RECORD, Volume 31, Number 1*

New Orleans Life Spring Meeting
May 22–24, 2005

Session 54TS
Introduction to Research Methods for Actuaries

Track: Education & Research

Instructors: KRZYSZTOF M. OSTASZEWSKI
ANTHONY WEBB

Moderator: KRZYSZTOF M. OSTASZEWSKI

Summary: Research in some areas of interest to actuarial practitioners is plentiful, while in others little research is published in the actuarial literature. In today's competitive environment, actuaries are becoming increasingly dependent upon staying abreast of the latest developments in applicable actuarial knowledge. Understanding the significance of existing research results, as well as having the skill to conduct research in areas where little is known, will likely be critical to the future of the actuarial profession. This session examines the importance of actuarial research, how and what to do, researching the literature, statistical tests of significance, writing a research report and how to get published.

MR. KRZYSZTOF M. OSTASZEWSKI: We will be both presenting some ideas about research in general and the relationship of research in general to what happens to actuaries. I am the actuarial program director at Illinois State University. I'm also involved in studying mortality improvement. There are some birth cohorts that have significantly better mortality improvement than other cohorts. That's our topic, which will be the second part of my presentation. Anthony Webb will speak after me about his ideas on research. But first I will discuss some general ideas about research with you. Albert Einstein said, "If we knew what it was we were doing, it would not be called research, would it?"

What is research in general? What is the meaning of the word? It actually means different things when you talk to different people. I'm trying to present the most

Note: The chart(s) referred to in the text can be downloaded at:
http://handouts.soa.org/conted/cearchive/NewOrleans-May05/054_bk.pdf.

general perspective and give you some ideas how different people can do this differently. Because if you engage in research in actuarial science, you're very likely to run into people who have very specific ideas about how research should be done—what we call scientific method. I remember participating in a discussion group on the Internet one day, when people were so fascinated with science that they proposed to change the timeline and date everything from the publication of *Novum Organum* where the scientific method was introduced. There was a very important change in attitudes toward knowledge and knowing the reality and understanding it at the point when we started using the scientific method openly.

When we're working using the scientific method, we always make our theories confront reality and be verified by reality. You might think that this is in a sense obvious, but we still don't do this in every activity, even not in every activity where we have scientists involved. The verification by reality is a very natural thing in all physical sciences. It gets much more difficult in other sciences, and it's really not as common. It gets to be problematic, because experiments in social science cannot be repeated, as I like to demonstrate to my students. I always ask my students, "Why is it that North Korea is poorer than South Korea?" People always volunteer answers about the economic structure and structure of the government.

But there are alternative theories. One is that all lazy Koreans live in North Korea, and the hard-working Koreans live in South Korea. I actually have empirical evidence to prove that. The productivity is significantly higher in South Korea, so they must be harder working, and that obviously explains everything. The other possibility is that there has been a series of about 40 or 50 years of bad crops in North Korea; it's just a natural random phenomenon that caused that country to have starvation. And again, I have empirical evidence, if you just look at the output of the agricultural sector. And that is my attempt to illustrate to them that in social sciences you have a much harder time performing an experiment and verifying everything. You can't go back and perform that experiment again under slightly changed circumstances, which is a very common thing to do in all physical sciences.

The one very important thing that I want you to understand when you have to deal with academics such as myself is that publication of research is the central part of the work of academics. And the tenure decision, which is life or death of an academic, is based mostly on research. I find it fascinating that when people talk about professors, they say that someone teaches at a particular university. I assure you they do some teaching, but that's a secondary activity. If you're an actuary, you don't always do actuarial work; you do other things at work, but the core of your work is actuarial. Most academics perceive the core of their work to be research, so if they "do" anything at their particular university, they do research. There are other secondary activities that hopefully help them pay the bills, but what's on their minds first and foremost is their research.

Research in an academic setting means publish or perish, which gets a bad name, but has some good things to it. The core principle of what they are supposed to do

is make a new contribution to the body of knowledge of whatever area they study. It has to be verified and tested by peers, and it has to be new. They have to, whether they like it or not, make some new contribution. It may be a very weird contribution sometimes, but it has to be new.

