The Lee-Carter Model for Forecasting Mortality Revisited

Siu-Hang Li* and Wai-Sum Chan†

Presented at The Living to 100 and Beyond Symposium
Sponsored by the Society of Actuaries

Orlando, Fla.

January 12-14, 2005

* Siu-Hang Li is a Graduate Student in the Department of Statistics and Actuarial Science at the University of Waterloo, Ontario, Canada, e-mail: shli@math.uwaterloo.ca.
† Wai-Sum Chan is an Associate Professor in the Department of Statistics and Actuarial Science at the University of Hong Kong, Hong Kong, PR China, e-mail: chanws@hku.hk.

Abstract

Interrupting phenomena are commonly encountered in time-series data analysis with the study of mortality trends being no exception. Nevertheless, previous demographic forecasts have paid little attention to the existence of such phenomena. In this study we use mortality data from the United States and Canada to perform time-series outlier analysis on the key component of the Lee-Carter model: the mortality index. We begin by employing a systematic outlier detection process to ascertain the timing, magnitude, and persistence of any outliers present in historical trends of the mortality index. We then try to match the identified outliers with important events that could possibly justify the vacillations in human mortality levels. At the same time, we adjust the effect of the outliers for model reestimation. The empirical results indicate that the outlier-adjusted model could achieve better fits and more efficient forecasts of variables such as the central rates of death and the life expectancies at birth. Finally, we conclude our study with possible extensions on the valuations of life annuities and the probabilistic distribution of the highest attainable age, incorporating the effect of mortality improvement portrayed by the revised model.

1. Introduction

The Lee-Carter model (Lee and Carter 1992a) is the current gold standard of mortality trend fitting and projection. Over the past 10 years, the model and its variants have been used by actuaries for a wide range of purposes, from the forecasting of mortality reduction factors (Renshaw and Haberman 2003) to the assessment of retirement income adequacy (Chia and Tsui 2003). Other applications in demographic science include population projections (Booth and Tickle 2003), the forecasting of sex differentials in mortality (Lee and Carter 1992b), and the projection of mortality patterns for the “oldest-old” (Buettner 2002). Intrinsically, the model assumes that the dynamics of death rates over time are driven by a single time-varying parameter, namely, the mortality index. The mortality forecast relies on the extrapolation of this index under an appropriate statistical linear time-series model.

With no exception to the mortality index, longitudinal data are often contaminated with various forms of discrepant observations, which, in the statistical literature, are commonly referred to as *outliers*. Outliers have a variety of sources: they may arise from mere recording or typographical errors, or from nonrepetitive exogenous interventions, such as pandemics or hostilities in a mortality series. Thus, the analysis of outliers may reveal invaluable information about the external shocks that affect the series, which may then enable actuaries to predict how the series will respond if these or similar interruptive events recur. When adequate data are available, outlier analysis may possibly shed some light on recent discussions on the impact of the SARS crisis (see Panjer 2003) and the effect of terrorist attacks (see Rowley and Bishop 2001; Wolak et al. 2003) on future mortality assumptions.

In addition, outlier detection and adjustment are critical to both model estimation and forecasting. Tsay (1986) noted that the existence of outliers might cause serious biases in both the sample autocorrelation function and partial autocorrelation functions, thus resulting in erroneous model specification. Hillmer (1984) and Ledolter (1989) found that outliers could have a disastrous effect when they occur near the forecast origin, due partly to the bias in model parameter estimation and partly to the carryover effect of the outlier on the forecast. Li and Chan (2005) further demonstrated that this carryover effect is exceptionally severe in the original Lee-Carter model (see Lee and Carter 1992a; Tuljapurkar et al. 2000), in which the mortality index is modeled by a random walk with drift.

Yet the existence of outliers often has been overlooked in extrapolative mortality forecasts. Even when the presence of outliers was acknowledged, the identification process was exclusively based on prior knowledge, rather than a sound statistical detection method. This is exemplified by Lee and Carter (1992a), who viewed subjectively the known influenza epidemic in 1918 as an anomaly and dealt with it by means of an intervention model with a dummy variable. Although not commonly used in mortality studies, outlier analysis is widely accepted by actuaries in stochastic investment modeling (see, e.g., Chan and Wang 1998; Chan 1998). It has also been applied in fields as diverse as macroeconomic analysis (Balke and Fomby 1994), data mining (Liu and Hudak 2001), and clinical studies (Thomas et al. 1992).

In this study we perform a systematic time-series outlier analysis of the mortality index encompassed in the Lee-Carter model, using mortality data of both the American and the Canadian populations. We begin our analysis by identifying exogenous events that might significantly affect the mortality dynamics. We do this via Chen and Liu's (1993) method to search iteratively for outliers present on the mortality index, and then we match the identified outliers with events that may have caused them. Noting that the masking effect of outliers might lead to biases in parameter estimates (Tsay 1986) and even erroneous model identification (Chan 1995), we also implement the additional iterative cycle proposed by Li and Chan (2005), aiming at unveiling the true model underlying the outlier-free series of the mortality index. Along the way, we scrutinize the efficiency of the outlier-adjusted model by comparing the interval forecasts of the mortality index and related variables such as the central rates of death, the complete life expectancies at birth, and the actuarial present values of a life annuity at pensionable age.

So far actuaries and demographers have concentrated on the estimation of population aggregates, such as life expectancies at birth or at some other ages, but have paid far less attention to extreme values. Nonetheless, the rapid emergence of centenarians and supercentenarians has highlighted the importance of the "tail" in different types of life tables (Fishman 2002) and has motivated some actuaries to look for appropriate ways for closing off the life tables instead of the prevailing practice of using the value of 1 in an arbitrarily chosen limiting age (Johansen 2002). Thatcher (1999) made the first attempt to derive the probabilistic distribution of the highest attainable age, actuarially denoted as ω, assuming that the age pattern of mortality follows a logistic distribution. Thatcher's study, however, focused on period mortality only and made no effort to consider the effect of changing mortality levels on the maximum life span. Following in Thatcher's footsteps, we attempt to compute the probabilistic distribution of ω for various cohorts of Canadians, incorporating the effect of mortality improvement portrayed by the outlier-adjusted Lee-Carter model.

2. The Data

To fit the Lee-Carter model and to perform the outlier analysis, we require both the central rates of death and the exposures to risk. Below we list, for each population to be investigated, the sources of data, sample period, and modifications made (if any).

The United States

The matrix of age-specific central rates of death from 1900 to 2000 is available from the National Center for Health Statistics (NCHS 2004a, 2004b). Unfortunately the mortality data are presented in an abridged form: that is, values of death rates are shown at age 0, age group 1–4, decadal age groups 5–14, 15–24, and so on up to 75–84, and the open age group 85 and over. Such a layout is not sufficient for the computation of various monetary functions involving life contingencies, which requires probabilities of death for every single year of age. To overcome this problem, we apply the disaggregation method proposed by Pollard (1988) to derive full life tables from the mortality data tabulated in 10-year age groups, assuming that the force of mortality (μ_x) varies in an exponential manner and that the population is stable within each of the age intervals.

Under the assumption of uniform distribution of death within each year of age, the exposure to risk is approximated by the midyear population estimate, which is available form the U.S. Census Bureau (2004). From 1990 to 2000 the estimates can be retrieved online, and from 1900 to 1989 the estimates are obtained by written request.

Canada

The required data are not available from a single source. From 1921 to 1997 the death count and the exposure to risk (midyear population estimate) for every single year of age up to 90 are obtained from the Human Mortality Database (HMD 2004). From 1992 to 2000 the death count and the exposure to risk are obtained, respectively, from Statistics Canada by written request and from CANSIM (an online socioeconomic database provided by Statistics Canada; see Statistics Canada 2004). By comparing the overlapping portion (1992–1997), we confirmed that the two data series commensurate with each other.

At very high ages the acquired central rate of death defined by the ratio of death counts to the exposure to risk may not be trustworthy, due partly to the inaccuracy of reported age at death and partly to the sampling error of estimated death rates when numbers are small. The significance of these problems in the United States and Canada is highlighted by Tuljapurkar and Boe (1998) and Boureau and Desjardins (2002). To have a satisfactory termination of the life tables, we apply the method suggested in Coale and Guo (1989, pp. 614–615) to extend the central rates of death from age 85 to 109. The original rates beyond 85 are discarded.

3. The Lee-Carter Model

The Lee-Carter model essentially describes the logarithmically transformed age-specific central rate of death ($m_{x,t}$) as a sum of an age-specific component that is independent of time (a_x), and the product of a time-varying parameter (k_t, also known as the mortality index) that summarizes the general level of mortality and an additional age-specific component (b_x) that represents how rapidly or slowly mortality at each age varies when the mortality index changes. Mathematically,

$$\log(m_{x,t}) = a_x + b_x k_t + \varepsilon_{x,t}\,. \qquad (1)$$

The final term, $\varepsilon_{x,t}$, is the error term, which reflects the age-specific influences not captured by the model. Mortality forecasting is carried out using the model of the mortality index time series, on which the outlier analysis in this study is performed.

The equation underpinning the Lee-Carter model is known to be overparameterized. To stipulate a unique solution, we take a_x as the arithmetic mean of the $\log(m_{x,t})$ over time, while the sums of b_x and k_t are normalized to unity and zero, respectively. As all parameters on the right-hand side of equation (1) are unobservable, fitting the model by the ordinary least squares method is impossible. To overcome the situation, we employ Lee and Carter's (1992a) two-stage estimation procedure, which gives exact solutions. In the first stage, singular value decomposition (SVD) is applied to the matrix of $\{\log(m_{x,t}) - a_x\}$ to obtain estimates of b_x and k_t. The SVD procedure can be implemented by using various standard mathematical/statistical packages such as GENSTAT (Payne et al. 1993), MATLAB (Brose 1997), and IMSL MATH/ LIBRARY (IMSL 1997). In the second stage, the time series of k_t is reestimated by solving for k_t such that

$$D_t = \sum_x \{\exp(a_x + b_x k_t) N_{x,t}\}, \qquad (2)$$

where D_t is the total number of deaths in time t, and $N_{x,t}$ is the exposure to risk of age x in time t. This is to ensure that the mortality schedules fitted over the sample years will reconcile the total number of deaths and the population age distributions.

An autoregressive integrated moving-average (ARIMA) model is then used to model the dynamics of k_t. The orthodox Box and Jenkins (1976) approach often is employed to obtain a fitted ARIMA model from the empirical k_t data.

