

THE LEE-CARTER METHOD FOR FORECASTING MORTALITY, WITH VARIOUS EXTENSIONS AND APPLICATIONS

Ronald Lee*

ABSTRACT

In 1992, Lee and Carter published a new method for long-run forecasts of the level and age pattern of mortality, based on a combination of statistical time series methods and a simple approach to dealing with the age distribution of mortality. The method describes the log of a time series of age-specific death rates as the sum of an age-specific component that is independent of time and another component that is the product of a time-varying parameter reflecting the general level of mortality, and an age-specific component that represents how rapidly or slowly mortality at each age varies when the general level of mortality changes. This model is fit to historical data. The resulting estimate of the time-varying parameter is then modeled and forecast as a stochastic time series using standard methods. From this forecast of the general level of mortality, the actual age-specific rates are derived using the estimated age effects. The forecasts of the various life table functions have probability distributions, so probability intervals can be calculated for each variable and for summary measures such as life expectancy. The projected gain in life expectancy from 1989 to 1997 matches the actual gain very closely and is nearly twice the gain projected by the Social Security Administration's Office of the Actuary. This paper describes the basic Lee-Carter method and discusses the forecasts to which it has led. It then discusses extensions, applications, and methodological improvements that have been made in recent years; considers shortcomings of the method; and briefly describes how it has been used as a component of more general stochastic population projections and stochastic forecasts of the finances of the U.S. Social Security system.

INTRODUCTION

Over the past ten years, a number of approaches have been developed for forecasting mortality using stochastic models (see Alho 1990, 1992; Alho and Spencer 1985, 1990; McNown and Rogers 1989, 1992; Bell and Monsell 1991; and Lee and Carter 1992). For a discussion of these methods, see Lee (1999) and Lee and Skinner (1996). In this paper, I will describe the Lee-Carter method; discuss extensions and applications; consider shortcomings of the method; and briefly describe how it has been used as a component of more general stochastic population projections and stochastic forecasts of the finances of the U.S. Social Security system.

Lee and Carter have developed a new method for the extrapolation of trends and age patterns in mortality. While it has some advantages over other extrapolative methods, it also shares the fundamental weaknesses of extrapolation: historical patterns may not hold for the future, and structural changes may therefore be missed. No attempt is made to incorporate current knowledge about actual and prospective advances in medicine, changing lifestyles, or new diseases such as AIDS (see Manton, Stallard, and Singer 1992 for an alternative approach that explicitly introduces such information). Experimental analysis and forecasts within the period of data availability show that the procedure performs quite well over the period 1900–1990: age patterns have been quite stable and the trend in our fundamental time-varying parameter has been surprisingly linear. However, if we were to go back much before 1900, the linearity of change would

*Ronald Lee, Ph.D., is Professor in the Departments of Demography and Economics at the University of California, 2232 Piedmont Ave., Berkeley, California 94720, e-mail, rlee@demog.berkeley.edu.

be violated: sustained mortality decline is primarily a twentieth century phenomenon in the developed countries.

Some methods forecast life expectancy (e_0) directly, and then explicitly build into their forecasts some degree of deceleration. For example, assuming an upper limit to attainable life expectancy has this effect. Given the forecasted level of life expectancy, one can then use assumptions to distribute mortality by age. It is true that life expectancy gains have tended to decelerate in many populations, including that of the U.S., and this must be a feature of plausible forecasts. However, the deceleration occurs for mechanical reasons deriving from the nature of life expectancy as a highly nonlinear summary measure of age-specific mortality. When mortality declines from a high level, many of the deaths averted are those of children who then gain many remaining years of life. When mortality declines from a low level, as in the U.S. today, most of the deaths averted are those of old people who then gain relatively few years of life. Even if each age-specific death rate declines at a constant exponential rate, life expectancy will increase at a decelerating pace.¹ The Lee-Carter approach models the rates of decline of individual death rates, and so deceleration of e_0 occurs naturally, with no imposition of an upper limit.²

MODEL

The model generates one-parameter families of age schedules for fertility and mortality, in the sense that variations in one parameter generate the entire range of schedules in the family. However, to express a single schedule requires a number of age-specific coefficients equal to twice the number of age groups. Different values of these coefficients define different fam-

ilies. Let $m_{x,t}$ be the central death rate for age x at time t .³ The model used for mortality is:

$$\ln(m_{x,t}) = a_x + b_x k_t + e_{x,t}.$$

Here the a_x coefficients describe the average shape of the age profile, and the b_x coefficients describe the pattern of deviations from this age profile when the parameter k varies. Parts of the Coale-Demeny-Vaughan (1983) life table system are built on a model somewhat like this (where k is taken to equal e_{10}), and preliminary versions of the U.N. model life table system also used an approach of this sort. Using exploratory data analysis, Gomez de Leon (1990) showed that a model of this sort gives the most satisfactory fit to the historical Norwegian data, out of a large class of simple models. It is our hope, of course, that the error term, $e_{x,t}$, is well-behaved and of relatively small variance. The hope is that most of the variance over time at any given age will be “explained” by the k parameter, and that the residual variance will be white noise. In the U.S. from 1933 to 1987, the model in fact accounted for over 97% of the temporal variance in mortality rates through ages 80–84.⁴

FITTING THE MODEL

This model cannot be fit by simple regression, because there is no observed variable on the right-hand side. Nonetheless, a least-squares solution exists and can be found using the first element of the singular value decomposition (or SVD; see Lee and Carter 1992 for details) or principal components (see Bell and Monsell 1991). On inspection, we can see that the solution cannot possibly be unique, however. To distinguish a unique solution, impose the further conditions that the sum of the b_x coefficients equals 1.0, and that the sum of the k_t parameters equals zero. Under these assumptions, it can be seen that the a_x coefficients must be simply the average values over time of the $\ln(m_{x,t})$

¹Keyfitz (1985:62–64) shows that if mortality at every age declines by the proportion δ , then e_0 rises by the proportion $H\delta$, where H is known as the entropy of the life table, and is 0.1 to 0.2 in modern low-mortality populations. As mortality falls, so does H , so that e_0 rises more and more slowly.

²We use the Coale-Guo (1989) method to model and forecast mortality from 85–89 to 105–109, and in that method it is assumed that nobody survives to 110. This assumption is not essential to our approach, however, and does not much affect the deceleration of life expectancy gains.

³By modeling the logarithm of $m_{x,t}$ we insure that the age-specific death rates themselves will never be negative, which otherwise occurs promptly in forecasts. Nothing insures that they will not exceed unity, however. In practice this is not a problem. While the possibility could be avoided by modeling the logit of the death rates, $\ln[m_{x,t}/(1 - m_{x,t})]$, in this case a linear trend in k would not imply a constant geometric rate of decline for each age-specific death rate. We prefer the alternative of modeling the force of mortality rather than m . The force of mortality can take on any positive value, and the implied m 's and q 's would always be between zero and unity. We have not done this yet, but it appears to be the most natural way to proceed.

