

Center for Retirement Research
Center for Retirement Research Working
Papers

Boston College

Year 2005

Life is Cheap: Using Mortality Bonds to
Hedge Aggregate Mortality Risk

Leora Friedberg

Anthony Webb

LIFE IS CHEAP: USING MORTALITY BONDS TO HEDGE

AGGREGATE MORTALITY RISK

Leora Friedberg and Anthony Webb*

CRR WP 2005-13

Released: October 2005

Draft Submitted: October 2005

Center for Retirement Research at Boston College
550 Fulton Hall
140 Commonwealth Ave.
Chestnut Hill, MA 02467
Tel: 617-552-1762 Fax: 617-552-0191
<http://www.bc.edu/crr>

*Anthony Webb, Ph.D., (corresponding author) is a research economist at the Center for Retirement Research at Boston College (CRR). He can be reached at 646-662-7254 (phone), at webbaa@bc.edu or, at the CRR (see information above). Leora Friedberg, Ph.D., is an Assistant Professor in the Department of Economics of the University of Virginia. She can be reached at 434-924-3225 (phone), 434-982-2904 (fax), at lfriedberg@virginia.edu, or P.O. Box 400182, Charlottesville, VA 22904-4182. The research reported herein was performed pursuant to a grant from the Actuarial Foundation. The opinions and conclusions are solely those of the authors and should not be construed as representing the opinions or policy of the CRR, Actuarial Foundation or Society of Actuaries. We are grateful to Marric Buessing, Francesca Golub-Sass, Jamie Lee, and Andrew Varani for excellent research assistance, and to Mauricio Soto, seminar participants at the Center for Retirement Research at Boston College, Steve Cooperstein and Bob Triest for their very helpful comments.

JEL Classifications: G12, G22, G23, J11, J14.

© 2005, by Anthony Webb and Leora Friedberg. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

About the Center for Retirement Research

The *Center for Retirement Research at Boston College*, part of a consortium that includes a parallel centers at the University of Michigan and the National Bureau of Economic Research, was established in 1998 through a grant from the Social Security Administration. The goals of the Center are to promote research on retirement issues, to transmit new findings to the policy community and the public, to help train new scholars, and to broaden access to valuable data sources. Through these initiatives, the Center hopes to forge a strong link between the academic and policy communities around an issue of critical importance to the nation's future.

Center for Retirement Research at Boston College

550 Fulton Hall
140 Commonwealth Ave.
Chestnut Hill, MA 02467
phone: 617-552-1762 fax: 617-552-1750
e-mail: crr@bc.edu
<http://www.bc.edu/crr>

Affiliated Institutions:

American Enterprise Institute
The Brookings Institution
Center for Strategic and International Studies
Massachusetts Institute of Technology
Syracuse University
Urban Institute

ABSTRACT

Using the Lee-Carter mortality model, we quantify aggregate mortality risk, the risk that annuitants might live longer than predicted by the model. We calculate that a markup of 4.3 percent on an annuity premium, or else shareholders' capital equal to 4.3 percent of the expected present value of annuity payments, would reduce the probability of insolvency resulting from uncertain aggregate mortality trends to five percent, and a markup of 6.1 percent would reduce the probability of insolvency to one percent. Using the same model, we find evidence that the projection scale that the insurance industry commonly refers to underestimates aggregate mortality improvements. Consequently, annuities that are priced on that projections scale without any conservative margin will be substantially underpriced.

Insurance companies could deal with aggregate mortality risk by transferring it to the financial markets through newly-available mortality-contingent bonds. We calculate the returns that investors would have obtained on such bonds had they been available previously, and the historical covariance between these bond returns and the growth in per-capita consumption. Using the Consumption Capital Asset Pricing Model (CCAPM), we determine the risk premium that investors would have required on such bonds. At plausible coefficients of risk aversion, investors should be able to hedge aggregate mortality risk via such bonds at very low cost.

INTRODUCTION

Annuities provide a means by which risk-averse households facing an uncertain lifespan can insure themselves against the risk of outliving their wealth. Annuitization is predicted to raise average expected utility under a variety of assumptions, but voluntary annuity purchases are rare. Dushi and Webb (2004) found that only 10.2 percent of a large sample of elderly households had voluntarily annuitized any of their wealth between 1993 and 2000.¹ Yet, interest in private annuities may jump in coming years as pre-annuitized defined benefit pensions provided by employers are largely replaced by lump-sum defined contribution pensions, and also if Social Security is similarly transformed.²

The meager demand for private annuities remains a puzzle. One reason for the small size of the private annuity market may rest on difficulties which insurers face in offering annuities. Annuity providers face two kinds of mortality risk: idiosyncratic risk, since any particular annuitant may live longer than anticipated, and aggregate mortality risk, since annuitants may, on average, live longer than expected.³ The former risk can be eliminated through diversification. The latter risk is non-diversifiable unless the insurer writes other classes of business, for example life insurance, with risks that are negatively correlated with that of its annuity business.

The considerable academic literature on the value of annuitization to households assumes that the degree of future mortality risk is known with certainty. In reality, there is considerable uncertainty about the possible course of mortality reductions. Forecasting life expectancy 50 or 100 years hence involves taking a controversial position on the potential for either scientific discoveries to slow the aging process or the emergence of new conditions to hasten it. In consequence, there is considerable disagreement about the potential for longevity gains over such time horizons. S. Jay Olshansky and Stephen Austad, two of the most famous researchers in the field, take divergent positions and have wagered \$500 million on whether someone alive today will live to 150 by 2150.⁴

¹ These statistics were computed using data from the Asset and Health Dynamics of the Oldest Old, a panel of households with a member born before 1924.

² Some Social Security privatization proposals include mandated annuities in various forms. Private Social Security investment fund providers in Chile are required to offer voluntary annuities. While the provision will end in April 2006, the U.K. has required that "individual pensions" be annuitized by the time that individuals turned 75.

³ To simplify the analysis, we ignore the possibility that the degree of selection experienced by the insurer may differ from that anticipated.

⁴ Due to the power of compound interest, the present value of the wager is considerably more modest!

Insurance companies and potential annuitants make their decisions over much shorter time horizons. Over these horizons, it is common to use models that combine extrapolations of past trends with some consideration of likely threats to health and likely medical advances, but which do not take a strong position on medical progress or the biology of aging. Extrapolative models fit the data closely, have been shown to perform well in sample, and importantly for our purposes, enable researchers to quantify the uncertainty surrounding the future course of mortality.

Past forecasts from these models indicate, nevertheless, that point estimates of expected mortality rates carry wide margins of error. Thus, insurance companies and other annuity providers basing their pricing on those estimates would be exposed to significant aggregate mortality risk. Annuity providers can respond to this risk in a number of ways. One approach is to add on a margin for error, though a substantial margin may further deter high-expected mortality households from annuitizing, exacerbating adverse selection. Other more efficient alternatives would be to transfer aggregate mortality risk to third parties through the reinsurance market or by issuing mortality-contingent bonds, with interest payments that depend on either population or the insurer's mortality experience.⁵

What is an appropriate risk premium for such bonds? This paper quantifies the magnitude of the aggregate mortality risk faced by annuity providers and then prices that risk. The remainder of the paper is organized as follows. In the first section, we discuss the literature on aggregate mortality risk and present the Lee-Carter (1992) model, which -- according to Deaton and Paxson (2004) -- has become the "leading statistical model of mortality in the demographic literature." We compare the predictions of that model with those of the Social Security Administration and published life tables. In Section Two, we use the model to quantify the magnitude of the aggregate mortality risk faced by financial institutions selling annuities to a single birth cohort. We calculate the combination of shareholders' capital and premium loadings that would be required to reduce to specified percentages the risk of that capital being exhausted by mortality-related losses.

