RECORD OF SOCIETY OF ACTUARIES
1985 VOL. 11 NO. 4B

CURRENT ACTIVITIES IN ACTUARIAL RESEARCH

Moderator:	STUART KLUGMAN
Panelists:	MARK G. DOHERTY*
	HARRY H. PANJER
	ARNOLD F. SHAPIRO
Recorder:	JAMES W. DALLAS

This session, sponsored by the Society of Actuaries Committee on Research, will focus on three projects funded by the Actuarial Education and Research Fund (AERF):

- Risk models
- Operations research applications for actuaries
- Probability modeling of losses

In addition, the Society's Director of Research will discuss current and prospective research activities at the Society office.

MR. STUART KLUGMAN: This discussion on actuarial research is sponsored by the Committee on Research. Our Committee has three ongoing tasks. The first is sponsorship of the annual Actuarial Research Conference. The most recent one was held in October 1984 in Berkeley, California, on the topic of Credibility. The next research conference will be November 21 and 22, 1985, at the University of Texas at Austin, and the topic for that one will be Financial Operations of Insurance Companies and Employee Benefit Plans.

Our second activity is the publication of <u>ARCH</u>, the Actuarial Research Clearing House, a vehicle for the timely distribution of interesting ideas. One issue each year contains the papers given at the Research Conference.

Our third activity is recommending candidates for the David Halmstad prize for the outstanding paper on an actuarial subject. This year's prize will be awarded at the general session of this meeting. We now hope that the sponsorship of a session at the annual meeting will become our fourth regular activity.

* Mr. Doherty, not a member of the Society, is Director of Research of the Society of Actuaries.

Of the four topics that will be presented today, three of them have some common elements. First, all were sponsored, at least in part, by the Actuarial Education and Research Fund (AERF). All of them either have or will result in publication, also sponsored in part by AERF.

Second, each of these presentations represents an application of some part of the material on the Society of Actuaries' Associateship syllabus. We hope to demonstrate that these applications are sophisticated enough so that they can be used in the kind of situations you routinely face, yet are sufficiently simple so as to be implemented with the background provided in the Associateship materials.

The four presentations represent four different views of research. The first view is the collection and distribution of data useful in the on-going activities of our discipline. The data may then be used as the basis of further research. Mr. Mark Doherty will discuss these activities, along with the general role of actuarial research. A second kind of research is the development of new methods and techniques. While perhaps based on previous work, this involves the solution of problems for which solutions had not previously been available. Harry Panjer will demonstrate that problems involving risk models can be solved and that useful models can be constructed leading to answers to practical problems. A third kind of research is one that demonstrates that techniques developed to solve problems in other fields can also be used to solve problems in our field. I will demonstrate that parametric modeling and maximum likelihood estimation can be used to resolve questions that are relevant to insurance pricing. A fourth type of research is the compilation of work already completed. This is done in a manner to make this work readily available to practitioners. Arnold Shapiro's presentation will review applications of operations research techniques.

MR. MARK G. DOHERTY: I view research as a process, not just simply the conducting of research as we often think about it. Research is a flow, and it is a cycle. It has some exogenous input at the beginning. It also has feedback. Often more questions are raised by the research effort than we answer in the research itself. To identify the process, we need to examine and to screen a variety of input. We try simply to take a look at a number of topics and best allocate a minimal amount of resources. The process is effectively conducting the research. Out of this, we hope to get some input that would be used for education purposes, to put in the syllabus, perhaps even to be published in a form that allows the general public to understand what it is that actuarial research is trying to accomplish.

Part of this output must be in some form of technology transfer; meetings such as this are a good example of that. The material should then be picked up and utilized, and from this, we'll get information and feedback, helping us refine what we are doing.

Input comes in a variety of ways. It is problems and controversies--a lot of this can be exogenous, for example, forced on the Society by request from regulatory bodies or the National Association of Insurance Commissioners (NAIC). We look at input as problems. The sources are many and varied. We have a good deal of interal sources, that is, the various committees of the Society. In addition, we have individual members. But getting ideas is perhaps one of the biggest problems we have for research to be done within the Society. Outside the Society, particular groups ask us to do things that we almost have to respond to. A good example from the NAIC work is the 1985 Commissioners Disability Table. It will replace the 1964 Table. The 1980 CSO Table is another good example.

We are looking at the problem of identifying ideas as they come in. Often, the problems we are told exist aren't really the problems. They are the symptoms. A good example is that when I took the job of Director of Research, I was told the problem is that the experience studies are published late. That is not the problem. That is the symptom. The problem was getting the data from the companies for the intercompany studies. Upon getting the data, we can effectively get the material out. We are going to screen this variety of input and narrow it down to a few manageable ones that we do have the resources to tackle.

Research has to have a purpose, regardless of whether that purpose is for the betterment or for the increase of the knowledge base of the Society. In managing projects in the Society, we tend to view this purpose as more practical. That is what can be done for our membership and for our constituents.

People say you can't schedule research. I have done research on a scheduling basis for about 15 years, and we did it on a contract basis for many years. Yes, research can be scheduled, and if it is not scheduled, it goes awry--you spend a lot of money, time and effort, and perhaps you are throwing good money after bad money, so to speak. It is clear that within the Society structure, we have to assign these research projects to a committee. If we do not assign them effectively, they'll just sit there. We've seen that happen in a couple of cases and had to move them to different committees.

A paper was written about four and a half years ago on the scope of research within the Society of Actuaries. This was done by a committee that began the Research Policy Committee. Research was divided into three areas: experience, practice, and theory. If anyone wants a copy of this paper, I could provide it from the Society office.

Output is some sort of service to the members. One example is the experience reports appearing in the TSA Reports. ARCH is another good example. We would like to see these also develop into seminar notes, study notes, and technical notes.

To be effective, we must transfer this information. We are looking at it from a perspective, often very biased, of use by the Society in our educational process for students and for continuing education for the membership. However, we must also be concerned with the fact that little is known about actuarial science and what actuaries do, and it is important that we transfer this information and develop it in a way that the general public can come to grips with some of what is being done.

We hope that the information and research we do is used in a form such as the development of principles and practices. Some of it becomes reference material. We hope to see it appearing in the examination syllabus. The NAIC will use it. Examples are the 1980 CSO Table and what I presume will be the 1985 CDT. States look at our research with respect to their needs and see what they can use and give us their feedback. As many questions as we have answered, we often are asked that many more.

The Society has the Mortality and Morbidity Experience Committees. They can be broken into four areas for individual products: life, annuities, health, and loss of time studies. The group area has been limited recently to weekly indemnity and to long-term disability. However, there are some plans to begin looking at a group life insurance study, and that will eventually feed into the Internal Revenue Service Uniform Table 1. We hope to begin some work in group health. It is a difficult area because a lot of people feel this information is proprietary. However, we have a centralized service out of which processes and data are compiled. Thus, the companies are no longer sending information to other insurance companies. They are, in effect, sending it to the Medical Information Bureau (MIB), and basically all the companies deal with the MIB.

We have other types of studies, such as the aviation statistics. We have a self-administered retirement-rate research project underway. We have special purpose projects, such as NAIC requests. The 1985 CDT is likely to be approved at one of the next couple of meetings of the NAIC and put in force. January 1987 is the scheduled date for that. It is a unique study in that the report itself doesn't tell you a lot, but the report really is a diskette containing software that develops all the valuation tables for you.

Research work can go into basic and/or continuing education. Basic education means that of the students. Continuing education is the updating of the skills needed by the members. A good example of the things that have been going on in that area is the work of the C-3 Risk Task Force. It started out with very theoretical work, developed it quickly into applied work, and then presented it in the form of a series of seminars. The Task Force is about to write study notes for students. This is a good example of something that has gone from a basic, theoretical research problem; moved into something very useful; and then moved into continuing and, we hope, basic education--completing the cycle. The C-3 Task Force work was viewed as extremely complicated. It didn't seem practical. The nature of some of the work is very complex, and not subject to a great degree of simplification.

For the work we do in the research areas, whether it gets into theoretical or methodological approaches, the principles are where we have to define the limits of practice. As far as we are concerned, the Society is a learned body that will foster principles. The American Academy of Actuaries will foster, in effect, standards of practice. There is a subtle difference, and the work we do will be aimed at principles. The reason for this is that legislative and regulatory bodies tend to come in when there is a void. If the research isn't there to support the work in the development of principles, then you may have regulatory bodies coming in and telling you what needs to be done.