⁴That is, the model accounted for over 97% of the variance in the $\ln(m_{x,t})$ matrix that remained after subtracting age group means.

values for each x . Wilmoth (1993) has developed a superior method for fitting the model, which will be discussed later.

The model was fit to U.S. data from 1933 to 1987 (Lee and Carter 1992), to Chilean data from 1952 to 1987 (Lee and Rofman 1994), and to Canadian data from 1922 to 1995 (Lee and Nault 1993). (For a treatment of male and female mortality differences, see Carter and Lee 1992. For an analysis of black-white differences by sex, see Carter 1996b). Figure 1 shows the estimated values of a_x and b_x coefficients for the U.S. The a_x coefficients, as noted, are just the average values of the logs of the death rates. Not surprisingly, the Chilean coefficients lie above those of the U.S. at all ages except the highest, reflecting the fact that mortality was higher, on average, in Chile from 1952 to 1987 than in the U.S. from 1933 to 1987.⁵ The b_x coefficients describe the relative sensitivity of death rates to variation in the k parameter. It is also not surprising that sets of coefficients for the U.S. and Chile look quite similar. Because of the normalization, their absolute levels have no particular meaning. With N age groups, if $b_x = b_y = 1/N$ for all x, y , then all the rates would move up and down proportionately, maintaining constant ratios to one another. However, it can be seen that, in fact, some ages are much more sensitive than others. Generally speaking, the younger the age, the greater its sensitivity to variation in the k parameter. The exponential rate of change of an age group's mortality is proportional to the b_x value: $d \ln(m_{x,t}) / dt = (dk_t / dt) b_x$. If k declines linearly with time, then dk_t / dt will be constant and each m_x will decline at its own constant exponential rate.

SECOND STAGE ESTIMATION OF K


At this point, we could proceed directly to the next step of modeling the k parameter as a time series process. Instead, we make a second stage estimate of k by finding that value of k which, for a given population age distribution and the previously estimated coefficients a_x and b_x , produces exactly the observed number of total deaths for the year in question. That is, we search for k_t such that:

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


where D_t is total deaths in year t , and $N_{x,t}$ is the population age x in year t .

⁵The Chilean data treat 65+ as the open age group. The Coale-Guo (1989) method was used to construct age-specific mortality estimates at older ages. In the U.S., 85+ is the open age group.

Figure 1
Estimated Values of a_x and b_x Coefficients for the U.S. and Chile
Comparison of Estimated a_x Coefficients for the U.S. and Chile



Comparison of Estimated b_x Coefficients for the U.S. and Chile



There are several advantages to making a second stage estimate of the k parameter in this way. First, this guarantees that the life tables fitted over the sample years will fit the total number of deaths and the population age distributions. Because the first stage estimation was based on logs of death rates rather than the death rates themselves, sizable discrepancies can occur between predicted and actual deaths. Second, in this way, the empirical time series of k can be extended to include years for which age-specific data on mortality are not available, because the second stage estimate of k yields an indirect estimate of mortality. For the U.S., this allows the base year of the forecast to be brought forward by two or three years, because of the long time lag in publishing age-specific rates. For Chile, it allows us to fill in several gaps in the middle of the time series of mortality. For other developing countries with less complete data (such as China), it permits many kinds of indirect estimation of mortality, even when data are restricted to a single initial population age distribution and annual totals of births and deaths.

While this two-stage method works quite well, Wilmoth (1993) has developed superior one-stage methods. The two stages are collapsed into one by using weighted SVD to fit the model, with the numbers of deaths at each age used as weights. In principle, this method can also provide estimates of the standard errors of the a_x and b_x coefficients. In the two-stage method, the standard errors for the b_x coefficients can be estimated using the bootstrap (see Lee and Carter 1992), but this is a cumbersome process. Wilmoth has also developed a maximum likelihood method.

VITAL RATES AS STOCHASTIC PROCESSES

The next step is to model k as a stochastic time series process. This is done using standard Box-Jenkins procedures. In most applications so far, k_t is well-modeled as a random walk with drift: $k_t = c + k_{t-1} + u_t$. In this case, the forecast of k changes linearly and each forecasted death rate changes at a constant exponential rate. However, sometimes a model of this general form, but with an added moving average term or autoregressive term, is superior. In this case the pattern of change is somewhat different.

Note that each of the m_x values is now itself modeled as a stochastic process driven by the process k . Note also that, if we ignore the error term $e_{x,t}$ (which we hope is relatively unimportant), the variations in the $\ln(m_x)$ values will be perfectly correlated with one

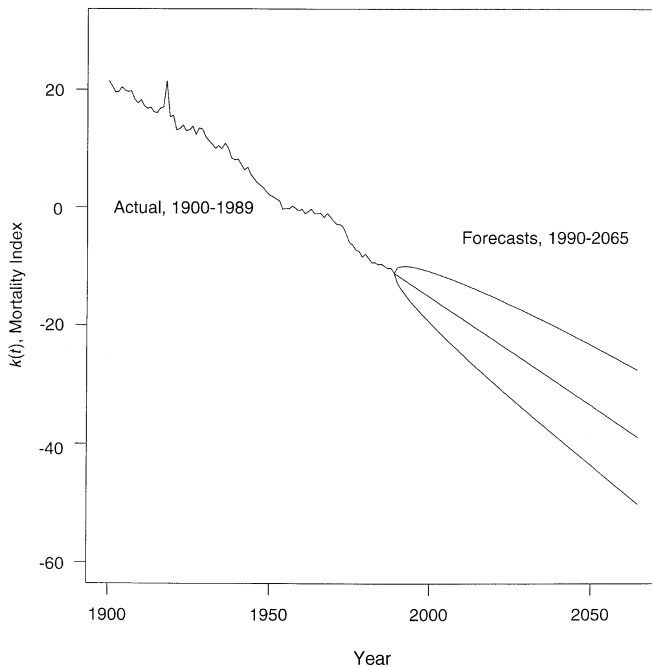
another, because all are linear functions of the same time-varying parameter k . This is a very convenient feature of the model, for it means that we can calculate the probability bounds on all (period) life table functions directly from the probability bounds on the forecasts of k , without having to worry about cancellation of errors. For a discussion of errors arising from $e_{x,t}$ and from the estimation of a_x and b_x coefficients, which are ignored here, see the appendix to Lee and Carter (1992).

