



DISCUSSION PAPER PI-0613

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

Leora Friedberg and Anthony Webb

May 2006

ISSN 1367-580X

The Pensions Institute
Cass Business School
City University
106 Bunhill Row London
EC1Y 8TZ
UNITED KINGDOM

<http://www.pensions-institute.org/>

LIFE IS CHEAP:
USING MORTALITY BONDS TO HEDGE
AGGREGATE MORTALITY RISK

Leora Friedberg
Anthony Webb

Anthony Webb, Ph.D., (corresponding author) is a research economist at the Center for Retirement Research at Boston College (CRR), located at 258 Hammond Street, Chestnut Hill, MA 02467-3808. 646-662-7254 (phone), 617-552-0191(fax). Webbaa@bc.edu. Leora Friedberg, Ph.D., is an Assistant Professor in the Department of Economics of the University of Virginia, P.O. Box 400182, Charlottesville, VA 22904-4182. 434-924-3225 (phone), 434-982-2904 (fax). lfriedberg@virginia.edu. We would like to thank Jeff Brown, Steve Cooperstein, Francesca Golub-Sass, Jim Poterba, John Ralfe, Mauricio Soto, Bob Triest, and seminar participants at the Center for Retirement Research, Boston College, for their very helpful comments. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

©2006 by Leora Friedberg and Anthony Webb. 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.

Life is Cheap: Using Mortality Bonds to Hedge Aggregate Mortality Risk
Leora Friedberg and Anthony Webb

ABSTRACT

Using the widely-cited Lee-Carter mortality model, we quantify aggregate mortality risk as the risk that the average annuitant lives longer than is predicted by the model, and we conclude that annuity business exposes insurance companies to substantial mortality risk. We calculate that a markup of 3.7% on an annuity premium (or else shareholders' capital equal to 3.7% of the expected present value of annuity payments) would reduce the probability of insolvency resulting from uncertain aggregate mortality trends to 5% and a markup of 5.4% would reduce the probability of insolvency to 1%. Using the same model, we find that a projection scale commonly referred to by the insurance industry underestimates aggregate mortality improvements. Annuities that are priced on that projection scale without any conservative margin appear to be substantially underpriced. Insurance companies could deal with aggregate mortality risk by transferring it to financial markets through mortality-contingent bonds, one of which has recently been offered. We calculate the returns that investors would have obtained on such bonds had they been available over a long period. Using both the Capital and the Consumption Capital Asset Pricing Models, we determine the risk premium that investors would have required on such bonds. At plausible coefficients of risk aversion, annuity providers should be able to hedge aggregate mortality risk via such bonds at a very low cost.

Leora Friedberg
Department of Economics
University of Virginia
114 Rouss Hall
P.O. Box 400182
Charlottesville, VA 22904-4182
and NBER
lfriedberg@virginia.edu

Anthony Webb
Research Economist
Center for Retirement Research at Boston College
258 Hammond Street
Chestnut Hill, MA 02467-3808
webbaa@bc.edu

INTRODUCTION

Annuities provide a means by which risk-averse households facing an uncertain lifespan can insure themselves against the risk of outliving their wealth in old age. Annuitization is predicted to raise average expected utility under a variety of assumptions (Brown, Davidoff, and Diamond 2005), yet voluntary annuity purchases are rare. Only 10.2% of a large sample of elderly households had voluntarily annuitized any of their wealth between 1993 and 2000 (Dushi and Webb 2004).¹ Immediate annuity sales were only \$10.2 billion in 2001 (NAVA 2002), and approximately half of that total may have been life-contingent (Brown and Poterba 2000).² Interest in private annuities may well jump in coming years, though, as pre-annuitized defined benefit (DB) pensions provided by employers are largely replaced by lump-sum defined contribution (DC) pensions, and also if Social Security is similarly transformed using private accounts.³

The meager market for private annuities remains a puzzle. One reason 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.⁴ The former risk can be eliminated through diversification. The latter risk is non-diversifiable unless the insurer can write other classes of business, for example life insurance, with negatively correlated risks.⁵ Notably, one explanation for the shift in private pension structure may be a growing reluctance of employers to take on aggregate mortality risk (ACA 2005, pp.13-14).

The academic literature on the value of annuitization to households assumes that average future mortality is known with certainty. In reality, there is considerable uncertainty about the possible course of mortality reductions. Forecasting life expectancy 50 or 100 years

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

² Sales of deferred annuities substantially exceed those of immediate annuities. A deferred annuity enables people to accumulate wealth while deferring taxes but lacks the essential feature of an immediate annuity, namely the irrevocable exchange of a capital sum in return for a lifetime income. Individuals may withdraw deferred annuity assets in the form of a lifetime income, but rarely do so.

³ Some Social Security privatization proposals involve annuitization. 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 turn 75.

⁴ Since adverse selection is not our focus, we simplify the analysis by ignoring uncertainty over the degree of selection experienced by the insurer.

⁵ The actual scope for doing so may be limited because the average ages of life insurance and annuity holders differ and mortality shocks differ in magnitude and occasionally even in sign across ages.

hence involves taking a controversial position on the potential for either scientific discoveries to slow the aging process or emergent conditions to hasten it. The potential for longevity gains over such time horizons generates considerable disagreement. 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 recently born will live to 150 by 2150.⁶

Insurance companies and potential annuitants make their decisions over much shorter time horizons. Over such horizons, it is common to use models that combine extrapolations of past trends with some consideration of likely medical threats or advances, while avoiding a strong position on medical progress or the biology of aging. Purely 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.

Yet, past forecasts from these models indicate that point estimates of expected mortality rates carry wide margins of error. Insurance companies and other annuity providers that base 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 doing so 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 through financial markets by issuing mortality-contingent bonds, with interest payments that depend on mortality experience.⁷

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. As far as we know, no other paper in the economics literature applies the tools of financial economics to price aggregate risks which beset the insurance industry. In Section 1 of the paper, 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.” The Lee-Carter model has been adopted by the U.S. Census Bureau, among others, and was viewed favorably by the Social Security Administration’s

⁶ Due to the power of compound interest, the present value of the wager is considerably more modest!