From an academic perspective, we are looking at the advancement of actuarial science. I have heard actuaries say all these other groups are infringing on our territory. If you think about it, the actuaries are also infringing on the others' territories. It is a two-way street. No one thinks of the actuarial community as a small group of people that does one specific thing. The actuarial community consists of academic people, consulting actuaries, and people working in insurance companies. There are a number of chief financial officers; people who are in management are actuaries. They are not practicing actuarial science per se. The actuarial community is moving out and doing other things.

We all need tools to work with, and through the advancement of actuarial science, one hopes to get this methodological approach. The other aspect of this kind of research which is important is education materials. There is excellent work being done by actuaries in academia.

At least 90 percent of the Society's research work is in experience studies. The experience studies have been the mainstay of the Society's efforts for years. It's what the Society members have come to know as research. We are attempting to improve these studies, so that outside bodies, other professional organizations, and the more educated general public begin to appreciate what is going on here. For example, we are looking at new studies. We are trying to add new aspects to studies, such as looking at new products--universal life and variable life--or in terms of life insurance, looking at smoker/nonsmoker and unisex, or trying to update some of the studies that are extremely old. We are going to look at different analyses, not only for an actuarial audience, but for a broader audience. We want to get more graphics and introduce other ways of looking at the same picture, not just tables. We are going to phase out or terminate studies we feel are no longer any use. We recognize that a lot of companies are large enough to do their own studies, but many small companies need information such as the experience studies. We are hoping to expand our analyses. We are looking again to broaden our audience to the insurance and to the more general public. We hope to increase the content and the analyses in the studies so that we can make some statements about what we are doing. And, again, we are going to try to move this faster. We are going to make available to the membership copies of reports as they are done. That is, there will be a notice in The Actuary or in a general mailing telling you what reports are available and how much they cost. The charge would be simply the duplication and mailing cost. These will be done on a much more timely basis in the future.

We are going to be working in the research area to define the limits of principles. We want to establish the profession clearly before we have to react to legislative and regulatory aspects. Again, part of this is educational needs. Students tell us they need things. We hope to anticipate legislation and regulations and technological change. We want to incorporate more things, such as work with personal computers.

Every science, every profession, every industry, in effect, needs to expand and advance the theoretical knowledge base. If you look at how the steel business has declined, research effects were minimal, at least in the United States. If you move over to the electronics industry, you can see how much they pour into it. One business is stagnating, the other is growing. We need to work towards developing tools and techniques. I have a project ongoing with respect to futures, options and financial risks, and the management of financial risks. We hope to be coming out with a monograph or book, which can be used for continuing education or for basic education of students.

Another thing we need to do is to work at the development of educators, and we feel very committeed to this at the Society, as well as working with the AERF on this area. We need to push these research projects into the syllabus as reference material, if nothing else.

Practical research is an area that if one could define it, I would be happy to try to manage the projects that are out there. The problem is simple. We hear in a survey from the membership that practical research is important. When we said to the membership, "What do you want in terms of practical research?", we received little response. We know we have to be on top of current topics. We need to have a method to get information out quickly, so we are working toward that.

Why are we worried about all of this? The Society has a definite role in research. One of the two stated purposes of the Society of Actuaries is research. The other is education. We often act as the facilitator of research. That is, we don't do the research per se in the Society office, but rather, we cause it to be done somehow. Because people can work in a volunteer effort and committees under the auspices of the Society, we are able to attract a lot of talent and put that toward specific work. The research on the valuation actuary responsibilities is a good example of something that would not be done if people were trying to get together as just representatives of their companies. They are working at a professional level. This makes it a little easier because they are not worried about conflicts of interest.

Research work will feed into that education process. It will keep the new actuaries current. It will help develop and hone the skills of the practicing actuaries right now. We hope the research area will be a standard setter, and eventually, the Society office will grow and have a number of actuaries on the staff so that research projects such as experience studies will be done in-house. The volunteer effort is fading away. We see it on the education side, and we definitely see it on the research side. We hope that this doesn't happen terribly fast, but we can see the day when the volunteer effort will almost totally disappear, at least in the research area.

Research, education, and the practice and principles of working in an actuarial field are all interrelated. One needs the other. Without research, we are going to have a problem. Without education, we obviously have a problem. Research plays a role, and we need to foster it.

We need input from the membership with respect to research ideas. We need people to participate on a variety of good mortality and morbidity committees. There are other committees that might be of interest to you in terms of theory of risk and other areas. Requirements may be forced on us from the outside. It is hoped that we can anticipate these and not have to worry about reacting. We hope to support the educational process through the research, that is, feeding information into the syllabus, into the examination process itself. We are going to have to understand better what the membership means by practical research. We want to broaden our horizons. We have a task force that is looking into the expansion of morbidity and mortality data. We have to find other professional groups we can interact with so that we can use their information and they can use the talents and skills that actuaries can give. We are going to provide financial support. We are underwriting the projects on futures, options, and the management of financial risks. It is the first time the Society has done that in a research mode. It should develop into a sufficient monograph to bring other people up to speed. It will be good for the education of students. We are going to have to augment the volunteer system. The time of people is not being given as freely by their companies. As a result, we are going to have to bring more research in-house. Finally, we are working to improve the experience studies. We hope to use those as the vehicle to express more of the research of the Society to a broader public. We hope to have it set up so that these studies which are sound for actuarial use also can attract the outside public. It may seem like it may be impossible, but I am sure we can get around some of the tables and move into graphics to make these studies more attractive.

MR. CHARLES A. ORMSBY: Do you have on your list of projects to be done in the near future studies of the cost/benefit analysis of underwriting tools?

MR. DOHERTY: No, not at this time.

MR. HARRY H. PANJER: I am a university professor. Mr. Klugman and Mr. Shapiro are university professors. But we are also professional actuaries. As university professors, we have the responsibility and the opportunity of working on problems that excite us. As professional actuaries, we have a responsibility to work on problems that are useful to actuaries. The problems that we choose to work on are sometimes viewed as being too theoretical. Some of you will think that the work discussed in this presentation is theoretical, but I believe it to be totally practical in all lines of insurance and it can be used to address a wide range of problems. The C-2 risk is one of the obvious problems that can be addressed using the models that I will discuss.

I am currently writing a book entitled Insurance Risk Models with Gordon Willmot, my colleague at the University of Waterloo. I plan to reflect on some of the content of that book. I will not use a lot of mathematical symbols, but rather make more general references to principles and applications.

The topic of risk models deals with the rational development of parametric models to describe the aggregate claims or losses of an insurance company for a particular block of business. It is not restricted to life and health insurance and reinsurance, but it is also applicable in property and casualty insurance at least as much as in the other areas. In developing models, we try to develop rational explanations for variations in observed results, and then, having developed those models, we develop practical tools for applying these models.

Risk can be described as being associated with uncertainty of results. For an insurance company, risk is associated with the uncertainty of results from a variety of areas. I will be talking about the claims area only, not about the investment or the expense areas.

Adverse claims can result from two sources. The company can have too many claims, or it can have claims that are too large. I will be discussing the problem of too many claims, explaining the frequency of claims on a rational basis. Mr. Klugman will discuss the problem of the size of claims. When the two are put together, you have a model for total claims (or aggregate claims).

Risk is associated with the right-hand tail of the distribution of aggregate claims. We are interested in the probability that total claims will exceed some particular amount.

There are several sources of adverse deviation. A company can have a bad year for reasons of pure randomness; reasons over which it has no control. Some years will be good, and some will be bad. You hope in the long run that they will average out. On the other hand, adverse deviation could be a result of mispricing the policy, that is, choosing the wrong assumptions. Both of these sources of adverse deviation apply to the frequency and size of claims--the size, in particular, in health insurance; the frequency, in all types of insurance.

Adverse deviation in frequency might occur because the actuary used the wrong morbidity rates in pricing health insurance.

Frequency models should reflect both these sources of variation. We would like our models to be reasonably simple mathematically so that the results are in a tractable form to be used for the calculation of the distribution of total claims. When we discuss parametric models we want models that have a small number of parameters. Another important aspect of choosing a model is that it should have a rational interpretation. We should be able to interpret the results, or explain the behavior, of claims given the model. A third practical attribute of a particular model is we would like to develop models which we can easily work with numerically.