FORECASTS

We can now use the fitted time series model for k to forecast it over the desired time period. Figure 2 shows past values of k for the U.S. from 1900 to 1989 and their forecasts from 1990 to 2065. Note that the estimated values of k over the base period change in a linear fashion. In fact, the change in k over the first half of the period almost exactly equals its change in the second half. This contrasts with our usual understanding that the pace of mortality decline has decelerated over time, an understanding based on the trends in life expectancy at birth, which we will examine in a moment. The approximate linearity of k in the base period is a great advantage from the point of view of forecasting. Long-term extrapolation is always a hazardous undertaking, but it is less so when supported in this way by the regularity of change in a ninety-year empirical series. Also note that, aside from the influenza epidemic of 1918, the variability of the series is similar throughout the period. This is also a desirable feature for forecasting purposes.

Figure 2 shows the point forecasts, which are essentially linear extrapolations of the base period series. Our analysis of Chilean data showed a similar result, as did Gomez de Leon's (1990) analysis of Norwegian mortality data, which took a slightly different approach. 95% probability intervals are also shown for the forecast of k . The figure shows intervals which reflect only the error term in the random walk, which expresses the innovation error. However, these could be augmented to include uncertainty about the rate of decline in k , which is itself estimated (see Lee and Carter 1992). The statistical model included a dummy variable to remove the influence of the 1918 influenza epidemic. One might prefer to treat this epidemic as representative of continuing epidemic risks in the future. In more recent forecasts, we have incorporated a 1 in 97 chance each year of a shock to k of the size of the shock in 1918, where 97 is the number of years

Figure 2
Mortality Forecast from 1900–1989 to 2065,
With 95% Probability Interval
Model is (0,1,0) with a dummy for flu.



in the sample period. The influence on the forecasts is barely discernible.

The next step is to convert the forecasts of k into forecasts of life table functions, given the previously estimated age specific coefficients a_x and b_x , using the earlier equation for $\ln(m_{x,t})$. Once the implied forecasts of $m_{x,t}$ have been recovered in this way, any desired life table function can be calculated. For period life table functions, the probability intervals can be found directly from the intervals on k . To find probability intervals for cohort life table functions, we would have to take into account the autocovariance structure of errors in k , which is less straightforward.


Figure 3 displays actual (fitted) base period values of life expectancy at birth along with the forecasts derived from the forecasts of k in Figure 2. The figure also shows the high, medium, and low projections by the U.S. Bureau of the Census (1989) and the Social Security Administration's Office of the Actuary (Bell, Wade, and Goss 1992) which have similar base line dates. There are several points to note. First, we predicted an increase in life expectancy of 10.4 years, from 75.7 to 86.1, between 1989 and 2065 (Lee and Carter 1992). This is nearly twice as large as the gain forecast by the Office of the Actuary and the Bureau

of the Census (from 75.08 to 80.61 for the Actuary, or 5.5 years). The difference translates into substantially higher forecasts of the elderly population. Second, although we projected a linear trend in k , the forecast for life expectancy is for growth at a slowing pace. The difference is due to the decreasing entropy of the survival curve, as discussed earlier. Third, the probability interval is fairly tight, even though the plot includes uncertainty in the estimated drift term in addition to the innovation error. We see here that time series methods need not lead to very wide probability intervals, with little information after a couple of decades; indeed, the Lee-Carter intervals have been criticized as being implausibly narrow. Note, however, that the intervals on the forecasts of k were not particularly narrow. We believe that the narrowness of the forecasts of life expectancy arise from the low entropy of the survival curve when life expectancy reaches high levels: even sizable variations in k have little effect on the level of life expectancy, because mortality has become so concentrated at the older ages.

How has this forecast performed compared to actual life expectancy in the intervening years? First, it is now clear that an attempt to estimate indirectly the level of life expectancy for the jump-off year of 1989, based on methods described earlier in this paper, led to an error of 0.58 years.⁶ The actual e_0 was 75.08 in 1989 (U.S. NCHS 1998b), while Lee and Carter (1992) estimated it to be 75.66. Given the random walk specification, the influence of this baseline error of 0.6 years persists in the forecast (a procedure for avoiding this problem is described later in the paper). If we concentrate on the projected gain in life expectancy rather than on the projected levels, we find that Lee and Carter projected e_0 to rise from 75.7 in 1989 to 77.0 in 1997, a gain of 1.3 years. The actual gain has been 1.4 years, from 75.1 to 76.5 in 1997 (U.S. NCHS 1998a). The agreement with the Lee-Carter gain forecast is very close. By contrast, the Office of the Actuary (Bell, Wade, and Goss 1992) projected a gain of 0.725 years over the same period, or about half the actual; 76.5 is not reached in this projection until 2002 or 2003. Similarly, the Lee-Carter forecasts of the age pattern of decline over this period are very close to the actual, although over earlier parts of the comparison period this was not so. In summary, the basic Lee-Carter method successfully predicted e_0

⁶At the time of our forecast, we had age-specific mortality data available through 1987, and we used our model to estimate indirectly the level of mortality for 1988 and 1989 based on the aggregate number of deaths and the population age distribution.

Figure 3
 e_0 Forecasts with Intervals: Lee-Carter, Census Bureau, and Social Security Administration for the U.S., Sexes Combined, 1990–2065



Note: The plot shows the median and 95% confidence interval for Lee-Carter and the high-middle-low forecasts of the Census Bureau and the SSA Office of the Actuary.

gains over the period since it was published, although an erroneous jump-off life expectancy level introduced a persistent error in the level of e_0 that Lee-Carter forecast.

