Session 4A: New Models of Advanced Age Mortality

Discussant: Tom Edwalds, FSA, ACAS, MAAA
Munich American Reassurance Co.

Presented at the Living to 100 and Beyond Symposium
Orlando, Fla.
January 7-9, 2008

Presented Papers:

Inference for Logistic-type Models for the Force of Mortality
Louis G. Doray, PhD, ASA

Predictive Modeling for Advanced Age Mortality
Lijia Guo, PhD, ASA, MAAA

The two papers presented in this session both extend our knowledge about modeling advanced age mortality, but do so in very different ways. Therefore, this discussion will address each paper separately.

Predictive Modeling

In her paper, "Predictive Modeling for Advanced Age Mortality," Lijia Guo introduces predictive modeling to the study of advanced age mortality. Predictive modeling is the term for a suite of statistical techniques used to extract useful information from large, complex databases, often including loosely linked data from disparate sources. These techniques have been used with great success in health insurance and in property and casualty insurance. Their application to the study of mortality holds great promise, but is relatively new. One predictive modeling technique, projection pursuit regression (PPR), was used recently by the Valuation Basic Table Subteam of the Individual Life Experience Studies Committee in its analysis of the data underlying the Society of Actuaries' 2002-2004 Individual Life Experience Report.

In the draft paper presented at the Living to 100 Symposium, Ms. Guo described two predictive modeling techniques, decision trees and generalized linear models (GLM), and

illustrated their application to mortality analysis. While her exposition of the techniques was good, there were serious flaws related to the use of the data in her analysis that invalidated her conclusions. The linkage of the underlying data sets and the interpretation of the data elements in these data sets were both flawed, and these led to findings that were misleading at best. These flaws have been remedied in her final paper, but it is useful to explore them in greater detail in this discussion.

Ms. Guo used two data sets in her analysis in the draft paper: the data collected by the Society of Actuaries for the development of the RP-2000 mortality tables; and the Surveillance, Epidemiology and End Results (SEER) data set from the National Center for Health Statistics (NCHS) of the Centers for Disease Control and Prevention (CDC). The RP-2000 data is mortality experience data from 1990–1994 collected from self-administered pension plans. Plans contributing data were classified by union status (union, non-union or mixed), pay type (hourly, salaried or mixed) and industry. Each plan submitted data grouped into cells with exposures and deaths by age group, gender and participant status (active employees, retirees, beneficiaries or those receiving disability pensions), and most plans further split these cells by annuity amount group for those receiving pensions (less than $6,000 per year, more than $14,400 per year or in between). The SEER data set contains cancer treatment data. It has individual records of patients diagnosed with cancer, including date of diagnosis, age at diagnosis, gender, current status (alive or dead) and, if dead, the cause of death and the survival time from diagnosis to death.

Ms. Guo did not explain her method for linking these data sets, and there does not appear to be a reasonable basis for linking them. Gender and age are the only two variables these data sets have in common, and the age variable is defined differently in the two data sets. It seems obvious that, at any given age and gender, the mortality of persons known to have cancer will be greater than the mortality of any group of pension plan participants. The RP-2000 data set has no information about the incidence of cancer among covered plan participants. The SEER data set has no information about the employment status of the patient. There is no basis for determining or estimating the extent to which mortality differences in the RP-2000 data might be related to differences in the incidence of cancer by subgroup. There is also no basis for determining or estimating the extent to which differences in survival times of cancer patients recorded in the

SEER data set might be related to differences in employment status. Although she did not state her assumption, it appeared that Ms. Guo assumed that the drivers that influence mortality in the RP-2000 data set and the drivers that influence mortality in the SEER data set will operate independently and consistently throughout both data sets. This assumption does not seem reasonable to this reviewer. In her final paper, Ms. Guo does not use the SEER data set at all, and focuses her analysis on the RP-2000 data set without attempting to link to another data source.

The claims Ms. Guo made in her draft paper about the information contained in the data were also inaccurate. In her construction of a "Senior Mortality Risk Score," she claimed that the RP-2000 and SEER data sets provide a differentiation by health status: those with cancer, those with some other non-cancer disease and those who are healthy. The fact is that neither of these data sets provides such information. All of the records in the SEER data set are of people who have cancer. It is inaccurate to claim that the SEER data sheds any light on the mortality of those who are healthy or of those who have a disease other than cancer. The only distinction in the SEER data set is between those cancer patients who are still alive, those cancer patients who died of cancer and those cancer patients who died of some other cause. The RP-2000 data set has precious little information about health status and none about health status after retirement. The distinction between "disabled" and "retired" in the RP-2000 data refers only to the health status of the participant at the time plan benefits were first collected. Mortality differences between these groups disappear by about age 85, about 20 years after pension benefits were first collected. In her final paper, the Senior Mortality Risk Score no longer includes a component for health status.