In addition to the two-stage estimation procedure, many alternative methods have been proposed in the actuarial literature. For example, Wilmoth (1993) and Brouhns et al. (2002) considered a regression-type model under the Poisson assumption on the number of deaths (D_t). In this way standard maximum likelihood estimates for the a_x's, b_x's, and k_t's can be obtained. This parametric approach has excellent statistical properties if the assumed Poisson density is the correct one, but could lead to grossly incorrect inferences if the D_t (or k_t) series is contaminated with outliers. To alleviate the situation, the penalized likelihood approach may be used. A penalty function $P(k_t)$ is subtracted from the original log-likelihood function, where $P(k_t) \geq 0$ is a roughness penalty that decreases as k_t becomes smoother. Maximum penalized likelihood estimates can be computed by maximizing the modified log-likelihood function. This method has the advantage of smoothing the estimated k_t's and allowing for an extrapolation directly from the penalty structure. However, it requires a subjective choice of $P(k_t)$, and different choices of the penalty function will yield different estimators. Furthermore, it may smooth out useful information concerning occasional disturbances to the general mortality evolution pattern (i.e., outliers in k_t). Therefore, we shall concentrate only on the original fitting methodology in this paper.

Figures 5 and 6, respectively, plot the fitted values of k_t based on the American and Canadian mortality (they also plot forecasts, which should be ignored for now). The trends are roughly linear despite the existence of several aberrant observations, which will be discussed in later sections. For completeness, corresponding fitted values of a_x and b_x for both countries are shown in Table 1.

TABLE 1
Fitted Values of a_x and b_x at Selected Ages

Age x	United States a_x	United States b_x	Canada a_x	Canada b_x
0	−3.3629	0.0176	−3.6219	0.0226
5	−6.4513	0.0224	−7.1893	0.0212
10	−7.1535	0.0183	−7.5977	0.0189
15	−6.7123	0.0153	−7.1418	0.0141
20	−6.2711	0.0124	−6.5359	0.0111
25	−6.1198	0.0127	−6.5745	0.0128
30	−5.9686	0.0131	−6.4752	0.0131
35	−5.7186	0.0119	−6.2481	0.0123
40	−5.4685	0.0107	−5.9261	0.0108
45	−5.1183	0.0091	−5.5443	0.0091
50	−4.7680	0.0075	−5.0929	0.0076
55	−4.3822	0.0066	−4.6864	0.0059
60	−3.9963	0.0057	−4.2096	0.0059
65	−3.6006	0.0054	−3.7491	0.0058
70	−3.2048	0.0052	−3.3306	0.0057
75	−2.7790	0.0052	−2.8741	0.0056
80	−2.3531	0.0052	−2.3993	0.0056
85	−1.9278	0.0052	−1.9082	0.0056
90	−1.5134	0.0048	−1.4287	0.0059
95	−1.1134	0.0042	−0.9981	0.0055
100	−0.7279	0.0032	−0.6165	0.0044
105	−0.3567	0.0018	−0.2838	0.0026

4. Outlier Analysis

Having fitted the Lee-Carter model, we are now ready to proceed to the outlier analysis of the mortality index, k_t. Outlier analysis of time-series data comprises two key issues:

- To search for the location and type of outliers in a contaminated time series (known as the outlier detection problem in the time-series literature).

- To obtain better estimates of parameters in the underlying time-series model through the incorporation of outlier effects within a model (known as the outlier adjustment problem).

Time-series outlier analysis was first studied by Fox (1972), who employed the likelihood ratio test for outlier detection, but it has also been considered by many other researchers. Chang et al. (1988) developed an iterative procedure for outlier detection and employed intervention models to incorporate outlier effects. Chen and Liu (1993) proposed an augmented iterative procedure for the joint estimation of model parameters and outlier effects. Other references on this topic include Abraham and Chuang (1989), Ljung (1993), Muirhead (1986), Tsay (1986, 1988), Vaage (2000), and Wei (1990).

The first phase of an outlier detection procedure is to specify the underlying stochastic structure, that is, the appropriate ARIMA model, for the outlier-free series. However, the true underlying stochastic structure is seldom known a priori in practice, and this brings about the following problems, as noted by Chen and Liu (1993):

- The existence of outliers may result in erroneous model selection. Chan (1995) showed that outliers create spurious autocorrelations, which can lead to erroneous model specification. This cannot be avoided even when the sample size is enlarged.

- Even if the model is correctly specified, outliers may lead to biases in the estimation of parameters. This affects the detection of outliers and ultimately obscures the reestimation of parameters; the whole process repeats itself indefinitely.

In this study we employ both the Chen and Liu (1993) method and the additional iterative cycle proposed by Li and Chan (2005) to ameliorate the problem of model selection. For brevity we will restrict our discussion to points necessary for describing the applications in this paper.

4.1 Outlier Models

We assume that the outlier-free time series Z_t follows an autoregressive integrated moving average, ARIMA(p,d,q) model,

$$\phi(B)(1-B)^d Z_t = \theta(B)a_t , \qquad (3)$$

where B is the backshift operator such that $B^s Z_t = Z_{t-s}$,

$$\varphi(B) = 1-\varphi_1 B - \cdots - \varphi_p B^p ,$$
$$\theta(B) = 1-\theta_1 B - \cdots - \theta_q B^q ,$$

and $\{a_t\}$ is a sequence of white noise random variables, iid with mean 0 and constant variance σ^2.

By definition, outliers are nonrepetitive exogenous interventions. Quantitatively, an outlier-contaminated time series Y_t consists of an outlier-free time series Z_t plus an exogenous intervention effect, denoted as $\Delta_t(T,\omega)$, that is,

$$Y_t = Z_t + \Delta_t(T,\omega) , \qquad (4)$$

where T is the locality of the outlier and ω is the magnitude of the outlier.

We consider four common types of outliers, namely, an additive outlier (AO), innovational outlier (IO), level shifts (LS), and temporary change (TC).

- An AO affects only the level of a single observation, that is,

$$\Delta_t(T,\omega) = \omega D_t^{(T)}. \tag{5}$$

- An IO affects all observations beyond T through the memory of the underlying outlier-free process, that is,

$$\Delta_t(T,\omega) = \frac{\theta(B)}{\phi(B)(1-B)^d}\omega D_t^{(T)}. \tag{6}$$

- An LS affects a series at a given time, and its effect is permanent, that is,

$$\Delta_t(T,\omega) = \frac{\omega}{1-B} D_t^{(T)}. \tag{7}$$

- A TC affects a series at a given time, and its effect decays exponentially according to a dampening factor, say, δ, that is,

$$\Delta_t(T,\omega) = \frac{\omega}{1-\delta B} D_t^{(T)}. \tag{8}$$

In practice, the value of δ often lies between 0.6 and 0.8 (Liu and Hudak 1994, p. 76). In our study we take $\delta = 0.7$ as recommended by Chen and Liu (1993).

Here $D_t^{(T)}$ is an indicator variable that equals 1 when $t = T$ and 0 otherwise. A graphical illustration of each type of outlier is given in Figure 1. For more detailed discussions of these types of outliers, see Chen and Tiao (1990), Fox (1972), and Tsay (1988). In general, a time series may contain more than one, say, m, outliers, and we have the general time-series outlier model,

$$Y_t = Z_t + \sum_{i=1}^{m} \Delta_t(T_i, \omega_i). \tag{9}$$

Figure 1
Different Types of Time-Series Outliers

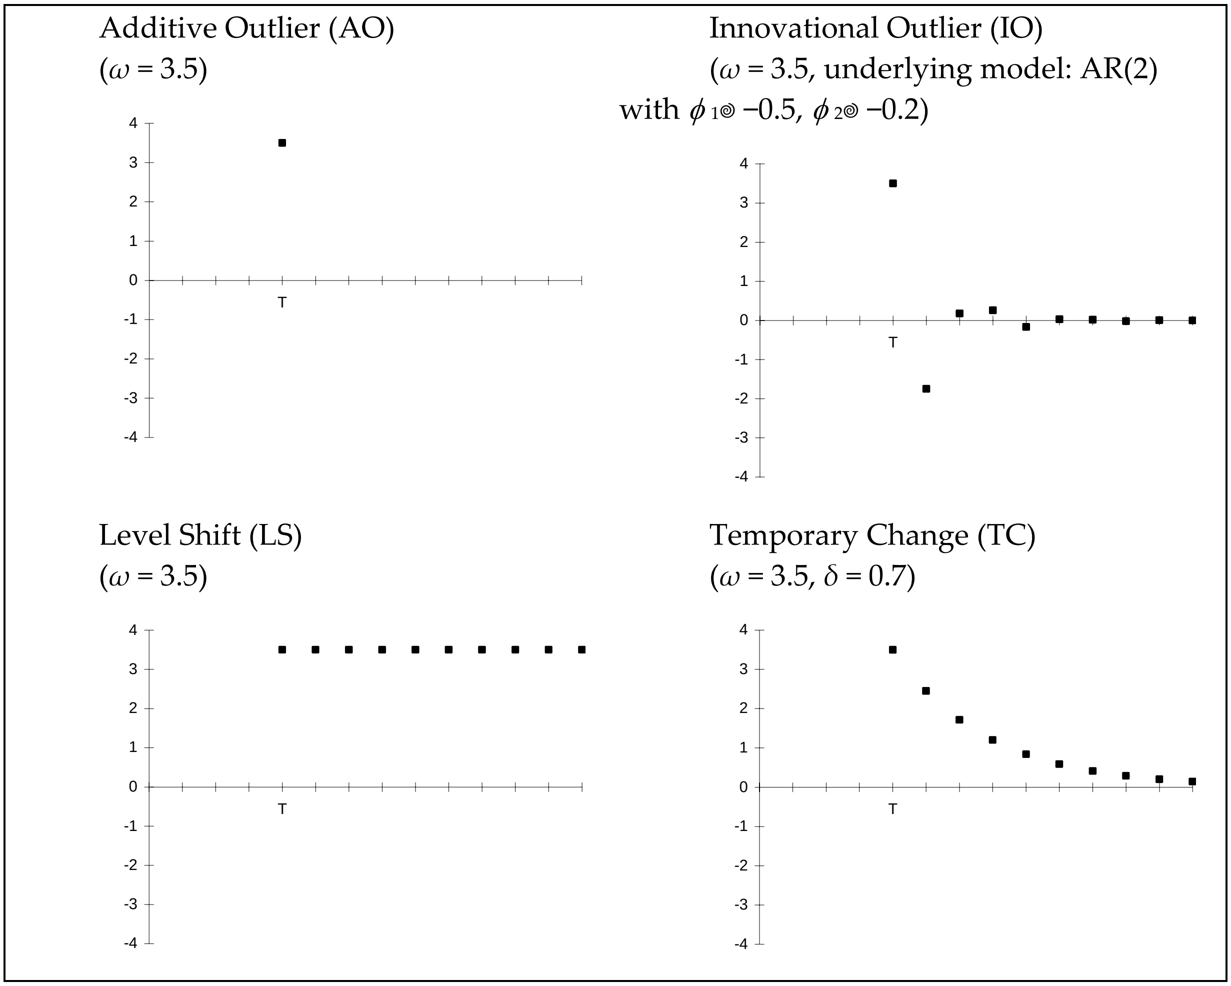

4.2 Detection and Adjustment

The detection process proposed by Chen and Liu (1993) is primarily based on the effect of outliers on the estimated residuals. To begin, we define a polynomial $\pi(B)$ as

$$\pi(B) = \frac{\varphi(B)(1-B)^d}{\theta(B)} = 1 - \pi_1 B - \pi_2 B^2 - \cdots . \tag{10}$$