SOME PROBLEMS WITH THE METHOD

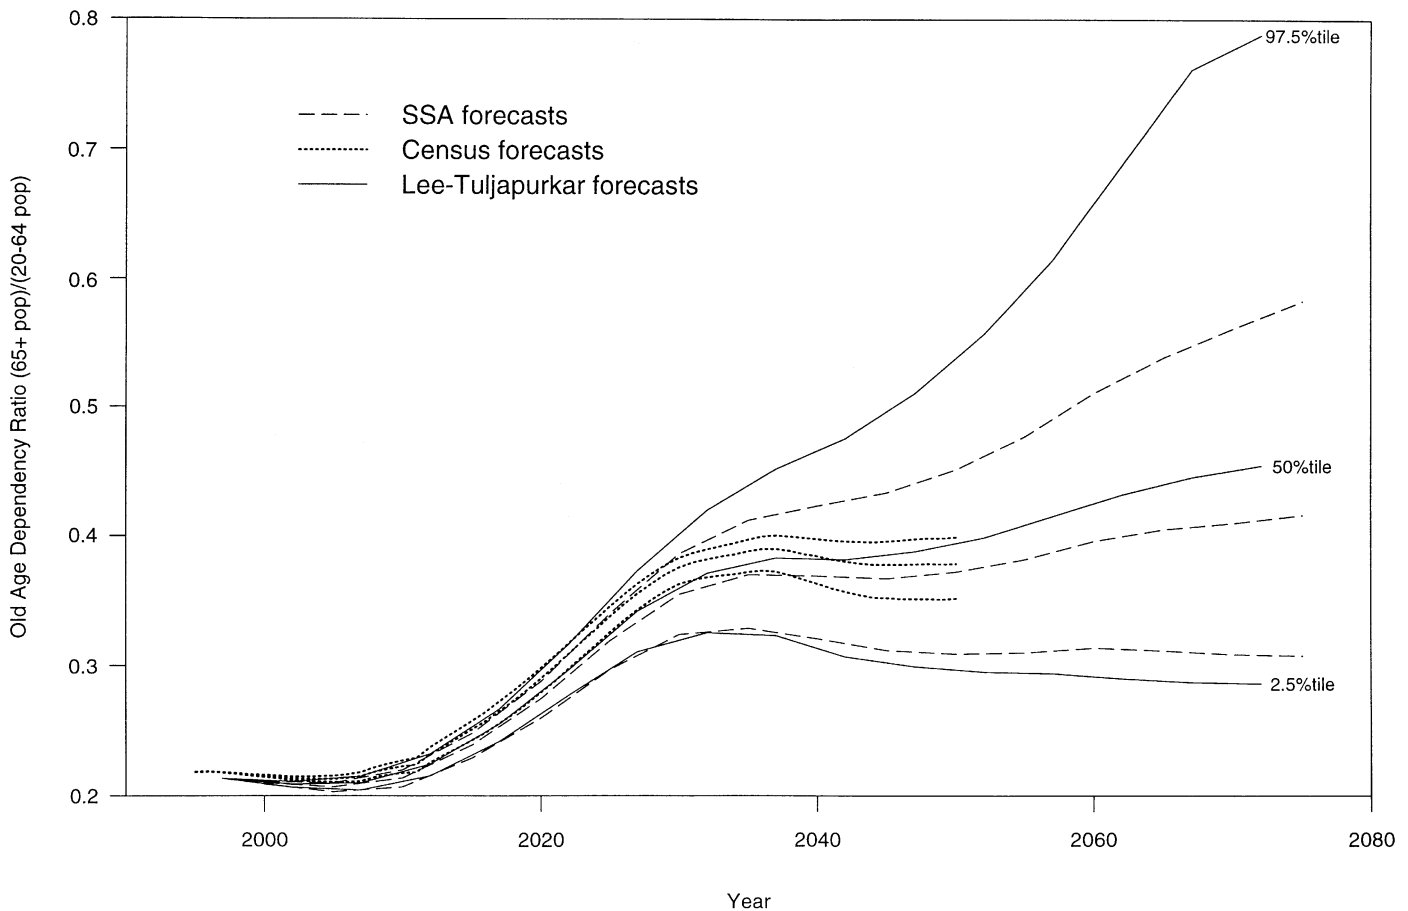
The Lee-Carter method is a useful and appropriate way to extrapolate historical trends in the level and age distribution of mortality. However, extrapolation may not always be a sensible procedure to employ. Perhaps it will be useful to discuss some of the problems and limitations of the method.

- As mentioned earlier, mortality in the U.S. has not always declined along the path represented by the plot of k in Figure 2. If it had, then by retrojection we would arrive at impossibly high levels of mortality early in the nineteenth century. So we know that the time series behavior of k during the period of observation from 1900 to 1996 is not typical of the whole historical experience and cannot possibly reflect a fundamental property of mortality change over time. Why, then, should we expect this pattern to hold over the next century? Perhaps the institutionaliza-

tion of biomedical research provides some reason to believe that twentieth-century trends will continue in the future.

- The method assumes a certain pattern of change in the age distribution of mortality, such that the rates of decline at different ages (given by $b_x(dk_x/dt)$) always maintain the same ratios to one another over time. But in practice, the relative speed of decline at different ages may vary. For example, Horiuchi and Wilmoth (1995) point out that in Sweden, mortality rates at old age used to decline more slowly than at other ages, but that in recent decades they have come to decline more rapidly than at other ages. In the U.S., it appears that there has been a slowdown in mortality declines for ages 5 to 50 relative to the older and younger ages. The method cannot take such shifts in pattern into account. It is possible to modify the method to accommodate changes in age pattern, but it is not known whether it would perform well as a forecasting method when so modified.
- The method does not readily accommodate extraneous information about future trends. Perhaps it is best viewed as providing a kind of baseline forecast of what would happen if present trends were to continue. Then, if there were a compelling reason to expect long-term future trends to be more or less rapid than in the past, the baseline forecast could be suitably modified. In our view, however, such compelling reasons are seldom available.
- The method most easily incorporates forecast uncertainty arising from uncertainty in the forecast of k , the mortality index. It is also possible without much difficulty to incorporate uncertainty about the estimated trend in mortality, arising from the uncertain drift coefficient in the Lee and Carter implementation (1992). However, uncertainty arising from errors in the estimation of the b_x coefficients is quite difficult to estimate and to incorporate, as is uncertainty arising from errors in the fit of the basic model to the actual matrix of age-specific mortality rates over time. Further errors arise from violation of the assumption that fitting errors are uncorrelated across age. Some, but not all, of these issues are discussed in the appendix to Lee and Carter (1992).
- Because the probability intervals do not reflect uncertainty about whether the model specification is correct, nor uncertainty about whether the future will look like the past, some people believe that they are too narrow and that they understate the uncertainty about future levels of life expectancy. While

Figure 4
Stochastic Forecasts of the Old-Age Dependency Ratio, U.S., 1997–2075



Note: Based on 1,000 stochastic simulations, using Lee-Carter stochastic mortality model fit on U.S. data, 1900 to 1996, and Lee-Tuljapurkar stochastic fertility model (see Lee and Tuljapurkar, 1998a) with long run mean constrained to 1.9 children per woman, with deterministic immigration as in Office of the Actuary's intermediate assumption. Census forecasts are based on U.S. Bureau of the Census (1996). SSA forecasts are based on F. Bell (1997).

some people expect major breakthroughs in medical technology that will accelerate mortality decline, others expect drug-resistant strains of diseases to proliferate and cause a slowing of decline.


- Implausible sex differentials: In the application to Canadian data, the sex difference in forecasted mortality grew in an implausible way, whether the sexes were forecasted independently or with identical k series. Presumably this reflected divergent trends in the historical data, for example, the earlier adoption of smoking by males. Most analysts expect sex differences to narrow in the U.S. and Canada.