⁷ The pricing and availability of reinsurance may depend on reinsurers’ ability to transfer the same risk yet more broadly. Later, we discuss relevant evidence and describe how mortality bonds can be used to effect this transfer.

1999 technical advisory panel.⁸ We compare its predictions with those of the Social Security Administration and published life tables. We also consider the sensitivity of these and other calculations in the paper to a large temporary shock, like a repetition of the 1918 flu pandemic, and to discrete medical gains like curing cancer and heart disease. In Section Two, we use the model to quantify the magnitude of aggregate mortality risk faced by a financial institution 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 aggregate mortality-related losses.

In Section Three, we investigate whether a 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 such forecasts. We calculate that if the Lee-Carter model provides an unbiased estimate of the pace of mortality improvement among annuitants, then insurance companies that use Projection Scale AA to price annuities without any conservative margin will underestimate that improvement substantially. Such annuities will be underpriced by amounts that range from 8.7 to 11.2% of the premium paid, relative to annuities with prices based on the Lee-Carter model. These findings of potential underpricing deepen the puzzle of low voluntary annuitization rates.

In Section Four, we explain the structure of mortality-contingent bonds, which could transfer aggregate mortality risk from borrowers to bondholders, and we price them using the Capital Asset Pricing Model (CAPM) and the Consumption Capital Asset Pricing Model (CCAPM). The CCAPM extends the CAPM by relating the risk premium that investors require to the covariance between the risky return on an investment and the growth rate of consumption. According to the CCAPM, investors will prefer assets that offer high returns just when consumption happens to be low, relative to assets such as stocks that tend to offer high returns when consumption is high. They will demand an expected return that exceeds the risk-free rate to hold the latter and will accept a return of less to hold the former.

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 of those returns with the stock market was close to zero, while the correlation with per capita

⁸ The Social Security Technical Panel on Assumptions and Methods (1999), p.64.

consumption growth turns out to be slightly negative, since an unexpected mortality decrease may stretch 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 mortality-risk discount on a longevity bond that was recently proposed by EIB but not issued as planned. Although small, the proposed discount is greater than that which is predicted by the CCAPM – so investors would have had to pay more than the model predicts for the opportunity to hedge aggregate mortality risk. We discuss reasons why it may be cheaper to hedge aggregate mortality risk in financial rather than reinsurance markets, as EIB intended.

Section Five summarizes. We conclude that insurance companies face substantial aggregate mortality risk. Moreover, we find that a projection scale commonly referred to by the insurance industry underestimates aggregate mortality improvements, so insurance companies may be underpricing annuities substantially. Were insurance companies instead to hedge the aggregate mortality risk predicted by Lee-Carter using mortality bonds or similar instruments, we calculate that they could do so at extremely low cost. These findings are also relevant 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 DB pension plans that are typically annuitized continue to be displaced by 401(k) and other DC plans, in which annuitization is almost never mandatory and usually not even an option, and as such concerns lead to calls for government action to cover this risk (Blake and Burrows 2001).⁹

1. QUANTIFYING AGGREGATE MORTALITY RISK

1.1 The Lee-Carter Model

Wong-Fupay and Haberman's (2004) review of academic research on population mortality forecasts has a valuable discussion about using mortality models to quantify the uncertainty surrounding mortality forecasts. Actuarial models are often unsuited to this purpose because they focus on point estimates, and many of the intervals bounded by high and low scenarios in official mortality projections appear much too narrow. Blake and Burrows (2001) show that plausible assumptions about forecasting errors can lead to quite imprecise estimates of

⁹ Brown et al (2004) note that government entry into reinsurance for terrorism risk in 2002 appears to have crowded out private market responses.

actuarially fair annuity yields. With mortality uncertainty rising over longer time horizons, he finds that confidence intervals are wider for escalating annuities whose benefits are back-loaded.

However, studies of that 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. We employ the Lee-Carter (1992) model, which is, as we described earlier, the most widely-used model of population mortality improvements. In the Lee-Carter model, mortality risk m at age x in year t is

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

The parameters a and b vary with age. Lee-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. They estimate that a random walk with drift fits the time path of k , as follows:

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

where flu is the impact of the 1918 epidemic. According to these results, a one standard-deviation shock to k translates into a roughly two-month change in age-65 life expectancy. Lee-Carter conclude that the fitted model explains well over 90% of within-age group variances in mortality rates.

We will use the parameters estimated from this model to forecast mortality rates at various ages for all future years. Importantly for our purposes, we can calculate both unconditional and conditional forecasts of future mortality rates, plus associated confidence intervals. For example, we can calculate trends in life expectancy – the model predicts life expectancy at birth in 2065 of 86.05 years, with a 95% 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 potential weaknesses which have been noted by others, we do not generally need to address them for our application. First, the parameter values depend on

the period over which they are estimated. They do not vary much, though, so our results will be relatively unaffected by our choice of base period. Second, the error term may not capture large but infrequent mortality shocks – for example, 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, such transitory shocks are asymmetric. Third, the model does not fit age-specific mortality data exactly in the jump-off year. Bell (1997) shows that setting a_x to fit the initial conditions exactly 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 errors associated with estimating the parameters a and b dominate 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) find a substantial age-time interaction that improves the fit of the model to Australian data. We do not incorporate such a term, as its importance has not been demonstrated for the U.S. and as we find below that, in any given year, k varies little with age. Sixth, the Lee-Carter model was estimated on combined male and female data. Although 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. Seventh, Ruhm (2004) describes evidence of a statistically significant correlation between U.S. business cycle conditions and aggregate mortality rates. However, the estimated effects of cyclical fluctuations are small and temporary, whereas the larger shocks that Lee-Carter identify affect not only current but also future mortality.

We note in passing that some authors have proposed other enhancements to the Lee-Carter model (Renshaw and Haberman 2003a, 2003b, Brouhns, Denuit, and Vermunt 2002). Alternately, other stochastic models besides Lee-Carter are available (Lee 2000, Yang 2001, Milevsky and Promislow 2001, Dahl 2004, Currie, Durban, and Eilers 2004, Cairns, Blake, and Dowd 2005a). In contrast to Lee-Carter and most other models, Sanderson and Scherbov (2004) project no deceleration in the rate of mortality improvement, based on mortality data across fourteen countries.¹⁰ We retain the original model, though, since it is widely accepted and has published parameter values for U.S. data, and the differences in the predictions of other models (besides Sanderson-Scherbov) are not substantial.