In Section Three, we investigate whether the projection scale that the insurance industry commonly refers to when forecasting mortality improvements appears to take account

of the predictions of the Lee-Carter model. The Society of Actuaries (SOA) recommends that insurance companies use Projection Scale AA as a basis for forecasting mortality improvements.⁶ We calculate that if the Lee-Carter model provides an unbiased estimate of the pace of mortality improvement, then insurance companies that use Projection Scale AA to price annuities without any conservative margin will, on average, underestimate that improvement substantially. Such annuities will be underpriced by amounts that range from 8.7 to 11.2 percent of the premium paid, relative to annuities with prices based on the Lee-Carter model. These findings of potential underpricing deepen the puzzle that we observe such low rates of voluntary annuitization.

In Section Four, we explain how mortality-contingent bonds, which could transfer aggregate mortality risk from borrowers to bondholders, are structured, and we price them using the Consumption Capital Asset Pricing Model (CCAPM). The CCAPM recognizes that the premium that investors must receive to take on an investment is determined by the correlation between the risky return on the investment and the growth rate of overall consumption. According to the CCAPM, investors will prefer assets that offer high returns just when consumption happens to be low and the marginal utility of consumption is correspondingly high, relative to assets such as stocks that tend to offer high returns when consumption is high and the marginal utility of consumption low. They will demand an expected return that exceeds the risk-free rate to hold the former and will accept a return of less than the risk-free rate to hold the latter. So, by the CCAPM, the risk premium on mortality-contingent bonds should be determined by the covariance between the bond's expected returns and overall consumption growth.

We calculate the annual returns that would have been earned on mortality-contingent bonds, had such instruments been available during the period 1959-1999. The correlation between those returns and per capita consumption growth is negative, since an unexpected mortality decrease stretches resources which are relatively fixed in the short-term over a larger-than-expected population. Consequently, under the CCAPM investors will accept a risk discount to hold a mortality-contingent bond. We compare the predicted risk discount to the actual

⁵ The pricing and availability of reinsurance should depend on reinsurers' ability to transfer the same risk yet more broadly. We discuss later how Swiss Re, a major reinsurer, has issued mortality bonds, which transfer its risk to the financial markets.

mortality-risk discount on the only longevity bond that has yet come to market. Although small, the observed discount is greater than that which is predicted by the CCAPM – so investors seem to be paying too much for the opportunity to hedge aggregate mortality risk. This “mortality premium puzzle” represents the flip side of the so-called “equity premium puzzle”, by which investors enjoy what appear to be excess returns from holding equity, relative to its observed riskiness.

Section Five summarizes. Our conclusion is that insurance companies face substantial aggregate mortality risk. Moreover, insurance companies may be mispricing this aggregate mortality risk in the annuities that they offer. While we show how the risk could be transferred to bondholders, a viable alternative would be to transfer it to annuitants. A workable scheme for such a transfer has been proposed by Piggott, Valdez, and Detzel (2005), and the annuity contracts issued by TIAA-CREF already involve some element of mortality risk-sharing. Although insurance companies would suffer a reduction in expected returns were they to hedge aggregate mortality risk by investing in mortality bonds or similar instruments, our calculations indicate that the magnitude of the reduction, and its effect on annuity prices, would be extremely modest even at prevailing mortality bond yields. These findings have broader implications for employers and for the government, who assume aggregate longevity risk through defined benefit pensions and Social Security, respectively. The issues will assume growing importance as defined benefit pension plans that typically provide benefits in annuitized form continue to be displaced by 401(k) and other defined contribution plans, in which annuitization is almost never mandatory and usually not even an option.

1. QUANTIFYING AGGREGATE MORTALITY RISK

1.1 The Lee-Carter Model

Wong-Fupay and Haberman (2004) review academic research into population mortality forecasts. The paper has an extremely useful discussion about using mortality models to quantify the uncertainty surrounding mortality forecasts. They point out that actuarial models are often unsuited to this purpose because they focus on point estimates. They also show that many of the intervals bounded by high and low scenarios in official projections of mortality

⁶ This and other tables can be viewed in the table manager that can be downloaded from the SOA website: www.soa.org.

improvements appear much too narrow. Blake (2001) shows that plausible assumptions about forecasting errors can lead to quite imprecise estimates of actuarially fair annuity yields. Mortality uncertainty rises over longer time horizons and he finds that confidence intervals are wider for escalating annuities whose benefits are back-loaded.

However, studies of this type provide essentially ad-hoc calculations based on authors' estimates of the likely level of mortality-table forecasting error. What is required is a means of quantifying the forecasting error, so we employ the Lee-Carter (1992) model, the most widely-used model of population mortality. It has been adopted by, among others, the U.S. Census Bureau and it was viewed favorably by the Social Security Administration's 1999 technical advisory panel.⁷ The actuaries forecast mortality by combining age-specific trend extrapolation with information collected from medical experts.

In the Lee-Carter model, mortality risk m at age x in year t is as follows:

$$\ln[m(x,t)] = a_x + b_x k_t + e_{x,t} \quad (1)$$

The parameters a and b vary with age. Lee and Carter find that the k trend declined roughly linearly over the period 1900-1989, which translates into a decreasing rate of increase in life expectancy, consistent with mortality data for the United States. They find that a random walk with drift fits the time path of k and estimate the following model:

$$k_t = k_{t-1} - 0.365 + 5.24 flu + e_t \quad (2)$$

flu is the impact on death rates of the 1918 flu epidemic. According to these estimates, a one standard-deviation shock to k translates into a change in age-65 life expectancy of about two months.

We will use the parameters estimated from this model to produce forecasts of mortality rates at various ages for all future years. Importantly for our purposes, it can be used to calculate both unconditional and conditional forecasts of future mortality rates, plus associated confidence intervals. Conditional on information at time t , we can calculate the probability distribution of survival rates at each age for all future years. For example, we can say that we expect the age-85

⁷ See p. 64 of the 1999 report of the Social Security Technical Panel on Assumptions and Methods, which is available at <http://www.ssab.gov/Publications/Financing/tech99.pdf>.

survival probability in 2020 to be some value $q_{85,(t=2020)}$, and we can report the standard error of the forecast. We can also calculate trends in life expectancy – for example, it predicts life expectancy at birth of 86.05 years in 2065 with a 95 percent confidence interval of [80.45, 89.95] years.

Lee and Miller (2001) found that the model under-predicted gains in in-sample forecasts but only by small amounts, and that the confidence intervals were a little too wide over the time horizons that might concern insurance companies and bondholders. From the viewpoint of an insurance company selling annuities, the latter type of error provides insurance against the former. Tuljapurkar, Li, and Boe (2000) found that the model applied not only in the United States, but also in the other G7 countries with different estimated parameters.

While the model suffers from several potential weaknesses which have been pointed out by other authors, we do not generally need to address them for our application. First, the parameter values depend on the period over which they are estimated. The parameter values do not vary much, though, so our results are unlikely to be affected by our choice of basis period. Second, the error term may not capture large but infrequent mortality shocks – for example, the risk of a repetition of the 1918 flu epidemic. We find that allowing for such low-probability events would have little effect on our results, especially if, as may be likely, mortality rates experience asymmetric shocks. Third, the model does not fit age-specific mortality data exactly in the jump-off year, and an alternative is to set a_x to fit the initial conditions exactly. Bell (1997) shows that this approach produces somewhat superior forecasts. We do not use his modification, though, as most of our calculations are little affected by it. Fourth, over short forecasting intervals, the error associated with estimating the parameters a and b dominates the variability of the k_t forecasts; however over the life of an annuity, the reverse is true and errors in forecasting k , which are our focus, dominate. Fifth, Booth, Maindonald, and Smith (2002), analyzing Australian data, find a substantial age-time interaction that improves the fit of the model. We do not incorporate such a term as no one has demonstrated whether the same is true for the United States. Moreover, we find below in our own analysis that, in any given year, k varies little with age. Sixth, the Lee-Carter model was estimated using combined male and female mortality tables. Although patterns of mortality declines differ by gender, the reasons for

the differences are not well understood. We assume that both male and female mortality declines are subject to the same process and the same mortality shocks.