EXTENSIONS

There are many extensions of this basic method. There will be space here for only a brief description of each.

- Disaggregation by sex (see Carter and Lee 1992): There are a number of possible approaches. One approach, of course, is simply to treat the male and female forecasts as two separate applications of the basic approach. It is tempting, however, to try to take advantage of similarities between the two to reduce the dimensionality of the model, and to impose some coherence on the two sets of forecasts. One way to do this is to treat the male and female rates as if they were just different portions of the same set of age-specific rates for a single population,

Figure 5
Stochastic Forecasts of the Total Dependency Ratio, United States, 1997–2075

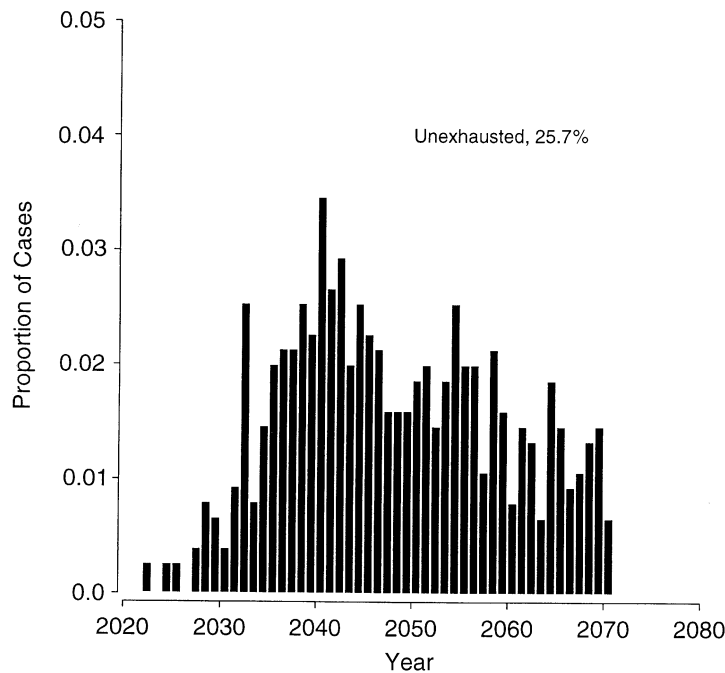


Note: Based on 1,000 stochastic simulations, using Lee-Carter stochastic mortality model fit on US data, 1900 to 1996, and Lee-Tuljapurkar stochastic fertility model (see Lee and Tuljapurkar, 1998a) with long run mean constrained to 1.9 children per woman, with deterministic immigration as in Office of the Actuary's intermediate assumption. Census forecasts are based on U.S. Bureau of the Census (1996). SSA forecasts are based on F. Bell (1997).

with $2N$ age groups, if the actual number of age groups is N . In this way, one estimates and forecasts a single time series for k which is used to drive both male and female mortality. The two can have different age-shapes, and can decline at different rates, because they have different a_x and b_x coefficients. Further simplifications of structure can be achieved by constraining the a_x coefficients to be identical while estimating different b_x coefficients, or conversely. These constraints on the basic model should be carefully tested and evaluated.

- Geographic disaggregation (see Lee and Nault 1993): Exactly the same issues and options arise as for disaggregation by sex when dealing with a set of regional populations comprising a national total. An application was made to the 12 provinces in Canada.
- Disaggregation by cause: Similar issues arise in Wilmoth (1995), who carried out a cause-disaggregated application to Japanese data (1996). He points out that a cause-disaggregated forecast will always yield higher mortality in the future than will an aggregate forecast, when trend extrapolation is used as the method. This is because the most slowly declining cause of death will come to dominate in the long run (see Wilmoth 1995). Often, some causes of death will actually have increasing rates, and they will be forecast to increase indefinitely.
- Lower bounds for death rates: Wilmoth (through personal communication with the author) has also

Figure 6
Histogram of Dates of Exhaustion with an Immediate Tax Increase of 2%



suggested that it would be a simple matter to have each age-specific death rate decline toward a lower bound greater than zero, by subtracting the bound before modeling and forecasting the death rates, and then adding it back in.

- Matching latest death rates: Lee and Carter (1992) mention that rather than estimating α_x coefficients as the average of the $\ln(m_{x,t})$, it would be possible to set the α_x coefficients equal to the logarithms of the most recent death rates. This assures that the first year of the forecast will match up smoothly and closely with the most recently observed death rates. William Bell (1997) tests a variety of methods for forecasting mortality, including both the original Lee-Carter method and this variation, and concludes that this variation outperforms all other methods. Perhaps this is more likely to be true over shorter horizons, such as in Bell's analysis. The advantage of the original method is the possibility of avoiding peculiarities of any particular year's mortality rates. My recommendation is that the model be fitted in the way described, but that forecasts should take the most recently observed age-specific death rates as initial values. Then forecasts would be given by $\ln(m_{x,t+s}) = \ln(m_{x,t}) - b_x(k_{t+s} - k_t)$. This assures that the baseline values exactly match the observed age-specific death rates and life expectancy, and that the

forecasts of age-specific rates progress smoothly from these values.

- Other model specifications: It is possible to complicate the model in various ways, for example by including higher order effects (see for example Bell and Monsell 1991). Gomez de Leon (1990) tests other possibilities using Norwegian data and concludes that this specification works best. However, by adding higher-order effects and cohort-specific effects, it will generally be possible to improve the fit within the sample. The question is whether such additional effects would represent fundamental and enduring aspects of mortality patterns such that the forecasts would be improved.
- Variable trends: Carter (1996a) develops a method in which the drift or trend term in the forecasting equation for k is itself allowed to be a random variable. This is done using standard state-space methods for modeling time series. Somewhat surprisingly, the forecast and probability intervals remained virtually unchanged under this approach.
- Leader countries: Wilmoth (1996) has suggested that lagging countries, with life expectancy below that of the leading country (now Japan), can achieve more rapid gains by borrowing more advanced technology or other health-related practices. However, once they have caught up with the leader they would

then be expected to progress more slowly. Therefore, it might be a mistake to extrapolate a mortality trend which would take a population past the projected mortality of the leader population. He applies this approach to forecasts for Japan and Sweden.