Then equation (3) can be rewritten as

$$\pi(B) Z_t = a_t , \tag{11}$$

and the fitted residuals $\hat{e}_t$ in equation (4), which may be contaminated with outliers, can be readily obtained and expressed as

$$\hat{e}_t = \pi(B)Y_t . \tag{12}$$

So, for each type of outlier, we have the following:

$$\text{AO} \qquad \hat{e}_t = \omega D_t^{(T)} + a_t, \tag{13}$$

$$\text{IO:} \qquad \hat{e}_t = \omega \pi(B) D_t^{(T)} + a_t, \tag{14}$$

$$\text{TC:} \qquad \hat{e}_t = \omega \frac{\pi(B)}{(1-\delta B)} D_t^{(T)} + a_t, \tag{15}$$

$$\text{LS:} \qquad \hat{e}_t = \omega \frac{\pi(B)}{1-B} D_t^{(T)} + a_t. \tag{16}$$

Alternatively, we can rewrite equations (13)–(16) as a general time-series regression,

$$\hat{e}_t = \omega d(j,t) + a_t. \tag{17}$$

where $j \in J = \{\text{AO,IO,TC,LS}\}$; $d(j,t) = 0$ for all j and $t < T$; $d(j,T) = 1$ for all j; and for all $k \geq 1$,

$$\begin{aligned}
&d(\text{AO},T+k) = 0, \\
&d(\text{IO},T+k) = -\pi_k, \\
&d(\text{TC},T+k) = \delta_k - \sum_{i=1}^{k-1} \delta^{k-j} \pi_j - \pi_k, \\
&d(\text{LS},T+k) = 1 - \sum_{i=1}^{k} \pi_k.
\end{aligned}$$

Hence, for a given T (suspected locality of the outlier) and j (suspected type of outlier), the standardized t-statistic $\tau(j,T)$ for the effect of the outlier, that is, the slope parameter ω, can be readily computed using the principle of least squares. The final test statistic is the maximum value of this t-statistic over all possible T and j, that is,

$$\mathcal{T} = \max_{1 \leq T \leq n} \max_{j \in J} \{\tau(j,T)\} . \tag{18}$$

For a given j, the test statistics follow approximately a normal distribution. An outlier of type j is detected if the final test statistic is greater than a critical value of C. We employ $C = 2.5$ in this paper as recommended by Liu and Hudak (1994) for a reasonable level of sensitivity. With the type and the location of an outlier, we can jointly reestimate the model parameters and the outlier effect. After the estimation is complete, we can adjust the original series for the effect of the detected outlier. The process is repeated until no other outlier is found.

The outlier detection-estimation-adjustment process can be implemented by using various standard statistical software packages for time-series analysis, such as AUTOBOX (Ord and Lowe 1996), SAS/ETS (SAS Institute 2004), and SCA (Liu and Hudak 1994).

To prevent problems with potential erroneous model selection, we also employ the following external iteration cycle proposed by Li and Chan (2005) to incorporate outlier effects in model specification.

Step 1—Tentative Model Identification

Use the orthodox Box and Jenkins approach (Box and Jenkins 1976), or any other appropriate method, to tentatively identify the order of the underlying outlier-free ARIMA (p,d,q) model.

Step 2—Outlier Detection and Adjustment

Use Chen and Liu's (1993) iterative procedures for a joint estimation of outlier effects and ARIMA model parameters. After the incorporation of outlier effects, an outlier-adjusted data series is obtained.

Step 3—Reidentification of the Model

Using the method employed in Step 1, identify the ARIMA model underlying the adjusted data series. If the reidentification makes a difference in p, d, and/or q, go back to Step 2 using the original unadjusted data series under the reidentified order (p,d,q). Otherwise, terminate the iteration cycle, and the ultimate estimates of outliers and ARIMA model parameters are those obtained in the immediately previous Step 2.

5. Empirical Results

5.1 Detected Outliers

Table 2 displays all the outliers found in the mortality indices derived from the American and Canadian data. Note that a negative outlier stands for improvement in mortality, whereas a positive one means deterioration. The outliers mostly are found in the first half of the century, and they can be roughly classified as pandemic-related and war-related.

World War I best explains the mild positive level shift in 1916. After a long stretch of isolationism, President Woodrow Wilson requested that the U.S. Congress declare war, which it did on April 6, 1917. The American contribution to the war was substantive: total casualties of the American armed forces amounted to 327,010 (Wikipedia 2004a). This deadly war ended in 1918, and the fact that young men were no longer being killed could possibly account for the negative temporary change in 1921.

The positive additive outlier in 1918 is most likely a consequence of the "Spanish flu" pandemic, which particularly affected healthy young adults. The earliest known case in the United States was at Fort Riley, Kansas, on March 11, and it soon spread swiftly to other locations. During the next few months, the virus infected about 28 percent of the U.S. population and killed approximately 500,000 Americans (Wikipedia 2004b). This pandemic of unprecedented virulence spread around the world and killed between 20 and 40 million people over the next six months (Holmes 2004); it was thought to have been the most deadly pandemic so far in human history.

Figure 2 shows the death rates from influenza and pneumonia in the United States during this time of turmoil—the Great Depression and war. Abrupt jumps are observed in 1926, 1928–29, and 1936–37. This suggests that some small-scale and probably not widely reported influenza epidemic may have occurred during this time of economic insecurity, which may have given rise to outliers in the late 1920s and 1930s.

In 1950 North and South Korea went to war. The Korean War killed about 44,000 U.S. servicemen (Wikipedia 2004c) and temporarily reduced the population of healthy and robust Americans. The cessation of young men being killed after the war's end might account for the negative temporary change in 1954.

Finally, the negative level shift in 1975 may be associated with better nutrition, the use of antibiotics, or some other improvements in social conditions.

TABLE 2
Summary of Outliers Detected in the Mortality Index

Year	Size	*t*-Value	Type
United States			
1916	7.182	3.26	LS
1918	24.770	11.91	AO
1921	−9.574	−4.25	TC
1928	6.679	2.97	TC
1936	9.438	4.21	TC
1954	−6.280	−2.8	TC
1975	−8.278	−3.78	LS
Canada			
1926	7.635	3.36	LS
1937	6.736	3.75	AO

Figure 2
Death Rates from Influenza and Pneumonia, United States, 1924–1939

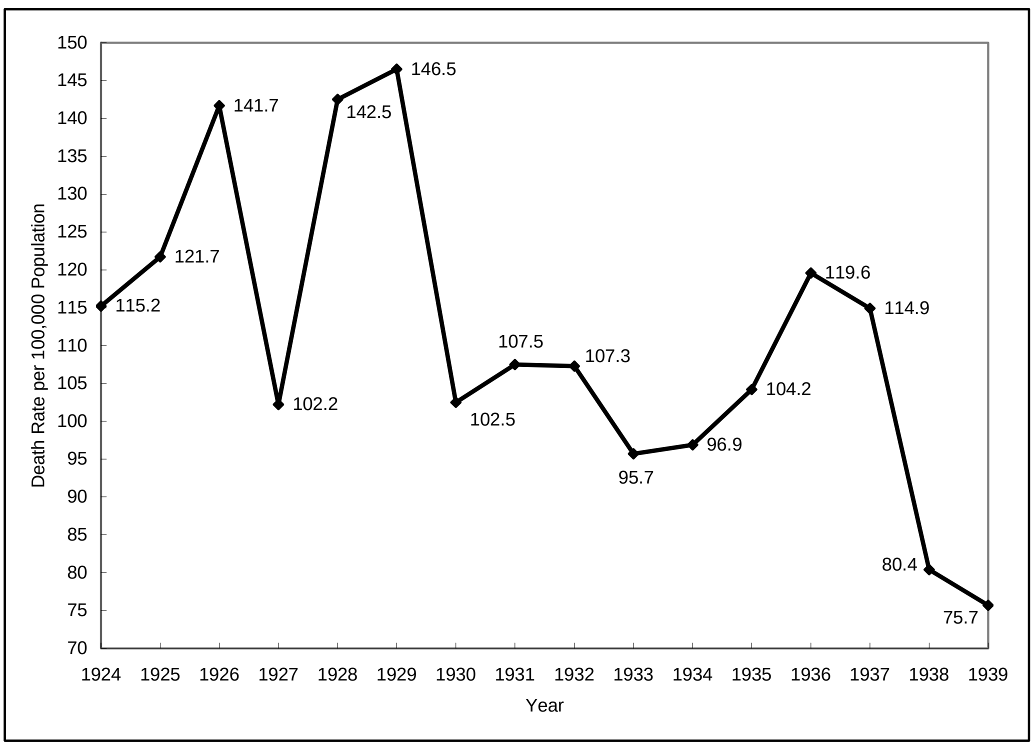

Source: National Center for Health Statistics (2004a).

5.2 The Outlier-Adjusted Model

Below we evaluate the model performance in each mortality series under investigation. The identified ARIMA order does not make a difference in Li and Chan's (2005) iterative cycle: (0,1,0) fits both indices well throughout the entire process. Ljung-Box's Portmanteau statistics (Ljung and Box 1978) are computed up to the tenth lags for testing serial correlation of the residuals. The results (not shown) do not suggest any inadequacy of the fitted models, either with or without outlier adjustment.

To provide a tentative impression of the fitting performance, we report the residual standard errors in Table 3. Upon outlier detection and adjustment, values of residual standard error are significantly reduced. However, under the principle of parsimony, one should employ the least possible number of parameters for adequate representations, and it is therefore premature to make conclusions solely from the reduction in residual standard errors as the total number of model parameters is increased, as additional intervention components for outlier effects are included. To compare the fitting performance of models with a different number of parameters, Akaike (1974) introduced the Akaike Information Criterion (AIC), which is defined as

$$\text{AIC} = n\ln(\hat{\sigma}^2) + 2M, \tag{19}$$

where n is the number of effective observations, $\hat{\sigma}$ denotes the residual standard deviation, and M represents the number of parameters in the model. The criterion rewards the reduction of the residual standard deviation, but at the same time harshly penalizes the increased number of parameters. In other words, a smaller AIC value is more favorable. In Table 3 we also report the AIC values both before and after outlier adjustment. The results support the view that the outlier-adjusted models provide a better fit.