The formula for the probability that n claims will occur, based on the Poisson distribution, is

POISSON DISTRIBUTION

$$\Pr\{n \text{ claims}\} = \frac{e^{-\lambda}\lambda^{n}}{n!}, \quad n=0,1,2, \ldots\ldots$$

This is a reasonable model to use if you assume that claims occur randomly through time in some way. This has been used in all kinds of areas, including life insurance, and it has a convenient formula associated with it for the development of total claims. This is a simple formula which related the distribution of total claims with the individual size of claim distribution.

$$\Pr\{\text{total claims} = x\} = f_S(x)$$

$$= \frac{\lambda}{x} \sum_{y=1}^{x} y\, f_X(y)\, f_S(x-y), \qquad x = 1,2,3,\ldots..$$

This distribution has been used by actuaries since about 1903 when the Scandinavian actuaries first published some of these models, and generally it has worked very well. However, there are many instances when it does not work well. For example, the data in table 1 are from a study of the number of injuries among workers over a period of time. Out of the 647 workers over the particular time period, 447 had no injuries, 132 were injured once, 42 were injured twice, 21 injured three times, 3 injured four times, and 2 were injured five times.

TABLE 1

Number of Injuries	Actual	Poisson
0	447	406.31
1	132	189.03
2	42	43.97
3	21	6.82
4	3	.79
5	2	.07
	647	646.99
Mean	.465	.465
Variance	.691	.465

The first question is whether or not the Poisson model is an appropriate model. We have some observations, and we have a model. We should test the fit of the model, to see if it makes sense for the data. When we fit the Poisson distribution using the method of maximum likelihood we end up with the expected frequencies given in table 1. Rather than 447 with no injuries, we would have predicted 406 using the Poisson distribution. Rather than 132 persons with one injury, we would have predicted 189, and so on. Although this may look reasonably good, when we conduct a statistical test of fit, we find that the model fails. It does not explain enough of the variation in these data. One way of looking at the variability is to look at the first few moments of the number of claims in the data and the number of expected claims in the model. The expected number of claims is .465 for both the actual and the model. Because the mean and variance in the Poisson distribution are equal, the variance for the Poisson model is .465 as well. However, if we calculate it for the actual data in table 1, we find that the variance is .691. It is considerably larger than the mean, which suggests that this model is not appropriate.

The next question one would ask when confronted with these data is, "What is a rational explanation for the deviation from the Poisson assumption?" An explanation that was put forward many years ago is that all risks are not alike. There is heterogeneity in the portfolio. All workers are not working in identical jobs or in the same work area, and so some are more prone to accidents than others. This suggests that the individuals in the population have different expected numbers of claims. Yet, we are putting all persons in one package and treating them as identical. That, of course, happens in many group health or group life situations. So, rather than modeling the expected frequency as a fixed number, we could model it as a random variable. That is, in our population, the measure of risk for the individuals varies among them. Some people are more risk-prone than others. We will describe that risk-proneness by a distribution, which might look something like the one in figure 1.

FIGURE 1

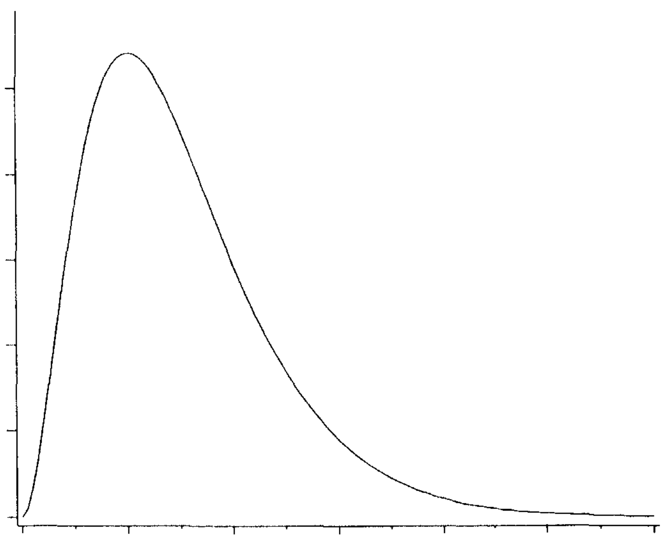

GAMMA DISTRIBUTION

Now Λ is our risk parameter. It has the same mean as before, in our example .465, but we will put a distribution about that mean because we know that, on average, some people will have few claims while others will have more claims.

The question arises, "What is an appropriate model for a distribution describing the heterogeneity in the population?" The only one that has received much attention to date is the gamma distribution. The gamma distribution is a well-known distribution which has the general shape of the curve in figure 1. One of its characteristics is that it is light-tailed. That means that the right-hand end of the distribution tails off relatively rapidly. In other words, we don't have a lot of really bad risks. When we take into consideration the distribution of the number of claims, we end up with a negative binomial distribution as the distribution for claims frequency.

All actuarial students should be familiar with a negative binomial distribution. It is similar to the Poisson but is more dispersed. Table 2 uses the same data as seen earlier, but with the negative binomial fitted.

TABLE 2

Number of Injuries	Actual	Poisson	Negative Binomial
0	447	406.31	445.89
1	132	189.03	134.90
2	42	43.97	43.99
3	21	6.82	14.69
4	3	.79	4.96
5	2	.07	1.69
	647	646.99	646.12
Mean	.465	.465	.465
Variance	.691	.465	.715

With the negative binomial distribution, instead of having one parameter, we have two parameters. So, we would expect to do a better job of fitting. As it turns out, the expected frequencies from the negative binomial distribution are very close to the actual, and we are doing this with only two parameters. Now, the question again arise, "Is this an adequate representation of the observed phenomenon?" If we conduct a statistical test of fit, it turns out that it is, in this case. When we compute the means and variances for these distributions, we see that for the negative binomial distribution that is fitted, using a method of maximum likelihood, the variance is .715, which is relatively close to .691, so this distribution is much more like the observed data. To conclude this example, the Poisson plus the gamma distribution as a heterogeneity distribution gives us the negative binomial distribution

Using the negative binomial rather than the Poisson distribution adds no complexity to whatever problem you are trying to solve. It adds no numerical difficulty because there is a formula that looks like this, which shows you how to calculate the distribution of total claims given the claim size distribution.

This formula is easy to apply and gives answers quickly on a personal computer, even for large portfolios.

COMPOUND NEGATIVE BINOMIAL DISTRIBUTION

Table 3 shows the observed number of claims from automobile accidents in California, in a particular zip code area, over a particular time period, and it is based on the California Driver Record Study conducted by the State of California. It turns out that the negative binomial fits very well indeed. It suggest that there is a heterogeneity among drivers. Some drivers are better drivers than others, so that the expected number of claims for some is larger than for others. When an insurer is using this distribution to model claims frequencies, the distribution will indicate more risk than under the Poisson assumption in the sense that the probability of total claims exceeding a certain amount will be larger for this distribution than for the Poisson.

TABLE 3

Number of Claims	Actual	Negative Binomial
0	129,524	129,527
1	16,267	16,261
2	1,966	1,955
3	211	232
4	31	27
5	5	3
6	1	.4
7	1	.04
	148,006	148,006

This has all been review. I am now going to introduce some alternatives. The negative binomial may be viewed as being too light-tailed in the sense that the gamma distribution did not include enough really bad risks. If you are offering insurance in situations where there is a significant risk of antiselection, you may wish to make sure that you use a model that reflects the possibility of antiselection.

There are a number of models that one could propose instead of a gamma distribution. One might be a simple two-point distribution in which there are only two homogeneous types of risks--good and bad. This model has been used in some insurance applications but is rather simple-minded because one would expect that there is a continuum of accident-proneness in the population. Some people are better than others. They are not simply in two classes: good and bad. One might also propose a discrete distribution like the Poisson, the negative binomial, or the geometric to describe the risk levels in the population. None of these models adds any more complexity. The computational techniques can easily be generalized to include these models. But these models have the drawback that they are discrete. This may be criticized since intuitively risk-proneness is not a multiple of some number. It makes more sense that it is continuous over the population in some way.

Another distribution which has not been used in an actuarial context before for this purpose is the inverse Gaussian distribution. Its distribution has the following form.

