optimal estimates in a number of senses of that word. It works, in an economic sense, because there is a competitive advantage for insurance companies to pursue this approach.

The mathematical technique of Bayesian credibility theory is a least squares regression of the unknown underlying class rates (which we are trying to find) on the actual class experience. Imagine a graph of the unknown underlying class loss cost which we are trying to estimate on the independent axis and the observed class experience on the dependent axis. The method is to fit a straight line to whatever points we can observe in our insurance statistics. Because it's just a linear regression, it's the kind of thing we can do on even our small computers fairly quickly. One just summarizes the data according to whatever parameter he wants to test, applies the regression routine and out come the credibility coefficients that are the parameters of the least squares regression and the class rates.

Because it can be done quickly, this gives us a lot of flexibility in designing class plans. The procedure that I am suggesting today is sometimes called stagewise regression. Basically the approach is to take a particular variable, one that is the most socially acceptable for a line of insurance (such as merit rating), and allocate the burden of losses according to the loss experience of that rating variable using the Bayesian approach. The result will be a new set of rates for each individual in the class, rates that reflect the classification according to that one rating variable.

Then take the residuals (the differences between each individual's experience and his rate as estimated using this single variable) and explain those residuals using the second most acceptable variable. We can continue in this way through the available rating variables until we have run out of variables or until we are no longer able to explain anything significant in the residuals.

In practice, it appears (based on a limited amount of experience) that an important thing happens. One stops getting much credibility in the new variables pretty quickly. For example, the SRI report indicated that the first two or three rating variables that were included in a private passenger automobile classification plan would explain almost all of the variance that was going to be explained by even the ISO's 260 class plans. When moving through this analysis, the variables that are introduced later in the stagewise regression get very little credibility. This means that the class's experience on the variable is not going to have much affect on the class's rate. For example, if sex is the fourth or fifth variable in this routine in the computation of private passenger automobile rates, it has a lot less affect on the rate than if it is the first variable studied. The advantage of this, in terms of designing socially acceptable class plans, is farily obvious. If there isn't much rating difference, either that small difference will quiet the public furor about the use of the variable or, if it's hardly significant, one can leave the variable out of the rate plan altogether.

Now it may be that private passenger automobile insurance is not going to be amenable to leaving out the variable of sex altogether, but just to reduce the amount of rate variation that depends on sex would be a big step towards social acceptance.

At this point, let us quickly run through the criteria for social acceptability of rating variables. Of course, all these criteria interact. Any rating variable may be good in terms of one of these criteria, bad on another and some place in between on the others.

A general concern of the public right now is incentive value. Incentive value includes both the controllability of the rating variable by the insured and the casualty, or closeness of the casual relationship of the rating variable to the particular insured loss. There are two other areas slightly related to the topic of incentive value. One is that the incentive must be financially significant to have any hope of accomplishing the actual change in driving pattern; the other is that the incentive must be simple. That is, the class definition must be simple enough that people can understand what it is they must do in order to realize the incentive value.

A second general area of social acceptability is in the meaningful classification of risks. I'm surprised, in my discussions with people outside of the actuarial profession, that they agree with actuaries on this list. A lot of the items have been talked about in other sessions in terms of being handy or administratively easy, but in fact, attributes like simplicity are socially desirable, at least by the laymen that I've been talking to. So the list would run something like: simplicity; casuality or at least, correlation with loss; reliability, both in terms of preventing fraud and reducing administrative error; objectivity; and, then, four types of statistical information. The first of these is homogeneity of the risk characteristics within a class. That is to say, people want to be grouped with people like themselves. They don't think a class plan is fair if there are people not at all like themselves in their class. The second is the homogeneity of risk experience in the class. As actuaries we want classes that identify groups of people that have similar loss experience. Laymen want this too. The third is the heterogeneity of risk characteristics between classes. Both the public and actuaries believe that the classes should identify different types of people. The fourth is the heterogeneity of the risk experience between classes. If all the classes have the same average, the class plan hasn't accomplished much.

The homogeneity of the risk experience within a class, the heterogeneity of the risk experience between the class and the volume of the experience in the classification plan are the three factors that go into the computations discussed earlier. Hence, this particular approach for developing expected values tends to produce rates that are socially acceptable, in the sense of recognizing the homogeneity of risk experience within a class and the heterogeneity of experience between classes.

The last group of criteria for social acceptability is consistency with social goals. This would include such items as protection of privacy, stability of rates from one policy period to the next, and the affordability of rates. I do not want to pretend that we should put these above and beyond rating on an actuarially sound basis, but if one has a choice between two class plans, one of which makes it easier for poor people to afford insurance than does the other, and they're equally valid statistically, it seems clear that rate changes will be accepted more quickly using a class plan that aids the affordability of insurance.

In summary, then, there is a mathematical approach that is relatively straightforward, that provides best guess rates in the traditional actuarial sense and that is responsive to social values in the design of class plans.