TABLE 3
Fitting Diagnostics, ARIMA Model for the Mortality Indices

	Residual Standard Error		**AIC**	
	Before Adjustment	**After Adjustment**	**Before Adjustment**	**After Adjustment**
United States	4.4262	2.4323	301.51	195.77
Canada	2.6218	2.2318	156.29	134.84

TABLE 4
Interval Forecast of the Central Rate of Death at Selected Ages at 2025

Age	Without Outlier Adjustment		With Outlier Adjustment		Percentage Reduction in Width of Interval
	Lower Limit	Upper Limit	Lower Limit	Upper Limit	
United States					
20	0.000240	0.000479	0.000285	0.000406	49.29%
50	0.002436	0.003703	0.002705	0.003352	48.90
80	0.039981	0.053458	0.042990	0.049890	48.80
Canada					
20	0.000152	0.000307	0.000237	0.000359	20.93
50	0.002062	0.002976	0.001777	0.002362	35.96
80	0.059350	0.071268	0.036401	0.044906	28.65

Chatfield (1993) pointed out that it is important to provide interval forecasts so that forecasts obtained by different methods may be compared more thoroughly. In Table 4 we report the interval forecast for the central rates of death at selected ages; and in Figures 5–10 we plot both the mean and the interval forecasts of the mortality index, the complete life expectancy at birth, and the actuarial present value (APV) of a life annuity at age 65. The confidence intervals are substantially narrower after the adjustment for outliers; they are coherent with the empirical results of Ledolter (1989), which found that estimated prediction intervals are sensitive to outliers, as they strongly inflate the estimate of the innovation variance. For mathematical derivations of the interval forecast for the APV of life annuities, see the Appendix.

Readers should note the interpretation of the reduction in width of the interval forecasts. There are a number of reasons for the increased narrowness. First, as mentioned earlier, the detection and adjustment process effectively mitigates the problem of biases in parameter estimation due to outliers. Second, if an outlier occurs at the forecast origin, an assumption on the type of that outlier is required, and the uncertainty of the type of the outlier is not reflected the interval forecast. Third, the narrowness essentially reflects the optimistic nature of the outlier-adjusted forecasts. Under the outlier-adjusted models we presume that the same event will not recur in the future, and so the question arises which model a practitioner should pursue during the forecasting exercise. Chan (2002, p. 559-560) offers the following suggestion for stochastic investment modeling:

Whether or not it is appropriate to adjust the data for outliers depends on the purpose to which the model so derived will be used. If the model will be used in an application for which extreme stochastic fluctuations are less important (e.g. to ensure that premiums are adequate in most, but not extreme, scenarios), then it may be preferable to use a model based on outlier-adjusted data. If, however, the model will be used in an application for which extreme stochastic fluctuations are important (such as pricing catastrophe risks or ensuring that investment guarantee reserves are sufficient to keep an insurance company solvent in all but the most extreme scenarios), then a model which is sympathetic to outliers in the data ought to be used.

This rationale also applies to mortality forecasting. The outlier-adjusted version may be more preferred if the model is used in an application where the major mortality trend is the primary focus, for example, in forecasting cohort life tables for the assessment of pecuniary loss in personal injury litigations (Sarony et al. 2003). Some other measures that allow for the recurrence of outliers should be undertaken if the model is used in an application where extreme fluctuations are important. Further discussions are provided in the next section on predicting the highest attainable age.

6. Extension—Predicting the Highest Attainable Age of Canadians

6.1 Extension of Period Life Tables

The age pattern of mortality of the oldest-old is of paramount importance in the estimation of the highest attainable age. Yet, as noted earlier, mortality data at higher ages are subject to problems in data quality, and because of the lack of reliable observation at advanced ages, there has been no consensus on which mathematical technique to use to generate mortality rates for advanced ages as an extension of earlier mortality.

A mutual shortcoming of the method proposed by Coale and Guo (1989, pp. 614–615) and the method of extinct generations suggested by Vincent (1951) is that they require an assumed age limit at around 110, which is contradictory to our goal of predicting the maximum life span. In the following analysis, we revise the estimates of death rates beyond age 85 by relying on the relational mortality model (Himes et al. 1994), which imposes no assumption on the limiting age. The relational model, recommended by the Working Group on Projecting Old-Age Mortality and Its Consequences of the Population Division in the United Nations (United Nations 1997), consists of a "standard" age pattern of mortality by sex and by single year of age from 45 to 99, calibrated from 82 different mortality schedules (for each sex) observed in a variety of low-mortality countries. (The "standard" mortality schedule from age 45 to 99 can be found in Himes et al. [1994], p. 273.) The "standard" is then made useful at higher ages by fitting a straight line to the logits of the age-specific death rates, beginning at age 80, as recommended by Himes et al. (1994). Mathematically,

$$\text{logit}(m_x^s) = \alpha + \beta x, \tag{20}$$

where $\text{logit}(m_x^s) = \ln\left(\frac{m_x^s}{1-m_x^s}\right)$, and m_x^s denotes the "standard" central rate of death at age x. It is then used to produce "standard" age-specific death rates from age 100 to 150. Finally, we relate the "standard" schedule to each of the historical Canadian period life tables by regressing on their logit transformations, that is,

$$\text{logit}\, m_x = \delta + \gamma\, \text{logit}\, m_x^s, \tag{21}$$

where m_x's are the death rates in the life table where values for the higher ages are to be estimated, and δ and γ are the regression coefficients.

Buettner (2002) showed that the linear extension in the logit domain represented by equation (20) is equivalent to the old-age term of the Heligman-Pollard model (Heligman and Pollard 1980). Furthermore, it is noteworthy to see the limiting behavior of this extension. Equation (20) implies that

$$\lim_{x\to\infty} m_x^s = \lim_{x\to\infty} \frac{e^{\alpha+\beta x}}{1+e^{\alpha+\beta x}} = 1, \tag{22}$$

for our estimates of α and β, which satisfy the condition of α ◎◎0 and β ◎◎0. Then it follows that

$$\lim_{x\to\infty} m_x = \frac{e^{\delta+\gamma}}{1+e^{\delta+\gamma}}. \tag{23}$$

In other words, as x tends to infinity, the value of m_x tends to a certain limit that is less than 1. This is parallel to the property of the logistic model of mortality that was used in Thatcher's (1999) study on the highest attainable age.

6.2 The Probability Model on the Highest Attainable Age

Suppose that N members of a single birth cohort survive to age x_0; then when all of these N members have died, there will be a highest value, say, ω_N, among their ages at death. Prior to realization, ω_N can be regarded as a random variable, and by basic actuarial mathematics, its distribution function is given by

$$\Pr(\omega_N < x) = (1 - {}_{x-x_0}p_{x_0})^N. \tag{24}$$

Thatcher (1999) showed that this distribution function demonstrates strong robustness with respect to both N and x_0. In this study we shall take $x_0 = 70$ and N as the size of a cohort when it reached age 70 years, which can be approximated by the appropriate midyear population estimate. In addition, Thatcher fitted the logistic mortality model to some period life tables and worked out their corresponding distribution function, which has the following closed form:

$$\Pr(\omega_N < x) = \left[1 - \exp[\gamma(x_0 - x) + \beta^{-1}\ln\{1+\alpha\exp(\beta x_0)\} - \beta^{-1}\ln\{1+\alpha\exp(\beta x)\}]\right]^N, \tag{25}$$

where $\alpha, \beta, \text{and } \gamma$ are the parameters in the logistic mortality model, given by

$$\mu_x = \frac{\alpha\exp(\beta x)}{1+\alpha\exp(\beta x)} + \gamma. \tag{26}$$

Unfortunately, Thatcher's layout failed to acknowledge the improvement in human mortality, and it is therefore likely to underpredict the maximum life span that a specific cohort could achieve. To illustrate, we consider the cohort of Canadians born in 1886, that is, the supercentenarians at age 118 in the year 2004. This cohort attained age x_0 ◎◎70 in 1956. Following Thatcher's methodology, we compute the distribution function of the highest attainable age for this particular cohort using the period life tables in 1956, and the results are depicted in Figure 3. The modes, far less than 110, are seemingly too low. Moreover, the 99th percentiles fail to capture the currently living supercentenarians in Canada, such as Julie Winnefred Bertrand (currently at the age of 113 in 2004; see Gerontology Research Group 2004), and even the Canadian supercentenarians who died in the last few years, be they Jeanne Clament (dead at the age of 112 years in 1997) or Louise Meilleur (dead at the age of 117 years in 1998; see Robine and Vaupal 2002).

Here we attempt to incorporate the effect of mortality improvement on the highest attainable age by fitting the probability model in equation (24) to cohort life tables instead of the period ones. To construct the required cohort life tables, we first fit the Lee-Carter model to the extended sex-specific period life tables obtained in Section 6.2. Then we apply the procedures described in Section 4 to the time series of k_t for estimates of outliers and outlier-adjusted model parameters. Based on the outlier-adjusted model, we obtain tentative extrapolative forecasts of k_t. However, as noted earlier, these extrapolative forecasts essentially are based on the assumption that the outliers will not recur in the future, which is inappropriate in the application of the forecast to the prediction of the highest attainable age in which extreme fluctuations are important. In light of this, we relax this assumption by using a bootstrapping procedure: outliers in the postsample forecasts are obtained by sampling with replacement from the array of detected outliers, separately for each sex. Informal experiments indicate that the resultant distribution of the highest attainable age is rather insensitive to the combination of outliers we added in the postsample forecast of k_t. Having obtained the revised forecasts, future period life tables can be readily computed, and the appropriate cohort life tables can be constructed with basic mathematics from demography. Detailed steps can be found in Brown (1993).

Figure 3
Distribution Function of the Highest Attainable Age (ω) for the Cohort of Canadians Born in 1886, Using Period Life Tables in 1956

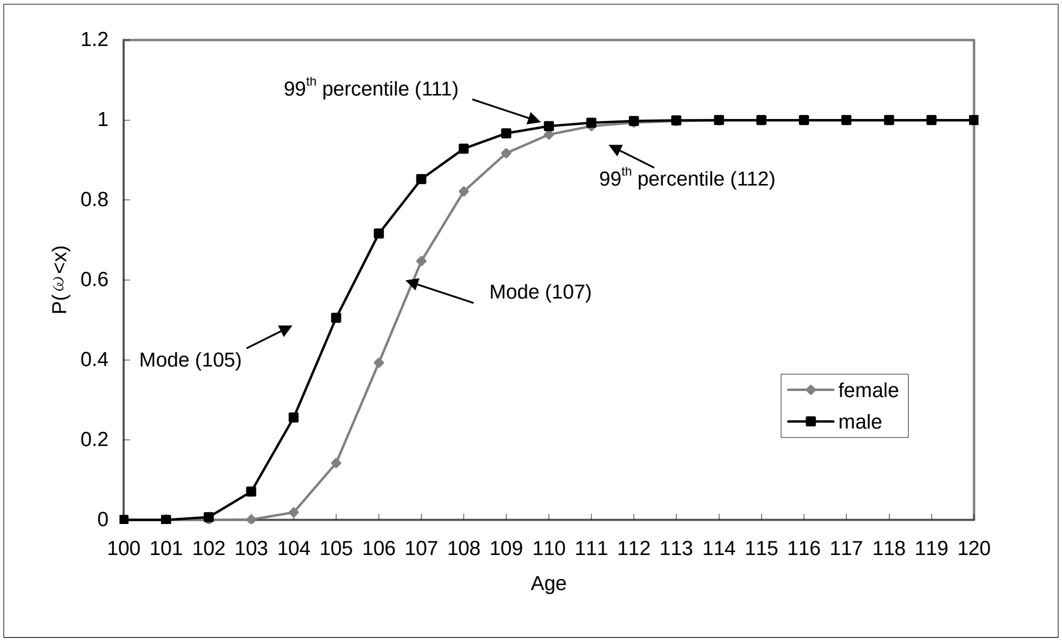

Table 5 summarizes the predicted distributions of the highest attainable age for cohorts born in 1886 (the prevailing supercentenarians), 1900, 1910, 1920, and 1930. Note that these cohorts attained age x_0 ◎◎70 in 1956, 1970, 1980, 1990, and 2000, respectively. For the information of readers, results derived from the unadjusted Lee-Carter model are also presented in tandem. In the first column the predicted distributions for the cohort of current supercentenarians agree reasonably well with the reality. The mode for females is pretty close to the current age of Julie Winnefred Bertrand, which is 113; and apparently the 99th percentile may be able to capture all possible values of the highest age that she will attain. Across the row, we can observe that the effect of mortality improvement on the maximum life span is substantial. On average, the mode increases by approximately three years in every 10-year period, and the trend is seemingly accelerating.