INVERSE GAUSSIAN DISTRIBUTION

$$f(x) = \mu(2\pi\beta x^3)^{-\frac{1}{2}}\exp\{-\frac{(x-\mu)^2}{2\beta x}\}, \quad x>0$$

You never need to use this form of the distribution. You never need to actually quantify this function. You could ignore this function entirely. The inverse Gaussian distribution has the advantage that it is more spread out in the sense that it is more skewed than the gamma distribution. So it includes more bad risks.

If you are worried about antiselection, you might wish to include a heavier tail to describe the risk parameter in the population. Table 4 gives the results of a study in which both the Poisson and the negative binomial fail to describe the observed data.

TABLE 4

Number of Claims	Actual	Poisson-inverse Gaussian
0	103,704	103.710.03
1	14,075	14,054.65
2	1,766	1,784.91
3	255	254.49
4	45	40.42
5	6	6.94
6	2	1.26

The distribution, using the inverse Gaussian, is referred to as a Poisson-inverse Gaussian distribution. When it is fitted by the method of maximum likelihood, the fit is adequate. We achieve this fit at no more cost, in terms of the number of parameters, than was the case for the negative binomial distribution. For the negative binomial distribution there were two parameters; for this distribution there are two parameters. So, it is really another mathematical function with a somewhat different shape that can be used to describe the behavior which is exhibited in the data. Computationally, this distribution also can be used very effectively in calculating the distribution of total claims.

Another distribution is the Stable distribution. Again you never have to use this form of the distribution. The ultimate functions that we are interested in are very simple functions. This distribution is a three-parameter model, and it is a good alternative when a two-parameter model is inadequate to explain the variation exhibited in the data. The distribution that results is called the generalized Poisson-Pascal. The density looks like this:

STABLE DISTRIBUTION

$$u(x) = \frac{e^{\lambda}}{\mu\lambda^{1/\alpha}} e^{-\frac{x}{\mu}} f_{\alpha}\left(\frac{x}{\mu\lambda^{1/\alpha}}\right), \quad x>0$$

where

$$f_{\alpha}(x) = \frac{1}{\pi}\sum_{k=1}^{\infty} \frac{\Gamma(k\alpha+1)}{k!}(-1)^{k-1} x^{-(k\alpha+1)} \sin(\alpha k\pi), \quad x>0$$

It is a distribution that has not appeared in literature before, but we have found it useful for some data sets. (When you develop models you have to find data sets to justify them.) Table 5 contains a large set of data on claims from a particular line of insurance in Australia. In this large volume of data (the study was done a few years ago), it turned out that the Poisson did not fit, the negative binomial did not fit, and the Poisson inverse Gaussian did not fit, so we fit the generalized Poisson-Pascal distribution. According to statistical criteria, it fit very well. The statistical criteria require consideration of the number of parameters. So, according to the statistical test, this fits adequately at a high level of significance.

TABLE 5

Number of Claims	Actual	Generalized Poisson-Pascal
0	565,664	565,661.24
1	68,714	68,721.23
2	5,177	5,171.68
3	365	362.86
4	24	29.61
5	6	2.98

So these are alternative models to describe how one can incorporate heterogeneity into a risk portfolio. I imagine one could use this kind of model, at least as an alternative to the basic model, to look at the effect on the required surplus for cafeteria-type group insurance, when one allows for a significant variation in the expected numbers of claims among the individuals in the portfolio.

There is another modeling problem which has been addressed by actuaries relating to a catastrophic risk; that is, multiple claims arising from a single claim-causing event. Claim-causing events are referred to as "accidents" and the number of claims arising from those as simply "the number of claims." There is dependence between risks in many types of insurance, for example, in group life and health. In group life, you might expect that when one employee is killed in an accident, others may also be killed in the same accident. If an elevator crashes with several employees in it, you are going to have several claims as a result of a single accident.

To model multiple claims, we model the number of accidents separately from the number of claims per accident. We then combine these and develop the number of claims. Using the Poisson distribution to represent the number of accidents through time is a reasonable assumption. We are simply assuming that accidents occur randomly through time. They occur now and then, but they are totally random. For some accidents, we will have more than one claim, and for some accidents we will have only one claim. The logarithmic distribution may represent the number of claims per accident. If we put these together, we also get the negative binomial distribution!

Table 6 goes back to the data exhibited originally in table 1, that is, the number of injuries among a group of workers over a particular time period. When discussing table 1, I provided one explanation for why those results might look like the negative binomial rather than the Poisson. I am now going to give you a second explanation.

TABLE 6

Number of Claims	Actual	Negative Binomial
0	447	445.89
1	132	134.90
2	42	43.99
3	21	14.69
4	3	4.96
5	2	1.69

POISSON ACCIDENTS λ = .372274

LOGARITHMIC CLAIMS PER ACCIDENT

n	Probability	n	Probability
1	.813	5	.0024
2	.142	6	.00071
3	.033	7	.00021
4	.0087	8	.000065

Suppose the accident had been purely in accordance with the Poisson distribution, and the numbers of claims per accident has a logarithmic distribution with the probabilities given in table 6. That is, 81 percent of accidents gave rise to a single claim; 14 percent to two claims; 3 percent to three claims; and so on. If we put these two distributions together, we get the same negative binomial distribution and the same negative binomial predicted frequencies. So we have quite a different explanation for how the claims arose. There are other explanations. But, if you are an actuary making a decision with respect to some business where you think the negative binomial may be a reasonable distribution, you should investigate whether or not these claims were from different accidents or whether some accidents could have given rise to multiple claims. Each model will lead to its own decision.

There are many other models that one could use. I have listed a few here:

Number of Accidents	Number of Claims Per Accident	Number of Claims
Poisson	Two-point on 1,2	Hermite
Poisson	Poisson	Neyman Type A
Poisson	Negative Binomial	Poisson-Pascal
Poisson	Extended Truncated Negative Binomial	Generalized Poisson-Pascal

All of these models make use of the same type of recursive formulas for developing the distribution of total claims. So there is a wide range of models which you can use to explain the behavior of how claims can arise and which will still allow you to accurately model and complete the total claims distribution so that you do not have to rely on some simple inappropriate model.

MR. THOMAS P. EDWALDS: Was your test of fit the chi-square test?

MR. PANJER: Yes it was.

MR. THOMAS G. WALSH: You said that the Poisson model is used in life insurance claims as well?

MR. PANJER: Yes.

MR. WALSH: The binomial model strikes me as a better one.

MR. PANJER: Well, it depends. Consider yourself a group insurer. The binomial is really the sum of Bernoulli random variables, which assumes that, when an employee dies, there is no replacement made. However, if you have a large group life portfolio, it is reasonable to assume that when a person dies, that employee is replaced by another employee. The Poisson assumption implies the person is replaced by a person with the same mortality characteristics (although this is not critical). The binomial is another distribution which could be used and which also presents no more computation difficulty. It assumes that all employees have the same mortality rate and that no replacement is made.

MR. STUART KLUGMAN: I hope to convince you of two things: (1) that parametric continuous distributions should be your choice for modeling losses and (2) that those parameters should be estimated by the method of maximum likelihood. Perhaps a third item is that variability, in particular the estimate of standard error, is just as important as the expected value. I'll use an example based on dental coverage, but the techniques clearly apply to any health insurance, or more generally, to any insurance in which the benefit amount is random. This would certainly be the case in the risk models just discussed by Mr. Panjer, and these methods would definitely be useful in specifying the distribution for the amount of one claim. This also would apply in the group life insurance context where the benefit paid on the next death is random, even though, of course, each insured's benefit is for a fixed amount. You don't know who that next death would be, so you have a random distribution for the amount of the next claim.