We note in passing that a number of authors have proposed other enhancements to the Lee-Carter model (Renshaw and Haberman 2003a, 2003b, and Brouhns, Denuit, and Vermunt 2002). We propose retaining the original model, as parameter values are available for U.S. data, the original model is widely accepted, and the differences between the predictions of the models are not substantial. Moreover, the Lee-Carter model is not the only one we could choose. Other stochastic models include Lee (2000), Yang (2001), Milevsky and Promislow (2001), Dahl (2004), and Cairns, Blake, and Dowd (2005a). In contrast to Lee-Carter and most other models that predict a deceleration in the rate of mortality improvement, Sanderson and Scherbov (2004) examine mortality across fourteen countries and project no deceleration in the rate of increase of life expectancy.⁸ As we noted earlier, we adopt the Lee-Carter model because it has achieved widespread acceptance among demographers and provides an excellent fit to U.S. data.

1.2 Comparing Lee-Carter with Social Security Administration Forecasts

The Lee-Carter model was originally estimated using mortality data up to 1989. We therefore examine forecasts over two periods: 1989 to 2001, when Lee-Carter and Social Security Administration forecasts can both be compared with actual mortality improvements; and 2001 onwards.

Figures One and Two compare the model's predictions of survival rates for the male and female birth cohort of 1924, starting at age 65, with those made by the Social Security Administration (SSA) in their 1989 intermediate forecast, and with actual survival rates from 1989-2001. The Lee-Carter model predicts a mortality-weighted annual mortality decline of 1.13 percent per year for the over 65s. In contrast, in the 1993 Trustees' Report, the Social Security Administration forecast male and female mortality declines of 0.54 and 0.50 percent for the entire period 1989 to 2017. It transpired that mortality declined much more rapidly among men than women from 1989 to 2001. Bell and Miller (2005) report male and female mortality among the over 65s declined at annual average rates of 0.76 and 0.17 percent respectively over the period 1981-2001. Our own analysis of mortality data for the period 1989-2001 shows an

⁸ They analyze "best practice" life expectancy by determining the average life expectancy in the country that, at a point in time, has the greatest life expectancy.

average mortality-weighted decline in mortality among men of 0.84 percent a year, but an average increase of 0.06 percent a year among women, so both forecasts in Figure Two lie above actual mortality. So, the decline for men proved to almost midway between the Lee-Carter and Social Security Administration forecasts, whereas neither the Social Security Administration nor Lee-Carter accurately predicted female mortality.

Preston and Wang (2005) investigated the recent narrowing of the male-female mortality differential, which dropped from around 7.8 years in 1972-79 to 5.3 years in 2003. They found evidence of a delayed response to increases in the prevalence in smoking among women. They project an acceleration in mortality reductions even if the prevalence of smoking remains unchanged. However, male mortality will decline faster than female mortality because male smoking rates have declined, whereas female smoking rates have not changed radically. We conclude from this research that the period 1989-2001 is merely a temporary pause in the decline in female mortality.

Looking forward from 2001, the SSA is less optimistic compared to Lee-Carter, however. SSA provides high, intermediate, and low mortality forecasts. Among the over 65s and taking 2001 as a baseline, they project annual mortality reductions of 0.24, 0.47, and 0.70 percent under the three forecasts, increasing to 0.29, 0.67, and 1.17 percent after 2029.⁹ This compares with historical average reductions of 0.72 percent over the period 1900 to 2001 and 0.47 percent more recently over 1979 to 2001.¹⁰ The intermediate forecast of 0.47 percent until 2029 and then 0.67 percent afterwards predicts that mortality will initially decline at the pace experienced between 1979 and 2001 and then jump to approximately the pace experienced between 1900 and 1979. The high forecast predicts that the somewhat higher pace of mortality declines experienced between 1900 and 1979 will prevail initially and then accelerate.

The Lee-Carter model results in forecasts of mortality reductions for the over 65s that are considerably higher than even the SSA's "high" forecast of 0.70 percent through 2029. Using the parameter values estimated by Lee-Carter from 1900-89 mortality data, we forecast a future age-sex weighted average mortality reduction of 1.13 percent a year.¹¹ It should be noted

⁹ The SSA do not break down their projected declines by age within this group.

¹⁰ 2005 Social Security Trustees Report, page 71. Historically, percentage reductions in mortality have been higher among the young. Among the whole population, mortality declined by an average of 1.06 percent per year from 1900 to 2001 and 0.72 percent per year from 1979 to 2001.

¹¹ As the Lee-Carter model is log-linear, the forecast percentage decline in mortality each period is constant at $b_x(E_t k_{t+1} - k_t)$. This varies with age, but not with calendar year.

that the SSA “high” forecast rises substantially after 2029 to 1.17 percent, almost the same as Lee-Carter. Also, by comparison, mortality declines exceeded the rates predicted by the model in other time periods. For example, Lee-Carter report that k actually declined by 0.548 a year in the 1970s, well above the estimated long-run average of 0.365.

1.3 Comparing Lee Carter with Actuarial Forecasts

Insurance and pension actuaries are understandably more concerned with mortality rates among annuitants and members of employer pension plans than with mortality rates among the population as a whole, as predicted by the Lee-Carter model. Rates of decrease in mortality may differ between annuitants and non-annuitants due to widening socio-economic mortality differentials, as documented by Willetts (1999) using United Kingdom data and Schalick et al (2000) using older U.S. data from 1967 to 1986. It is nonetheless informative to compare the mortality declines projected by the Lee-Carter model with the projection scales that actuaries commonly use when forecasting mortality declines among annuitants and members of pension plans.

Table One compares, by age group, the mortality declines projected by the Lee-Carter model with those by the SOA in their projection scale AA, the most recent such scale that SOA has issued. The SOA projection scale is a blend of Federal Civil Service and Social Security mortality improvements from 1977 to 1993, subject to various adjustments that smoothed and placed upper and lower bounds on the forecasted improvements. When preparing the RP2000 employer pension life tables, the SOA considered whether scale AA should be updated but concluded that it was consistent with recent data on mortality improvements.¹² As expected given the source of the SOA data, the SOA forecasts a slower pace of mortality reductions than ours based on the Lee-Carter model. The age-sex weighted average for the SOA is 0.71 percent for the whole population and 0.64 percent for the over 65s.

From age 45 to 79, the SOA and the Lee-Carter forecasts are remarkably close. From age 80 upwards, the SOA projects successively smaller percentage reductions in mortality, approaching zero at age 100 plus. In contrast, the Lee-Carter model projects somewhat larger reductions in mortality over age 80 than among those aged 45-79.

¹² “The RP 2000 Life Tables” are available from the SOA website.

The disagreement above age 80 is particularly important. Mortality rates are much higher at these ages, so given percentage reductions in mortality have a disproportionate effect on overall life expectancy. Mortality rates at old ages are even more important to insurance companies because most annuities are purchased by older households.¹³ It will be important to keep in mind this issue as we present the rest of our results.

2. QUANTIFYING THE MORTALITY RISK FACED BY ANNUITY PROVIDERS

In this section, we use simulations to quantify the aggregate mortality risk faced by annuity providers selling annuities to a single birth cohort. We calculate the combination of capital and premium loading that would reduce the risk of insolvency to specified percentages.

So that we can focus on the effect of aggregate mortality risk, we impose a number of simplifying assumptions. We assume that the annuity provider sells a “large” number of annuities of a single type to people in a single birth cohort who have not only population average mortality but also population average changes in mortality rates as predicted by the Lee-Carter model. In reality, as we noted in the last section, there is evidence that higher socio-economic classes are experiencing more rapid declines in mortality. We ignore these differences in the likely pool of potential annuitants, which would amplify our findings, and we ignore considerations of adverse selection, which may already affect this market. We assume that the annuity price is set at a level that will enable the annuity provider to break even at expected mortality rates. We further assume that the provider sells a real annuity, invests in a risk-free asset offering a real 3 percent return, and has zero administrative expenses.¹⁴

There are two ways in which mortality outcomes may differ from expectations as given by the Lee-Carter model. First, the time trend of -0.365 in equation (2) is estimated rather than known for certain, and the uncertainty associated with that estimate should be taken into account when forecasting mortality trends. Second, in each period there is a mortality shock – the e_t term in equation (2). While the a and b terms in equation (1) are also estimated, Lee-Carter

¹³ The DB pension market is much larger than the market for immediate annuities, and DB pensions, as deferred annuities, are “bought” by people of all ages.