Here we would like to sound a few cautionary notes in the interpretation of the results. First, the 99th percentile of the highest attainable age is completely different from that of the lifetime: the former is the extreme of the *maximum* life span, whereas the latter is merely the extreme of the life span. Although there may be some extent of positive correlation, these two values are significantly different in magnitude and are by no means comparable with each other. Second, the probabilistic distribution of the highest attainable age is conditional on the mortality experience of the specific cohort from which the distribution is derived, and thus it is not applicable to any other cohort whose mortality experience is different. Third, the variance of the distribution takes no account of the uncertainty that arises from the forecast of future period life tables. The widening of the prediction intervals in Table 5 and the thickening of the "tails" in Figure 4 should be attributed to the rectangularization and the expansion of the survival functions arising from the rapidly improving mortality. The actual variance, taken into account the forecast error, could be higher.

Note also that in the prediction of maximum age, we require a very long-term forecast of future period life tables. To illustrate, we consider the cohort of male Canadians born in 1930 and who attained age $x_0 = 70$ in 2000. Table 5 shows that the 99th percentile is 135, which implies that more than 65 successive future life tables have been used in the computation of the probability distribution corresponding to this cohort. This leads to the question of whether the inertia of mortality improvement still continues in such a long forecast horizon. In Tables 6 and 7 we present a number of scenario forecasts to provide readers with some ideas on how the predicted distributions will vary if mortality improvement ceases at some specific future dates.

TABLE 5
Predicted Distribution of the Highest Attainable Age for Selected Cohorts of Canadians, Assuming That Future Mortality Follows the Lee-Carter Forecast

	Year of Birth				
	1886	**1900**	**1910**	**1920**	**1930**
	Male, with Outlier Adjustment				
Mode	110	114	117	120	123
99th percentile	119	124	128	132	135
	Female, with Outlier Adjustment				
Mode	116	120	123	127	133
99th percentile	126	132	137	144	>150
	Male, without Outlier Adjustment				
Mode	110	114	117	120	123
99th percentile	119	124	128	132	136
	Female, without Outlier Adjustment				
Mode	116	121	126	132	138
99th percentile	127	135	142	>150	>150

Figure 4

Predicted Probability Mass Function of the Highest Attainable Age for Selected Cohorts of Male Canadians, Based on the Outlier-Adjusted Lee-Carter Forecast

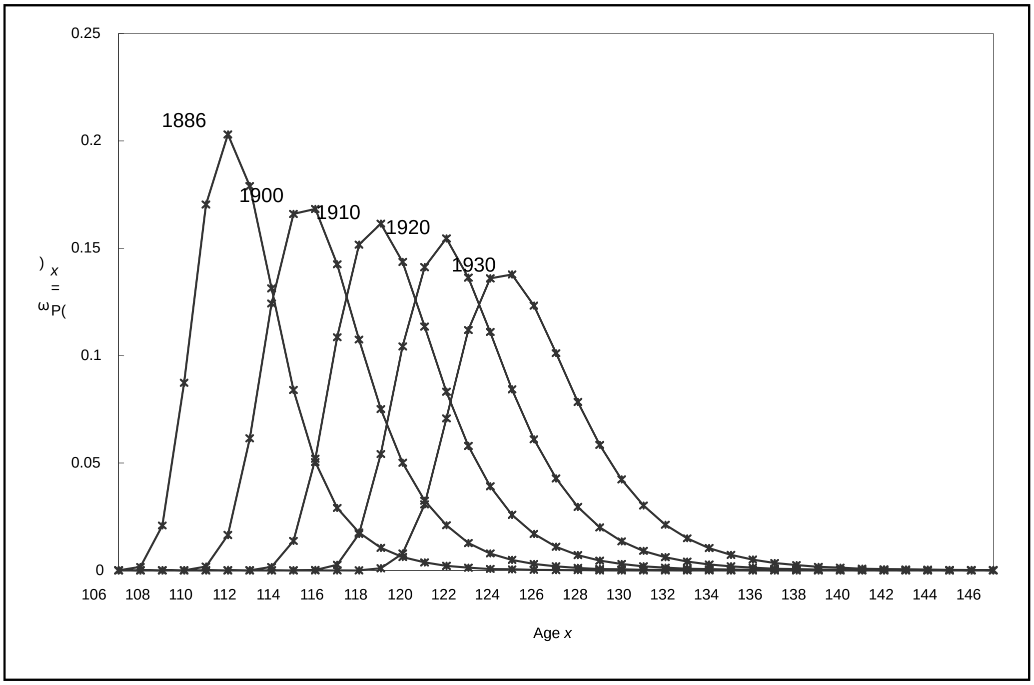

TABLE 6

Predicted Distribution of the Highest Attainable Age for Selected Cohorts of Canadians, Under Different Scenarios of Mortality Improvement, Male

	Year of Birth			
	1900	**1910**	**1920**	**1930**
	Scenario 1: Mortality Improvement Ceases in 30 Years			
Mode	114	117	120	121
99th percentile	124	128	130	132
	Scenario 2: Mortality Improvement Ceases in 20 Years			
Mode	114	117	119	120
99th percentile	124	127	129	130
	Scenario 3: Mortality Improvement Ceases in 10 Years			
Mode	114	116	117	118
99th percentile	124	126	127	128
	Scenario 4: Mortality Improvement Ceases after 2000			
Mode	113	115	116	116
99th percentile	122	124	125	125

TABLE 7
Predicted Distribution of the Highest Attainable age for Selected Cohorts of Canadians, Under Different Scenarios of Mortality Improvement, Female

	Year of Birth			
	1900	**1910**	**1920**	**1930**
	Scenario 1: Mortality Improvement Ceases in 30 Years			
Mode	120	124	128	130
99th percentile	133	138	141	143
	Scenario 2: Mortality Improvement Ceases in 20 Years			
Mode	120	123	126	126
99th percentile	132	135	137	138
	Scenario 3: Mortality Improvement Ceases in 10 Years			
Mode	120	121	123	123
99th percentile	130	132	133	133
	Scenario 4: Mortality Improvement Ceases after 2000			
Mode	119	120	121	121
99th percentile	129	130	131	131

7. Discussion and Conclusion

In this paper we performed a systematic outlier analysis of the mortality index built in the Lee-Carter model. Through outlier detection we found that mortality levels in the United States and Canada are vulnerable to events like pandemics and wars. By incorporating the effect of outliers, we created an outlier-adjusted Lee-Carter model. Given the better fit and the enhancement in forecast efficiency, the outlier-adjusted version seems to be an attractive alternative to the original model.

To simplify the analysis, we assumed that the mortality index follows a linear time-series model throughout the process of outlier detection and adjustment. Nevertheless, the detected outliers could be signs of nonlinearity in the mortality index, and it is therefore warranted to explore the use of nonlinear time-series models, for example, the Threshold model (Tong 1983), for a more comprehensive picture of the underlying stochastic structure.

The availability of data is a critical factor in the analysis of old-age mortality. In this study the original data are disaggregated (historical life tables are available only in abridged form in the United States) and extrapolated to very advanced ages. In this way the data are unavoidably contaminated, and this might influence the ultimate conclusion. To avoid this problem, we might perform the analysis using the Canadian Pension Plan/Quebec Pension Plan data (for Canada) and the Social Security Administration (OASDI) data (for the United States), which are more detailed.

As noted earlier, the prediction of the highest attainable age requires a long-term mortality forecast, whereas the period on which the forecast is based is relatively too short. In addition, as pointed out by Smith (1999) and Ledford and Robinson (1999), the prediction might be improved by considering some modern statistical methods based on exceedances over a high threshold (see, e.g., Davison and Smith 1990; Pickands 1975). These methods, however, require detailed mortality data over the threshold selected. All in all, for a better understanding of the maximum life span, we require more accurate mortality data for the supercentenarians, and we wholeheartedly look forward to future developments of the International Database on Longevity.

Acknowledgments

The authors are grateful to the U.S. Census Bureau, the National Center for Health Statistics, the Human Mortality Database (www.mortality.org/www.humanmortality.de), the Gerontology Research Group, and Statistics Canada for providing the data used in this study. This work was supported by a grant from the Research Grants Council of the Hong Kong Special Administrative Region (Competitive Earmarked Research Grant Project No. HKU 7111/01H).

The authors would also like to acknowledge stimulating discussions with the discussant (Prof. Robert Brown) and the participants in "Living to 100 and Beyond: Survival at Higher Ages Symposium, 2005," organized by the Society of Actuaries.