¹⁰ 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.

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 both Lee-Carter and Social Security Administration (SSA) forecasts from the 1993 Trustees' Report can be compared with actual mortality improvements; and 2001 onwards.¹¹

Figures 1 and 2 compare the model's predictions of survival rates starting at age 65 for the male and female birth cohort of 1924 with those made by the 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% per year for the over-65s. In contrast, the SSA forecasts male and female mortality declines of 0.54 and 0.50% for the entire period 1989 to 2017. It transpired that mortality declined much more rapidly among men than women from 1989 to 2001, with an average mortality-weighted decline in mortality of 0.84% a year among men, but an average increase of 0.06% a year among women.¹² Therefore, the decline in female mortality in Figure 2 was lower than either the Lee-Carter or the SSA forecasts, whereas the decline in male mortality in Figure 1 proved to be almost midway between the two forecasts.

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 increased smoking prevalence among women earlier on. They project an acceleration in mortality reductions even if smoking rates do not change further, with male mortality declining yet faster because male smoking rates have declined considerably, whereas female smoking rates have not. We conclude that the 1989-2001 period represents a temporary pause in the decline in female mortality.

Looking forward from 2001, the SSA continues to be less optimistic than Lee-Carter. The 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% under the three forecasts, increasing rapidly to the 0.29, 0.67, and 1.17% assumed for 2029 and later.¹³ The intermediate forecast of 0.47% initially and 0.67% from 2029 predicts that mortality will initially decline at the pace experienced between 1979 and 2001 and then increase rapidly to

¹¹ The SSA actuaries forecast mortality by combining age-specific trend extrapolation with information collected from medical experts.

¹² Bell and Miller (2005) found a similar differential for male and female mortality over the period 1981-2001.

¹³ The SSA does not break down their projected declines by age within this group.

almost the pace of 0.72% experienced between 1900 and 2001.¹⁴ The high forecast predicts that the somewhat higher pace of mortality declines experienced between 1900 and 2001 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% initially. 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% a year.¹⁵ It should be noted that the SSA "high" forecast of 1.17% from 2029 on is almost the same as Lee-Carter. By comparison, mortality declines exceeded the rates predicted by the model over some fairly lengthy earlier periods as well. 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

Next, we 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, while noting that these populations may differ. Dushi and Webb (2004) show that annuity purchasers in the U.S. are much wealthier than average. 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, but it is difficult to infer whether mortality differentials have widened since 1986.¹⁶

Table 1 compares, by age group, the mortality declines projected by the Lee-Carter model and by the SOA in their Projection Scale AA, the most recent that the 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

¹⁴ 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% per year from 1900 to 2001 and 0.72% per year from 1979 to 2001.

¹⁵ 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.

¹⁶ Table 7.1 in "The RP-2000 Mortality Tables" available from www.soa.org 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 can discern no systematic differences among them

and lower bounds on the forecasted improvements. When preparing the RP-2000 life tables, the SOA considered whether scale AA should be updated but concluded that it was consistent with recent data on mortality improvements.¹⁷

The SOA forecasts a slower pace of mortality reductions than the Lee-Carter model does, which is unsurprising since the SOA data is based on a limited time period when mortality declined relatively slowly. The age-sex weighted average for the SOA is 0.71% for the whole population and 0.64% 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. If the annuitant population continues to enjoy better than average mortality reductions, even the Lee-Carter forecasts may prove to be too conservative.

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 annual survival probabilities. Mortality rates at old ages are even more important to insurance companies because most annuities are purchased by older households.¹⁸ It will be important to keep this issue in mind 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 an annuity provider selling level or increasing 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.¹⁹ To 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 population-average

that are related to their presumed average socio-economic status. However, as the RP-2000 report points out in Chapter 4, “the measurement of mortality improvement requires voluminous consistent data covering many years.”

¹⁷ “The RP-2000 Mortality Tables” page 73. Projection Scale AA is also reported in this publication in Table 7-3.

¹⁸ This distinction is less important for pension funds, since DB pensions, as deferred annuities, are “bought” by people of all ages.

¹⁹ It would not be feasible to take the alternative approach of using price data and backing out expected mortality improvements from that. A single price is consistent with many possible rates of improvement depending on the assumptions made regarding expenses, profits, current period mortality, and cost of capital.

mortality and changes in mortality rates as predicted by the Lee-Carter model.²⁰ 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 adverse selection, which likely affects this market already (Finkelstein and Poterba 2004). We assume that the annuity price is set at a level that will enable the provider to break even at expected mortality rates. We further assume that the provider sells either a real level annuity or a real annuity increasing at 3% per year, invests in a risk-free asset offering either a real 3 or 5% return, and has zero administrative expenses.²¹

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, and there is in each period the idiosyncratic mortality shock $e_{x,t}$ in equation (1), Lee-Carter show in their Appendix B that it is reasonable to ignore those 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, the values associated with Lee-Carter's estimated time trend of -0.365. We add this draw to -0.365 in order 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, the values associated with Lee-Carter's estimation of e_t .²² 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. Based on these 10,000 simulations, we construct separate male and female mortality tables, assuming that the ratio of mortality rates is equal to the gender-weighted average of population mortality rates reported by SSA by age and

²⁰ We chose to take current period population rather than annuitant mortality as our starting point as the Lee-Carter model was estimated on population rather than annuitant data. We obtain very similar results if we start out by using annuitant mortality instead.

²¹ A 3% 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. Insurance companies are starting to offer real annuities.

birth cohort.²³ 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 it experiences results in payments that exceed the premium net of mark-up. 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 as capital to reduce the probability of insolvency to those same percentages, or the combination of markup and capital that is required.

