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

Leora Friedberg Anthony Webb*

October 20, 2006

ABSTRACT. Insurance companies, employers, and the U.S. government all provide annuities and therefore assume aggregate mortality risk. 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 determine that annuities expose providers to substantial risk. We also focus on other recent actuarial forecasts, some of which lie at the edge or outside of the 95% confidence interval of Lee-Carter.

We focus here on the implications of aggregate mortality risk for insurance companies; this analysis can be extended to private pension providers and Social Security. Given the forecasts of the Lee-Carter model, we calculate that a markup of 3.9% on an annuity premium (or shareholders' capital equal to 3.9% of the expected present value of annuity payments) would be required to reduce the probability of insolvency resulting from aggregate mortality shocks to 5%, and a markup of 5.7% would reduce the probability of insolvency to 1%. Based on the same model, we find that a projection scale commonly referred to by the insurance industry underestimates aggregate mortality improvements and would leave annuities underpriced.

Annuity providers could deal with aggregate mortality risk more efficiently by transferring it to financial markets through mortality-contingent bonds. We calculate the returns that one recently-proposed mortality bond would have paid had it 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 to hold the bond. At plausible coefficients of risk aversion, annuity providers should be able to hedge aggregate mortality risk via such bonds at very low cost.

* Anthony Webb, Ph.D., (corresponding author) is a research economist at the Center for Retirement Research at Boston College, located at 258 Hammond Street, Chestnut Hill, MA 02467-3808, 646-662-7254, 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, 434-982-2904 (fax), lfriedberg@virginia.edu. We would like to thank Jeff Brown, Steve Cooperstein, Irena Dushi, 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. JEL Codes: G12, G22, G23, J11, J14.

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 7.4% 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 jump in coming years, though, as pre-annuitized defined benefit pensions provided by employers in the United States and United Kingdom are largely replaced by lump-sum defined contribution pensions, and even more so if Social Security is similarly transformed using private accounts.[3]

Thus, it is important to understand why the market for private annuities is so meager. One reason may involve difficulties faced by annuity providers in managing aggregate mortality risk. Annuity providers face two kinds of mortality risk: idiosyncratic risk, since any particular annuitant may live longer than anticipated; and aggregate risk, since annuitants may on average live longer than expected.[4] The business of insurers is to eliminate idiosyncratic risk through risk-pooling. The latter risk is non-diversifiable unless the insurer can write other classes of business, for example life insurance, with negatively correlated risks.[5] The growing reluctance of employers to take on aggregate mortality risk has been offered as one of the explanations for the shift in private pension structure (ACA 2005, pp.13-14).

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 emergent

---

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

[2] Sales of deferred annuities substantially exceed those of immediate annuities. A deferred annuity offers tax deferral but, when sold to individuals, lack 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. Investment funds in Chile must offer voluntary annuities. Until April 2006, the U.K. required that "individual pensions" be annuitized by the age of 75.

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

[5] Milevsky and Promislow (2005) discuss this possibility. The scope for doing so is limited by the difference in average ages of life insurance and annuity holders. Mortality shocks differ in magnitude and occasionally even in sign across ages.

crises to hasten it. S. Jay Olshansky and Stephen Austad, two of the most famous researchers in the field, have wagered $500 million on whether someone recently born will live to 150 by 2150.[6] Providers of immediate annuities and also potential annuitants make their decisions over much shorter time horizons.[7] Over shorter 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 either. Purely extrapolative models fit the data closely, have been shown to perform well in-sample, and importantly for our purpose, quantify the uncertainty surrounding the future course of mortality.

Yet, past forecasts from extrapolative models carry wide margins of error in their point estimates of expected mortality rates. Using such forecasts will expose annuity providers to substantial aggregate mortality risk. We consider a number of ways in which annuity providers could manage this risk. We assess by how much insurance companies would need to mark up the annuity premium or increase their capital reserves so as to gain protection against specified levels of risk. A more efficient alternative would be for annuity providers to transfer aggregate mortality risk to a third party by purchasing mortality-contingent bonds, with payments that depend on mortality experience of a reference population. A few such bonds have recently been proposed or issued. We calculate the returns that one such bond would have paid had it been available over a long period, and we determine the risk premium that investors would have required to hold such a bond. As far as we know, no other paper in the economics literature has applied the Capital Asset Pricing Model (CAPM) and Consumption CAPM to price aggregate risks which beset the insurance industry.[8]

This paper is organized as follows. 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 1999 technical

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

[7] Although an annuitant aged 60 might live fifty more years, the mortality risk over near periods holds prominence as distant periods are subject to substantial time and mortality discounting.