While the above-mentioned flaws in the draft paper have been remedied in the final paper, a few still remain. The Senior Mortality Risk Score presented in the paper includes a factor for occupation, split into professional, non-professional and high risk. As mentioned above, the SEER data set has no information about employment status or history. The RP-2000 data set defined plans covering predominantly salaried, non-union employees as "white collar," while union plans or plans covering predominantly hourly employees were defined as "blue collar." Conceivably this might be used as a proxy for professional versus non-professional status, but the identification of the "high risk" occupation category is completely mystifying.

The paper also claims to investigate the link between wealth and senior mortality risk, and uses the terms “financial wealth” and “income level” interchangeably. The only available data element related to wealth is the RP-2000 annuity amount group, which is a rough indication of the amount of the employer-provided pension benefit received from one employer. There is no information about other sources of income, and no information about accumulated assets. This data element is a crude proxy for overall wealth, and is probably a poorer proxy for female wealth than for male wealth for the RP-2000 retiree data, since all of the retirees included in this data were born before 1930. A woman who entered the labor force relatively late in life would have earned a small pension of her own, but might have substantial financial resources due to her spouse’s earnings over his career. Therefore, the lack of a correlation between annuity amount group and mortality for females between the ages of 70 and 85 is not adequate evidence for Ms. Guo’s conclusion that financial wealth does not affect female mortality at these ages.

Ms. Guo’s claim that the mortality risk score presented in the paper provides more than a 4,000 percent differentiation in mortality is both flawed and inflated. Most of the difference that can actually be documented is due to age, gender, and participant status (disability pensioner vs. normal retiree). These differences are already included in the RP-2000 mortality tables and should not be claimed as a new benefit of the mortality risk score. As mentioned above, the differences due to health status and occupation are not substantiated by the data Ms. Guo claims to have used. This leaves the differences measured by annuity size group and union status as valid benefits of the mortality risk score. These factors provide about a 160 percent differentiation in mortality.

While the mortality risk score presented in the paper is not as robust as claimed, the exposition of predictive modeling tools is valuable. The illustration of the application of GLM to mortality modeling is also useful, as are the senior mortality risk score factors for the type of pension plan (union, non-union or mixed) and annuity amount group. Predictive modeling techniques are already essential tools for actuaries studying health or property and casualty risks. The advantages of the use of these tools to improve both underwriting and pricing are now becoming evident to actuaries studying mortality and longevity risks as well. This paper is a

good first step in the application of predictive modeling tools to the analysis of mortality at advanced ages, and Ms. Guo is to be congratulated for this contribution to the actuarial literature.

Logistic-Type Models

Although the title of this session is "New Models of Advanced Age Mortality," the logistic-type models addressed by Louis Doray in his paper "Inference for Logistic-Type Models for the Force of Mortality" are firmly established. The advantage of these models over the more traditional Makeham model is that the force of mortality is bounded, leading to a maximum value of q_x at advanced ages that is less than certain. This is consistent with the perception that mortality levels off at very high ages.

Even though the models aren't new, Prof. Doray's method of calculating the parameters is new. His exposition is clear, and it is easy to understand how to apply his method in an Excel spreadsheet. He presents in detail two estimators for logit (μ_x): ordinary least squares (OLS) and weighted least squares (WLS). He represents the OLS estimator as $\bar{\theta}$ and the WLS estimator as θ^*. He also briefly mentions a third estimator he represents as $\tilde{\theta}$, which is a modification of the WLS estimator. This paper makes a significant contribution to the actuarial literature with the development of the formulas for the OLS and WLS estimators and the formulas for calculating the variances of the estimators for p_x.