While nothing is observed on a continuous basis (we know that any measuring device will always have a limit on its accuracy), we are accustomed to using continuous models. For example, we use the normal distribution to model persons' heights, and we often use the exponential distribution to model the time to the next claim. There are advantages in using a continuous model. First, irregularities in the data are smoothed. Of course, we should verify that those observed irregularities do not have a good reason for being there. Second, data are often collected in frequency table form. A continuous model will provide for interpolation in the frequency table. Third, data are often censored or truncated. For example, if there is an upper limit on the coverage, the actual amount of losses above the limit may not be recorded, only the fact that such a loss occurred. This is called censoring. At the other end, if there is a deductible, losses below that amount will not be recorded at all, or even the fact that the event producing that loss occurred. This is called truncation. If the current coverage is to be changed by either lowering the deductible or raising the limit, a continuous model can provide the probabilities over those regions on which data had not been observed. Additionally, in the presence of a deductible or limit, the effect of inflation is clearly not the same as it would be if there were not deductible or limit, and these effects can be measured with a continuous model. My goal will be to model the amount of loss prior to the imposition of any coverage modifications. Then we can easily study the effects of any proposed modifications.

There are many continuous probability distributions to choose from. One way to narrow the field is to restrict attention to parametric families. By this I mean situations where the probability density function depends on a small, say four or fewer, number of parameters that vary over some subset of the real line, so within one parametric family there will still be a large number of possibilities. Having a few such families to work with will probably be sufficient. One reason for using parametric models is that it is easier to compare a few parameters than to compare two probability density functions. We may want to see if loss patterns have changed over time, or if one group, for example, adults, has a different loss experience than another group, for example, children. A few parameters may make that comparison possible.

A second reason is that for using parametric models standard errors can be obtained. We will see in the discussion of maximum likelihood estimation that confidence intervals can be constructed for the parametric estimates and the quantities of interest. Consider the following losses on a basic dental coverage.

DENTAL LOSSES, DEDUCTIBLE OF $50

Amount of Loss	Number
$ 50- 99	281
100- 199	314
200- 299	147
300- 499	162
500- 699	81
700- 999	52
1,000-1,499	30
1,500-1,999	7
2,000-4,999	4
5,000-7,499	1
	1,079

These data were from actual losses on adults, but the number of groups has been reduced, and a deductible of $50 artificially imposed. So the total number of observed losses is 1,079, and the average loss was $294. If we further assume that the group to which we will be selling this coverage is expected to generate 200 claims per year, the annual net premium would be estimated as 200 claims times $244, after the removal of the deductible, or $48,800.

What would the net premium become if the deductible was removed? What would the net premium become if a limit of $1,000 per claim was imposed? Suppose inflation increases losses uniformly by 10 percent. What is the effect on the net premium? None of these questions can be answered directly or accurately from the data as given. We need either information on values outside the observed range, or the precise distribution within an interval. For example, the second question requires knowledge of the average amount paid on losses over $1,000.

A first step in the analysis is an examination of the histogram that describes these observations. It is important to note that when constructing a histogram based on unequal interval widths, the heights are not proportional to frequencies, but rather it is the areas of each rectangle that are proportional to the observed frequency. The histogram in figure 2 is typical of the type of distribution encountered in insurance losses. The dominant feature is the positive skewness, also known as the presence of a heavy tail. One thing that cannot be determined from these dental losses, however, is the location of the mode. It seems reasonable to expect that any loss experience is from a unimodel distribution, but all we can tell from this histogram is that the mode is less than 50. The distributions I am about to list all are unimodel. Some of them always have a mode at zero. Some have it at a positive location.

Our objective now is to create an inventory of parametric families of heavy-tailed distributions, and I begin with two well-known distributions. Their basic forms are as follows.

FIGURE 2

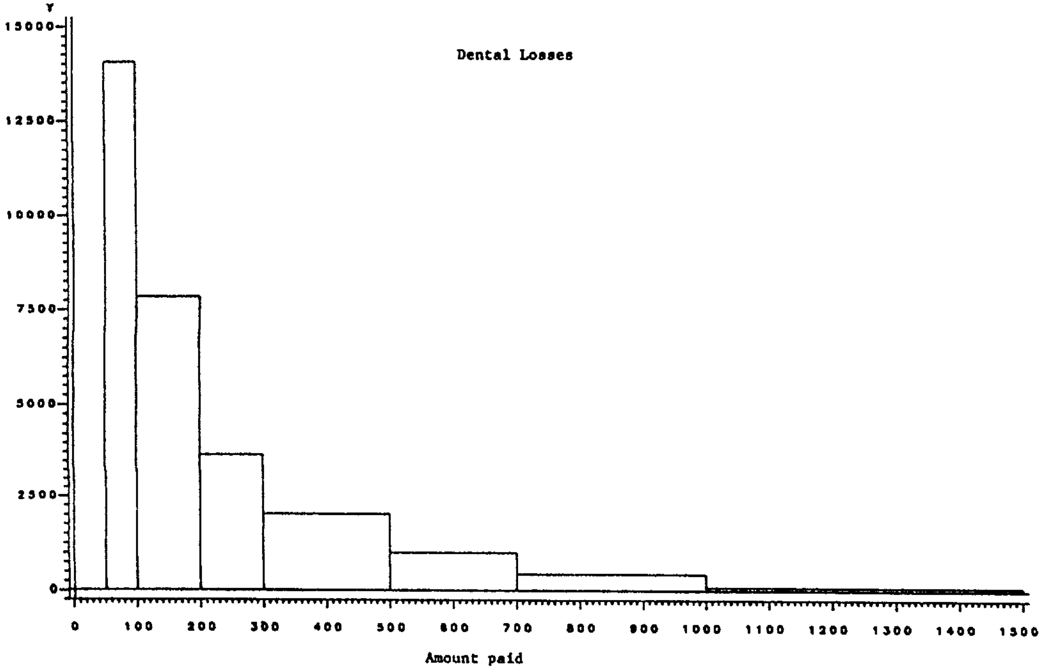
Dental Losses
Y
15000
12500
10000
7500
5000
2500
0
0
100
200
300
400
500
600
700
800
900
1000
1100
1200
1300
1400
1500
Amount paid

$$\text{gamma} \quad \text{--} \quad f(x) = \frac{e^{-x/\lambda}(x/\lambda)^{\alpha-1}}{\lambda\Gamma(\alpha)}$$

$$\text{Pareto} \quad \text{--} \quad f(x) = \frac{\alpha}{\lambda(1 + x/\lambda)^{\alpha+1}}$$

The gamma distribution is a familiar one from our Part 5 syllabus. One that is not as prevalent in the actuarial literature is the Pareto distribution. There is not general agreement on exactly what a Pareto distribution is, and this is but one version of a Pareto distribution.

The gamma distribution places less probability in the tail than does the Pareto. Evidence of that is that all moments exist for the gamma distribution, but for the Pareto distribution, the expected value of X to the k power exists only for those values of k that are less than alpha.

There are two easy methods of getting additional distributions from the distributions you already possess. The first is called mixing. Assume that one of the parameters, in this case θ_1, is random, assign a distribution to it, and then construct the marginal distribution after the conditioning has taken place. The formula for the marginal probability density formula (PDF) is

$$f(x|\theta_2,\ldots,\theta_m,\beta_1,\ldots,\beta_n) = \int f(x|\theta_1,\ldots,\theta_m)f(\theta_1|\beta_1,\ldots,\beta_n)d\theta_1.$$

An example is

$$X|\alpha,\lambda \sim \text{gamma}(\alpha, 1/\lambda)$$

$$\lambda|k,\beta \sim \text{gamma}(k, 1/\beta)$$

then

$$f(x) = \frac{\Gamma(\alpha+k)(x/\beta)^{\alpha-1}}{\Gamma(\alpha)\Gamma(k)\beta(1 + x/\beta)^{\alpha+k}}$$

the generalized Pareto distribution.

The example has X being a gamma distribution. Then let the parameter lambda also have a gamma distribution, depending on parameters k and beta. After doing the mixing, the resulting distribution is a three-parameter distribution, which has been called the generalized Pareto distribution. A special case of this distribution is the well-known F distribution used in analysis of variance.

One justification for using mixing is parameter uncertainty. Suppose you believe that a particular family adequately describes the loss process. However, the parameters you have been using have been based on old data. You further believe that next year's losses will be similar in pattern but inflated by an unknown amount. If the distribution of this amount can be modeled, it can be used as a mixing distribution on the scale parameter, lambda in the case illustrated here. Such mixing would usually increase the variance, thus reflecting that the future is less certain than the past.

A second justification for mixing is heterogeneity, which has already been discussed.