¹⁴ A three percent real return is a common assumption in the academic literature, although the long-run risk-free rate as measured by TIPS is currently well below that level. Insurance companies are starting to offer real annuities, although the Wall Street Journal of 25 May 2005 reports Vanguard describing the consumer response as fairly lukewarm.

show in their Appendix B that it is reasonable to ignore these other sources of error when making life expectancy forecasts.

We run 10,000 Monte-Carlo simulations of the evolution of aggregate mortality. In each simulation, we make a single draw from a normal distribution with mean zero and standard deviation 0.069; these are the values from Lee-Carter associated with the estimated time trend of -0.365. We add this draw to -0.365 to obtain the value of the k trend which will stay fixed for that particular simulation. In the next step of each simulation, we obtain the error term e_t for each year by making a series of draws from a normal distribution with mean zero and standard deviation 0.655; these are the values from Lee-Carter associated with the estimation of e_t .¹⁵ We fill in the value of k year by year using these draws and that particular simulation's k trend and use equation (1) to calculate the associated annual mortality risk. Using the results of these 10,000 simulations, we construct separate male and female mortality tables, assuming that the ratio of male to female mortality rates is equal to the gender-weighted average of the population mortality rates reported by SSA for the appropriate age and birth cohort.¹⁶ We calculate annual survival probabilities, discount the resulting cohort annuity payments by a rate of interest, and sum the payments to arrive at a present value. We calculate the return earned by the annuity provider by comparing this present value with the premium paid, calculated under the same assumptions.

The insurance company makes a loss if, in a particular simulation, the mortality draw experienced by the insurance company results in payments that exceed the premium net of markup. We calculate the percentage markup over an actuarially fair premium that the insurance company must impose in order to reduce the probability of loss to specified percentages. This markup can also be interpreted as the percentage of the expected present value of annuity payments that the annuity provider must hold by way of capital if it wishes to reduce the probability of insolvency to those same percentages, or the combination of markup and capital that is required.

¹⁵ As noted earlier, a one standard-deviation shock to k translates into a change in age-65 life expectancy of about two months.

¹⁶ It would be preferable to use annuitant life tables for this purpose, but these are only available for periods, not cohorts.

Table Two reports our results for annuities issued to married couples (with survivor benefits of 50 and 100 percent), single women, and single men aged 65, 70, 75, 80, and 85.¹⁷ For an insurance company selling joint life and 50 percent survivor annuities to couples aged 65, discounted at 3 percent, some combination of shareholders' capital and premium loading equal to 3.74 percent of the premium is required to reduce the probability of loss or insolvency to five percent. Shareholders' capital or loading of 5.39 percent is required to reduce the probability to one percent. For the same annuity sold to couples aged 85, the corresponding numbers are lower, at 3.67 and 5.24 percent.

A higher interest rate reduces the required mark-up. At an interest rate of five instead of three percent, the premium loadings are somewhat lower – 2.94 and 4.29 percent for couples aged 65. At a higher interest rate, relatively less weight is attached to the payments made at older ages, which are increasingly risky.

There is no clear age-related pattern in the required mark-up. The standard error of the forecast of k in t periods time is $\sqrt{t \text{ var}(k)}$ so aggregate mortality risk increases with the forecasting horizon t . In consequence, annuities that are heavily back-loaded – those sold to married couples, to younger persons, or with large survivor benefits – carry more aggregate mortality risk. This is offset by two factors. First, back-loaded risk is subject to greater time discounting. Second and more importantly, annuities sold to younger persons include a substantial period during which mortality rates are extremely low. Even if mortality during that period turns out to differ substantially in percentage terms from that predicted, it will have very little effect on the present value of the annuity, and this tends to reduce the overall riskiness of annuities sold to younger individuals. These factors more or less counterbalance the effect of increasing the annuitant's age and hence the forecasting horizon.

3. DO INSURANCE COMPANIES CORRECTLY PRICE AGGREGATE MORTALITY RISK?

When we calculated the required premium loading for insurers to offset mortality risk, we assumed that insurers price annuities in accordance with Lee-Carter. In the following section, we investigate whether this is really the case. We find that, by the Lee-Carter benchmark, insurance companies systematically underprice annuities if they use the Projection

¹⁷ For couples, we assume that both spouses are of the same age. We obtain almost identical results when we assume that the wife is three years younger than the husband.

Scale AA without making any compensating adjustment elsewhere in their pricing formulas. This finding is a corollary of our earlier results showing that actuarial life tables appear to understate aggregate mortality risk.

3.1 Methodology

The problem in determining whether aggregate mortality risk is correctly priced is that we only observe prices charged, and not necessarily the insurance company's mortality assumptions. Moreover, prices depend not only on mortality assumptions, but also on assumptions regarding expenses and asset returns.

While we take a different approach, Mitchell, Poterba, Warshawsky, and Brown (1999) also dealt with some of the problems that arise because we do not observe insurers' mortality assumptions. They focused on the value of annuities to potential annuitants, not on whether insurance companies were pricing their products correctly. As a result they did not concern themselves with expenses or asset returns, but their treatment of mortality risk is relevant for our purposes. They calculated the expected present value (EPV) of annuities held by people with either population or annuitant mortality.¹⁸ To calculate the EPV for someone with population mortality, they used SSA cohort life tables. Annuitant life tables are period tables, though, and they could not similarly use them to calculate the EPV of an annuity to someone with annuitant mortality. Furthermore, annuitant tables were not available for the period they were studying. Instead, they first interpolated between the 1983 Individual Annuity Basic A and the Annuity 2000 annuitant life tables, constructing a set of annuitant life tables for 1995. Then, they constructed annuitant cohort life tables by converting the period table through multiplication of the resultant mortality rates by the ratio of the relevant mortality probabilities from the 1995 Social Security cohort table for the appropriate birth cohort to the 1995 Social Security period table.

However, their calculation rests on two questionable assumptions. First, it assumes that the SSA has correctly forecasted the pace of mortality reductions. As we have seen, the SSA may be underestimating the pace of reductions, particularly among the oldest old. Second, as

¹⁸ Another difference between their approach and ours is that they assumed that annuitants discount future annuity payments by either the Treasury strip or the corporate bond interest rate. We simply assume a three percent real interest rate. Assumptions about interest rates have only a second-order effect on our calculations of the consequences of underestimating the rate of improvement in mortality.

Mitchell et al point out, it assumes that the pace of mortality reductions among annuitants will equal that among the population as a whole. But, as mentioned earlier, there is evidence that socio-economic disparities in mortality were widening through 1986.¹⁹

Given the difficulty that we do not observe the mortality data that annuity providers use, we proceed as follows. We assume that the Annuity 2000 mortality table correctly describes current period mortality of the persons buying the annuity. We further assume that insurance companies make use of projection scale AA, discussed previously, to forecast mortality improvements.²⁰ Thus, we focus on annuity providers and not just annuitants, as Mitchell et al did, and we use information commonly available to annuity providers to gauge their mortality projections without making assumptions about expenses and asset returns. We assume that insurance companies face zero administrative costs and can invest in a risk free asset yielding a real three percent interest rate.

In contrast, when we calculate the expected present value (EPV) of the annuity to the annuitant, we use the Lee-Carter model to project mortality. It follows that the EPV will exceed 100 percent of the premium paid if and only if the Lee-Carter mortality forecasts are more optimistic than the projection scale. As we are assuming symmetry of beliefs regarding current mortality, this result holds irrespective of our choice of current period life table that the insurance company may use.