Figure 5
Mean Forecast and 95 Percent Pointwise Confidence Interval for the Mortality Index, k_t, United States

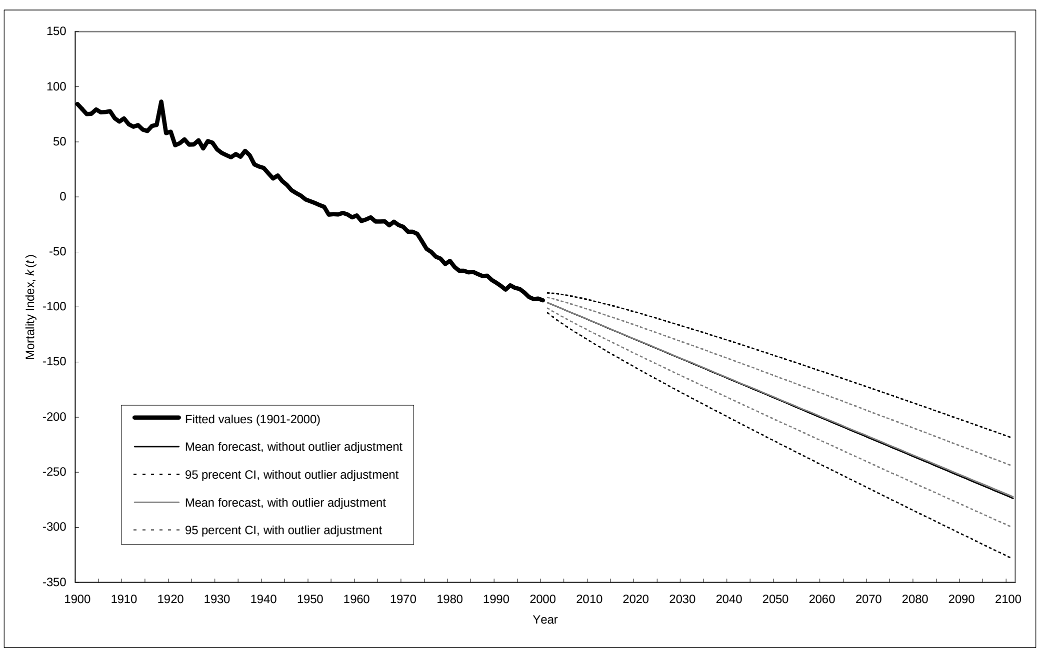

Figure 6
Mean Forecast and 95 Percent Pointwise Confidence Interval for the Mortality Index, k_t, Canada

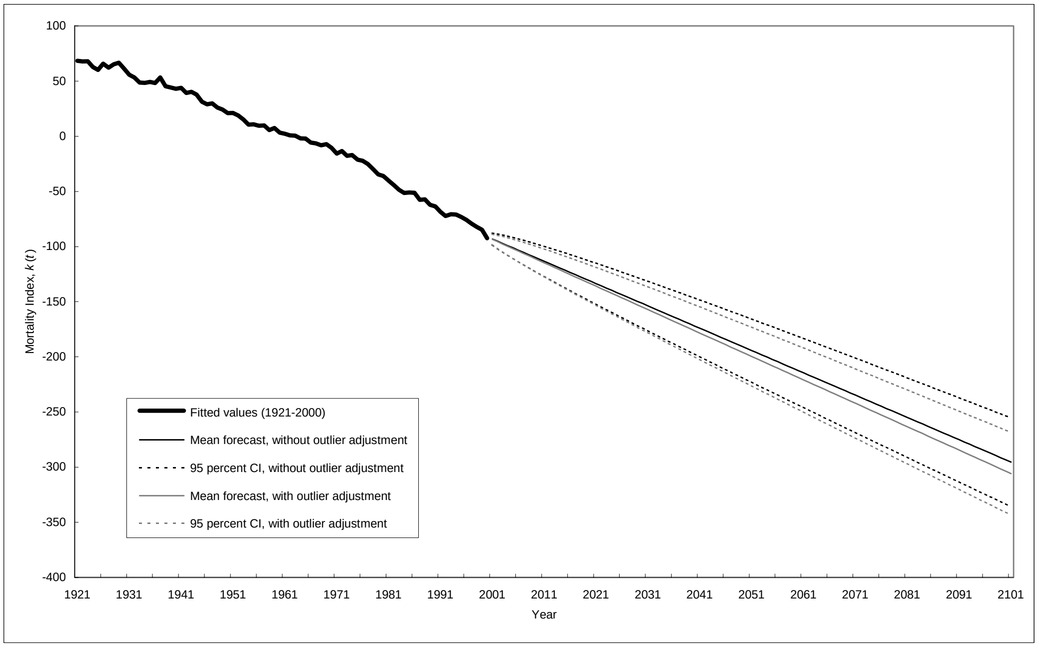

Figure 7
Mean Forecast and 95 Percent Pointwise Confidence Interval for the Complete Life Expectancy at Birth, United States

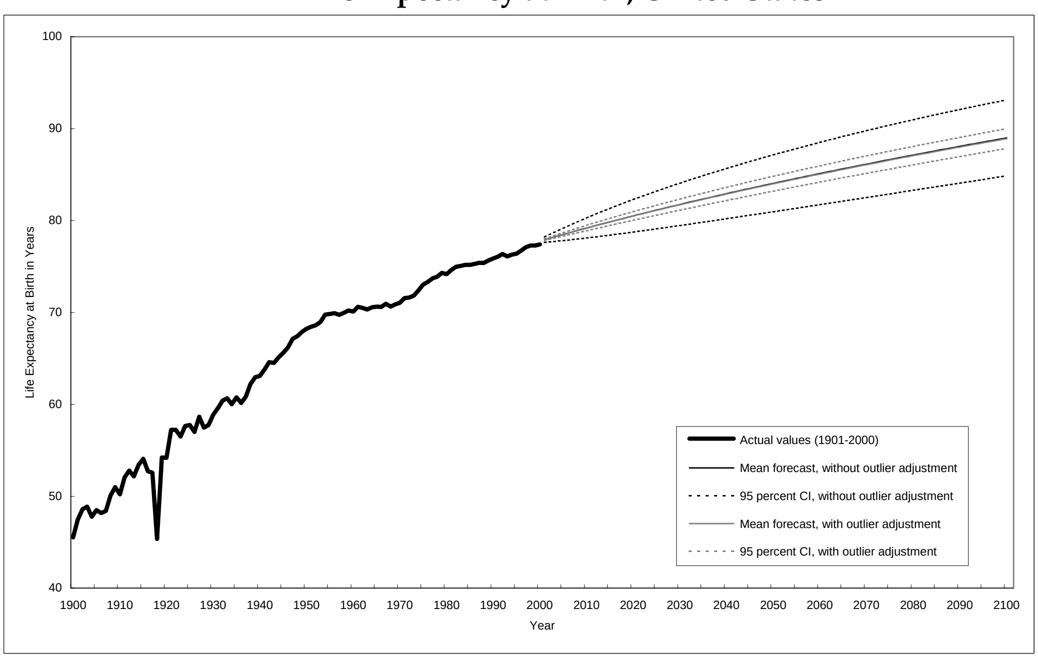

Figure 8
Mean Forecast and 95 Percent Pointwise Confidence Interval for the Complete Life Expectancy at Birth, Canada

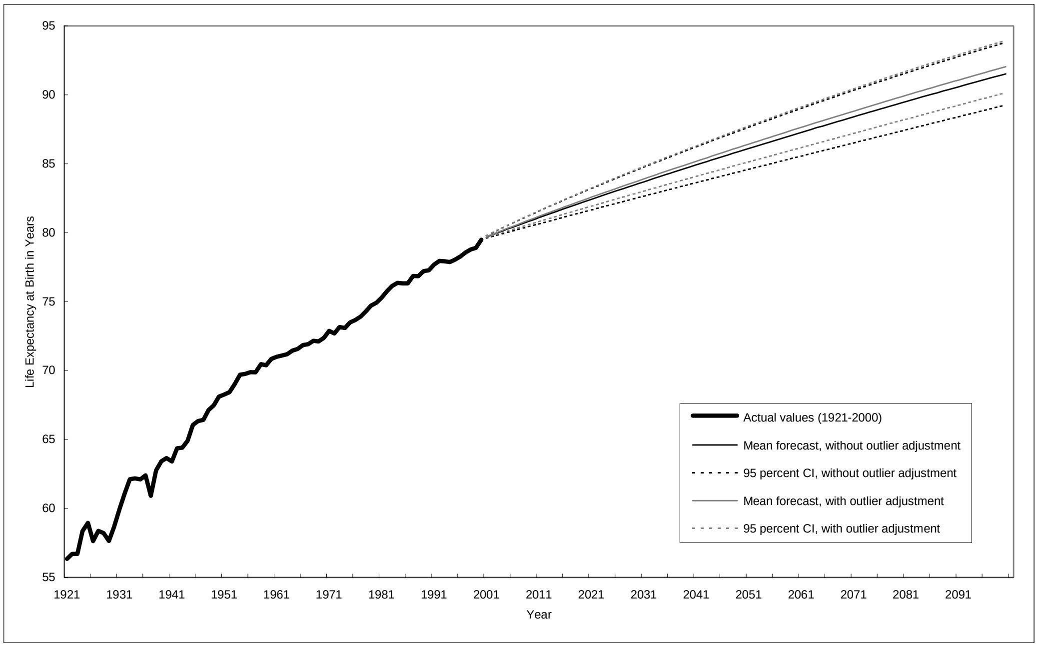

Figure 9
Mean Forecast and 95 Percent Pointwise Confidence Interval for the Actuarial Present Value of Life Annuity at Age 65, *i* = 2.5 percent, United States

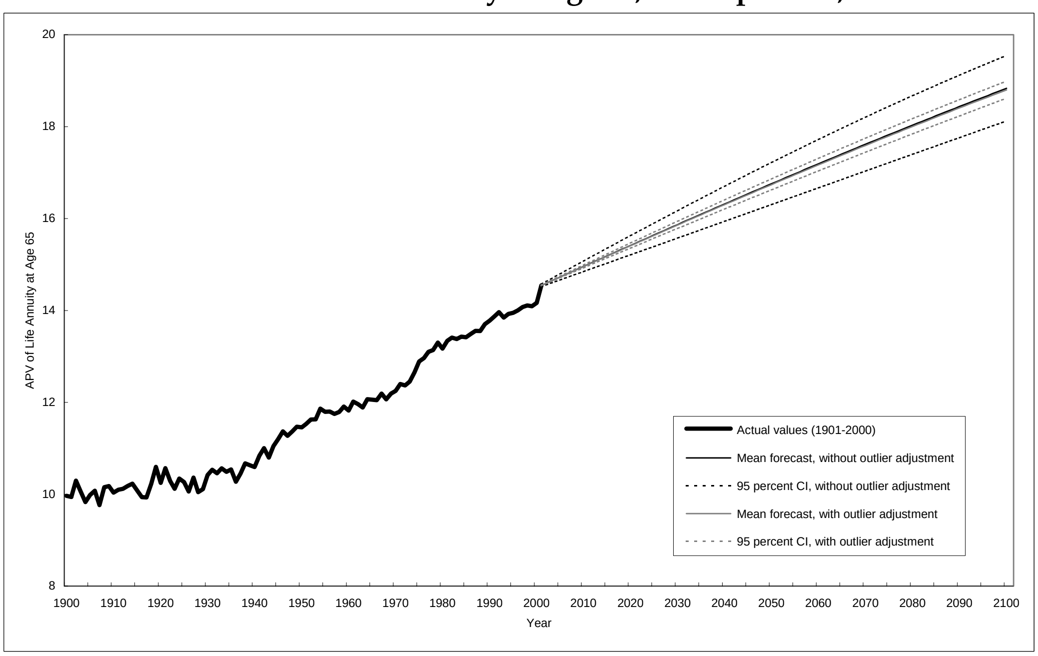

Figure 10
Mean Forecast and 95 Percent Pointwise Confidence Interval for the Actuarial Present Value of Life Annuity at Age 65, *i* = 2.5 percent, Canada