Another way to obtain a new distribution is through a transformation. One is the transformation Y equals e to the X power, and the usual use for that is transforming the normal distribution to the lognormal distribution. The one that I like to use is as follows:

$$Y = \lambda(X/\lambda)^{1/\tau}$$

To obtain the PDF of Y replace x/λ with $(x/\lambda)^{\tau}$

and multiply by $\tau(x/\lambda)^{\tau-1}$.

The main reason that lambda is there is to preserve its role as a scale parameter. If you apply the transformation to the Pareto distribution given earlier we obtain a three-parameter distribution, which has been called the Burr distribution and has proved useful for fitting a great variety of losses.

$$\text{Burr} \quad -- \quad f(x) = \frac{\alpha\tau(x/\lambda)^{\tau-1}}{[1 + (x/\lambda)^{\tau}]^{\alpha+1}\lambda}$$

Another feature of the Burr distribution, unlike the gamma distribution, is that the distribution function as well as the density function can be written in a closed form. One special case of this transformation is worth mentioning. When tau equals minus one, we have what are called the inverse distributions. For example, the inverse gamma distribution is the distribution of the random variable one over a gamma distributed random variable.

$$\text{inverse gamma} \quad -- \quad f(x) \quad \frac{e^{-\lambda/x}(\lambda/x)^{\alpha+1}}{\lambda\Gamma(\alpha)}$$

Finally, an ultimate extension of all of these processes is the transformation of the generalized Pareto distribution.

$$f(x) = \frac{\tau\Gamma(\alpha+k)(x/\beta)^{\tau\alpha-1}}{\beta\Gamma(\alpha)\Gamma(k)[1 + (x/\beta)^{\tau}]^{\alpha+k}}$$

This distribution has been given a number of names, "generalized beta of the second kind" and "transformed beta" among them.

All of the distributions mentioned prior to this one are either special or limiting cases of this four-parameter distribution.

So, with a large number of families at our disposal, we proceed to the next two questions. First, for a given family, how we can find the one member of that family that best describes the data? And second, once we have the best member selected from a number of families, how do we select the one to be the model? Both questions can be answered through the method of maximum likelihood. In brief, the estimate is that value of the parameter or parameters that maximizes the probability of observing what was actually observed. Among various families, we can select the one with the largest such probability overall.

I will be taking a look today only at data collected in groups, so let $d = c_0 < \cdots < c_k = u$, with the c_1 in increasing order, be the class boundaries, and let f_i be the frequency for class i. If we then let $F(x)$ be the distribution function and recall that it will depend on the parameter values, we can describe the probability of observing a loss in the i-th class.

$$P_i = \frac{F(c_i) - F(c_{i-1})}{F(u) - F(d)}$$

This is in the presence of an upper limit u and a deductible d. The likelihood function is the product of the probabilities of observing f_i in each class. That is,

$$\Pi(P_i)^{f_i}$$

Then rather than maximizing the likelihood function, it is usually easier to minimize the negative of its logarithm. The main advantage is that we now have a sum instead of a product. Unfortunately, in most cases, it is difficult to directly minimize this function. Usually the likelihood function is a fairly complex function of the parameters, so you need a numerical approach. A number are available, many in standard computer packages. The one that I have had success with is the Procedure NLIN in the Statistical Analysis System (SAS). It requires that you turn the minimization into the least squares problem.

Minimize

$$\Sigma(y_i - f(x_i,\theta))^2$$

where each $y_i = 0$ and $\qquad f(x_i,\theta) = \sqrt{-f_i \log(P_i)}$

Another approach is to use the method of scoring.

Let

$$P_i(r) = \partial P_i / \partial \theta_r$$

and

$$S_r = \Sigma f_i P_i(r)/P_i .$$

The matrix H has (rs)th element

$$n \Sigma f_i P_i(r) P_i(s)/P_i .$$

The iterative step is then

$$\theta_{new} = \theta_{old} - H^{-1} S$$

where S is the vector $(S_1, \ldots, S_k)$.

This approach is a generalization of the Newton-Raphson method for finding the root of a system of equations. You begin by setting the partial derivatives of the likelihood equal to zero, and then find the solution. The reason I like this method is that you need the partial derivatives of P with respect to each parameter. That is equivalent to taking partial derivatives of the distribution function with respect to the parameters, and in most cases, it is easy to do. There is one difficulty, of course, in using any iterative approach to obtain a solution. You need to have a good starting value. Generally you can use a simple estimation technique, such as the method of moments or percentile matching, to get starting values.

A final useful item from the method of scoring is the matrix H^{-1}. When you are finished with the iteration procedure, the contents of that matrix is an estimate of the co-variance matrix of the estimator.

I then fit three distributions to the dental data. The results were:

Pareto	$a = 3.14027$	$\lambda = 484.309$		$L = 1911.510$
Burr	$a = 14.3359$	$\lambda - 5593.07$	$r = 0.71511$	$L = 1908.763$
Loglogistic	$\tau = 1.57713$	$\lambda = 143.929$		$L = 1924.816$

The loglogistic distribution is the inverse Burr distribution with the parameter alpha set equal to one. At first it seems obvious that the Burr distribution is the best choice from these three since it has the smallest value of L (the negative of the loglikelihood). However, the Burr distribution must provide a better fit than the Pareto, since the Pareto is a special case of the Burr distribution. To determine if the third parameter of the Burr distribution is of value, we examine twice the difference of the loglikelihood, or 5.5, which is well past the 97.5

percentile of the chi-square distribution with one degree of freedom. It appears that the Burr distribution does indeed provide a good fit for what we have observed. We can also compare the observed frequencies with the fitted frequencies. We can also compare the Burr PDF to the histogram (see figure 3):

BURR MODEL

Loss	Observed	Fitted
50- 99	281	281
100- 199	314	305
200- 299	147	163
300- 499	162	162
500- 699	81	73
700- 999	52	50
1,000-1,499	30	29
1,500-1,999	7	9
2,000-4,999	4	6
5,000-7,499	1	0
Average	294	295

So, it looks reasonable to accept the Burr distribution as an appropriate model for dental losses, and now we can try to answer the three questions posed previously.

What would the net premium become if the deductible was removed? First of all, the probability of a loss less than \$50 is given by the cumulative distribution function evaluated at 50, which is 0.38, so if there were 200 losses at \$50 or more, we can expect 200 divided by 0.62, or 324 losses in all. This Burr distribution has a mode of zero, so it is not surprising that it produced a large number of losses under \$50. Second, this Burr distribution has a mean of \$189. Therefore, if the deductible of \$50 were to be removed, the annual net premium would increase to 324 times \$189, or \$61,418. Part of this increase is due to paying the extra \$50 on each of the present 200 claims. The remaining \$2,600 is from the additional 124 losses. While extrapolation is always dangerous, we have no other source of information about losses less than \$50. As a matter of fact, this data set did contain information on losses below \$50. It turned out that the average loss was \$33 and the frequency 0.16. So, my extrapolation was not close at all, and the net premium should have been \$60,000 instead of the \$61,400 as observed.

What would the net premium become if a limit of \$1,000 per claim was imposed while retaining the deductible of \$50? The net premium can be found from the PDF as follows:

$$324\left\{\int_{50}^{1000} (x - 50)f(x)dx + 950[1 - F(1000)]\right\} = 44,813 .$$

It represents the expected losses over the range, plus the \$950 paid any time a loss exceeds \$1,000, leading to a net premium of \$44,800. Imposing a limit leads to a reduction of \$4,000 in the net premium.