When you submit your work to a research journal of any kind, you will automatically be in the company of academics, because what you're doing is their core activity. They live or die by submission of papers. When you submit a paper to a research journal, the paper is submitted for review by an editor or two or more referees. The referees are almost automatically people who have published in the area. If you write a paper, at the end you should refer to a previous work published. Any name among those references is a possible candidate for being a referee for the paper. In addition, anybody who is known to have done work in the area is a possible referee for that paper. It doesn't reflect well on you if somebody is well known for work in the area and that name does not appear in your list of references. Clearly you ignored the work of somebody who's important in the area, and you may pay a price for that.

The two central issues that are always supposed to be checked are the accuracy of the work and whether the work is relevant for the existing body of knowledge. Any inaccuracy will probably disqualify the work for publication. You also have to put your work in the context of what is out there. You have to fit it into the picture to some degree.

So what do you do if you are an aspiring researcher and want to get published in one of those journals with the swarm of academics who view the world their own way? It should be easy to check accuracy, which means what you're doing should be transparent, should be obvious, and should be very easy to check. How do you do this? You explain what you're doing as much as possible. That's actually not always how you write practical papers. If you ever write a paper that is eventually going to give you a job in advising or consulting for somebody, you don't want to give all the details and all the explanations. You want to hold something back. But that's a very bad thing for anything that's supposed to be published. Every single detail is supposed to be verifiable by others. And if it's not, then it's not a scientific paper. It's a very different perspective.

Your paper should not redo what has already been done. That's the kiss of death for publication, because you're very likely dealing with people whose job is to know everything about the area. If they see that your topic was done by so-and-so in 1984 in this and that publication, then that's the end of the story. They will wonder why you are even bothering submitting this when it's been done already.

Your paper should address the existing research on the subject. If you're not putting it in the context of what's been done, you're not doing what is perceived by academics in general to be theoretical research, or even any kind of research. Research is supposed to make a contribution, so you have to make it clear what the

contribution is by placing it in the context of what exists. And clearly define what your contribution is. That's one of the principles of communication: at the end of the presentation, you make a point. Don't try to ask people to guess what your point is. You have to make it clear what your contribution is and how you're adding to the body of research.

So you're a practicing actuary and you want to do some kind of research and get published. It's not that there's any perception out there that actuaries don't know what they're doing in terms of setting up their models, doing calculations, or doing their experiments. I like this as a description of actuaries: actuaries are persnickety. There are other words to describe it, but "persnickety" is perfect. Actuaries pay attention to detail and make sure that things are correct. That's a very big plus if you're working with anything research-oriented. But knowledge of what has been done already and placing things in the context of existing research sometimes may be a problem for actuaries, because you have to know what that body of research is. You may say, "Well, I know how things are done with the people that I've worked with," but it's not the same thing from the perspective of people who publish in research journals, because you're supposed to have read all the papers on the subject. That's what academics do.

I have a friend who's an economist who says that a true researcher spends 95 percent of his time reading and only 5 percent writing. That's partly because you're supposed to and partly because there's so much to read out there. You have your assistants who do some analysis for you. For academics, the most important thing is reading, reading, reading about things that have been done, and they have graduate students doing this. If you don't have anybody doing this for you, it really becomes much more difficult to compete with those people. The first part of their job is to find all existing papers, not just the ones written in the last couple of years, but also the ones written in the 1950s and 1890s, if necessary.

How do you find existing papers? There are databases of research publications. The world is getting much better in terms of how this works. It's now possible to retrieve many papers online. If you are a member of the American Risk Insurance Association, for example, it's now become possible to retrieve papers from old issues of *Journal of Risk and Insurance.* Of course, you can retrieve papers published by the SOA online from their Web site. There are research databases that actually allow you to get PDF files of papers. It's very common for universities to have memberships in those databases as part of a membership in their libraries. So if you become a member of a library at any larger university, you're likely to have access to those online databases of research papers. And if you don't have access to those, find some academic who can help you with that, because they're supposed to help the general society. They should help you out if they're interested in actuarial research.