[8] An alternative is to price mortality bonds using a Wang transform (Wang 2000), as in Lin and Cox (2004), but the results would depend heavily on assumptions about the insurer's unobserved level of expenses. Milevsky and Promislow (2001) and Milevsky, Promislow, and Young (2005) derive theoretical formulas that could be used to price aggregate mortality risk based on prices of other securities. Our calculations are based on historical data using models that take utility-maximizing behavior as their starting point.

advisory panel.[9] We compare its predictions with those of the Social Security Administration and published life tables. While the Lee-Carter estimates are based on a long time-series of data, the degree of divergence across expert forecasts suggests that even the Lee-Carter confidence intervals may fail to capture the true magnitude of aggregate mortality risk. To this end, we consider the sensitivity of these and other calculations in the paper to a large, temporary, negative mortality shock and to the possibility of discrete medical gains.

In Sections 2 and 3, we focus on aggregate mortality risk faced by insurance companies; much of this analysis can be extended to annuities provided by private pension plans (as in Dushi, Friedberg, and Webb 2006) and by Social Security. In Section 2, we use the Lee-Carter model to quantify the magnitude of aggregate mortality risk faced by an insurer selling annuities to a single birth cohort. We calculate the shareholders' capital reserves and/or premium loadings that would be required to reduce to specified percentages the risk that capital is exhausted by aggregate mortality-related losses. Under typical assumptions, the capital reserve/mark-up needed to reduce the probability of insolvency to 5% lies in the range of 2.7-3.3% for a joint life and 50% survivor annuity sold to couples between the ages of 65 and 85.

The Society of Actuaries recommends that insurance companies use Projection Scale AA as a basis for forecasting mortality improvements. We find that Projection Scale AA underestimates mortality improvements, compared to the Lee-Carter model. We show in Section 3 that annuities based on Projection Scale AA without any conservative margin would be underpriced by amounts that typically range between 2-3% of the premium paid, relative to annuities with prices based on Lee-Carter, for single or married annuitants aged 65 and over.

In Section 4, we consider how annuity providers of any type might efficiently shed aggregate mortality risk by purchasing mortality-contingent bonds. We explain the structure of mortality bonds that have recently been issued or proposed, and we price one such bond using the CAPM and the Consumption CAPM. According to the CAPM, investors will prefer assets with returns that negatively covary with financial market returns. They will accept an expected return that is smaller than the risk-free rate to hold such an asset and will demand an expected return that is greater to hold assets that positively covary with the market. 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.

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

We calculate the annual returns that would have been earned on mortality-contingent bonds modeled on one that was recently proposed, though not issued as planned, by EIB, had they been available in the U.S. during the period 1959-1999. The covariance of those returns with the stock market was close to zero, while the covariance with per capita consumption growth was slightly negative. Consequently, under the CAPM investors will require little or no premium to hold such a bond, and under the CCAPM they will accept a discount to do so, thus paying a little to shed longevity risk. Although it was small, the proposed discount on the EIB bond would have required investors to pay more than either model predicts for the opportunity to shed aggregate mortality risk. We discuss reasons why it may have been cheaper to hedge such risk in financial markets rather than reinsurance markets, as EIB intended.

Section 5 summarizes. We conclude that annuity providers – including insurance companies that sell voluntary annuities, employers that offer defined benefit pensions, and the U.S government through Social Security and the Pension Benefit Guaranty Corporation – face substantial aggregate mortality risk. Were annuity providers to hedge aggregate mortality risk model by purchasing mortality-contingent bonds, we calculate that they could do so at extremely low cost. These issues will assume growing importance as pension plans that are typically annuitized continue to be displaced by 401(k) and other plans that offer lump-sum distributions and only rarely offer annuitization, and as such concerns lead to calls for government action to cover this risk (Blake and Burrows 2001).[10]

1. QUANTIFYING AGGREGATE MORTALITY RISK

1.1 The Lee-Carter Model

Wong-Fupay and Haberman's (2004) review of research on population mortality forecasts has a valuable discussion about the use of mortality models to quantify uncertainty surrounding mortality forecasts. Actuarial models often produce only point estimates, and many of the intervals bounded by high and low scenarios in official mortality projections appear somewhat narrow. Blake and Burrows (2001) show that plausible assumptions about forecasting errors can lead to quite imprecise estimates of actuarially fair annuity yields. With mortality

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

uncertainty rising over longer time horizons, they find 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

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

The parameters *a* and *b* vary with age. Lee-Carter find that *k* 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:

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

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. The fitted model explains well over 90% of within-age group variances in mortality rates.