FIGURE 3

Fitted Burr distribution

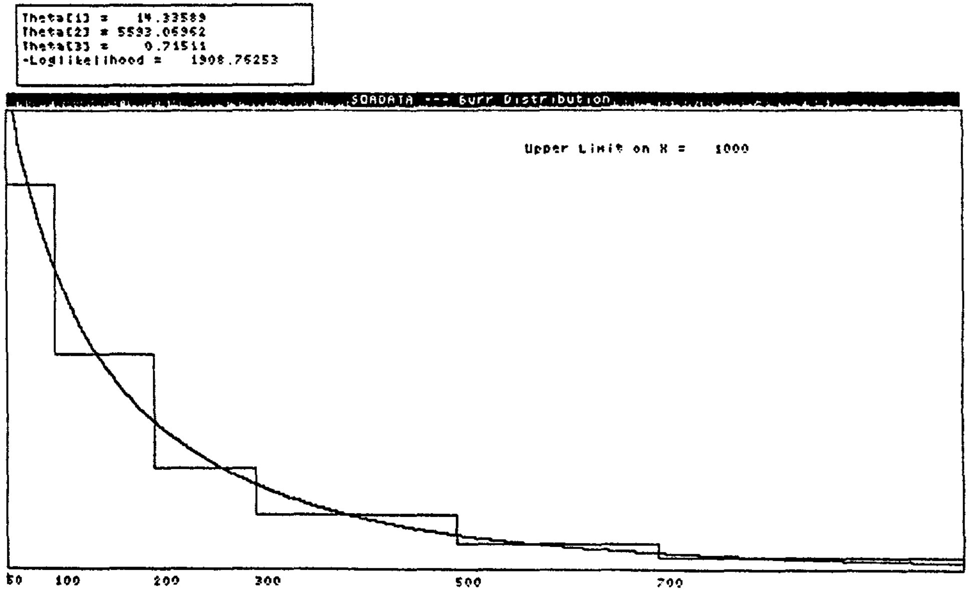

Suppose inflation increases losses uniformly by 10 percent. What is the effect on the net premium? This new net premium is:

$$324(1.1)\int_{50}^{\infty}(x - 50)f(x/1.1)dx = 54,923,$$

This net premium would be an increase of 11.2% over the previous net premium, a rate slightly greater than the rate of inflation.

My final comment is on the use of maximum likelihood estimates to make inferences about the estimates that we obtained.

The inverse of H from the method of scoring is an estimate of the covariance matrix of the maximum likelihood estimator. In the example, it is:

513.225	348,368	-2.67607
348,368	236,873,000	-1,839,03
-2.67607	-1,839.03	.0154663

Suppose I would like to estimate the standard error of the estimate of the number of additional claims under $50, if that deductible is to be removed. Recall that the formula by which that number 124 was obtained is 200/[1-F(50)]-200 which is a function of the three parameters I had to estimate. So consider this as a function of three parameters, and therefore, since the maximum likelihood estimates were inserted, this quantity can be considered as a random variable. Its variance can be estimated using a particular formula involving derivatives of this quantity with respect to the parameters, giving us the vector h. The variance is $h'H^{-1}h=2,549$. We can then construct the 95 percent confidence interval as 124±1.96(50). With 95 percent confidence, we can expect that when the deductible is removed, the number of extra claims will be anywhere from 25 to 223.

This gives an excellent indication of our ability to extrapolate. In this case, it points out how little faith we can place on it. The true number of extra claims, since I actually had that information, was 38. It is within the confidence interval, which gives me some reason to believe that we are operating correctly in this example. Also, with this accompanying confidence interval, recognize the limited value of the data set we collected. Perhaps such a confidence interval would encourage us to get more information from other sources to help refine that estimate.

I want to acknowledge the support provided by the AERF for this work and the considerable help provided to me by my colleague at the University of Iowa, Robert Hogg, and to reinsurance actuaries Charlie Hewitt and Gary Patrik. Our combined efforts resulted in the text, Loss Distributions, which contains much of the material included in this presentation. Also, my thanks to the anonymous donor of the dental data that has proved to be a useful illustration of these techniques.

MR. ARNOLD P. SHAPIRO: My topic Applications of Operations Research (OR) Techniques in Insurance, emphasizes tools and techniques to solve practical insurance problems.

Some would regard this methodology-oriented view of OR as too narrow. Jewell (1980, p. 113), for example, would prefer to stress the system-building opportunities and areas for constructive interaction within insurance, rather than tools and techniques. However, there has been a concern voiced by actuarial students and many practicing actuaries that they have a problem conceptualizing practical applications for OR techniques in their daily work, particularly applications that are unique to insurance. In response to this concern, the focus of this AERF-sponsored project has been a comprehensive review of the applications of OR techniques in insurance.

The purpose of this presentation is to share with you some of the applications found in the literature reviewed. The discussion is not meant to be comprehensive and is intended primarily to stimulate your interest.

It may be of interest to mention the journals that were consulted for the study. In this regard, an attempt was made to cover the major journals in operations research, insurance, and related fields of business. Table A shows the journals reviewed.

TABLE A
JOURNALS REVIEWED

American Economic Review
ASTIN Bulletin
Bell Journal of Economics
CLU Journal
Computers and Operations Research
CPCU Annuals
Decision Sciences
European Journal of Operations Research
The Geneva Papers on Risk and Insurance
Insurance: Mathematics and Economics
Journal of Applied Probability
Journal of Business
Journal of Financial and Quantitative Analysis
Journal of Political Economy
Journal of Risk and Insurance
Management Science
Mathematical Programming Study
Omega: The International Journal of Management Science
Operations Research Quarterly
Operations Research
Opsearch
Scandinavian Actuarial Journal
Skandinavisk Aktuarietidskrift

Turning now to the particular applications, consider first game theory. Game theory involves competition or conflict between two or more decision makers and is concerned with prescribing best strategies. Specific insurance applications include topics such as insurance purchases, management, and expense allocation.

Williams (1960) discussed the use of pure strategies in game theory for the evaluation of insurance consumption alternatives. The analysis was based on loss in utility associated with the decision of whether or not to buy fire insurance. This is one of the traditional OR applications in insurance that has found its way into insurance and OR textbooks.

Further insight into the development of this model was contained in Williams and Dickerson (1966), and an empirical investigation of the model was reported in Neter and Williams (1973).

One of the earlier models of a game for a property and liability insurance company was the executive game for officers and middle management suggested by McGuinness (1960). The inputs to the model included basic decisions, assumptions, and data while the outputs of the model were reports to the players and an analysis of the effect of their decisions. Assumptions included those that were known to the players and those which described the environment and were not known to the players. The data included such information as economic activity, underwriting experience fluctuations, and the cost of training agents.

Lemaire (1984) discussed the application of game theory to the problem of allocating expenses among the departments of an insurance company when cooperation leads to economies of scale. He first showed that the cost allocation problem was identical to the value of a game with transferable utilities and then discussed the attributes of four cost allocation methods based on game theory. The criteria advocated are collectively rational, in the sense that no departments subsidize another; monotonic costs, in that all departments contribute to an increase in global costs; and additivity, in the sense that a subdivision of a department does not affect the cost allocation.

Stochastic dominance provides another means for preference ordering when uncertain alternatives are involved. Gandhi, Saunders, and Sugars (1981), discussed a simple reinsurance application where the manager of a stock company must choose between two portfolios: one containing reinsurance and the other not. The paper described in detail the characteristics of first-, second-, and third-order stochastic dominance and discussed the superiority of stochastic dominance over mean-variance, coefficient of variation, and expected utility models.

Consider next the area of linear programming. This has long been recognized as one of the most important techniques of OR because of its versatility and power in resolving problems involving the allocation of scarce resources.

An interesting example deals with a perennial problem in the sale of life insurance, the optimal combination of various types of life insurance policies and alternate investments. Schleef (1980) showed how a linear programming model could be used to help resolve this problem.

The objective of the model was to maximize the present value, adjusted for the marginal tax rate, of future cash flows due to cash value recovery, loans, and other investments. The constraints of the model included a budget constraint, which provided for the payment of premium and loan interest and recognized alternative investment cash flows; a death benefit constraint, which provided for the desired level of death benefits; and nonnegative constraints on the alternative investment fund, the cash value, the loan balance, the face amount, and all decision variables.

An additional dimension is added by quadratic programming. Here, the concern is with optimizing a quadratic objective function subject to linear constraints. One of the most common uses of quadratic programming is in the resolution of questions related to portfolios. The origin of this approach was the work of Markowitz (1952), who suggested using probability estimates of future security performance to develop an efficient set of portfolios which could then be matched with an investor's preference.

Sharpe (1963) discussed an efficient resolution of the problem which greatly simplified the analysis.

Markle and Hofflander (1976) applied the Markowitz model to nonlife insurers under the assumption that the goal is to maximize returns for given levels of risk subject to institutional solvency constraints. The objective function is the underwriting and investment profits. For this purpose, the overall return for a given line was assumed to be an average return for that line over the period of investigation, as were the expected returns from securities.