STOCHASTIC FORECASTS OF POPULATION AND OF SOCIAL SECURITY SYSTEM FINANCES

One advantage of this approach is that it forecasts the joint probability distributions of age-specific death rates, and these can be viewed as components of stochastic Leslie matrices. When combined with stochastic fertility forecasts (see Lee 1993), it is then possible to produce fully stochastic population forecasts. This is done for the U.S. in Lee and Tuljapurkar (1994). Figure 4 plots more recent stochastic forecasts of the old-age dependency ratio, showing the median forecast along with the 95% probability interval. For comparison, the high, medium, and low forecasts by the Office of the Actuary (Felicite Bell 1997) and the U.S. Census Bureau (1996) are also shown. Note that the median Lee-Tuljapurkar forecast is somewhat higher than that of the Actuary due to the greater longevity gains forecast by the Lee-Carter method. Note also that while the lower 2.5% bracket is close to the Actuary's low-cost projection, the upper 2.5% bracket is far higher. Also note that the Census high-low range is very narrow. Census combines high fertility and low mortality in their high-cost scenario, and low fertility and high mortality in their low-cost scenario, to generate a wide range for population size and population growth rates. However, these combinations minimize the range for old-age dependency, because high fertility and low mortality work in opposite directions on the ratio. The Census forecasters realize this, of course, and provide alternative scenarios to generate a plausible range of uncertainty for the old-age dependency ratio. The Office of the Actuary, seeking a plausible range of uncertainty for this ratio, constructs its scenarios differently, combining low fertility and low mortality in their high-cost scenario, and high fertility and high mortality in their low-cost scenario.

However, it is not possible to avoid the inherent inconsistencies in the scenario-based approach to assessing uncertainty, as shown by Figure 5 for the total dependency ratio. Note first that the Census middle forecast for the ratio is substantially higher than the others, because Census assumes the total fertility rate will rise to 2.245 children per woman, whereas Lee-

Tuljapurkar follow Social Security in assuming a long-term mean of 1.9 children per woman (see Lee and Tuljapurkar, 1998c, for a discussion of these assumptions). The point to note in the present context, however, is that the Census high-low interval is fairly wide, while now the Office of the Actuary's interval is implausibly narrow. The stochastic forecast automatically provides consistent estimates of uncertainty, regardless of the variable examined. Some readers may believe the Lee-Tuljapurkar intervals to be too wide. In assessing their width, it must be kept in mind that they are designed to contain annual variations, not just long-term trends. For a discussion of this issue, see Lee and Tuljapurkar, 1998c.

These stochastic population forecasts can then be used as the basis of stochastic forecasts of the finances of the Social Security system, or of other aspects of the federal budget. For example, the Congressional Budget Office, in its long-term forecasts published in 1996, 1997, and 1998 (Congressional Budget Office 1998), uses the Lee-Tuljapurkar stochastic population projections to prepare stochastic projections of the federal budget balance out to 2050. Lee and Tuljapurkar (1998a, b, c) use their stochastic population projections, together with time series models of the real interest rate and productivity growth rates, to generate stochastic forecasts of the finances of the Social Security system under varying policy assumptions. As an illustration, Figure 6 displays the probability distribution of the dates of exhaustion of the OASDI Trust Fund, conditional on the assumption that the payroll tax rate for OASDI is immediately raised by 2 percentage points, from 12.4% to 14.4%. According to our forecasts, there would still be a 75% chance of fund exhaustion by 2070. Lee and Tuljapurkar (1998c) present probability distributions for outcomes under other assumptions, such as current policy, rising retirement ages, and investment of differing fractions of the trust fund in equities.

SUMMARY AND CONCLUSION

The Lee-Carter method is based on a simple representation of year-to-year variations in the set of age-specific death rates in terms of a single time-varying parameter. The time series of estimated values of this parameter is then modeled and forecasted using statistical time series methods. From these forecasts, the probability distributions for forecasts of age-specific death rates and related variables such as life expectancy are calculated. The published forecast of a 1.3-year gain in life expectancy from 1989 to 1997

matches closely the 1.4 years actually realized, and over this period the pattern of declines by age is also very closely matched by the forecasts. Alternative methods have been developed for estimating the basic model, and for modeling and forecasting the time series. There have been a number of extensions of the method, including the development of coherent forecasts by sex and by race, and forecasts for regions comprising a national system. Problems and shortcomings of the method have been discussed.

The mortality model has been used together with a similar fertility model and deterministic migration assumptions to generate stochastic forecasts of the population and its components. These stochastic population forecasts, in turn, have been used as the key component of stochastic forecasts of the finances of the U.S. Social Security system. The stochastic forecasts avoid some of the problems inherent to the use of the classic scenario method for representation of forecast uncertainty.