In practice, actuaries develop “prudent best estimates” of mortality by constructing life tables based on both their own experience and published life tables, and then adding a conservative margin reflecting data uncertainty. This conservative margin extends to forecasting mortality improvements. For example, the American Academy of Actuaries Variable Annuity Reserve Working Group proposes that actuaries be required to take account of mortality

¹⁹ Table 7.1 in “The RP-2000 Mortality Tables” lists rates of mortality declines among various types of lives (Federal Civil Service, Social Security, Railroad Retirees, Group Annuity Lives, and Group Annuity Amounts) from roughly 1980 to the late 1990s, and we see no discernable relationship between the presumed average socio-economic status of the lives and the rate of mortality decline. However, as the RP 2000 report points out in Chapter Four, “the measurement of mortality improvement requires voluminous consistent data covering many years,” so it is difficult to infer whether mortality differentials have continued to widen since 1986.

²⁰ The Annuity 2000 table is an individual annuitant mortality table, while Projection Scale AA is reported in “The RP 2000 Life Tables” issued by the Society of Actuaries (www.soa.org). We chose the Annuity 2000 table over others, like those published by the Social Security Administration, because our focus is on the individual annuity market. Our results would not be substantially affected by using an alternative period table or by projecting the Annuity 2000 table to 2005 because they flow from the discrepancy between Lee-Carter and actuarial projections of declines in mortality. Lee-Carter compared their predictions with those of the SSA, but to the best of our knowledge, no one has compared the Lee-Carter predictions with those of life tables.

improvements when that would err on the side of caution, and permits them to do so otherwise.²¹ The assumption relating to such improvements must be based on “current relevant data” with an unspecified margin for error in the appropriate direction. Our calculations therefore represent an upper bound of the extent of the underpricing.

Table Three reports our results. EPVs are in all cases greater than 100 percent of the premium paid, which indicates underpricing if annuities are priced according to Scale AA. The extent of the underpricing ranges from 8.7 to 11.2 percent. It is, at all ages, greater for women than for men, but this simply reflects the fact that the Scale AA projected mortality improvements are greater for women than for men, whereas the Lee-Carter model only allows us to calculate unisex improvement factors, as we mentioned earlier.

If the Lee-Carter model is indeed an unbiased forecast of mortality improvements, then we would conclude that insurers are underestimating the pace of mortality reductions. Moreover, Dushi and Webb (2004) show that annuity purchasers are much wealthier than average households. Since, as we pointed out earlier, wealth and mortality rates are strongly correlated, and if socio-economic disparities in mortality risk continue to widen, then the above calculations will understate the cost of recognizing future mortality reductions.

Another implication of our finding, again assuming that the Lee-Carter projections are unbiased, is that Mitchell, Poterba, Warshawsky, and Brown (1999) substantially overstate the actuarial unfairness of annuities, since their cohort tables, which are based on Social Security Administration forecasts of mortality improvements, probably underestimate the rate of improvement in mortality. The low rate of voluntary annuitization is therefore even more puzzling than it previously appeared. It also cannot be explained by individuals mistakenly overestimating their own mortality risk. Gong and Webb (2005) construct subjective life tables for respondents in the Health and Retirement Study (HRS) based on their assessments of their own subjective survival probabilities. They conclude that the HRS households generally hold reasonable beliefs since the aggregated data resemble published mortality tables and co-vary appropriately with education and ethnicity. Although their calculations are not precise enough to distinguish between small differences in forecasts of mortality declines, they found no evidence of systematic overestimating of mortality risk.

²¹ http://www.actuary.org/pdf/life/varwg_0305.pdf

4. PRICING AGGREGATE MORTALITY RISK

Given our evidence of substantial aggregate mortality risk and the possibility that insurers systematically underprice such risk, this section discusses how insurers might reduce their exposure. We describe new mortality-contingent bonds and calculate the returns that investors would have experienced on such bonds had they been available over the period 1959-1999. To price such bonds, we outline the consumption capital asset pricing model (CCAPM), present data on the relationship between mortality-contingent bond returns and consumption growth, and price aggregate mortality risk.

4.1 Mortality-Contingent Bonds

As of August 2005, only two mortality-contingent bonds had been issued. One, issued by Swiss Re, is a \$400 million three-year bond paying LIBOR plus 135 basis points. The bond contains a provision that, if a five-country weighted mortality index exceeds 130 percent of the 2002 level, then the principal will be reduced. If it goes above 150 percent, the principal will be exhausted. The other bond, issued by the European Investment Bank (EIB) is for £540 million. It has a life of 25 years and is interest-free with no return of principal, but it has mortality-related payments. The payments on the EIB bond decline proportionately with the annual survival rate of the UK male population reaching 65 in 2003, subject to a short time lag. Life insurers should go short on the Swiss Re bond and annuity providers should go long on the EIB bond in order to hedge their exposure to aggregate mortality risk. Given that our focus is longevity risk, we will make calculations to price this latter bond.

Cairns, Blake, and Dowd (2005b) report that the EIB bond traded at a yield that was some 20 basis points lower than those at which similar non-mortality related EIB bonds traded – in other words, the bond attracted a mortality risk discount, not a premium.²² One explanation for this is that annuity providers exposed to aggregate mortality risk have driven up the price of mortality-contingent bonds relative to regular bonds. Yet, while risk-averse annuity providers would be willing to pay a risk premium to transfer aggregate mortality risk to a third party, the market price of that risk will be determined by its covariance with other sources of risk, which determines its value to potential buyers – a point which will drive our discussion of asset pricing.

Before proceeding with that, we will mention some issues related to the possible nature of mortality shocks. It seems unlikely that market prices reflect concerns about transitory shocks to mortality. In particular, the long duration of the EIB bond means that one-off shocks have relatively little effect on its value. The only one-off shock incorporated into the Lee-Carter model was the 1918 flu epidemic. Most or all one-off shocks will be negative, increasing but not decreasing mortality, so we do not consider the possibility of large positive transitory shocks.²³ We calculate that an extreme event like a repetition of the 1918 flu epidemic, which is large and increased k temporarily by 5.24 relative to an annual trend decline of -0.365, would reduce the value of the EIB bond by only 0.30 percent at a three percent discount rate.

A much smaller one-standard deviation permanent shock to k of 0.655 has a much greater – but still modest – effect on the value of the EIB bond because forward-looking investors will recognize the change in future mortality. Table Four shows the impact on the expected present value of an EIB-type bond, based on the survival rates of the U.S. 1938 birth cohort, of a shock to k of this magnitude. The impact of the shock varies from 0.63 to 1.17 percent of the value of the bond, depending on the age of the cohort, increases with age, and is lower at higher discount rates.

It should be noted that Beelders and Colarossi (2004) have priced the Swiss Re bond using extreme value theory, which attempts to estimate the extreme tail of the probability distribution. This technique is less relevant for pricing the EIB bond because, as Table Four showed, the extreme tail of the distribution of one-off shocks has little impact on the price.

4.2 Historical Returns

We calculate the returns that investors would have experienced had the EIB bond been made available in the U.S. market and had those investors believed that the Lee-Carter model correctly described mortality improvements. In the Lee-Carter model, the expected mortality decline for people of age x equals $0.365b_x$, and the percentage deviation of actual from expected mortality in period $t+1$ can be expressed as

²² The calculation of the mortality discount was based on the assumption that the Government Actuary Department's forecasts of mortality improvements are unbiased. It is also possible, as this is the first bond of its type to be issued, that the market has not yet reached an equilibrium, and that the mortality yield discount will narrow over time.

²³ It is difficult to think of any improvements in mortality that are not long-lasting.

$$\ln(m_{x,t+1}) - E_t[\ln(m_{x,t+1})] = [a_x + b_x(k_{t+1})] - [a_x + b_x(E_t(k_{t+1}))] \quad (3)$$

or $b_x e_{x,t+1}$.