A related area is goal programming. The essential feature of goal programming is that it provides an opportunity to assign priorities to conflicting objectives and then minimizes deviations from those objectives. Gleason and Lilly (1977) examined goal programming as an agency decision-making tool in the context of property and casualty insurance agency decisions regarding number of insurers to represent, cost reduction efforts, and expanded commercial lines. Rather than deal with all possible goals that might confront such an agency, they limit themselves to common goals.

The common approach to the analysis of alternative investment opportunities is based on quadratic programming models because the objective function involves the variance of distributions, a quadratic. Brodt (1983) showed how to develop a linear programming alternative based on the mean absolute deviation of returns, rather than the variance of returns. The objective was to minimize risk, as measured by mean absolute deviation, subject to intra-temporal and inter-temporal constraints. Since the objective function was nonlinear, goal programming was required.

Linear programming is a one-stage process. Dynamic programming extends the single-stage assumption of linear programming to a multistage environment and is concerned with the overall effectiveness of sequential, interrelated decisions over the planning horizon. One of the early attempts to bring dynamic programming to bear on the problem of choosing an optimal life insurance program was a study by Belth (1964). The study was naive by current standards, but the interrelationship of the variables is as relevant today as when it was first done.

Another OR technique is inventory models. They seek the optimal balance between the cost of holding inventory and the cost of procuring it. There are, of course, numerous examples of applications involving general inventory problems in business, including the insurance business. However, an interesting specific insurance application deals with proceduring insurance coverage.

Smith (1968) envisioned the optimal insurance coverage in the context of an optimal inventory stockage under uncertainty. Specifically, if an insurable loss (demand) is exceeded by the insurance coverage (inventory level), excessive insurance cost (inventory holding cost) is incurred. Conversely, if insurable loss is greater than insurance coverage, unrecoverable losses must be absorbed by the insured. The problem, then becomes one of choosing optimal insurance (inventory) levels.

Similarly, the numerous applications of queuing theory, used to resolve waiting-time problems of business, are directly applicable to the insurance industry. An application, however, which is unique to insurance has to do with the ruin of an insurance enterprise. Under a proper scenario, the solvency of an insurer may be viewed as a queuing problem, where the probability of not being ruined by some time t, u(u,t), is essentially the same probability as that involving a customer waiting less than some time U, given that the customer joined the queue t periods after the server was free. A discussion of this aspect is given by Seal (1978, Chapter 2).

Another technique is the Markov process, which provides a dynamic system under which only the immediate past is relevant to the prediction of future behavior. Its application to life tables is obvious. Since working life tables evolve from the dynamics of labor-force participation, Hoem (1977) advocated that such tables be produced using the theory of continuous-time Markov chains. The transitions in this case are due to death, accession to the labor force, and separation from the labor force. His discussion of an application of such a model is based on a previous study by him and Fong (1976), which contains the formal details of the model.

Similarly, Braun (1978) emphasized stochastic stable population theory where the forces of fertility and mortality depend on age, parity, and place of residence.

Because of the easy access to computers, the use of simulation within insurance and related industries has become commonplace. Paralleling this application of computers and simulation has been the need for the development of models which adequately encompass particulars of specific areas within the industry. Representative examples of areas of application include variable life insurance and reinsurance.

Brennan and Schwartz (1979) used simulation to explore risk-reducing investment strategies associated with equity-linked life insurance policies. There are implications for the development of a Bayesian approach for finding estimates of the model's parameters.

Galitz and Brown (1981) discussed the qualitative nature of a simulation model for insurance and reinsurance operations. Although the relationships were not specifically defined, important overall considerations were delineated. The basic components of the model were surplus, capital, unearned premium reserve, and loss reserve.

The foregoing examples stress the application of OR techniques. However, OR authorities view the "systems approach," which coordinates overall relationships and interdisciplinary teams, as the fundamental thrust of OR. Areas where the systems approach have been used are population planning and workers' compensation.

Reinke (1970) discussed the role of model building in population planning in underdeveloped countries. The problems were a conflict between national and family goals, sparce relevant information, and limited resources. Within this framework, the role of OR was to analyze the decision process and organize professional activities in this area.

Jewell, Johnston, and Leavitt (1974) discussed the multidisciplinary nature of a comprehensive project involving workers' compensation insurance. Each phase of the project is discussed, as are the interrelationships of the phases. The study embodies the spirit of model building which OR authorities stress.

This presentation provides only a cursory overview of the literature pertaining to applications of OR techniques in insurance. Nonetheless, it is hoped that it stimulates discussion and provides direction and insight into further research in this area. To the extent it has met these criteria, it will have served its purpose.

REFERENCES

Belth, J. M., 1964, "Dynamic Life Insurance Programming," Journal of Risk and Insurance, 31 No. 4, 539-56.

Braun, H., 1978, "Stochastic Stable Population Theory in Continuous Time," Scandinavian Actuarial Journal, No. 4, 185-203.

Brennan, M. J. and Schwartz, E. S., 1979, "Alternate Investment Strategies for the Issuers of Equity Linked Life Insurance Policies with an Asset Value Guarantee," Journal of Business, 52 No. 1, 63-93.

Brodt, A. I., 1983, "Min-Mad Life: A Multi-period Optimization Model for Life Insurance Company Investment Decisions," Insurance: Mathematics & Economics, 2, 91-102.

Galitz, L. and Brown, M., 1981, "ASIR: The Advanced Simulation Model of Insurance and Reinsurance Operations," Geneva Papers on Risk and Insurance, 19, 85-96.

Gandhi, D. K., Saunders, A., and Sugars, E. G., 1981, "Stochastic Dominance: An Application to the Insurance Portfolio Decision," Geneva Papers on Risk and Insurance, 21, 51-62.

Gleason, J. M., and Lilly, C. C., 1977, "A Goal Programming Model for Insurance Agency Management," Decision Sciences, 8 No.1, 180-90.

Hoem, J. M., 1977, "A Markov Chain Model of Working Life Tables," Scandinavian Actuarial Journal, No. 1, 1-20.

Hoem, J. M. and Fong, M. S., 1976, A Markov Chain Model of Working Life Tables, Copenhagen University Laboratory of Actuarial Mathematics, Working Paper No. 2.

Jewell, W. S., 1980, "Models in Insurance: Paradigms, Puzzles, Communications and Revolutions," Transactions of the 21st International Congress of Actuaries, S, S87-S141.

Jewell, W. S., Johnston, T. L., and Leavitt, S. S., 1974, "Operations Research in the Insurance Industry: II. An Application in Claims Operations of Workman's Compensation Insurance," Operations Research, 22 No. 5, 929-41.

Lemaire, J., 1984, "An Application of Game Theory: Cost Allocation," ASTIN Bulletin, 14 No. 1, 61-81.

Markle, J. L., and Hofflander, A. E., 1976, "A Quadratic Programming Model of the Non-life Insurer," Journal of Risk and Insurance, 43 No. 1, 99-120

Markowitz, H. M., 1952, "Portfolio Selection," The Journal of Finance, 12 (March), 77-91.

McGuinness, J. S., 1960, "A Managerial Game for an Insurance Company," Operations Research, 8 No. 2, 196-209.

Neter, J. and Williams, C. A., Jr., 1973, "Performance of the Expected Utility Method and Two Other Normative Methods in Insurance Decision Making," Decision Sciences, 4, 517-32.

Reinke, W. A., 1970, "The Role of Operations Research in Population Planning," Operations Research, 18 No. 6, 1099-1111.

Schleef, H., 1980, "Using Linear Programming for Planning Life Insurance Purchases," Decision Sciences, 11 No. 3, 522-34.

Seal, H. L., 1978, Survival Probabilities, New York: John Wiley & Sons.

Shapiro, A. F. and McAdon, J. M., 1982, Operations-Research Applications for Pension Consultants and Administrators, Washington, D.C.: American Society of Pension Actuaries.

Sharpe, W. F., 1963, "A Simplified Model for Portfolio Analysis," Management Science, 9 No. 2, 277-93.

Smith, V. L., 1968, "Optimal Insurance Coverage," Journal of Political Economy, 76 No. 1, 68-77.

Williams, C. A., Jr., 1960, "Game Theory and Insurance Consumption," Journal of Risk and Insurance, 27 No. 4, 47-56.

Williams, C. A., Jr. and Dickerson, O. D., 1966, "Game Theory and Insurance Consumptions: The Worry Factor Revisited," Journal of Risk and Insurance, 33 No. 3, 411-25.