There's a very interesting Google search engine, scholar.google.com, but also just general search on the Internet is a very good idea. Google is getting in trouble with

universities for scanning too many things, but it is probably the future that more and more materials will be available online. Of course, there is a great positive benefit from this, that we will have easier access to all those materials. Some of them are owned by somebody, and there will be ownership issues. Hopefully we'll resolve them fairly. Nevertheless, easy access to historical data for the researcher is very important.

Research conferences that people attend are also useful for the purpose of exposure to existing research and talking to people about it. The Actuarial Research Conference looks very different than what you're used to in SOA meetings. You'll hear people discussing their research details and who has done what in this year or that.

Now I'll tell you a little bit about the project on mortality improvement that we're working on. We're actually not studying mortality improvement per se, but whether there are birth cohorts that have better mortality improvements than other birth years. We're looking for this globally. The partners in this research are Richard MacMinn, me, Ranee Thiagarajah and FrederikWeber.

Here is another illustration of how the Internet is changing the world for us. Our data source is a human mortality database that's been created over the last several years as a joint project of the University of California at Berkeley and the Max Planck Institute for Demographic Research in Germany. They regularly collect updated data on mortality worldwide, at least once a year. At the very least, you can get information about all cohorts for which they could get data. For Sweden and England, for example, from roughly 1750 until now, you can see population size for every birth cohort and see how they change over time. It's quite a fascinating project, and it's getting easier to collect the data, so we will have more and more of it over time. We also used some interesting data for the United Kingdom from the Bank of England.

The slide on page 20 shows a perspective on the history of the cost of life annuities at 3 percent, 5 percent and market rates from age 65 to age 100 in England and Wales. The X-axis is the birth year, and the Y-axis is the actual cost per 10,000 annual income. The top line is at 3 percent; the bottom line is at 5 percent; and the yellow line that oscillates is the going nominal market rate for the longest term existing bond as of the date when the person reaches 65. It's quite a fascinating graph that tells us why we would be worried about mortality improvement, because there's a point where these graphs take off. The wild oscillations in the yellow curve are not caused at all by mortality; they're caused by the oscillations in interest rates that we experienced in the first part of the 20th century, essentially through 1980.

For people born around 1840, we have a large increase in terms of the cost of retirement. There is an unprecedented improvement in longevity that creates problems, because someone is going to have to pay for it. We have social insurance

pension systems, private retirement plans and life annuities where costs will dramatically change. In addition to the increased cost of retirement annuity, we also found that the oscillations between two consecutive birth years for the cost of retirement annuity could be as high as 3 percent of the cost of the annuity, with 1 percent also occurring quite often. We estimated, for comparison, the duration of life annuity at aged 65 to be around seven for the cohorts that we had in the database. It's probably higher for the cohorts retiring now, but seven was a reasonable estimate. So if we used seven, the convexity of the life annuity would be about 120, and we could estimate the change in price that corresponds to a change in interest rate. Based on these estimates of duration and convexity, we find that a 1 percent annuity price increase is roughly is equivalent to a 14 basis points decline in interest rates.

The warning here is that we found that a 1 percent oscillation between two birth years in the history of these birth cohorts that we studied is actually a very common occurrence. That could be up or down. You have cohorts that are 1 percent more costly to retire than the two neighboring cohorts, and that's in terms of interest rates roughly equivalent to 14 basis points. So then the question is: What is the profit margin on life annuities that you have? If you're anywhere in this range where 14 basis points is significant, there's real trouble out there, because this oscillation occurs frequently.

The first equation in the slide on page 23 shows the official definition of mortality improvement that's been used by people who study it. Z is the birth year. We compare the birth year, the birth year before, and the central death rates. Central death rates are easier to work with. They're more directly connected to force of mortality, so it makes things a little more convenient. We are not convinced that this is the best definition. For example, we could compare longevity improvement the same way. Then there's the question of what it means to be a select cohort or a birth year cohort. So we studied these definitions where a given birth year has more improvement than the neighboring birth years at a majority of ages. You compare these mortality improvements at every age separately. We also studied better improvement than the 80^{th} percentile of a given age improvement at a majority of ages or 80 percent of ages, or better longevity improvement. And you do get different definitions; you get different outcomes based on these definitions.