The Berkeley Human Mortality database holds U.S. period life tables for 1959-1999. Based on these tables, we calculate the yearly percentage change in mortality at each age from 65 up. We compare this with the change predicted by Lee-Carter and recover, for each year, a vector of mortality shocks, $e_{x,t+1}$. We confirm that the Lee-Carter specification fits the data well. First, an important assumption of the Lee-Carter model is that the k shocks are assumed to affect all ages in the same way. As a result of this assumption, insurance companies would not be able to reduce their aggregate mortality risk much by diversifying across birth cohorts. In accordance with this assumption, we find that the error term $e_{x,t}$ does not vary significantly with age. In any given year, the mortality shocks are almost invariably of the same sign, and usually of a similar magnitude. An adverse shock would thus affect the company's obligations across all age groups more or less proportionately to the amounts in Table Four, discussed later. Second, we would expect $e_{x,t+1}$ to have mean zero and a standard deviation of 0.655 as found by Lee-Carter. The sample mean is 0.034, not significantly different from zero, and the sample standard deviation is 0.5391, not significantly different from 0.655.

To continue, the financial impact of the persistent mortality shocks which we have just calculated equals not only the additional or reduced payment that the bondholder receives in the current period, but also the present value of the additional (or reduced) payments that are now expected in all future years. Having expressed the mortality shock for each of the years 1960 to 1999 in terms of the innovation to k , we now calculate the impact of that innovation on the numbers expected to survive each year and consequently on the EPV of the income stream from the bond, assuming that investors use the Lee-Carter model to price the bond. The magnitude depends on the duration of the bond, and for each year, we assume that bond payments are based on a population that is currently age 65. We express the impact of the innovation as a percentage of the market value of the bond prior to the shock, using a three percent interest rate. The resulting mean is 0.04 percent and the standard deviation 0.64 percent. This standard deviation is very small – the general consensus is that the standard deviation of stock returns is about 17 to 20 percent, or roughly thirty times greater. It is important to note that this low level of risk

reflects the design of this particular bond. The return on a bond designed to pay out on mortality in excess of some floor would be much more volatile. We conclude that, if investors used the Lee-Carter model to forecast mortality and price mortality bonds and if the Lee-Carter model provides unbiased estimates, then a bond structured like the EIB bond would be a relatively low risk investment.

4.3 The Consumption Capital Asset Pricing Model

Mehra and Prescott (2003) provide an overview of the Consumption Capital Asset Pricing Model (CCAPM). The intuition behind the standard Capital Asset Pricing Model is that the expected return on a risky asset equals the risk-free interest rate plus a premium for bearing risk. Idiosyncratic risk does not command a premium because it can be eliminated through portfolio diversification, but the same is not true for systematic or general market risk that cannot be diversified.

In the CCAPM, what determines the premium for systematic risk is the relationship between the expected return on the asset and the marginal utility of consumption. When consumption suffers in a particular period, then utility in that period will be correspondingly low and marginal utility high relative to other periods. Extra income from investment is relatively more valuable at that time compared to others, while it is relatively less valuable in times of unexpectedly high consumption. Thus, investors require a risk premium for holding assets whose returns are positively correlated with shocks to overall consumption because those assets provide the biggest payoffs in states of the world in which consumption and utility are high and the marginal utility of consumption low. Conversely, investors place a high value on assets that offer high returns when consumption is low (so the marginal utility of consumption is high) and will buy such assets even when the expected return is less than the risk-free rate.

Mathematically,

$$E_t(R_{e,t+1}) = R_{f,t+1} + Cov_t \left\{ \frac{-U'(C_{t+1}), R_{e,t+1}}{E_t(U'(C_{t+1}))} \right\} \quad (4)$$

so the expected return on the risky asset in period $t+1$ equals the risk-free rate plus the covariance of the asset return with the marginal utility of consumption. A little algebra and a few assumptions²⁴ results in:

$$\text{Log}E_t \left[(1 + R_{b,t+1}) / (1 + R_{f,t+1}) \right] = \gamma \sigma_{b,c} \quad (5)$$

where γ is the coefficient of risk aversion and $\sigma_{b,c}$ is the log of the covariance of the asset return in question with consumption growth. Thus, a positive covariance requires a premium above the risk-free rate, so $R_b > R_f$, and the required premium increases with the coefficient of risk aversion.

We note in passing that an alternative to the CCAPM is to price the mortality-contingent bond using a Wang transform (Wang 2000). This latter method was adopted by Lin and Cox (2004). We favor the CCAPM approach, as the risk premium obtained using the Wang transform depends on assumptions about the insurer's level of expenses.

4.4 The relationship between mortality bond returns and consumption growth

Our hypothesis is that there should be a negative covariance between the growth in per-capita consumption and returns on EIB-type bonds. In other words, we expect that the contemporaneous correlation between aggregate consumption shocks and mortality shocks is negative. The key point is that most of the variation in mortality rates occurs at older ages, since that is when most deaths occur. Suppose an unexpectedly small number of older people die in a particular year – then total national income is unlikely to increase significantly in response because the capital stock is largely fixed at short time horizons, and few older people work. In consequence, per-capita income will decline because output is relatively constant and must be shared among a larger-than-expected number of individuals.²⁵ Moreover, to the extent that the mortality shock is global, then one country would not be able to borrow from another to offset

²⁴ The assumptions, which are standard, are that equity and stock returns are jointly log-normally distributed and that consumption growth and risky asset returns are both i.i.d.

²⁵ Mortality shocks may affect aggregate interest rates by altering the shares of national income taken by labor and capital. A major adverse mortality shock makes capital abundant and labor scarce. Ralph Higden, a contemporary English chronicler, documented how wages jumped and rents fell in the aftermath of the Black Death, which killed one-third of the European population during 1347-1352. The focus of our paper, though, is on the additional risk premium that a mortality bond should command relative to a similar non-mortality linked bond, with both bonds affected in the same way by a shift in the risk-free interest rate.

the consumption shock. A final point is that reductions in mortality may result from increased medical spending, leaving less money available for general consumption.

We indeed find that per-capita consumption and the returns on the pseudo-EIB bond we simulated above negatively co-vary. We follow convention by focusing on the consumption of non-durables and services. We use consumption data for 1959-1999 from the National Income and Product Accounts. Consumption growth is a relatively smooth series with a mean of 2.25 percent and a standard deviation of 1.19 percent.

The correlation between consumption growth and mortality bond returns is -0.1958 and is significantly different from zero. Thus, over the period we examine, mortality bonds would have provided their holders with the highest returns in periods when the rate of consumption growth was low and the marginal utility of consumption was correspondingly high.

It follows that investors should be willing to accept a risk discount for holding mortality bonds.²⁶ In contrast, they should require a modest risk premium for holding equities. The literature has found a large positive correlation between stock returns and consumption growth of about 0.5 – in other words, stocks offer the highest returns when the marginal utility of consumption is low. The covariance is much smaller, though, as consumption growth is a relatively smooth series, and under the CCAPM, it is the covariance between asset returns and consumption growth that determines the risk premium. This modest covariance has resulted in estimates of an appropriate equity premium that are much lower than that observed in historical data, which has in turn led to extensive discussion of the “equity premium puzzle.”

Although the correlation between mortality bond returns and consumption growth is some 40 percent of that between stock returns and consumption growth, the standard deviation of mortality bond returns is much less than that of equities. As a result, the covariance between mortality bond returns and consumption growth is extremely small at -0.0015 percent, just as the positive covariance between equities and consumption growth is small. Applying the CCAPM, the risk discount is only two basis points when the coefficient of risk aversion equals ten.

Recall that the EIB bond trades at a discount of 20 basis points. Yet, we expect that the same relationship between consumption growth and mortality shocks holds in the United

²⁶ This point assumes that agents are aware that aggregate mortality has improved. Yet, the government only publishes the mortality data on which EIB bond payments are based after the year's end, and the bond is currently structured so that the payments in any given year reflect lagged survival rates. Were such bonds to become

Kingdom as in the United States. Thus, if the British Government Actuary Department's mortality estimates are unbiased, or believed by investors to be unbiased, then the market appears to be overpaying for the EIB bond.²⁷ Alternatively, investors could believe that the Actuary Department is too conservative in its mortality forecasts. We calculate that the 20 basis point discount would be eliminated if investors believed that the actuaries were underestimating mortality declines by 0.44 percent a year. This represents about two-thirds of the gap between the intermediate SSA forecast (0.47 percent per year through 2029) and our Lee-Carter forecasts for U.S. data (1.13 percent per year) – so bondholders may believe that official forecasts are underestimating the rate of decrease in future mortality. To sum up, in just the same way that there is an equity premium puzzle, with stocks appearing to offer a return that is excessive in relation to their risk, there also appears to be a “mortality premium puzzle,” with investors overpaying for the insurance provided by the EIB bond.