Table 2 reports our results for annuities issued to married couples (with survivor benefits of 50 and 100%), single women, and single men aged 65, 70, 75, 80, and 85.²⁴ For an insurance company selling joint life and 50% survivor annuities to couples aged 65, discounted at 3%, some combination of shareholders' capital and premium loading equal to 3.78% of the premium is required to reduce the probability of loss or insolvency to 5%. Shareholders' capital or loading of 5.37% is required to reduce the probability to 1%. For the same annuity sold to couples aged 85, the corresponding numbers are lower, at 3.73 and 5.26%. A higher interest rate reduces the required mark-up, since relatively less weight is attached to the payments made at older ages, which are increasingly risky. At an interest rate of 5 instead of 3%, the premium loadings are somewhat lower – 2.95 and 4.13% for couples aged 65

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 people, 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, as Milevsky (2003) points out, annuities sold to younger persons include a substantial period during which mortality rates are extremely low.

²² As noted earlier, a one standard-deviation shock to k translates into a change in age-65 life expectancy of about two months.

²³ It would be preferable to use annuitant life tables, but these are only available for periods, not cohorts.

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.

Both inflation-linked and increasing annuities are backloaded relative to level annuities in that they pay a greater proportion of their benefits at older ages. Both are therefore relatively more attractive to long-lived households but also expose the insurance company to additional aggregate mortality risk – as demonstrated by the results for 3% increasing annuities in the lower portion of Table 2. Brown, Mitchell, and Poterba (2002) report that inflation-linked annuities are sold at more actuarially unfair prices than level annuities. Finkelstein and Poterba (2004) obtain a similar result when comparing level with increasing annuities and attribute this to the effects of adverse selection, since increasing annuities are particularly attractive to low-mortality types. It is equally possible that the higher prices reflect the greater amount of aggregate mortality risk to which sellers of increasing annuities are exposed.

Our calculations are similar in spirit to those in Dowd, Cairns, and Blake (2005) of value-at-risk and expected shortfall for U.K. annuity providers.²⁵ They estimate a 90% value-at-risk for long-dated coupon-paying mortality bonds of 4-5% of the premium, depending on whether they allow for parameter uncertainty. While somewhat higher than our results, it is of the same order of magnitude.

Lastly, we consider the possibility of a sudden medical breakthrough, although even major innovations are likely to have incremental effects on mortality. The Lee-Carter model does not tell us about the likelihood of such a breakthrough at a point in time. Olshansky, Carnes, and Cassel (1990) calculate that eliminating cancer and ischemic heart disease, while a highly unrealistic prospect, would raise life expectancy at age 50 by 6.57 years for females and 7.83 years for males, while eliminating cancer, all circulatory diseases, and diabetes would raise life expectancy at 50 by 15.30 and 15.02 years. These latter gains correspond to a 75% reduction in mortality from all causes, which has a less than 0.01% probability of occurrence over a 30

²⁴ 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.

²⁵ They use the two-factor mortality model developed in Cairns, Blake, and Dowd (2005a). As we only have price data for a single mortality bond, it is not possible to determine the relative weights that financial markets might place on each factor. They calculate risk premiums for mortality bonds of various maturities and initial cohort ages, assuming alternative weightings.

year horizon under the Lee-Carter model. Such a shock would increase the expected present value of the payments on a joint life and 50% survivor annuity by 50.7%. Insurance companies would not actually suffer losses of that magnitude, though, as mortality would decline gradually and hence payments to existing annuitants would increase by lesser amounts, and insurance companies would raise prices to new purchasers.

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 this issue. To the best of our knowledge, no one has compared the Lee-Carter predictions with those of life tables. We find that, by the Lee-Carter benchmark, insurance companies systematically underprice annuities if they use Projection Scale AA without making compensating adjustments elsewhere in their pricing formulas. This finding is a corollary of our earlier results that actuarial life tables appear to understate expected aggregate mortality improvements.

3.1 Methodology

The problem in using price data to determine whether aggregate mortality risk is correctly priced is that we do not necessarily observe the insurance company's mortality assumptions. Prices depend as well on unobservable assumptions regarding expenses and asset returns. Given these difficulties, we proceed as follows. We assume that the Annuity 2000 mortality table correctly describes current mortality of people buying the annuity.²⁶ We further assume that insurance companies use Projection Scale AA, discussed previously, to forecast mortality improvements. We apply a 3% real interest rate and zero administrative costs and calculate the expected present value of a one dollar-a-year annuity, assuming in the first instance that mortality declines according to Projection Scale AA, and in the second instance that it declines as forecast by the Lee-Carter model. A comparison of these expected values tells us the extent to which using Projection Scale AA would contribute to underpricing of annuities.

²⁶ The Annuity 2000 life tables can be read in a table manager downloadable from 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.

It follows that the EPV will exceed 100% of the premium paid if and only if the Lee-Carter mortality forecasts are more optimistic than Projection Scale AA. As we are assuming that the Annuity 2000 table correctly describes current mortality among annuity purchasers, this result holds irrespective of the current period life table that the insurance company may use, although the magnitude of the difference would be slightly affected.

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 various types of data uncertainty – over not only the future path of mortality declines but also, for example, the degree of adverse selection that they may experience. The American Academy of Actuaries Variable Annuity Reserve Working Group proposes that actuaries be required to take account of mortality improvements when that would err on the side of caution, and permits them to do so otherwise.²⁷ 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 will therefore overstate the degree of underpricing to the extent that a particular insurer’s “conservative margin” is in fact protecting against population level aggregate mortality risk. On the other hand, our calculations will understate it if wealth and mortality rates are strongly correlated, as we pointed out earlier, and if socio-economic disparities in mortality risk are continuing to widen.

Table 3 reports our results. EPVs are in all cases greater than 100% of the premium paid, which indicates underpricing of annuities compared to annuities priced according to scale AA. The extent of the underpricing ranges from 8.7 to 11.2%. It is, at all ages, greater for women than for men, reflecting the fact that the scale AA projected mortality improvements are greater for men than women, 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.

Mitchell, Poterba, Warshawsky, and Brown (1999) undertake a related calculation to determine the value of annuities to potential annuitants, rather than evaluating the pricing of

²⁷ http://www.actuary.org/pdf/life/varwg_0305.pdf

annuities by insurance companies.²⁸ Again assuming that the Lee-Carter projections are unbiased, our results imply that Mitchell et al substantially overstate the actuarial unfairness of annuities. Their cohort tables are based on SSA forecasts of mortality improvements and may thus 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 survival probabilities. They conclude that 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 small differences in forecasted mortality declines, they found no evidence of systematic underestimation.