The interesting thing is that though mortality improvement is studied quite widely and is very important for both governments and private companies, there's surprisingly little research on these select birth cohorts. Most of it was done by Richard Willets, who made the point that there is a very strong collection of birth cohorts in the United Kingdom of people born sometime in the 1930s. In his paper "Insights and Explanations," he suggests that the major contributing factor was that there was a generation who went to World War II who didn't get killed. He didn't put it that way, but I can't resist putting it this way, because I grew up in a country where military service was mandatory, in Poland, which was a Communist country. The one thing that happened to you when you were drafted was that you

started smoking. It was automatic. You went into the military and started smoking, then you came out and you tried to quit. One thing that happened during World War II was that all the countries that went to war started smoking. You know, it was a bit stressful to be out there.

Willets makes the point that the generation that came of military draft age immediately after World War II ended smoked less. In fact, there's a decline in rates of smoking that he illustrates. He believes this to be the major contributing factor to this cohort effect in Britain that those people who didn't get to start smoking during World War II will start retiring soon.

It's a very interesting hypothesis. We are not completely convinced of it, because we find cohorts in other places too. But it's still an interesting idea; that's why I have this perspective that the war made them smoke, and the following generation didn't. We found that there seems to be in many countries a very strong birth cohort born in 1919 and in 1946. Any way you look at it, these two years have only one thing in common: they are the years after the respective wars ended. Our initial research on these birth cohorts indicates that it's a small effect relative to the general trend of mortality improvement. That's good news.

The bad news is that it's probably big enough to kill somebody's profit margins. While general mortality improvement is something people are afraid of and study a lot, this is new. The higher cost for one birth year and then possibly lower the next year is something that may come as a surprise. In addition, when you combine a group of cohorts for various birth years that are together, if those turn out to have significantly better mortality improvements than some neighboring group of cohorts, this may be a problem for governments. We have long-term projections that are in aggregate, but then when a group like this starts retiring there may be a very sharp increase in the cost of benefits that will be unexpected, because in the other projections we have to average in everything else.

If you look at the projections of Social Security and Medicare, you will see that there's a dramatic increase happening starting around 2008 in the costs of Social Security and in 2011 for Medicare. I think that's to be expected, because when will the babies born in 1946 be 62? That's really the question, because that's the earliest retirement age for Social Security.

We are studying additional scenarios and asking the basic question of whether it's possible to do a better pricing of annuities based on your birth year. Interestingly, there was a pension reform in Poland in 1999, where the government now pays the social insurance pensions based on the life expectancy of your birth cohort. So when you retire, you will receive a benefit that is based on the social insurance authority's estimate of your birth cohort. I always joke that the more mathematically astute kids in high schools will now be handing out cigarettes to other kids in the class to ensure themselves a higher retirement benefit.

There was a first issue of survival bonds, which pay a coupon proportional to the number of survivors in a given birth cohort, in 2004 by European Investment Bank. It's a very interesting question, again purely theoretically, but it may have the long-term practical implication of who should issue such bonds, because effectively governments that issue social insurance benefits are issuing them. But we don't trade them. If we traded them, maybe we would have a better market signal about what the real cost of retirement is and what the risks are. This is still a very, very small market, but it's a very interesting financial innovation to have mortality-based securities out there.

In closing, the slide on page 28 shows an amazing picture of the data that's available in this mortality database. In virtually every country, the percentage of the population over 65 is 4 percent for a very long time, and then we have a lift off. This shows the calendar year they become 65, so we have a lift off in the 20th century. What happened in Japan is just astonishing, while in the United States and in England and Wales this process has actually somewhat turned around, in the United States especially. There was a study at Boston University that asked the question of why the United States was the leader in terms of expected longevity in 1980 at age 65, and that's no longer the case. The hypothesis proposed that it's our eating habits. They don't know for certain. But this graph shows that this process has at least slowed down in the United States. It doesn't show any signs of that in Japan. It's also going to be interesting to watch China in this respect. Just recently, the Chinese government started a campaign about exercise and good health so that you can work in your older years. In 20 years, the Chinese working population will start declining. It's one thing to have a lot of people, but another thing to have a lot of people who can't possibly work, because they have reached the age when it's very difficult. It's going to be very interesting to watch this.