4.5 Pricing Aggregate Mortality Risk

Our last goal is to determine the impact of aggregate mortality risk on annuity prices by recalculating annuity prices on the assumption that the insurance company earns not the standard bond rate of return, but that rate of return minus the mortality risk discount. In the previous section, we calculated the risk discount for mortality bonds at the time of issue when the cohort on whose lives the payments are based is aged 65, and the duration of the bond is 25 years. The discount may vary, though, as both the cohort ages and the number of remaining payments decreases.

We therefore repeat the calculations for 25-year bonds issued at the same time for each succeeding age beyond 65, computing the mortality bond risk discount at each age until expiration. We find that it continues to be extremely small at all ages. We recalculate the annuity's expected present value, assuming that the insurance company discounts each period's annuity payment not at the risk-free rate, but at the lower age-related rates that we just obtained. We then calculate the markup that the insurance company must charge to compensate for the lower rates. Given the very small mortality risk discounts, the cost of an annuity increases only

widespread, investors should find it worthwhile to collect contemporaneous mortality data, and so bond prices would reflect current mortality shocks.

very slightly. For men and women of all ages, the increase is less than one percent at a coefficient of risk aversion of ten. So, if annuity providers were to hedge their aggregate mortality risk on the capital markets, and if those capital markets were to price that risk in accordance with the predictions of the CCAPM, the effect on annuity prices would be extremely small.

Even if annuity providers were to pay the much higher 20 basis point mortality risk discount observed on the EIB bond, the effect on the price of a joint life 100 percent survivor annuity at age 65 would be only 2.1 percent, assuming a three percent real interest rate. At age 85, the difference is only one percent. This is a considerably smaller amount than the results reported in Table 2, where we found that markups of three to four percent would be required to reduce the probability of making a loss to just five percent.

5. CONCLUSIONS

We conclude that annuity business exposes insurance companies to significant aggregate mortality risk and that insurance companies need to maintain substantial capital against mortality shocks. An alternative is for the insurance companies to hedge that risk through the use of mortality contingent bonds. Our calculations indicate that this might be accomplished at what is, for all practical purposes, zero cost, according to the CCAPM. Even if annuity providers were to pay the mortality risk discount observed on the EIB bond, the effect on annuity prices would be considerably less than the markups required to reduce the probability of making a loss to just five percent.

The United States annuity market is currently extremely small. The National Association for Variable Annuities (2002) reports that immediate annuity sales were only \$10.2 billion in 2001.²⁸ Brown and Poterba (2000) point out that approximately half of all immediate annuity sales are period certain rather than life-contingent, so sales of immediate life contingent annuities may have been about \$5.1 billion in 2001. In the near future, the demand for annuitization may rise as unannuitized defined contribution plans displace annuitized defined

²⁷ The effect of a mortality shock on bond returns is amplified by the fact that a current period shock affects not only current but also future mortality. If investors failed to understand this, the required risk premium would be even smaller.

²⁸ Sales of deferred annuities substantially exceed those of immediate annuities. A deferred annuity enables people to accumulate wealth with the benefit of tax-deferral but lacks the essential feature of an immediate annuity, namely

benefit pension plans. Whether this demand is satisfied will be affected by the risk premium insurance companies require for accepting aggregate mortality risk, and their willingness to commit additional capital to this line of business.

the irrevocable exchange of a capital sum in return for a lifetime income. Individuals may subsequently withdraw deferred annuity assets in the form of a lifetime income, but rarely do so.

REFERENCES

- Beelders, Owen and David Colarossi. 2004. Modeling Mortality Risk with Extreme Value Theory: The Case of Swiss Re's Mortality-Indexed Bonds. *Global Association of Risk Professionals* (19): 26-30.
- Bell, William R. 1997. Comparing and Assessing Time Series Methods for Forecasting Age Specific Demographic Rates. *Journal of Official Statistics* (13): 279-303.
- Bell, Felicitie and Michael Miller 2005 Life Tables for the United States Social Security Area 1900-2100 Actuarial Study No. 120 Social Security Administration, Office of the Chief Actuary
- Blake, David. 2001. Survivor Bonds: Helping to Hedge Mortality Risk. *Journal of Risk and Insurance* 68(2): 339-348.
- Booth, Heather, John Maindonald, and Len Smith. 2002. Applying Lee-Carter Under Conditions of Variable Mortality Decline. *Population Studies* (56): 325-336.
- Brouhns, Natacha, Michel Denuit, and Jeroen K. Vermunt. 2002. A Poisson Log-Bilinear Regression Approach to the Construction of Projected Lifetables. *Insurance: Mathematics and Economics* (31): 373-393.
- Brown, Jeffrey R. and James Poterba. 2000. Joint Life Annuities and Annuity Demand by Married Couples. *Journal of Risk and Insurance* 67(4): 527-53.
- Cairns, Andrew, David Blake, and Kevin Dowd. 2005a. A Two-Factor Model for Stochastic Mortality with Parameter Uncertainty. Working Paper, School of Mathematic and Computer Sciences, Heriot-Watt University, Edinburgh.
- Cairns, Andrew, David Blake, and Kevin Dowd. 2005b. Pricing the Risk on Longevity Bonds. Preprint, School of Mathematic and Computer Sciences, Heriot-Watt University, Edinburgh.
- Dahl, Mikkel. 2004. Stochastic Mortality in Life Insurance: Market Reserves and Mortality-Linked Insurance Contracts. *Insurance: Mathematics and Economics* 35(1): 113-136.
- Deaton, Angus and Christina Paxson. 2004. Mortality, Income, and Income Inequality Over Time in Britain and the United States. National Bureau of Economic Research Working Paper No. 8534.

- Dushi, Irena and Anthony Webb. 2004. Rethinking the Sources of Adverse Selection in the Annuity Market. In *Christian Gollier* ed. Forthcoming.
- Gong, Guan and Anthony Webb 2005 Mortality heterogeneity and the distributional consequences of mandatory annuitization *Unpublished Working Paper*
- Lee, Ronald and Lawrence R. Carter. 1992. Modeling and Forecasting U.S. Mortality. *Journal of the American Statistical Association* 87(419): 659-671.
- Lee, Ronald and Timothy Miller. 2001. Evaluating the Performance of the Lee-Carter Method for Forecasting Mortality. *Demography* 38(4): 537-549.
- Lee, Ronald. 2000. A General Framework for Stochastic Investigations of Mortality and Investment Risks. Presented at the Wilkiefest, Heriot-Watt University.
- Lin, Yijia, and Samuel Cox. 2004. Securitization of Mortality Risk in Life Annuities. Presented at Managing Retirement Assets Symposium, Sponsored by the Society of Actuaries.
- Mehra, Rajnish, and Edward Prescott. 2003. The Equity Premium in Retrospect. In *Handbook of the Economics of Finance*. Edited by George Constantinides, Milton Harris, and Rene Stulz. JAI Press.
- Milevsky, Moshe and David Promislow. 2001. Mortality Derivatives and the Option to Annuities. *Insurance: Mathematics and Economics* 29(2001): 299:318
- Mitchell, Olivia, James Poterba, Mark J. Warshawsky and Jeffrey J. Brown. 1999. New Evidence on the Money's Worth of Individual Annuities. *American Economic Review*. 89(5): 1299-1318.
- National Association for Variable Annuities (NAVA). 2002. Annuity Fact Book. Reston: NAVA.
- Piggott, John, Emiliano Valdez and Bettina Detzel. 2004. The Simple Analytics of a Pooled Annuity Fund. *Journal of Risk and Insurance*, forthcoming.
- Preston, Samuel and Haidong Wang 2005 "Sex Mortality Differentials in the United States: The Role of Cohort Smoking Patterns" *Boettner Center for Pensions and Retirement Research*, The Wharton School, Working Paper 2005-1
- Renshaw, Arthur E. and Steven Haberman 2003a "Lee-Carter Mortality Forecasting With Age-Specific Enhancement" *Insurance: Mathematics and Economics* 33 pp 255-272.