4. PRICING AGGREGATE MORTALITY RISK

Given our evidence of substantial aggregate mortality risk and possible underpricing of that risk by insurers, 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 Capital and the Consumption Capital Asset Pricing Models, present data on the relationship between mortality-contingent bond returns, stock market returns, and consumption growth, and price aggregate mortality risk accordingly.

4.1 Mortality-Contingent Bonds

As of 2005, only one mortality-contingent bond had been issued and another one proposed. In December 2003 Swiss Re issued a \$400 million three-year bond paying LIBOR plus 135 basis points. The bond provides that, if a five-country weighted mortality index exceeds 130% of the 2002 level (i.e., mortality rises substantially), then the principal will be reduced. If it goes above 150%, the principal will be exhausted. In November 2004, the European Investment Bank (EIB) announced a £540 million bond. It was to have a life of 25

²⁸ They compute the EPV of an annuity under the assumption that either population or annuity tables, with the latter converted into cohort tables, accurately describe the mortality of the birth cohort purchasing the annuity.

years and offer mortality-related payments instead of interest or return of principal. The payments on the EIB bond would 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.

The Swiss Re bond exploits the feature that financial markets may be willing to retain non-diversifiable risks at lower cost than insurance markets – a point that we will come back to later. A non-diversifiable risk may attract only a small premium on the financial markets if that risk is both small in relation to the market and is uncorrelated with market returns.

Cairns, Blake, and Dowd (2005b) report that the EIB bond was to be issued 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 was priced at a mortality risk discount, not a premium. While risk-averse annuity providers should be willing to pay a 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 drives our discussion of asset pricing.

Before proceeding with that, we will mention some issues related to the possible nature of mortality shocks. According to our calculations, it is unlikely that market prices will be influenced by expectations about transitory mortality shocks. In particular, the long duration of the EIB bond means that one-off shocks have little effect on its value. Most or all one-off shocks will be negative, increasing but not decreasing mortality temporarily, so we do not consider the possibility of large positive transitory shocks.²⁹ The only one-off shock incorporated into the Lee-Carter model was the 1918 flu epidemic. An extreme negative event like a repetition of the epidemic, which 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% at a 3% discount rate when the population on whose lives the payments are based is age 65. Similarly, it should be noted that Beelders and Colarossi (2004) use of extreme value theory, with its focus on the extreme tail of

²⁹ It is difficult to think of any substantial improvements in mortality (i.e., something more than an unexpectedly mild flu season) that are not long-lasting.

the probability distribution, to price the Swiss Re bond is less relevant for the much longer-horizon EIB bond.

By way of comparison, a much smaller but permanent one-standard deviation shock to k of 0.655 has a greater – but still modest – effect on the value of the EIB bond because forward-looking investors will recognize the change in future mortality. Table 4 shows the effect on the expected present value of an EIB-type bond, assuming that the shock took place in 1989, the last year of data used to estimate the Lee-Carter model. The shock reduces the value of the bond by 0.63 to 1.17%, depending on the age of the cohort. The reduction increases with age and decreases with the discount rate. These calculations demonstrate the magnitude of typical fluctuations that EIB bondholders might experience.³⁰

4.2 Historical Returns

We calculate the returns that investors would have experienced had the EIB bond been 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

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

or $b_x e_{t+1}$.

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. We compare this with the change predicted by Lee-Carter and recover, for each year, a vector of combined mortality shocks, $E_{x,t+1}$. We confirm, first, that the Lee-Carter specification fits the data well. An important assumption it makes is that the k shocks affect all ages in the same way. If this is true, then insurance companies would not be able to reduce their aggregate mortality risk much by diversifying across birth cohorts. We find that the error term $E_{x,t}$ does not, in fact, vary much with age x . In any given year, these mortality shocks are almost invariably of the same sign across ages, 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 4, discussed earlier. Second, we would expect $E_{x,t+1}^-$ to have a mean close to zero, and the sample mean of 0.034 is not significantly different from zero.

To continue, the financial impact of the persistent mortality shocks e_t 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 to k as a percentage of the market value of the bond prior to the shock, using a 3% interest rate. The resulting mean of this percentage is 0.04%, very close to zero as expected, and the standard deviation 0.64. This standard deviation is also very small – the general consensus is that the standard deviation of stock returns is about 17 to 20% by comparison, or roughly thirty times greater.³¹

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 Capital Asset Pricing Models

The Capital Asset Pricing Model (CAPM)

The CAPM decomposes risk into its systematic and specific components. Specific risk does not command a risk premium because it can be diversified away. Systematic risk cannot be similarly diversified away. The magnitude of the risk premium that an investor requires to hold a risky asset like a mortality bond equals its beta, defined as

³⁰ It would be difficult to make similar calculations of the impact on bond prices of cures for cancer and ischemic heart disease because we do not know the time interval over which such cures would take effect.

³¹ This low level of risk reflects the design of this particular bond. Just as derivatives are more volatile than the value of the assets underlying them, the return on a bond designed to pay out on mortality in excess of some floor would, for example, be much more volatile.

$$\beta_b = \text{cov} \frac{(R_b, R_m)}{\sigma_m^2} \quad (4)$$

the covariance of the bond with the market, divided by the variance of the market return. Thus, the expected return on the mortality bond equals

$$E(R_b) = R_f + \beta_b [E(R_m) - R_f] \quad (5)$$

the risk free rate plus beta multiplied by the excess of the expected market return over the risk free rate.

The Consumption Capital Asset Pricing Model (CCAPM)

Mehra and Prescott (2003) provide an overview of the CCAPM. In the CCAPM, what determines the premium for systematic risk is not its relationship to general market risk but rather to 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 is high and the marginal utility of consumption low. Conversely, investors place a high value on assets that offer high returns when consumption is low and the marginal utility of consumption high and will buy such assets even when the expected return is less than the risk-free rate. Because it prices systematic risk in relation to the rest of an individual's wealth and to the intertemporal nature of investment decisions, the CCAPM may be theoretically preferable to the CAPM when considering annuities, which are crucial for smoothing life-cycle marginal utility in the presence of lifespan risk.