This is the end of my discussion, so I'll let Tony have the floor.

MR. ANTHONY WEBB: I'm actually an economist rather than an actuary. I was previously working at the International Longevity Center in New York. I've very recently started work at the Center for Retirement Research at Boston College. I'd like to build on Krzysztof's comments and hopefully not repeat any of the very good points he's made.

My first observation is: Why bother? Academic research is very time consuming. A good and productive researcher who is spending most of his time undertaking research might publish at most one or two papers a year. One can realistically reckon that a quality paper takes an individual maybe four to six months of full-time work to write. Consequently, first you need an employer who's willing to have you spend that amount of time doing work that may not necessarily be directly related to the employer's work. You may be undertaking work in the course of your employment that you think might turn into publishable research. One of the problems you'll face is that work that might interest your employer may need a lot of modifications to make it publishable in an academic journal. Finally, the end

products of most academic research get rejected, so your six months of hard work may result in nothing at all to show for it.

The next issue that you might consider is whether you have the requisite skills. There is a set of skills that may make a good researcher: academic writing, which is not the same as professional writing, math and econometric skills, programming skills, presentation skills and, obviously, perseverance. If you don't have all of these skills, you should go out and find a coresearcher whose skills complement your skills.

The strongest piece of advice I would give, and building on what Krzysztof has said, is that it's important to get to know the relevant literature. Learn how the research in your chosen field has developed over the years. Learn what statistical or econometric techniques have been used, what data sets have been used, and who the important contributors to the literature are. I would recommend spending a period of months or even years simply reading. In the course of my doctorate, we were led into the process of doing independent research. We would first spend one term simply reading the literature and writing critical appraisals of what we read. I would strongly recommend that.

Regarding data sources, I obviously agree entirely with everything that Krzysztof said. There are three sources that I find most useful. First is Econlit, a search engine that's available through academic libraries. It's not available to the general public, so you'll need to find a friendly academic who is willing to help you. It is a wonderful search engine. You can search by author, by keyword, and by words that appear in the abstract. So one can go into it and type in Krzysztof's surname, if you can spell it, and it will bring up all of his publications. One can type in the key phrase "longevity improvement," and it will bring up all the publications that have "longevity improvement" in the abstract. It obviously has a focus on economics, but it also covers all the papers that appear in the *North American Actuarial Journal.* It's a little bit weaker in the field of health economics, and a number of the health economics papers tend not to get cited there.

The second source is the references cited in the related literature. All academic papers end with a list of references. You should read this list carefully to get to know who the important contributors are. Get the papers cited, then go and get the other papers cited by those persons.

The third source is academic conferences. Kris made reference to the actuarial conferences. The problem with a lot of conferences is they tend to be closed. As an example, the National Bureau of Economic Research (NBER) holds a series of summer meetings in Boston where the leading researchers in various fields get together. Unfortunately, this set of meetings is simply not open to the general public. However, you can do the next best thing and go to the NBER's Web site and download NBER working papers.

At the point when you think you have an interest in a particular field, you need to critically appraise the relevant literature. Think about where it's deficient, what questions are left unanswered and so on and so forth. A good exercise is to try writing either yourself or your employer a research grant proposal. Academics hate writing research grant proposals, but they are actually very helpful. It forces you to think about what you really want to do. If you want your employer to authorize you spending six months or so on a particular topic, you may want to convince him that it's worth doing.

For actually writing the paper, there tends to be a standard form. Papers have an introduction, which briefly describes the problem. They have a literature review explaining how your paper fits into the existing literature and contributes to it. Almost all academic papers have an economic or an actuarial model. You spend time describing that model. You usually have a data set, unless the paper is a purely theoretical paper. You then have a set of results and a conclusion. A typical academic paper runs to maybe 20 or 30 pages.

After you write your paper, prior to sending it off to a journal, it's important to get comments. The main source of comments is academic conferences. They present opportunities to get feedback, correct errors, find out what possible extensions could be made and so on. The people who attend these conferences are likely to be your referees. It's much better to get those comments from some referees prior to sending the paper off to an academic journal, rather than when the paper is in front of a referee and he has to decide whether or not it gets published.

