Session 2B: Mortality Modeling I-Modified Lee-Carter Methods Q&A

Presenters: Louis G. Doray
Xiaoming Liu
Sharon S. Yang

Presented at the Living to 100 Symposium
Orlando, Fla.

January 5-7, 2011

Copyright 2011 by the Society of Actuaries.
All rights reserved by the Society of Actuaries. Permission is granted to make brief excerpts for a published review. Permission is also granted to make limited numbers of copies of items in this monograph for personal, internal, classroom or other instructional use, on condition that the foregoing copyright notice is used so as to give reasonable notice of the Society's copyright. This consent for free limited copying without prior consent of the Society does not extend to making copies for general distribution, for advertising or promotional purposes, for inclusion in new collective works or for resale.

Paul Sweeting: I had a few questions for Xiaoming on the approaches used. First I wondered what criteria had been used to fit the Lee-Carter model, whether it was singular value decomposition or maximum likelihood method because that would have an impact on the standard deviation of k_t, which could tell you about the volatility of your outcomes.

I also wondered if there had been any investigation of whether there was a trend or difference stationarity in k_t. It looked like random drift was being assumed, and in Sharon's presentation as well. I wondered if there was any test used to see whether that was a correct assumption, particularly when looking at things like structural breaks in k_t. Rather than having a random walk with drift, it could be that you have a fixed level of k_t with the levels of trending and drift changing from time to time, and you just have volatility around those instead. So I wondered whether there had been any analysis done on that at all.

I just wanted to clarify something on the drift uncertainty. My understanding from what you're saying is that you're essentially assuming that the level of drift is drawn from a random distribution around the central value, so you're seeing a level of drift that has a distribution and then you just drew out a value from that. But given some of the comments that were made about the instability in the level of k_t, I was wondering whether you had used another approach of actually changing the mean value that you're taking your value of k_t from. So rather than just drawing from the same distribution, actually change the distribution that you take it from if it does change over time.

Xiaoming Liu: Thank you for your questions. So your first question is about clarifying the method of estimation, right? We use maximum likelihood methods based on the presumed assumption for the death number so I guess that's similar to Sharon's approach. For your second question, whether we have tried different methods for k_t, no, we haven't. We just assumed it's universally the random walk with drift model. Actually, that's one of the reasons we came up

with this project, because when you use different period of data to fit the model, you may get a different shape for k_t so it may suggest a different model, but then that's the model risk in predicting the future (mortality rates). One of the reasons is to test the simplest model and to check if it can do a good job.

The answer to your third question is that I use the median. You suggest if we use the mean of the distribution, right, is that your question?

Paul Sweeting: I was really looking at the way in which you were dealing with the drift uncertainty. My understanding is the way you'd be doing it, you'd say, "What is the distribution for the level of drift?" And you just randomly take a level out for the uncertainty rather than saying, "It's a different distribution, a different period of time."

Xiaoming Liu: Yes, thank you. That's a lot of detail on how we get the drift uncertainty, I didn't have much time to go into the technical details and I hope if you read the paper you can get an idea on how we derive that, but the approach is, it's distribution-free. We didn't make any assumption on what the distribution is about the drift.

The other papers, the paper by Cairns and Dowd and Blake, made a distribution assumption about the drift. We didn't use that approach. We used the bootstrap method, so it's distribution-free.

Thank you for all your questions and also I want to thank Ward for his discussion. There are lots of topics covered by the paper, but I didn't have time to discuss them in the presentation. I concentrated on the methodology and ideas but the details are in the paper. It's a quite long paper.

Will Mitchell: Question or maybe just a comment for Sharon, and that is on pricing the life annuity using the coherent model. It sounds like you're using population mortality data. So somebody doing pricing, I would think, of an individual product might want to use annuity mortality experience. I don't know if you would have any comments about what kind of industry mortality data you could get from the various countries in your study: Japan, Taiwan, Canada and the United States.

Sharon Yang: Thank you for this comment. Actually, if I can gather the data from the insurance company it would be better. However, I don't have a chance to gather this kind of data so that's why I use the population data. Everyone here, if you are able to provide me the data, I will be happy to do these kinds of analyses. Thank you.