- Renshaw, Arthur E. and Steven Haberman 2003b “On the Forecasting of Mortality Reduction Factors” *Insurance: Mathematics and Economics* 33 pp 379-401
- Sanderson, Warren and Sergei Scherbov. 2004. Putting Oeppen and Vaupel to Work: On the Road to New Stochastic Mortality Forecasts. *International Institute for Applied Systems Analysis*. Working Paper.
- Schalick, Lisa Miller, Wilbur C. Hadden, Elsie Pamuk, Vicente Navarro, and Gregory Pappas 2000 “The Widening Gap in Death Rates Among Income Groups in the United States from 1967 to 1986” *International Journal of Health Services* Vol. 30 No. 1 pp13-26
- Tuljapurkar, Shripad, Nan Li and Carl Boe. 2000. A Universal Pattern of Mortality Decline in the G7 Countries. *Nature* (405): 789-792.
- Social Security Advisory Board. 1999. *Technical Panel on Assumptions and Methods – Report to the Social Security Advisory Board*. (November).
- Wang, Shaun. 2000. A Class of Distortion Operators for Pricing Financial and Insurance Risks. *Journal of Risk and Insurance* 67(1): 15-36.
- Willets, Richard. 1999. Mortality in the next millennium. Paper presented in the *Staple Inn Actuarial Society*.
- Wong-Fupay, Carlos and Steven Haberman. 2004. Projecting Mortality Trends: Recent Developments in the United Kingdom and the United States. *North American Actuarial Journal* 8(2): 56-82.
- Yang, S. 2001. Reserving, Pricing, and Hedging for Guaranteed Annuity Options. Ph.D Thesis, Heriot-Watt University.

Figure 1

Actual vs. Forecast Survival
- Males -




Figure 2




Table 1 - Comparison of Lee-Carter with Projection Scale AA Annual Mortality Reductions

Age	0	1-4	5-9	10-14	15-19	20-24	25-29	30-34
Percentage reduction:								
Lee-Carter	3.36%	4.12%	3.41%	3.10%	1.75%	1.97%	2.20%	2.28%
Society of Actuaries	2.00%	2.00%	2.00%	1.97%	1.69%	1.63%	0.93%	0.70%
Age	35-39	40-44	45-49	50-54	55-59	60-64	65-69	70-74
Percentage reduction:								
Lee-Carter	2.18%	1.95%	1.64%	1.41%	1.24%	1.08%	1.06%	1.07%
Society of Actuaries	0.93%	1.25%	1.62%	1.66%	1.14%	0.98%	0.91%	1.04%
Age	75-79	80-84	85-89	90-94	95-99	100-104		65-104
Percentage reduction:								
Lee-Carter	1.19%	1.13%	1.13%	1.13%	1.13%	1.13%		1.13%
Society of Actuaries	0.99%	0.77%	0.52%	0.30%	0.16%	0.03%		0.80%

Sources:

Society of Actuaries - authors' mortality weighted average of Society of Actuaries projection scale AA data. This projection scale is a blend of Federal Civil service and Social Security mortality improvements from 1977 to 1993 subject to various adjustments discussed in the text.

Lee Carter - authors' calculations based on the Lee-Carter model. The average for 65-104 year olds is mortality weighted.

Table 2 - Potential Losses Arising From Aggregate Mortality Risk

3% interest rate				
5% loss probabilities				
	Single men	Single women	Married couples	
Survivor benefit			50%	100%
Age				
65	3.89%	3.67%	3.74%	3.91%
70	4.08%	3.74%	3.55%	3.69%
75	3.96%	3.87%	3.90%	4.10%
80	3.91%	3.81%	3.77%	4.12%
85	3.70%	3.64%	3.67%	4.01%
1% loss probabilities				
	Single men	Single women	Married couples	
Survivor benefit			50%	100%
Age				
65	5.45%	5.24%	5.39%	5.53%
70	5.87%	5.38%	5.05%	5.17%
75	5.97%	5.45%	5.62%	5.97%
80	5.61%	5.57%	5.51%	5.97%
85	5.25%	5.28%	5.24%	5.79%
5% interest rate				
5% loss probabilities				
	Single men	Single women	Married couples	
Survivor benefit			50%	100%
Age				
65	3.11%	2.82%	2.94%	2.87%
70	3.30%	3.00%	2.79%	3.12%
75	3.43%	3.15%	3.26%	3.34%
80	3.35%	3.23%	3.33%	3.56%
85	3.28%	3.26%	3.34%	3.59%
1% loss probabilities				
	Single men	Single women	Married couples	
Survivor benefit			50%	100%
Age				
65	4.46%	3.95%	4.29%	4.08%
70	4.80%	4.27%	3.87%	4.54%
75	4.94%	4.61%	4.70%	4.71%
80	4.82%	4.74%	4.71%	5.25%
85	4.87%	4.72%	4.73%	5.13%

Note: Analyses are for the 1924 birth cohort, assuming population mortality. They show the amounts by which the total payments by the insurance company on annuities sold to a single birth cohort will exceed the amounts forecast using the Lee-Carter model at the 95th and 99th percentiles of the distribution of payments, assuming that the only source of variation is aggregate mortality risk.

Table 3 - Percentage Underpricing Resulting From Use of Projection Scale AA

Survivor benefit	Female	Male	Couple	
			50%	100%
Age				
65	10.58%	8.73%	9.70%	10.47%
70	10.58%	9.07%	9.87%	10.85%
75	10.53%	9.65%	10.12%	11.22%
80	10.00%	9.59%	9.80%	11.01%
85	9.55%	9.03%	9.30%	10.56%

Note: Excess of expected present value over premium paid arising out of use of projection scale AA to price annuities when mortality improvement follows Lee-Carter model.

Table 4 - Impact on Bond Price of one Standard deviation Mortality Shock

Age of cohort	Interest rate	
	3%	5%
65	0.72%	0.61%
70	0.86%	0.73%
75	1.00%	0.88%
80	1.13%	1.02%
85	1.27%	1.17%

Note: Percentage impact on price of EIB type mortality bond of a one standard deviation permanent mortality shock.

**RECENT WORKING PAPERS FROM THE
CENTER FOR RETIREMENT RESEARCH AT BOSTON COLLEGE**

How Much Is the Working-Age Population Saving?

Alicia H. Munnell, Francesca Golub-Sass, and Andrew Varani, October 2005

Top Ten Myths of Social Security Reform

Jeffrey R. Brown, Kevin Hassett, and Kent Smetters, October 2005

What Replacement Rates Do Households Actually Experience in Retirement?

Alicia H. Munnell and Mauricio Soto, August 2005

Bequests, Inheritances and Family Traditions

Donald Cox and Oded Stark, August 2005

Local Labor Market Conditions and Retirement Behavior

Dan A. Black and Xiaoli Liang, May 2005

Design and Implementation Issues in Swedish Individual Pension Accounts

R. Kent Weaver, March 2005

Lashed to the Mast?: The Politics of Notional Defined Contribution Pension Reforms

Sarah M. Brooks and R. Kent Weaver, January 2005

Understanding Expenditure Patterns in Retirement

Barbara A. Butrica, Joshua H. Goldwyn, and Richard W. Johnson, January 2005

Changes in the Distribution of Long-Run Earnings and Retirement Incomes—Have Recent Cohorts Fallen Behind?

Peter Gottschalk and Minh Huynh, January 2005

The Age Profile of Income and the Burden of Unfunded Transfers in Four Countries: Evidence from the Luxembourg Income Study

Gary Burtless, December 2004

Projecting Immigration: A Survey of the Current State of Practice and Theory

Neil Howe and Richard Jackson, December 2004

All working papers are available on the Center for Retirement Research website (<http://www.bc.edu/crr>) and can be requested by e-mail (crr@bc.edu) or phone (617-552-1762).