Practice self-criticism. Look for weaknesses in your own arguments. So you have written a wonderful paper explaining the poverty of North Korea is as a result of 40 years of bad harvests. You've done all kinds of mathematical computations that show that it's entirely probable. What alternative explanations exist that a referee might find more plausible?

When you think you're ready for publication, choose your journal carefully. The different journals have different interests. The *North American Actuarial Journal* is interested in actuarial science. The *Journal of Pension Economics and Finance* is interested in pension economics and finance. The *American Economic Review* (*AER*) is interested in groundbreaking big ideas. Unless you have a groundbreaking big idea, the *AER* is not going to accept your paper. So don't aim too high, don't aim inappropriately, and take the advice of academics. Then read the paper again and think what comments you would make if you were asked to referee the paper.

I'd like to talk about an example of good academic research. I would like to talk about examples of bad academic research from my experiences of refereeing, but that is against the rules. I would also like to talk about my own research, because every researcher thinks that nothing is as fascinating as their own research. But my research is primarily in the field of savings behavior, which is maybe only of marginal interest to actuaries. So I'd like to talk about the Lee-Carter mortality

model, which is an example of good academic research, and about why the paper is so widely cited.

The authors address the very important issue of mortality improvements. There are lots of calculations that give point estimates of mortality improvements, but what actuaries and economists really want to know is how accurate those forecasts are and what confidence interval one can assign to the estimates of life expectancy in 2065. What is nice about this paper is it's accessible to a reader with a moderate level of math skills. It's not like some finance papers, which seem to be written for an audience of ten. What is even better is that the authors rigorously test the robustness of the results and anticipate the reader's questions and their concerns. So it's not only academically rigorous, it's also an act of salesmanship, and writing a paper is an act of salesmanship. You're marketing an idea.

The authors are honest about what the paper does and doesn't do. If you aren't honest it will become obvious right at the outset. There are various ways of tackling forecasts of mortality improvements. One is to gather a group of wise men. This is what the Social Security Administration (SSA) does. They ask various experts such as doctors what medical improvements they expect in the next 50 years. Will there be a cure for cancer? And based on medical expertise and other factors, they make a forecast of mortality improvements. The Lee-Carter model is not driven by medical science. It's purely extrapolative; they make an assumption that in some sense the future will be like the past. In addition, they give enough details to enable the reader to replicate their results, which can be a problem when people are using proprietary data. A lot of actuaries have access to confidential data, and it's obviously attractive to do research using that data. But there's been great concern in recent years that the findings be replicable. There have been a number of scandals—not in the field of economics or in the field of actuarial science—where researchers have simply made up results. There's been concern among the economics profession that published findings be replicable, and that means that the data must be available to other researchers.

They start off with a road map telling the reader where they are heading, which is very useful. They present the data and the demographic model, then they fit the model and evaluate its historical performance, and then give some forecasts. The data they use is U.S. mortality data 1900–1987. As you all know, there's been a 28 percent increase in life expectancy over that period. But there have been big differences in age-specific mortality decreases. There's been a 93.3 percent reduction in infant mortality rates, but only a 42 percent reduction for people aged 85 plus. Any mortality model has to explain those differences. There's also the 1918 influenza epidemic, and people have to decide whether that's a one-time event or whether it's something that might occur in the future.

The authors start off by discussing what alternative models might have been chosen and then make a reasoned case for why their particular model is appropriate. This formula

$$Ln[m(x,t)] = a_x + b_x k_t + e_{x,t}$$

shows the Lee-Carter model. Log mortality is obviously more attractive to use than actual mortality, because one doesn't want to have negative mortality rates. It's a function of age. It's also a function of the individual's birth cohort, but the birth cohort affects mortality in different ages differentially. The authors estimate this model and find that the k coefficient decreases almost linearly. That translates into a non-linear increase in life expectancy. This ties in with what we expect, that a lot of researchers think that there's a biological limit to life expectancy and that mortality will decrease at a decreasing rate as we approach that limit. They fit the model using data for 1900 to 1987 and find that fitted mortality is close to actual mortality. But that's not enough to convince a skeptic. If you take the example of North Korea, if you look at its gross national product (GNP) growth and harvests, there's probably also a close correlation. So a good researcher will always test the robustness of his model.