Jean-Claude Menard: I was surprised by your comment, Mr. Kingkade, about the usefulness of the HMD. For the first time we have done the actuarial report on the Canada Pension Plan(CPP) based on the human mortality database instead of the data from Statistics Canada, and there are two reasons for that. First, they developed a more recent mortality table than the one available from Statistics Canada. After we changed the methodology we have compared the results, which are the annual improvement rates calculated using the HMD for Canada and the Statistics Canada central death rates. We obtained the same results; so we don't have any evidence that the HMD is not correct, at least for Canada.

Ward Kingkade: I wasn't suggesting necessarily it's not correct. I was repeating that apparently it's been smoothed. Now, did the data look smooth to you, or did they have the kinds of kinks you would expect at late age?

Jean-Claude Menard: It depends what you mean by late age, but let's say that their mortality table stops at age 110 and that our model is going to age 120. So we changed slightly the mortality rates from age 105 to 120. But again, at these ages, the difference for us does not make a big deal of a difference because there are too few people still alive or even in the future

still alive. So we didn't see any problem. Another adjustment has been done in the report. We looked at the raw data of the beneficiaries database and we further adjusted the mortality rates of the general population.

Ward Kingkade: It's also not really that much of a problem if you're applying the Lee-Carter method to smooth the data; in fact, you may be getting a better projection if you have graduated data because you're not incorporating a stochastic disturbance that was there in your base here.

Jack Yue: Just two comments. The first one is if I remember this correctly, Dr. Wilmoth from Berkeley did some research on the estimation of Lee-Carter method. He found that the MLE (maximum likelihood estimation) and WLS (weighted least squares) can improve the estimates from the approximation method, but the estimates of these two methods are pretty similar to those derived from SVD (singular value decomposition). So, I don't think that's a problem if you use MLE. And also here, if I remember this correctly, MLE is using the iteration and gradually you will come up to the similarity as CBD (Cairns-Blake-Dowd) stochastic mortality model, so I guess that's not a problem. The other comment is about the HMD. I am also surprised to hear that the data from HMD has been graduated. I know the Taiwan data in the HMD probably is wrong. I definitely can say that because if you want to use Taiwan data you can go to the Web page. The Taiwan government has made the Taiwan data available ever since 1990. I don't know who uploaded Taiwan data into HMD, but I'm sure it's not correct. At least, before 1990, Taiwan didn't have the data for age 85 or over, so I'm curious who did the data uploading. I guess I need to go back and double check the data from the United States and Canada and see if the data has been graduated. Thank you.

Joseph Lu: My question has to do with the coherence model that you have presented. I think one of the potential uses of the Lee-Carter model is to generate best estimates. It also comes with some fan charts that indicates uncertainty about the future. So it can also be used for reserving, for example. In Europe we have this new Solvency II regulation that requires us to

hold reserves for one-in-200-years event over one year. It is hard enough to come out with a good best estimate let alone working out this one-in-200-years event. So a model like the one you have presented will allow us to pick a scenario with its attaching probability and say this is my one-in-200-years event. What is interesting here is that we have appeared to find a way to narrow this confidence interval, by merging data of different countries. This is good news for us because that would mean that the uncertainty has reduced and we can hold less capital for longevity risk. If I were to go to the regulator and say that, look, last year I told you if I used this model my reserve is 5 percent above the best estimate. Now I have found a new way to reduce the reserve by narrowing the confidence interval through merging my data with that in the United States and Canada. I'm just wondering whether I can get away with that. Do you have any comments on the way that I may potentially use or misuse your model? Thank you.

Sharon Yang: It's a tough question, but in terms of pricing, that's the idea of pricing the life of annuity. But I haven't thought about it for the reserving, if we reduce the uncertainty. So maybe we can discuss this one later.

Xiaoming Liu: I would like to make a comment to your question, although it is not addressed to me. My study with real data suggests that the confidence band used in practice is already narrower than it should be. Normally I find if you introduce more parameters, more factors, it can result in better-fitting results. It can result in a narrower prediction interval, but that doesn't mean the uncertainty in the future has been reduced. We don't know yet, but the observation is that the prediction interval used in the past is too narrow. It has been discussed many times. Now, you can have a better-fitted model and have a narrower prediction interval, so there will be more questions debating about that, whether this is right for risk management and reserving.