The biggest drawback of the WLS estimator Prof. Doray presents is that the calculation of the covariance matrix is cumbersome. The modified WLS estimator, $\tilde{\theta}$, obviates this problem by ignoring the covariances, which is equivalent to assuming that the estimators of p_x at different ages are independent. This reviewer contends that this assumption of independence is accurate for period data, and that the need for the covariance matrix is a consequence of using cohort data for the analysis. While it is clearly advantageous to use cohort data to estimate survival curves, assuming independence of the p_x estimators greatly simplifies the calculation of the WLS estimator for logit (μ_x).

The Canadian mortality data for the cohort born in 1888-1892 that Prof. Doray uses to fit the Kannisto logistic model is well-behaved, and the parameters he calculates are substantially the same for all four of the estimators he calculates (OLS, WLS, modified WLS and maximum likelihood). He concludes that the OLS estimator is optimal because it is the simplest to calculate. While this conclusion is accurate for the well-behaved Canadian data, this reviewer tested these estimators on a quirkier data set, the data collected by the Society of Actuaries Individual Life Experience Studies Committee for 2002-2004 (ILEC 2002-04). Thirty-five U.S. life insurance companies contributed over 75 million policy years of experience to this study, including nearly 700,000 deaths. This data set included nearly half a million life years of exposure and over 70,000 deaths at attained ages 90 and up. Table 1 lists the exposures and deaths at these ages in the ILEC 2002-04 study by gender, along with the raw mortality rates and the standard deviations of these mortality rates, calculated using Prof. Doray's method.

TABLE 1

ILEC 2002-04 High Age Mortality Data

Attained	Males				Females			
Age	**Exposure**	**Deaths**	**Raw qx**	**St Dev**	**Exposure**	**Deaths**	**Raw qx**	**St Dev**
90	65,112	9,399	0.14435	0.00138	59,503	6,772	0.11381	0.00130
91	51,301	7,940	0.15477	0.00160	46,969	5,796	0.12340	0.00152
92	39,838	6,544	0.16427	0.00186	36,620	5,014	0.13692	0.00180
93	31,017	5,190	0.16733	0.00212	27,943	4,135	0.14798	0.00212
94	24,021	4,223	0.17580	0.00246	21,110	3,268	0.15481	0.00249
95	18,022	3,313	0.18383	0.00289	15,268	2,531	0.16577	0.00301
96	11,412	1,785	0.15642	0.00340	10,194	1,674	0.16421	0.00367
97	4,608	1,094	0.23743	0.00627	6,084	1,103	0.18129	0.00494
98	2,938	695	0.23658	0.00784	4,088	783	0.19153	0.00615
99	1,981	420	0.21205	0.00918	2,701	493	0.18256	0.00743
100	1,160	265	0.22842	0.01233	1,759	280	0.15918	0.00872
101	527	145	0.27509	0.01945	897	153	0.17050	0.01255
102	244	62	0.25399	0.02786	415	89	0.21466	0.02016
103	133	34	0.25640	0.03792	231	53	0.22903	0.02762
104	89	30	0.33538	0.04992	131	23	0.17588	0.03329
105	63	15	0.23879	0.05379	84	16	0.18957	0.04266
106	45	6	0.13365	0.05078	57	7	0.12383	0.04381
107	28	2	0.07113	0.04847	34	4	0.11934	0.05600
108	25	0	0.00000	0.00000	17	4	0.24031	0.10473
109	18	1	0.05521	0.05366	11	1	0.09113	0.08688
110	17	2	0.11991	0.07954	15	3	0.20197	0.10417
111	14	1	0.07113	0.06855	5	1	0.20008	0.17895
112	10	1	0.10015	0.09500	2	0	0.00000	0.00000
113	9	0	0.00000	0.00000	3	0	0.00000	0.00000
114	10	1	0.10004	0.09490	2	0	0.00000	0.00000
115	4	0	0.00000	0.00000	0	0	0.00000	0.00000
Total	252,645	41,168			234,141	32,203		

Figure 1 illustrates two Kannisto logistic curves fit to the ILEC 2002-04 raw mortality rates for females at attained ages 90 to 103 using alternative methods described by Prof. Doray:

one using the ordinary least squares (OLS) method, and the other using the modified weighted least squares (WLS) method. The solid line in the graph represents the raw mortality rates, and the dashed lines represent the 95 percent confidence interval about the raw mortality rates as estimators of the underlying mortality rates. The line with the square markers is the Kannisto logistic curve fit to this data using the OLS method, and the line with the triangle markers is the curve fit to this data using the modified WLS method.