First, they forecast mortality reductions. They estimate a time series model for the k coefficient, and they then estimate confidence intervals for the forecast of mortality. This is what sets it this model apart from point estimates of mortality reductions. They then make an attempt to convince people that the model produces credible forecasts. One of the tests they do is to look at whether the forecast they get depends on the basis period used. So they can either forecast from 1900 to 1987 or they can use data from part of that period. They find that the forecast that they get doesn't significantly depend on the basis period used. The second test they do is in-sample forecasting. They have mortality data for 1900 to 1987. They use part of that data, for example 1900 to 1950, and generate a forecast for 1950 to 1987 based on that data. Then they check to see how that forecast compares with actual mortality 1950 to 1987, and they find that it forecasts actual mortality pretty well.

So what do the Lee-Carter forecasts tell you? First they tell you that infant and child mortality, which has decreased dramatically over the past 100 years, continues decreasing dramatically to very low, and what some people might regard as implausibly low levels. Because this is a log linear model and because k is decreasing linearly, they find that life expectancies are forecast to increase at decreasing rates. They forecast a 10.5 percent increase in life expectancy at birth over the period 1989 to 2065, the bulk of which occurs in the 65-plus age group.

Confidence intervals are what actuaries care about or ought to care about. The authors actually get fairly narrow confidence intervals. They estimate that the 95 percent confidence interval for the 2065 birth cohort, which is a long way ahead, is from plus 4.1 to minus 5.2 years. It's so narrow because k has historically decreased fairly close to linear trends, so there's not forecast to be much variation in k. The entropy of the life tables is decreasing as we're approaching what we think might be the natural limits to human longevity. In addition, the model ignores the 1918 flu epidemic and possible similar events in the future such as SARS or other

nasty bugs heading our way.

The authors anticipate the reader's questions. The obvious question that a reader might ask first is: How do these forecasts differ from SSA forecasts? The answer is that the Lee-Carter model forecasts substantially higher life expectancy than the SSA actuaries. In some sense that is comforting, because the SSA has historically underestimated improvements in life expectancy. In addition, this paper was written in 1990 when AIDS was just coming on the scene, and the authors obviously have to take a stand on AIDS. There's no right answer to this question, but they do the right thing by first acknowledging the issue and then justifying the choice that they've made.

The problem with the Lee-Carter model is that, as with all extrapolative forecasts, it assumes that the future will be like the past. If medical science does invent a cure for cancer or a longevity pill, then obviously the Lee-Carter forecasts will be invalid. But what this model does show is that if life expectancy at birth is to approach 100, it will require a radical break with previous trends.

So that's the Lee-Carter model. It's a very good paper, something that any of us would have been proud to have written. I would encourage you to read it as an example of good academic writing. It may get you started with ideas about where you could contribute to the literature. Thank you.

MR. OSTASZEWSKI: I believe there is also a paper by Lee alone in *North American Actuarial Journal* in approximately 1992 where he discussed the details of the model. You can go to the SOA Web site and do a search for *North American Actuarial Journal* to find it. It actually is a very nice explanation of the model itself as well.

I'd like to make one comment about research. I frequently speak to students to tell them how great it is to be an actuary. There are close to 900 high schools in Illinois, and every one of them gets a letter from me offering to tell the students how good it is to be an actuary. And about three or four a year ask me to do so. The students in those high schools ask about the profession. I do tell them roughly what the salaries are, so the first question is not about the salary. The first question is something else. They ask what kind of car I drive. The second question is if it's so great to be an actuary how come you're a professor? It's a good question. My answer is that I worked as an actuary for a while and found myself writing research papers at night. So I decided to find a job where that would be the main part of my job. I like doing this, and it's not always possible to do research on exactly the things you want to study, sometimes of a theoretical nature, if you have a practical job. And that's related to what we're talking about.