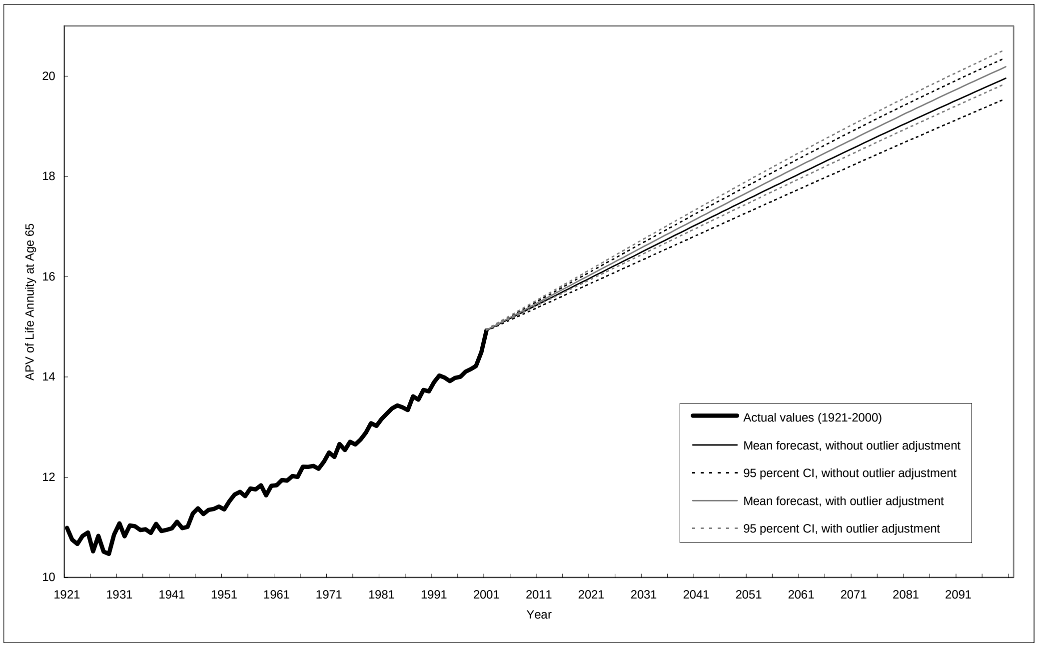

References

Abraham, B., and A. Chuang. 1989. "Outlier Detection and Time Series Modelling." *Technometrics* 31: 241–248.

Akaike, H. 1974. "A New Look at the Statistical Model Identification." *IEEE Transactions on Automatic Control* AC-19: 716–723.

Balke, N. S., and T. N. Fomby. 1994. "Large Shocks, Small Shocks, and Economic Fluctuations: Outliers in Macroeconomic Time-Series." *Journal of Applied Econometrics* 9: 181–200.

Booth, H., and L. Tickle. 2003. "The Future Aged: New Projections of Australia's Elderly Population." Department of Actuarial Studies Research Paper Series. Macquarie University, Sydney, Australia.

Bourbeau, R., and B. Desjardins. 2002. "Dealing with Problems in Data Quality for the Measurement of Mortality at Advanced Ages in Canada." *North American Actuarial Journal* 6(3): 1–13.

Bowers, N. L., Jr., H. U. Gerber, J. C. Hickman, D. A. Jones, and C. J. Nesbitt. 1997. *Actuarial Mathematics*. Schaumburg, IL: Society of Actuaries.

Box, G. E. P., and G. M. Jenkins. 1976. *Time Series Analysis Forecasting and Control*. 2nd ed. San Francisco: Holden-Day.

Brose, G. J. 1997. *Numerical Methods with MATLAB: A Resource for Scientists and Engineers* Boston: PWS Publishing.

Brouhns, N., M. Denuit, and J. K. Vermunt. 2002. "A Poisson Log-Bilinear Regression Approach to the Construction of Projected Lifetables." *Insurance: Mathematics and Economics* 31: 373–393.

Brown, R. L. 1993. *Introduction to the Mathematics of Demography*. 2nd ed. Winsted, CT: ACTEX Publications.

Buettner, T. 2002. "Approaches and Experiences in Projecting Mortality Patterns for the Oldest-Old." *North American Actuarial Journal* 6(3): 14–29.

———. 1992b. "Modelling and Forecasting U.S. Sex Differentials in Mortality." *International Journal of Forecasting* 8: 393–411.

Chan, W. S. 1995. "Outliers and Financial Time Series Modelling: A Cautionary

Note.” *Mathematics and Computers in Simulation* 39: 425–430.

———. 1998. “Outlier Analysis of Annual Retain Price Inflation: A Cross-Country Study.” *Journal of Actuarial Practice* 6: 149–172.

———. 2002. “Stochastic Investment Modelling: A Multiple Time-Series Approach.” *British Actuarial Journal* 8: 545–591.

Chan, W. S., and S. Wang. 1998. “The Wilkie Model for Retail Price Inflation Revisited.” *British Actuarial Journal* 4: 637–652.

Chang, I., G. C. Tiao, and C. Chen. 1988. “Estimation of Time Series Parameters in the Presence of Outliers.” *Technometrics* 30: 190–204.

Chatfield, C. 1993. “Calculating Interval Forecasts.” *Journal of Business and Economic Statistics* 11: 121–135.

Chen, C., and L. M. Liu. 1993. “Joint Estimation of Model Parameters and Outlier Effects in Time Series.” *Journal of the American Statistical Association* 88: 284–297.

Chen, C., and G. C. Tiao. 1990. “Random Level Shift Time Series Models, ARIMA Approximation, and Level Shift Detection.” *Journal of Business and Economic Statistics* 8: 170–186.

Chia, N. C., and A. K. C. Tsui. 2003. “Life Annuities and Compulsory Savings and Income Adequacy of the Elderly in Singapore.” *Journal of Pension Economics and Finance* 2: 41–65.

Coale, A., and G. Guo. 1989. “Revised Regional Model Life Tables at Very Low Mortality.” *Population Index* 55: 613–643.

Davison, A. C., and R. L. Smith. 1990. “Models for Exceedances over High Thresholds (with discussion).” *Journal of the Royal Statistical Society Series B* 52: 393–442.

Fishman, M. 2002. “Searching for Omega.” *The Actuary* 36(5): 2.

Fox, A. J. 1972. “Outliers in Time Series.” *Journal of the Royal Statistical Society Series B* 34: 350–363.

Gerontology Research Group. Validated Living Supercentenarians, as of September 29, 2004. Available at www.grg.org. Accessed October 1, 2004.

Heligman, L., and J. H. Pollard. 1980. "The Age-Pattern of Mortality." *Journal of the Institute of Actuaries* 107: 49–75.

Hillmer, S. 1984. "Monitoring and Adjusting Forecasts in the Presence of Additive Outliers." *Journal of Forecasting* 3: 205–215.

Himes, C. L., S. H. Preston, and G. A. Condran. 1994. "A Relational Model of Mortality at Older Ages in Low Mortality Countries." *Population Studies* 48: 269–291.

Holmes, E. C. 2004. "Enhanced: 1918 and All That." *Science* 303: 1787–1788.

Human Mortality Database. 2004. University of California, Berkeley (U.S.), and Max Planck Institute for Demographic Research (Germany). Available at www.mortality.org or www.humanmortality.de. Accessed January 25, 2004.

IMSL. 1997. *IMSL Math/Library: Fortran Subroutines for Mathematical Applications. Volume 1.* Houston: Vision Numerics.

Johansen, R. J. 2002. "Advancing the Horizon for Centenarian Survival Studies." *North American Actuarial Journal* 6(3): Guest Editorial.

Ledford, A. W., and M. E. Robinson. 1999. Discussion of A. R. Thatcher, "The Long-Term Pattern of Adult Mortality and the Highest Attained Age." *Journal of the Royal Statistical Society Series A* 162: 39.

Ledolter, J. 1989. "The Effect of Additive Outliers on the Forecasts from ARIMA Models." *International Journal of Forecasting* 5: 231–240.

Lee, R.D. and L.R. Carter.1992a. "Modelling and Forecasting U.S. Mortality." *Journal of the American Statistical Association* 87: 659–671.

Li, S. H., and W. S. Chan. 2005. "Outlier Analysis and Mortality Forecasting: The United Kingdom and Scandinavian Countries." *Scandinavian Actuarial Journal*, in press.

Liu, L. M., and G. B. Hudak. 1994. *Forecasting and Time Series Analysis Using the SCA Statistical System.* Chicago: Scientific Computing Associates.

Ljung, G. M. 1993. "On Outlier Detection in Time Series." *Journal of the Royal Statistical Society Series B* 55: 559–567.

Ljung, G. M., and G. E. P. Box. 1978. "On a Measure of Lack of Fit in Time Series Models." *Biometrika* 65: 287–303.

Muirhead, C. R. 1986. "Distinguishing Outlier Types in Time Series." *Journal of the Royal Statistical Society Series B* 48: 39–47.

National Center for Health Statistics. 2004. HIST290 Death Rates for Selected Causes by 10-Year Age Groups, Race, and Sex: Death Registration States, 1900–32, and United States, 1933–98. Available at www.cdc.gov/nchs/datawh/statab/unpubd/mortabs/hist290.htm. Accessed January 25, 2004.

National Center for Health Statistics. 2004. GMWK23R Death Rates by 10-Year Age Groups: United States and Each State, 1999–2001. Available at www.cdc.gov/nchs/datawh/statab/unpubd/mortabs.htm. Accessed January 25, 2004.

Ord, K., and S. Lowe. 1996. "Automatic Forecasting." *American Statistician* 50: 88–94.

Panjer, H. 2003. "SARS and Actuarial Modeling." *The Actuary* 37(6): 27.

Payne, R. W., P. E. Lane, P. G. N. Digby, S. A. Harding, P. K. Leech, G. W. Morgan, A. D. Todd, R. Thompson, G. Tunnicliffe Wilson, S. J. Welham, and R. P. White. 1993. *Genstat 5 Release 3 Reference Manual*. Oxford: Clarendon Press.

Pickands, J. 1975. "Statistical Inference Using Extreme Order Statistics." *Annals of Statistics* 3: 119–131.

Pollard, J. H. 1988. "On the Derivation of a Full Life Table from Mortality Data Recorded in Five-Year Age Groups." *Mathematical Population Studies* 2: 1–44.

Renshaw, A. E., and S. Haberman. 2003. "On the Forecasting of Mortality Reduction Factors." *Insurance: Mathematics and Economics* 32: 379–401.

Robine, J. M., and J. W. Vaupel. 2002. "Emergence of Supercentenarians in Low-Mortality Countries." *North American Actuarial Journal* 6(3): 54–63.

Rowley, M. C., and P. C. Bishop. 2001. "Impact of Terrorism on Future Mortality Assumptions." *Record of the Society of Actuaries* 28.

Sarony, N., W. S. Chan, and F. W. H. Chan. 2003. *Personal Injury Tables Hong Kong:*

Tables for the Calculation of Damages. Hong Kong: Sweet and Maxwell Asia.

SAS Institute. 2004. "What's New in SAS/ETS 9?" Available at http://support.sas.com/software/91x/etsugwhatsnew900.htm. Accessed August 4, 2004.