Figure 2 illustrates three Kannisto logistic curves fit to the ILEC 2002-04 raw mortality rates for males at attained ages 90 to 104: one using the OLS method, another using the modified WLS method, and the third using the modified WLS method, but excluding the outlier point for males age 96. The raw mortality rate for males age 96 in this data is substantially lower than the raw mortality rates at all other ages from 92 to 105, and yet it is supported by a large volume of data (1,785 deaths). The raw mortality rates at ages 97 and 98 could arguably be considered outliers also, as they are higher than the raw mortality rates for ages 99 and 100 and are remarkably higher than the raw mortality rates at ages 96 and under. For the purpose of this illustration, however, only the data point for age 96 was removed from the calculations for the "WLS (x-outlier)" fit for the Kannisto logistic curve.

As in Figure 1, the solid line in Figure 2 represents the raw mortality rates, and the dashed lines represent the 95 percent confidence interval about the raw mortality rates as estimators of the underlying mortality rates. The line with the square markers is the Kannisto logistic curve fit to this data using the OLS method, and the line with the triangle markers is the curve fit to this data using the modified WLS method. The broken line with open diamonds is the curve fit to the data excluding the point for age 96, using the modified WLS method, i.e., the WLS (x-outlier) curve.

FIGURE 1

ILEC 2002-04 Raw Mortality Rates

With Fitted Kannisto Logistic Curves

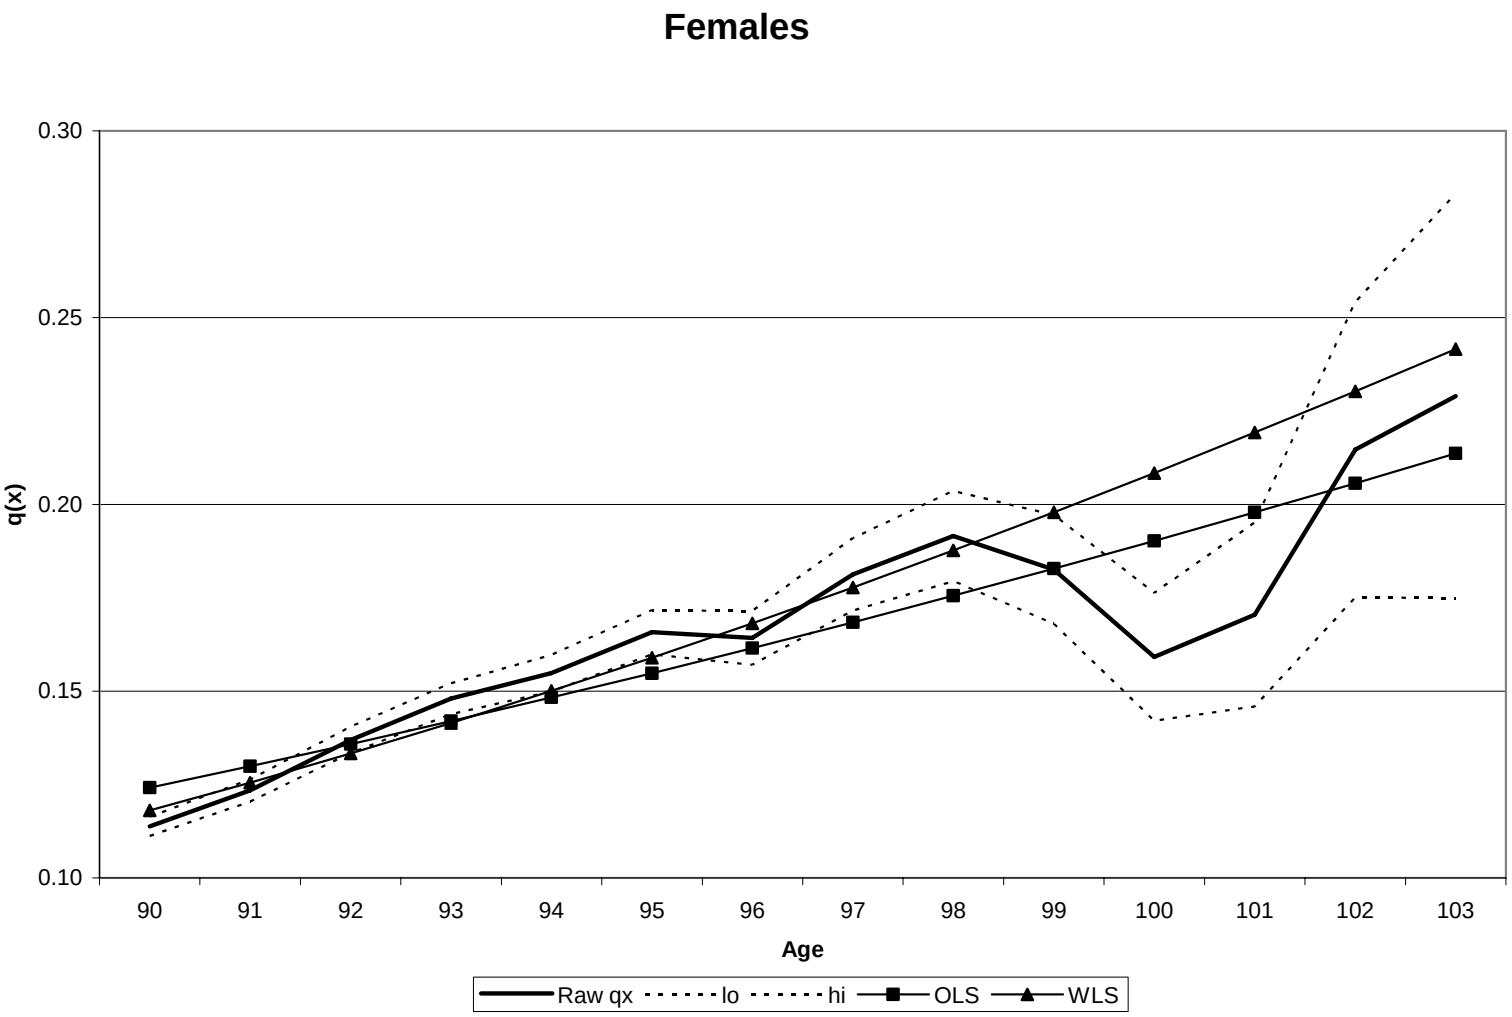

FIGURE 2

ILEC 2002-04 Raw Mortality Rates

With Fitted Kannisto Logistic Curves

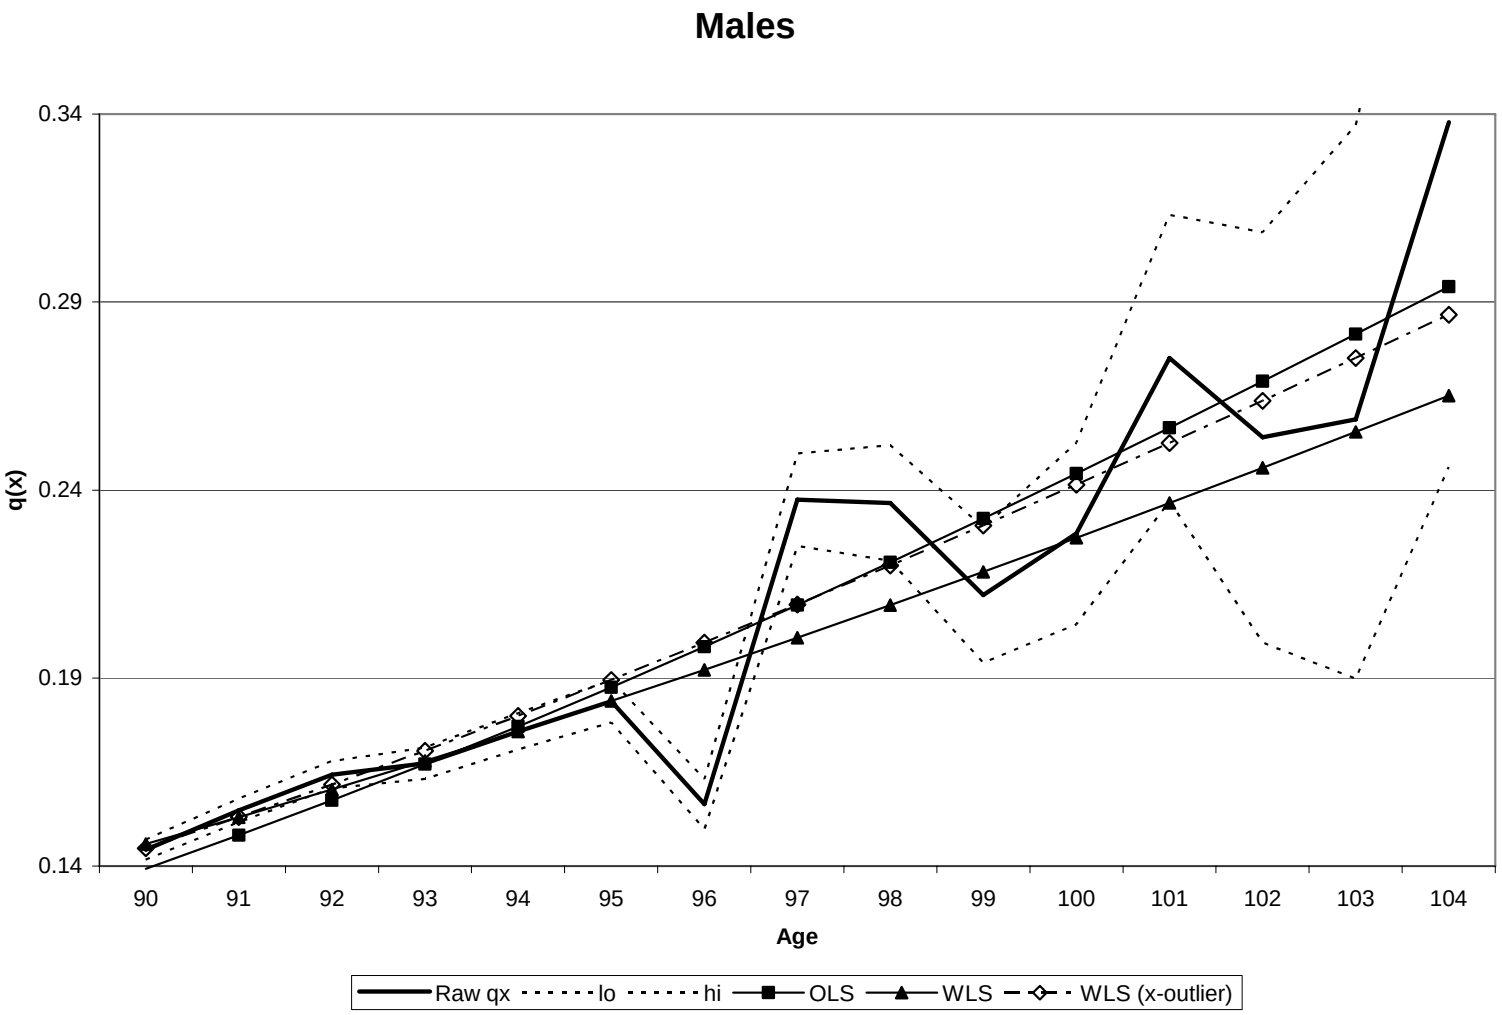

In both figures, the broadening of the confidence intervals about the raw mortality rates as age increases (and exposure decreases) is clearly illustrated. It is also apparent that the OLS and modified WLS estimates of the parameters are substantially different. Less obvious is the fact that the modified WLS curves are within the 95 percent confidence intervals more often. For females, the modified WLS curve is within the 95 percent confidence interval for seven of the 14 ages, while the OLS curve is within the 95 percent confidence interval for only five ages. For males, the modified WLS curve is within the 95 percent confidence interval for 10 of the 15 ages, while the OLS curve is within the 95 percent confidence interval for only five ages. While none of these fits are particularly good, the modified WLS method fits the data better for both males and females. Figure 2 also illustrates how sensitive the results are to outlier points. Excluding just one point from the calculations dramatically changes the estimated parameters of the logistic curve. The WLS (x-outlier) curve is very close to the OLS curve, while the modified WLS curve is significantly flatter. Interestingly, the WLS (x-outlier) curve fits the data as well as the modified WLS curve, based on being within the 95 percent confidence interval for 10 of these 15 ages. Based on these observations, this reviewer believes that the modified WLS method may be the best choice for fitting the Kannisto logistic curve to advanced age mortality data. Caution and judgment must still be used, however, since the results are sensitive to outliers.

Prof. Doray's paper provides some good tools for fitting logistic-type curves to advanced age mortality data. It is well-written, and his methods are easily applied. This paper is an excellent addition to the actuarial literature, and Prof. Doray is to be congratulated for it.