Mathematically, the CCAPM implies that

$$E_t(R_{b,t+1}) = R_{f,t+1} + \text{Cov}_t \left\{ \frac{-U'(C_{t+1}), R_{b,t+1}}{E_t(U'(C_{t+1}))} \right\} \quad (6)$$

so the expected return on the risky mortality bond in period $t+1$ equals the risk-free rate plus the covariance of the bond return with the marginal utility of consumption. Some algebra and assumptions result in:³²

$$\text{Log}E_t \left[(1 + R_{b,t+1}) / (1 + R_{f,t+1}) \right] = \gamma \sigma_{b,c} \quad (7)$$

where γ is the coefficient of risk aversion and $\sigma_{b,c}$ is the log of the covariance of consumption growth with the bond return. A negative covariance commands a discount below the risk-free rate, so $R_f > R_b$, and the discount increases with the coefficient of risk aversion.³³

4.4 Pricing the EIB Bond Using the CAPM and CCAPM

Pricing Using the CAPM

We assume that the S&P 500 represents market returns and estimate the beta on EIB-type mortality bonds, had they been available over the period 1959-1999.³⁴ The result is a beta of 0.15 with a 95% confidence interval of [-0.15, 0.46]. Thus, there is little or no correlation with the market.³⁵ Since the immediate annuity market is quite small, at roughly \$10 billion, then transferring all of the aggregate mortality risk to financial markets would only expose an investor holding the market portfolio to a very small change in risk. Opinion differs as to the level of the equity premium, but at a plausible level of 500 basis points, the beta would suggest a premium of 75 basis points (since $75 = 0.15 * 500$) above the risk-free rate, although the confidence interval is wide.

Pricing Using the CCAPM

We hypothesize that there should be a negative covariance between the growth in per-capita consumption and returns on an EIB-type bond. In other words, we expect that the contemporaneous correlation between aggregate consumption shocks and mortality shocks is negative. Moreover, the positive correlation of equity returns with consumption shocks will

³² 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.

³³ Another alternative is to price mortality bonds using a Wang transform (Wang 2000), as in Lin and Cox (2004), but the results depend heavily on assumptions about the insurer's unobserved level of expenses.

³⁴ The S&P 500 is a capitalization-weighted index of large capitalization stocks that is strongly positively correlated with other domestic stock returns and hence a good proxy for the market portfolio. The correlations of the S&P 500 with other international equity markets is lower, so they potentially offer an even better opportunity to hedge U.S. mortality risk.

³⁵ The lack of any substantial correlation here is consistent with the small estimated effects of business cycle conditions on aggregate mortality which Ruhm (2004) describes.

further widen the gap between the beta predicted by the CAPM and the risk discount predicted by the CCAPM.

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 rise significantly in response because the capital stock is largely fixed at short time horizons, as is the labor supply of older people, since few work. In consequence, per-capita income will decline because output is relatively constant and must be shared among a greater-than-expected number of individuals.³⁶ Moreover, to the extent that the mortality shock is global, then one country would not be able to borrow from another to offset 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 co-vary negatively. 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.³⁷ Consumption growth is a relatively smooth series with a mean of 2.25% and a standard deviation of 1.19%.

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, and given that our estimate of the negative correlation between consumption growth and bond returns is statistically significant, we can also say that we do not believe they should attract

³⁶ 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.

³⁷ The correct variable, although it is more difficult to measure, is the consumption of stockholders, who bear aggregate mortality risk by owning insurance companies and employers that offer DB pensions. Consumption of stockholders is more highly correlated with the stock market (Mankiw and Zeldes 1991), which would increase the risk discount on an asset that negatively covaries with the market.

a risk premium.³⁸ In contrast, investors 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% 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%. Applying the CCAPM, the risk discount is only two basis points when the coefficient of risk aversion equals ten.

Although it is unwise to make too much of differences between two small numbers – the risk discount predicted by CCAPM and the greater discount proposed by EIB – there are three reasons to expect that the issuers of the EIB bond would require a larger discount to the risk-free rate than that predicted by the CCAPM. First, the extensive literature on the so-called “equity premium puzzle” shows that the risk premium on equities, whose returns are positively correlated with consumption growth, is several times greater than that predicted by the CCAPM. It is possible that risk preferences are such that investors would also accept a larger discount to that predicted by the CCAPM on assets whose returns are negatively correlated with consumption growth – resulting in a “mortality premium puzzle.”

Second, we are assuming that the British Government Actuary’s Department’s mortality estimates are believed by investors to be unbiased. If the market believed those estimates to be too conservative, then a bigger discount would be justified.³⁹ We calculate that the 20 basis point discount would be eliminated if investors believed that the actuaries were

³⁸ This point assumes that agents are aware that aggregate mortality has improved. 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 widespread, investors should find it worthwhile to collect contemporaneous mortality data, and so bond prices would reflect current mortality shocks.

underestimating mortality declines by 0.44% a year. This represents about two-thirds of the gap in the U.S. between the intermediate SSA forecast (0.47% per year through 2029) and our Lee-Carter forecasts for U.S. data (1.13% per year).

Third, in contrast to the Swiss Re bond that transferred aggregate mortality risk out to financial markets, the proposed bond transferred aggregate mortality risk in to EIB, which anticipated transferring it to insurance markets through the purchase of reinsurance from Partner Re. Both Dowd (2003), discussing aggregate mortality risk, and Smetters (2004), analyzing similarly non-diversifiable terrorism risk, note that daily gains and losses in financial markets dwarf potential losses from the risks in question. Smetters also takes note of legal and regulatory barriers (like the double taxation of investment income earned by insurance companies) that limit the capacity of the insurance market to bear risk. Thus, financial markets may be better positioned than smaller and constrained insurance markets to absorb these non-diversifiable risks, so using reinsurance instead to hedge mortality risk may have driven up the risk premium which EIB required. In evidence of this is a report that Partner Re had “little appetite for additional [EIB-type reinsurance] deals” (Cass Business School and Pensions Institute 2005). An alternative design to transfer the risk to financial markets is possible, perhaps making bondholders residual claimants after survivor benefits had been met.