Smith, R. L. 1999. Discussion of A. R. Thatcher, "The Long-Term Pattern of Adult Mortality and the Highest Attained Age." *Journal of the Royal Statistical Society Series A* 162: 37.

Statistics Canada. 2004. "CANSIM Table 051-001—Estimates of Population by Age and Sex for Canada, the Provinces and the Territories, Annual." Available at http://cansim2.statcan.ca/cgi-win/cnsmcgi.exe?Lang=E&RootDir=CII/&ResultTemplate=CII/CII_pick&Array_Pick=1&ArrayId=510001. Accessed January 25, 2004.

Thatcher, A. R. 1999. "The Long-Term Pattern of Adult Mortality and the Highest Attained Age." With discussion. *Journal of the Royal Statistical Society Series A* 162: 5–43.

Thomas, G., G. Plu, and J. C. Thalabard. 1992. "Identification of Pulses in Hormone Time Series Using Outlier Detection Methods." *Statistics in Medicine* 11: 2133–45.

Tong, H. 1983. *Threshold Models in Non-linear Times Series Analysis*. Berlin: Springer.

Tsay, R. S. 1986. "Time Series Model Specification in the Presence of Outliers." *Journal of American Statistical Association* 81: 132–141.

———. 1988. "Outliers, Level Shifts and Variance Changes in Time Series." *Journal of Forecasting* 7: 1–20.

Tuljapurkar, S., and C. Boe. 1998. "Mortality Change and Forecasting: How Much and How Little Do We Know?" *North American Actuarial Journal* 2(4): 13–47.

Tuljapurkar, S., N. Li, and C. Boe. 2000. "A Universal Pattern of Mortality Decline in the G7 Countries." *Nature* 405: 789–792.

United Nations. 1997. "Projecting Old-Age Mortality and Its Consequences." Report on the working group, December 3–5, 1996. New York: United Nations.

U.S. Census Bureau. 2004. "Monthly Population Estimates, 1990–2000." Available at www.census.gov/popest/archives/1990s/nat_detail.html. Accessed January

25, 2004.

Vaage, K. 2000. “Detection of Outliers and Level Shifts in Time Series: An Evaluation of Two Alternative Procedures.” *Journal of Forecasting* 19: 23–37.

Vincent, P. 1951. “La mortalité des vieillards.” *Population* 6: 181–204.

Wei, W. W. S. 1990. *Time Series Analysis: Univariate and Multivariate Methods*. Redwood City, CA: Addison-Wesley.

Wikipedia. 2004a. “World War I.” Available at http://en.wikipedia.org/wiki/World_War_I#1916. Accessed May 30, 2004.

———. 2004b. “Spanish Flu.” Available at http://en.wikipedia.org/wiki/Aftermath_of_World_War_I#Influenza_pande mic. Accessed May 30, 2004.

———. 2004c. “The Korean War.” Available at http://en.wikipedia.org/wiki/Korean_War. Accessed May 30, 2004.

Wilmoth, J. R. 1993. “Computational Methods for Fitting and Extrapolating the Lee-Carter Model of Mortality Change.” Technical report. Department of Demography, University of California, Berkeley.

Wolak, D. L., D. Knapp, C. H. Meintel, and K. J. Trapp. 2003. “Pricing for Terrorism Risk in Group Life, Group Health and Group Disability Products.” *Record of the Society of Actuaries* 29: 1–30.

Appendix

Derivation of the Interval Forecast of the Actuarial Present Value of a Life Annuity Under the Lee-Carter Model

The s-period ahead forecast of the logarithm of each age-specific death rate is specified by

$$\hat{\ln}(m_{x,t+s}) = \hat{a}_x + \hat{k}_{t+s}\,\hat{b}_x, \qquad \text{(A1)}$$

where the carat indicates an estimate (for a_x and b_x) or a forecast (for k_t), and t denotes the base period. Assuming the model specification is correct, the true value of $\ln(m_{x,t+s})$ is given by

$$\ln(m_{x,t+s}) = (\hat{a}_x + \alpha_x) + (\hat{k}_{t+s} + u_{t+s})(\hat{b}_x + \beta_x) + \varepsilon_{x,t+s}, \qquad \text{(A2)}$$

where α_x and β_x are the errors in estimating a_x and b_x, u_{t+s} is the error in the s-period ahead forecast of k_t, and $\varepsilon_{x,t+s}$ is the error term that reflects all remaining age-specific influences not captured by the model. On the whole, the forecast error for the logarithm of the age-specific death rate can be expressed as

$$E_{x,t+s} = \ln(m_{x,t+s}) - \hat{\ln}(m_{x,t+s}) = \alpha_x + \varepsilon_{x,t+s} + (\hat{b}_x + \beta_x)u_{t+s} + \beta_x\,\hat{k}_{t+s}. \qquad \text{(A3)}$$

For a given x, the error terms are *assumed* to be independent of one another. In addition, following the suggestion by Lee and Carter (1992a), $E_{x,t+s}$ can be decomposed into two parts, namely, $\varphi_{x,t+s} = u_{t+s}\,\hat{b}_x$ and $\xi_{x,t+s} = \alpha_x + \varepsilon_{x,t+s} + \beta_x\,\hat{k}_{x,t+s} + \beta_x u_{t+s}$, such that $\varphi_{x,t+s}$ is perfectly correlated across x, and $\xi_{x,t+s}$ is *assumed* independent across x. Both $\varphi_{x,t+s}$ and $\xi_{x,t+s}$ are *assumed* to have mean 0. On the ground of $E_{x,t+s}$, the forecast error for the APV of a life annuity can be readily derived using the following three arguments. Actuarial notations here follow Bowers et al. (1997).

Claim 1:

The forecast error in $m_{x,t+s}$ is approximately $E_{x,t+s}\,\hat{m}_{x,t+s}$.

Proof of Claim 1:

$$\ln(m_{x,t+s}) = \hat{\ln}(m_{x,t+s}) + E_{x,t+s}$$

$$\Rightarrow m_{x,t+s} = \hat{m}_{x,t+s}\exp(E_{x,t+s}) \approx \hat{m}_{x,t+s}(1 + E_{x,t+s}) \Rightarrow m_{x,t+s} - \hat{m}_{x,t+s} \approx E_{x,t+s}\,\hat{m}_{x,t+s}.$$

Claim 2:

On a small change of m_x, say, Δm_x, q_x changes by $p_x \Delta m_x$.

Proof of Claim 2:

Assuming uniform distribution of death within each year of age, $m_x = \dfrac{q_x}{1-\frac{1}{2}q_x}$.

This implies that $\dfrac{\partial m_x}{\partial q_x} = \left(1-\dfrac{q_x}{2}\right)^2$.

Then, using a first-order approximation,

$$\begin{aligned}\Delta q_x &= \left(1-\frac{q_x}{2}\right)^2 \Delta m_x \\ &\approx (1-q_x)\Delta m_x \qquad \text{(for small values of } \Delta m_x\text{).} \\ &= p_x \Delta m_x\end{aligned}$$

Claim 3:

On a small increase in m_{x+y}, say, Δm_{x+y}, $\exists y \in \{0,1,2,\ldots\}$, $\ddot{a}_x$ deceases by $\Delta m_{x+y}\, {}_{y+1|}\ddot{a}_x$.

Proof of Claim 3:

$$\begin{aligned}\ddot{a}_x &= \ddot{a}_{x:\overline{x+y|}} + {}_{y+1}p_x v^{y+1} \ddot{a}_{x+y+1} \\ &= \ddot{a}_{x:\overline{x+y|}} + {}_{y}p_x (1-q_{x+y}) v^{y+1} \ddot{a}_{x+y+1},\end{aligned}$$

Suppose that m_{x+y} is increased by Δm_{x+y}; then by Claim 1, q_{x+y} will increase by $p_{x+y}\Delta m_{x+y}$. Let $\Delta \ddot{a}_x$ be the change of $\ddot{a}_x$ when m_{x+y} is increased by Δm_{x+y}; then

$$\begin{aligned}\ddot{a}_x + \Delta\ddot{a}_x &= \ddot{a}_{x:\overline{x+y|}} + {}_{y}p_x (1-q_{x+y} - p_{x+y}\Delta m_{x+y}) v^{y+1} \ddot{a}_{x+y+1} \\ &= \ddot{a}_x - \Delta m_{x+y}\, {}_{y+1}p_x v^{y+1} \ddot{a}_{x+y+1} \qquad\qquad \therefore \Delta\ddot{a}_x = -\Delta m_{x+y}\, {}_{y+1|}\ddot{a}_x\,. \\ &= \ddot{a}_x - \Delta m_{x+y}\, {}_{y+1|}\ddot{a}_x,\end{aligned}$$

By Claim 3, the forecast error of $\ddot{a}_x$ for time *t* + *s* can be written as $\sum_{y=0}^{\omega-x-1} - E_{x+y,t+s}\, m_{x+y,t+s}\, {}_{y+1|}\ddot{a}_y$.

Denoting **E** as expectation and using equation (29), the approximate error variance can be written as

$$\mathrm{Var}(\hat{\ddot{a}}_x) = \mathbf{E}\left[\sum_{y=0}^{\omega-x-1} \hat{m}_{x+y,t+s}\, u_{t+s}\, \hat{b}_{x+y}\, {}_{y+1|}\ddot{a}_x\right]^2 + \mathbf{E}\left[\sum_{y=0}^{\omega-x-1} (\hat{m}_{x+y,t+s}\, {}_{y+1|}\ddot{a}_x)^2 (\alpha_{x+y} + \varepsilon_{x+y,t+s} + \beta_{x+y}\hat{k}_{t+s} + \beta_{x+y}u_{t+s})^2\right]. \tag{A4}$$

Substituting variances, we get

$$\mathrm{Var}(\hat{\ddot{a}}_x) = \sigma^2_{u,t+s}\left[\sum_{y=0}^{\omega-x-1} \hat{m}_{x+y,t+s}\, \hat{b}_{x+y}\, {}_{y+1|}\ddot{a}_x\right]^2 + \sum_{y=0}^{\omega-x-1} (\hat{m}_{x+y,t+s}\, {}_{y+1|}\ddot{a}_x)^2 (\sigma^2_{\alpha,x} + \sigma^2_{\varepsilon,x,t+s} + \sigma^2_{\beta,x}\hat{k}^2_{t+s} + \sigma^2_{\beta,x}\sigma^2_{u,t+s}). \tag{A5}$$

Finally, the approximate 95 percent pointwise confidence interval is given by

$$\hat{\ddot{a}}_x \pm 2\sqrt{\mathrm{Var}(\hat{\ddot{a}}_x)}\,. \tag{A6}$$

For technical details on the computations of $\sigma^2_{u,t+s}$, $\sigma^2_{\alpha,x}$, $\sigma^2_{\beta,x}$, $\sigma^2_{\varepsilon,x,t+s}$, see Lee and Carter (1992a, p. 669).