REFERENCES

- ALHO, JUHA M. 1990. "Stochastic Methods in Population Forecasting," *International Journal of Forecasting* 6:521-30.
- ALHO, JUHA M. 1992. "The Magnitude of Error Due to Different Vital Processes in Population Forecasts" *International Journal of Forecasting* 8, no. 3 (November):301-14.
- ALHO, JUHA M., AND SPENCER, BRUCE D. 1985. "Uncertain Population Forecasting," *Journal of the American Statistical Association* 80, no. 390 (June):306-14.
- ALHO, JUHA M., AND SPENCER, BRUCE D. 1990. "Error Models for Official Mortality Forecasts," *Journal of the American Statistical Association* 85:609-16.
- BELL, FELICITIE C. 1997. "Social Security Area Population Projections: 1997," *Actuarial Study No. 112*. Washington, D.C.: Office of the Actuary, Social Security Administration.
- BELL, FELICITIE C., WADE, ALICE H., AND GOSS, STEPHEN C. 1992. "Life Tables for the United States Social Security Area: 1900-2080," *Actuarial Study No. 107*. Washington, D.C.: Office of the Actuary, Social Security Administration.
- BELL, WILLIAM R., AND MONSELL, BRIAN C. 1991. "Using Principal Components in Time Series Modeling and Forecasting of Age-Specific Mortality Rates," paper presented at the 1991 Annual Meetings of the Population Association of America, Washington, D.C. (March 21-23).
- BELL, WILLIAM R. 1997. "Comparing and Assessing Time Series Methods for Forecasting Age Specific Demographic Rates," *Journal of Official Statistics* 13:279-303.
- BOARD OF TRUSTEES, FEDERAL OLD-AGE AND SURVIVORS INSURANCE AND DISABILITY INSURANCE TRUST FUNDS. 1992. *The 1992 Annual Report of the Board of Trustees of the Federal Old-Age and Survivors Insurance and Disability Insurance Trust Funds*. Washington, D.C.: Government Printing Office.
- CARTER, LAWRENCE. 1996a. "Forecasting U.S. Mortality: A Comparison of Box-Jenkins ARIMA and Structural Time Series Models," *The Sociological Quarterly* 37, no. 1:127-44.
- CARTER, LAWRENCE. 1996b. "Long-Run Relationships in Differential U.S. Mortality Forecasts by Race and Gender: Non-Cointegrated Time Series Comparisons," presented at the Annual Meeting of the Population Association of America, New Orleans, May 9-11.
- CARTER, LAWRENCE, AND LEE, RONALD D. 1992. "Modeling and Forecasting U.S. Mortality: Differentials in Life Expectancy by Sex," *International Journal of Forecasting* 8, no. 3 (November):393-412.
- COALE, ANSLEY, DEMENY, PAUL, AND VAUGHAN, BARBARA. 1983. *Regional Model Life Tables and Stable Populations*, 2nd ed. New York: Academic Press.
- COALE, ANSLEY, AND GUO, GUANG. 1989. "Revised Regional Model Life Tables at Very Low Levels of Mortality," *Population Index* 55, no. 4 (Winter):613-43.
- CONGRESSIONAL BUDGET OFFICE OF THE UNITED STATES. 1998. *Long-Term Budgetary Pressures and Policy Options*. Washington, D.C.: U.S. Government Printing Office.
- GOMEZ DE LEON, JOSE. 1990. "Empirical DEA Models to Fit and Project Time Series of Age-Specific Mortality Rates," unpublished manuscript. Central Bureau of Statistics, Norway.
- HORIUCHI, SHIRO, AND WILMOTH, JOHN R. 1995. "The Aging of Mortality Decline," presented at the annual meeting of the Population Association of America, San Francisco, April 6-8.
- <http://www.census.gov/population/www/projections/nasrh.html>: U.S. Bureau of the Census, Population Division, "Population Projections of the United States by Age, Sex, Race, and Hispanic Origin: 1995 to 2050." Machine-readable data file; released May 1996.
- KEYFITZ, NATHAN. 1985. *Applied Mathematical Demography*. 2nd ed. New York: Springer-Verlag.
- LEE, RONALD D. 1993. "Modeling and Forecasting the Time Series of U.S. Fertility: Age Patterns, Range, and Ultimate Level," *International Journal of Forecasting* 9:187-202.
- LEE, RONALD D. 1999. "Probabilistic Approaches to Population Forecasting," in *Rethinking Population Projections*, Supplement to vol. 25 of *Population and Development Review*, eds. Wolfgang Lutz, James Vaupel, and Dennis Ahlburg, pp. 156-90.
- LEE, RONALD D., AND CARTER, LAWRENCE. 1992. "Modeling and Forecasting the Time Series of U.S. Mortality," *Journal of the American Statistical Association* 87, no. 419 (September):659-71.
- LEE, RONALD D., CARTER, LAWRENCE, AND TULJAPURKAR, SHRIPAD. 1995. "Disaggregation in Population Forecasting: Do We Need It? and How To Do It Simply," *Mathematical Population Studies* 5, no. 3:217-34.
- LEE, RONALD D., AND NAULT, FRANCOIS. 1993. "Modeling and Forecasting Provincial Mortality in Canada," presented at the World Congress of the International Union for the Scientific Study of Population, Montreal.

- LEE, RONALD D., AND ROFMAN, RAFAEL. 1994. "Modelacion y Proyeccion de la Mortalidad en Chile," *NOTAS* 22, no. 59: 182–213 (Also available in English from the authors, titled "Modeling and Forecasting Mortality in Chile").
- LEE, RONALD D., AND SKINNER, JONATHAN. 1996. "Assessing Forecasts of Mortality, Health Status, and Health Costs During Baby-Boomers' Retirement," in *Assessing Knowledge of Retirement Behavior*, ed. C. Citro. Washington, D.C.: National Academy of Sciences, pp. 195–243.
- LEE, RONALD D., AND TULJAPURKAR, SHRIPAD. 1994. "Stochastic Population Projections for the United States: Beyond High, Medium and Low," *Journal of the American Statistical Association* 89, no. 428 (December):1175–89.
- LEE, RONALD D., AND TULJAPURKAR, SHRIPAD. 1998a. "Stochastic Forecasts for Social Security," in *Frontiers in the Economics of Aging*, ed. David Wise. Chicago: University of Chicago Press, pp. 393–420.
- LEE, RONALD D., AND TULJAPURKAR, SHRIPAD. 1998b. "Uncertain Demographic Futures and Social Security Finances," *American Economic Review* (May):237–41.
- LEE, RONALD D., AND TULJAPURKAR, SHRIPAD. 1998c. "Population Forecasting for Fiscal Planning: Issues and Innovations," presented at a conference at the University of California at Berkeley, October 16–17, on Population Change and Fiscal Policy, sponsored by the Burch Center and the Center on the Economics and Demography of Aging.
- MANTON, KENNETH, STALLARD, ERIC, AND SINGER, BURT. 1992. "Projecting the Future Size and Health Status of the U.S. Elderly Population," *International Journal of Forecasting* 8, no. 3:433–58.
- McNOWN, ROBERT, AND ROGERS, ANDREI. 1989. "Forecasting Mortality: A Parameterized Time Series Approach," *Demography* 26, no. 4:645–60.
- McNOWN, ROBERT, AND ROGERS, ANDREI. 1992. "Forecasting Cause-Specific Mortality Using Time Series Methods," *International Journal of Forecasting* 8, no. 3:413–32.
- U.S. BUREAU OF THE CENSUS. 1989. Current Population Reports, Series P-25, no. 1018, *Projections for the Population of the United States, by Age, Sex, and Race: 1988 to 2080*, by Gregory Spencer. Washington, D.C.: U.S. Government Printing Office.
- U.S. NATIONAL CENTER FOR HEALTH STATISTICS. 1998a. "Births and Deaths: Preliminary Data for 1997," *National Vital Statistics Reports* 47:4.
- U.S. NATIONAL CENTER FOR HEALTH STATISTICS. 1998b. "Report of Final Mortality Statistics, 1996," *National Vital Statistics Reports* 47:9.
- WILMOTH, JOHN R. 1993. "Computational Methods for Fitting and Extrapolating the Lee-Carter Model of Mortality Change," Technical Report, Department of Demography, University of California, Berkeley.
- WILMOTH, JOHN R. 1995. "Are Mortality Projections Always More Pessimistic When Disaggregated by Cause of Death?" *Mathematical Population Studies* 5, no. 4:293–319.
- WILMOTH, JOHN R. 1996. "Mortality Projections for Japan: A Comparison of Four Methods," in *Health and Mortality*

among *Elderly Populations*, eds. Graziella Caselli and Alan Lopez. New York: Oxford University Press, pp. 266–87.

DISCUSSION

JUHA M. ALHO*

The Carter and Lee (1992) paper on forecasting age-specific mortality has been one of the most influential in recent times, and for a good reason. It provides a description of mortality change that is both easy to understand and empirically accurate. The review by Lee is a succinct summary of the work and shows how it can be applied to the stochastic forecasting of the population, and of the Social Security trust fund, for example. In the following I will try to put the work into a broader perspective.