To sum up, we have identified two reasons why the market may have accepted a bigger discount for the EIB bond than predicted by either CAPM or CCAPM. It is also possible that the discount which EIB required reflects mispricing due to the undeveloped nature of the market for mortality bonds or the limitations of insurance markets as opposed to general financial markets in absorbing this risk.

4.5 Pricing Aggregate Mortality Risk

Our last goal is to determine the impact of aggregate mortality risk by recalculating annuity prices on the assumption that the insurance company earns not the standard bond rate of return, but that rate 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

³⁹ We assume that a similar relationship between projected mortality bond returns and consumption growth holds in the U.K. as in the U.S.

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 EPV, assuming that each period's annuity payment is discounted 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 very slightly for men and women of all ages – with an increase of less than 1% at a coefficient of risk aversion of ten. So, if annuity providers were to hedge their aggregate mortality risk in 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 higher 20 basis point mortality risk discount proposed on the EIB bond, the effect on the price of a joint life and 100% survivor annuity at age 65 would be only 2.1%, assuming a 3% real interest rate. At age 85, the effect is only 1%. These are considerably smaller amounts than the markups (or required reserves) of around 4% that, as we reported in Table 2, would reduce the probability of making a loss to just 5%.

5. CONCLUSIONS

The U.S. annuity market is currently extremely small. In the near future, demand is expected to rise as unannuitized defined contribution plans displace annuitized defined benefit pension plans. The risk premium that insurance companies require for accepting aggregate mortality risk will influence annuity prices and therefore the extent to which the market expands.

We conclude that annuity business exposes insurance companies to substantial aggregate mortality risk and that insurance companies need to charge substantial markups or maintain substantial capital against mortality shocks. An alternative is for insurance companies to hedge that risk through the use of mortality-contingent bonds. Our calculations based on the CCAPM indicate that this might be accomplished at what is, for all practical purposes, zero cost. While we show how the risk could be transferred to bondholders, a viable (though probably less efficient) 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 sharing of aggregate mortality risk.

It is interesting to note that similar problems involving correlated risks may limit the scope of other insurance markets. Long-term care risks faced by individuals are large, yet the long-term care insurance market is very small. Cutler (1996) argues that aggregate uncertainty over the risk of utilization and cost increases may account for sharp limits on the coverage that is offered. The development of mechanisms to handle aggregate mortality risk may thus have applications in the long-term care insurance market. Smetters (2004) and Brown et al (2004) argue that private market responses to managing terrorism risk have been crowded out by public policy interventions. Our results support the argument that private markets may be able to manage these types of aggregate risks at little cost.

ACKNOWLEDGEMENTS

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 Marris Buessing, Francesca Golub-Sass, Jamie Lee, and Andrew Varani for excellent research assistance.

REFERENCES

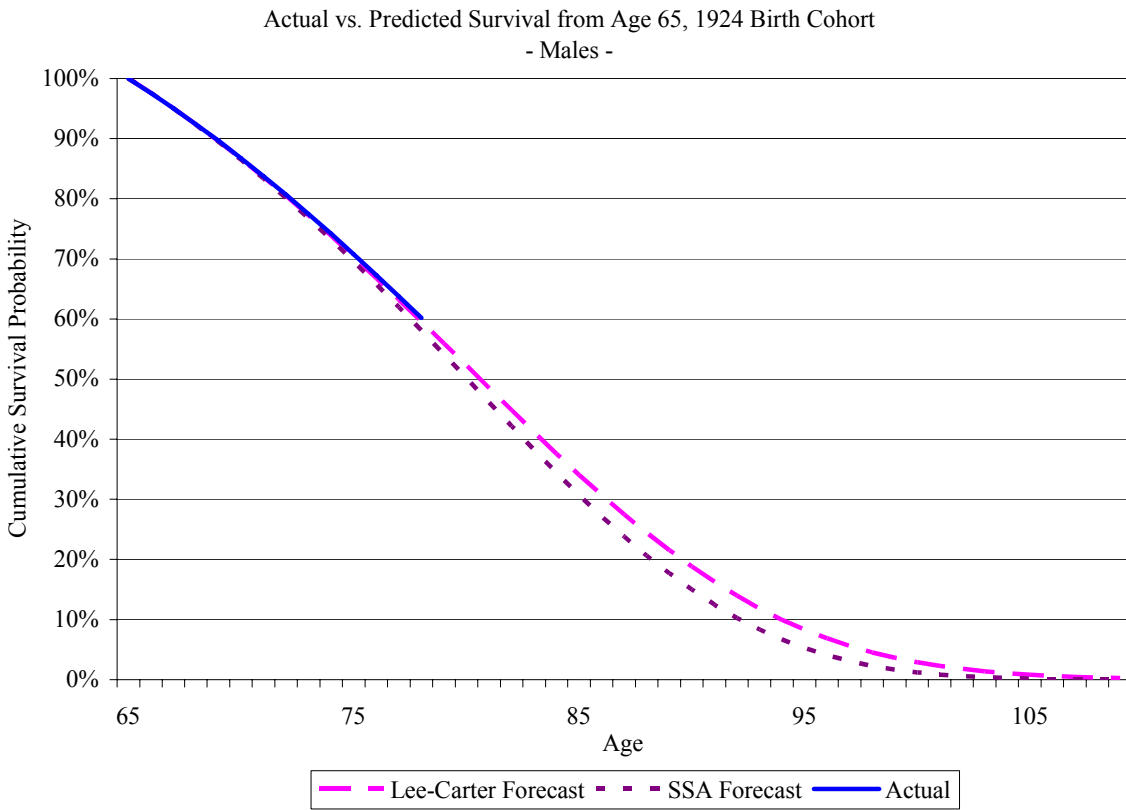
- Association for Consulting Actuaries. 2005. Pensions: Challenges and Choices. *Response to the First Report of the Pensions Commission*, 26 January.
- Beelders, Owen and David Colarossi. 2004. Modeling Mortality Risk with Extreme Value Theory: The Case of Swiss Re's Mortality-Indexed Bonds. *Global 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 and William Burrows. 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, J. David Cummins, Christopher Lewis, and Ran Wei. 2004. "An Empirical Analysis of the Economic Impact of Federal Terrorism Reinsurance." *Journal of Monetary Economics* 51: 861-898.
- Brown, Jeffrey, Thomas Davidoff, and Peter Diamond. 2005. "Annuities and Individual Welfare." *American Economic Review* 95(5): 1573-1590.
- Brown, Jeffrey, Olivia Mitchell, and James Poterba. 2002. Mortality Risk, Inflation Risk, and Annuity Products. In *Innovations in Retirement Financing*, University of Pennsylvania Press: Philadelphia, PA, 175-197.