Academics have jobs designed for doing things that are sometimes useless, or at least not immediately useful. Research very often is like that. The Lee-Carter model, for example, is very useful to the SSA, but how useful is it for practicing

actuaries and life insurance? Probably over time it will become useful, but I'm not sure if it's immediately applicable to you. The same goes for these mortality improvement select cohorts.

FROM THE FLOOR: Of all the actuarial science research papers that come out, how many of them come from actual employers who allow their employees time to write them, as opposed to coming out of universities with an actuarial science master's or doctorate program?

MR. OSTASZEWSKI: I never did a study, so I don't know. But I believe there has been some complaining that academics have too much influence at *North American Actuarial Journal*. If you look at the composition of the editorial body of *North American Actuarial Journal,* while they're not all academics, the academics are overrepresented versus the general population of actuaries. There's probably an unfair advantage to academics in terms of what they do on the job, so they probably publish far more than their fair share in *North American Actuarial Journal.* I'm sorry I don't know of any study though.

I would like to make one comment in relation to this. It's not always easy to define what actuarial science research is. The Lee-Carter paper was in the main dominant journal in statistics in the United States. They considered it to be research in statistics, but of course we think that it's actuarial science. There are large organizations of risk and insurance researchers: the American Risk and Insurance Association, Southern Risk and Insurance Association, Western Risk and Insurance Association. In fact, both American Risk and Insurance Association and Western Risk and Insurance Association have their own research journals where very few actuaries ever publish and probably are very rarely read by actuaries. *Journal of Risk and Insurance* is considered to be a very important and prestigious journal by insurance scholars, but is very rarely read by actuaries. *AER* has almost godlike status among economists, but is probably never read by actuaries. Did you have a comment, Tony?

MR. WEBB: Yes, I was going to say that there are a number of quasi-academic journals. So if you're asking what percentage of total publications comes from academics, it's obviously greater in the academic journals, but there are a number of respectable quasi-academic journals where practitioners certainly publish.

MR. OSTASZEWSKI: I understand that is, for example, the profile of the *Journal of Actuarial Practice*, that they want to emphasize practice, and there is a talk about creating a more practical journal in actuarial science. I'm a very big fan of writing case studies in actuarial science, but that's not done much. It's done in Course 7 right now. I consider Course 7 to be a fantastic educational experience for most candidates, but I'm biased because I'm one of the instructors. But I think that being put in a place where you have to solve a problem is something that employers look for. There's a bit of a conflict, because employers really look for solutions to business problems, not the underlying research. Very often, people who find an

ingenious solution would like to go back and present it as new research, but it's a slightly different objective.

MR. WEBB: Can I ask one question of the audience? There's been a lot of very good academic research that has made use of proprietary data sets. Are there any in the audience who have access to proprietary data sets, for example, on life expectancy or claims data or investment behavior, or things of that kind?

FROM THE FLOOR: Sure.

MR. WEBB: Can you tell us something about it without breaking confidences?

MR. OSTASZEWSKI: Probably not. This is a real issue for researchers. There are very significant limits on sharing the data. I always try to think whether it would be possible to perform some kind of transformation of the data or add white noise to it so that the data is not what it originally was, and then still be able to do research. Theoretically that's possible, but it's also partly a legal question.

MR. AL KLINE: To answer your question, we just did an older age mortality study, which we're selling. It's the most comprehensive older age mortality study that's been done. Another comment that I wanted to make is that the Society of Actuaries is about to put together a group. I don't know if it's filled yet or not, but for those who are interested, Ronora Stryker is the person to talk to at the SOA. They're going to do some research on the wearing off of the preferred criteria. It's going to be a futuristic study; they've hired a futurism group to help with it. It's going to be an interesting project.

MR. WEBB: I would like to point you in the direction of a good economist named Amy Finkelstein. She's published a lot of work using data on long-term care insurance claims and U.K. life annuity data. I would encourage you to read her research both because it's intrinsically interesting and also because it's an example of good research. It's available from the NBER Web site.

MR. OSTASZEWSKI: Any other questions? We welcome your comments or questions, and we are also available for any future questions. You can contact me if anything comes up. Let's go and do some research and some wonderful new creative things that will make the world better. Thank you very much.