STATISTICAL MODEL

Principal components techniques have been used in time-series analysis for quite some time. Brillinger's (1981, Ch. 9) discussion of the topic is based on lectures held in the late 1960s, for example. Bozick and Bell (1987) appear to have been the first to use the approach in the forecasting of age-specific fertility rates. The method can be described as follows. Let $\mathbf{M} = (M_{xt})$ be an $N \times T$ matrix of vital rates that have been observed for N ages during T years. Using the singular value decomposition we can write $\mathbf{M} = \mathbf{BK}^T$, where the columns of \mathbf{B} and \mathbf{K} are orthogonal. This representation can be used to approximate \mathbf{M} , in the least squares sense. The best one-dimensional approximation is $\mathbf{M} \approx \mathbf{B}_1\mathbf{K}_1^T$, where \mathbf{B}_1 and \mathbf{K}_1 are the columns of \mathbf{B} and \mathbf{K} that correspond to the largest singular value. In the method of Carter and Lee, the decomposition is first applied to the logarithms of the mortality rates that have been centered. The remarkable fact is that the one-dimensional approximation appears to be empirically adequate for mortality in many industrialized countries. For fertility data sets, a higher order approximation is often needed.

The second step is to develop a stochastic model for the components of $\mathbf{K}_1 = (k_1, \dots, k_T)^T$. Again, remarkably, a simple random walk with a drift appears to be adequate. In forecasting, \mathbf{B}_1 is kept fixed. The optimal forecast for k_{T+t} is obtained from a straight line that goes through the values k_1 and k_T . The forecast for the vector of the (transformed) rates is $\hat{k}_{T+t}\mathbf{B}_1$.

*Juha M. Alho is a professor in the Department of Statistics at the University of Joensuu, PO Box 111, 80101 Joensuu, Finland, e-mail, juha.alho@joensuu.fi.

In Alho and Spencer (1990) I have pointed out problems in the U.S. Office of the Actuary's method of forecasting cause-specific mortality. The method uses judgment in a way that is difficult to justify (Alho 1992). Simpler methods such as that of Carter and Lee (1992) would merit serious consideration. As always, there are details that might be refined; I will mention two.

The use of the least squares criterion is reasonable, but not optimal. Let D_{xt} be the number of deaths in age x during year t , and let V_{xt} be the corresponding person years. Maximum-likelihood estimates can be obtained under a Poisson model $D_{xt} \sim Po(\lambda_{xt}V_{xt})$, where the intensity is of the *log-bilinear* form $\log(\lambda_{xt}) = a_x + b_x k_t$. This is actually a version of the so-called association models developed in the 1980s but with a history that goes back to the 1930s (see Goodman 1991). Special software for their estimation has been developed, but point estimates can be obtained using, for example, the well-known package GLIM.

Since the bilinear model makes no assumption about the shape of the vector $\mathbf{B}_1 = (b_1, \dots, b_N)^T$, there is no guarantee that the forecasted mortality schedules in adult ages are monotone. In a University of Joensuu Master's Thesis (Eklund 1995) it was found that the one-dimensional approximation to mortality in Finland, by single years of age 65–99 during 1972–1989, fitted well but produced several non-monotonicities in the forecasts for ages over 90. The data may have been irregular because of Poisson variation. Yet, it would be of interest to have a model that would be constrained to produce monotone schedules.

APPLICATIONS TO FORECASTING

As pointed out by Lee, the high-low intervals of current official population forecasts suffer from serious logical problems that can be avoided by a probabilistic reformulation. Apparently, this was first noted by L. Törnqvist in 1949. Törnqvist was a statistics professor at the University of Helsinki, better known for his work in index number theory. In the late 1940s he helped Statistics Finland to improve their methods of population forecasting. Unfortunately, in those days effective means of computation were not available so he could not pursue a full stochastic analysis of a population forecast.

The extrapolation method of Carter and Lee provides a good point forecast for mortality. The handling of fertility in the Lee-Tuljapurkar forecast relies more

on judgment. The analysis of uncertainty is based on empirical standard errors from both models.

Even when forecasts are not based on formal statistical models, it is possible to provide an approximate probabilistic assessment of their uncertainty in countries with sufficiently long data series of mortality and fertility, such as the U.S. One can simply ask how large the error would have been in the past had the series been forecasted using some simple, baseline approach. In Alho (1998) I have done such a calculation for Finland. The baseline forecast assumes that age-specific mortality declines at the same rate that it has for the past 15 years. If the official forecasts are actually more accurate than those obtained with this simple assumption, then our empirical error estimates would be conservative (that is, too large). A similar calculation was done for fertility but with a constant baseline forecast. For migration, a mixture of time-series modeling and elicitation of expert judgment was used. The predictive distribution for future population was obtained via simulation, using the program PEP (Program for Error Propagation) we have written.

Lee's application to the Social Security trust fund is a beautiful example of how the use of simulation techniques allows us to take the various sources of forecast uncertainty seriously and to summarize the results in a predictive distribution of a complex "statistic" such as the balance of the fund, or its time of exhaustion. Similar calculations can be made for many other statistics. We have considered the system of state allocation of funds to municipalities in Finland, for example. These first calculations are based on many simplifying assumptions. A challenge for future research is to examine and model the various feedback mechanisms that may exist between the economic variables and population.

REFERENCES

- ALHO, JUHA M. 1992. "Estimating the Strength of Expert Judgment: the Case of U.S. Mortality Forecasts," *Journal of Forecasting* 11:157–67.
- ALHO, JUHA M. 1998. *A Stochastic Forecast of the Population of Finland*. Reviews 1998/4. Statistics Finland.
- ALHO, JUHA, M., AND SPENCER, BRUCE D. 1992. "Error Models for Official Mortality Forecasts," *Journal of the American Statistical Association* 85:609–16.
- BRILLINGER, DAVID R. 1981. *Time Series, Data Analysis and Theory*. San Francisco: Holden-Day.
- BOZICK, J. E., AND BELL, WILLIAM R. 1987. "Forecasting Age Specific Fertility Using Principal Components," *Proceedings of the American Statistical Association, Social Statistics Section* 369–401.

- CARTER, LAWRENCE, AND LEE, RONALD D. 1992. "Modeling and Forecasting U.S. Mortality: Differentials in Life Expectancy by Sex," *International Journal of Forecasting* 8, no. 3 (November):393-412.
- GOODMAN, LEO. 1991. "Measures, Models, and Graphical Displays in the Analysis of Cross-Classified Data," *Journal of the American Statistical Association* 86:1085-1111.

Additional discussions on this paper can be submitted until July 1, 2000. The author reserves the right to reply to any discussion. Please see the Submission Guidelines for Authors on the inside back cover for instructions on the submission of discussions.