- 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.*
- The Cass Business School and The Pensions Institute. 2005. Summary of the First International Conference on Longevity Risk and Capital Market Solutions. http://www.pensions-institute.org/conferences/longevity/conference_summary_18.02.05.pdf
- Currie, Iain, Maria Durban, and Paul Eilers. 2004. Smoothing and Forecasting Mortality Rates. *Working Paper, Heriot-Watt University.*
- Cutler, David. 1996. Why Don't Markets Insure Long-Term Risk? *Manuscript, Harvard University.*
- 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.*
- Dowd, Kevin. 2003. Survivor Bonds: A Comment on Blake and Burrows. *Journal of Risk and Insurance* 70(2): 339-348.
- Dowd, Kevin, Andrew Cairns, and David Blake. 2005. Mortality-Dependent Financial Risk Measures. *Centre for Risk and Insurance Studies Discussion Paper VII University of Nottingham.*
- Dushi, Irena and Anthony Webb. 2004. Rethinking the Sources of Adverse Selection in the Annuity Market. In *Christian Gollier* ed. Forthcoming.
- Finkelstein, A and Poterba, J. Adverse Selection in Insurance Markets: Policyholder Evidence from the U.K. Annuity Market *Journal of Political Economy*, 2004. Vol. 112(1): 183-208.

- 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*.
- Mankiw, N.Gregory, and Stephen Zeldes. 1991. The Consumption of Stockholders and Non-Stockholders. *Journal of Financial Economics* 29(1): 97-112.
- 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: 299-318
- Milevsky, Moshe and David Promislow. 2003. Can Life Insurance be used to Hedge Payout Annuities? Part 1: Modeling Longevity Risk. *York University, Toronto, Working Paper*.
- 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. 2002. Annuity Fact Book. *Reston: NAVA*.
- Olshansky, Jay, Bruce Carnes, and Christine Cassel. 1990. In Search of Methuselah: Estimating the Upper Bounds to Human Longevity. *Science, New Series*, 250(4981): 634-640
- 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: 255-272.
- Renshaw, Arthur E. and Steven Haberman. 2003b. On the Forecasting of Mortality Reduction Factors. *Insurance: Mathematics and Economics* 33: 379-401
- Ruhm, Christopher. 2004. Macroeconomic Conditions, Health and Mortality. *National Bureau of Economic Research Working Paper No. 11007*.
- 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* 30(1): 13-26
- Smetters, Kent. 2004. Insuring Against Terrorism: The Policy Challenge. *Brookings-Wharton Papers on Financial Services*.
- The Social Security Technical Panel on Assumptions and Methods. 1999.
<http://www.ssab.gov/Publications/Financing/tech99.pdf>.
- 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



Note: Actual mortality data for 1989 to 2001.

Figure 2



Note: Actual mortality data for 1989 to 2001.

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	4.17%	3.70%	3.78%	3.86%
70	4.03%	3.81%	3.30%	3.16%
75	4.03%	3.85%	3.91%	4.11%
80	3.88%	3.86%	3.87%	4.06%
85	3.32%	3.74%	3.73%	4.11%
1% Loss Probabilities				
	Single Men	Single Women	Married Couples	
Survivor Benefit			50%	100%
Age				
65	5.80%	5.31%	5.37%	5.47%
70	5.75%	5.30%	4.73%	4.44%
75	5.89%	5.54%	5.76%	5.88%
80	5.57%	5.57%	5.56%	5.92%
85	4.83%	5.27%	5.26%	5.94%
5% Interest Rate				
5% Loss Probabilities				
	Single Men	Single Women	Married Couples	
Survivor Benefit			50%	100%
Age				
65	3.06%	2.81%	2.95%	2.86%
70	3.25%	2.97%	2.83%	2.60%
75	3.41%	3.18%	3.28%	3.44%
80	3.43%	3.22%	3.25%	3.55%
85	3.34%	3.26%	3.32%	3.56%
1% Loss Probabilities				
	Single Men	Single Women	Married Couples	
Survivor Benefit			50%	100%
Age				
65	4.41%	4.04%	4.13%	4.14%
70	4.68%	4.28%	4.00%	3.65%
75	4.95%	4.48%	4.63%	4.94%
80	4.96%	4.69%	4.65%	5.15%
85	4.80%	4.71%	4.65%	5.12%

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 2 (Continued) - With Payments Growing at 3%, Mark-ups Required to Reduce Probability of Loss to Specified Percentages

3% Interest Rate				
Mark-up/Reserve Required to Reduce Probability of Loss to 5%				
	Single Men	Single Women	Married Couples	
Survivor Benefit			50%	100%
Age				
65	5.67%	5.64%	5.77%	5.99%
70	5.52%	5.40%	4.45%	4.36%
75	5.28%	5.14%	5.22%	5.65%
80	4.88%	4.90%	4.94%	5.34%
85	4.88%	4.90%	4.94%	5.06%

Mark-up/Reserve Required to Reduce Probability of Loss to 1%				
	Single Men	Single Women	Married Couples	
Survivor Benefit			50%	100%
Age				
65	8.26%	8.15%	8.11%	8.87%
70	7.98%	7.86%	6.53%	6.43%
75	7.77%	7.47%	7.44%	8.16%
80	7.11%	7.22%	7.18%	7.97%
85	7.11%	7.22%	7.18%	7.48%

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

	Female	Male	Couple	
Survivor Benefit			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.