We use this model and its estimated parameters to forecast mortality rates by age and future year. Importantly for our purposes, we can calculate forecasts and associated confidence intervals of future mortality rates at birth or conditional on reaching a given age. For example, the model predicts life expectancy at birth in 2065 of 86.05 years, with a 95% confidence interval of [80.45, 89.95] years, which appears relatively wide.

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 G-7 countries with different estimated parameters.

While the model suffers from potential weaknesses which have been noted by others, they do not generally concern 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 has little effect on our results, especially if, as may be likely, such transitory shocks are generally negative rather than positive.[11] 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 adopt this modification to the basic model where appropriate. 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. Its importance has not been demonstrated for the U.S., and we find below that *k* varies little with age in any given year. Sixth, the Lee-Carter model was estimated on combined male and female data. Although mortality declines for men have recently exceeded those for women, the difference has gone in the other direction during some periods. If annuities are bought by men and women in roughly equal numbers, so that insurance companies get to pool this particular risk, then gender differences in mortality trends only concern us when they help project population-level mortality.[12] We will discuss such instances as we proceed with our analysis. Seventh, Ruhm (2004) found a statistically significant correlation between U.S. business cycles 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 2006). In particular, Sanderson and Scherbov (2004), in contrast to Lee-Carter and most others, project no deceleration in the rate of mortality improvement, based on data from fourteen

[11] In the Lee-Carter model, the 1918 flu epidemic increased mortality rates among the over 65s by approximately 17.5%. This translates into a relatively small number of deaths today – 0.28% of 65 year olds today , for example – and a correspondingly modest impact on the expected present value of the insurance company's obligations.

[12] The reasons for recent differences are not well understood and may include different health behaviors, such as a relative increase in smoking by women. We choose not to consider the impact of specific behaviors, as an attraction of the Lee-Carter model is that it captures the uncertainty surrounding forecasts of behavior.

countries.[13] We retain the original model, though, since it is widely accepted and offers published parameter values for U.S. data, and the differences in the predictions of other models (besides Sanderson-Scherbov) are not substantial.

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 2002, when both Lee-Carter and Social Security Administration (SSA) forecasts from the 1993 Trustees' Report can be compared with actual mortality improvements; and 2003 onwards.[14]

Figures 1 and 2 shows actual mortality rates 1989 to 2002 for males and females aged 65-94, with some surprising patterns observed among the oldest women. They are compared with the predictions made by the SSA in their 1989 intermediate forecast and with an age-weighted forecast made by the Lee-Carter model. The Lee-Carter model forecasts an age-weighted annual mortality decline of 1.12% per year for the over-65s. In contrast, the SSA intermediate forecast predicts male and female mortality declines of 0.54 and 0.50% through 2017. It transpired that mortality declined much more rapidly among men than women, and more rapidly among the newly retired than the oldest old.[15] Male mortality declined by a little more than the Lee-Carter forecast and much more than the SSA intermediate forecast at ages 65-69, 70-74, and 75-79. This was true of none of the female age groups after 1992, and mortality actually rose among women aged 85-89 and 90-94 after the mid-1990s. While expected mortality from the Lee-Carter model did not closely match recent mortality trends, the annual age-weighted changes in mortality mostly fall within the 95% confidence interval around the Lee-Carter mean. Furthermore, these patterns highlight the inherent difficulty in forecasting mortality and hence in managing aggregate mortality risk.

Preston and Wang (2005) investigated the recent narrowing of the male-female life expectancy differential, which dropped from around 7.8 years in 1972-79 to 5.3 years in 2003. They found a delayed response to an earlier increase in smoking by women, which can explain

[13] 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.
[14] The SSA actuaries forecast mortality by combining age-specific trend extrapolation with information collected from medical experts.
[15] Bell and Miller (2005) found a similar differential for male and female mortality over the period 1981-2001.

recent anomalies for women. They project an acceleration in mortality reductions even if smoking rates do not change further, with faster reductions among males because their smoking rates have declined more than women's. Thus, the 1989-2002 period should represent a temporary pause in the decline in female mortality.

Figure 1 - Actual versus forecasted mortality declines
Males, 1924 birth cohort

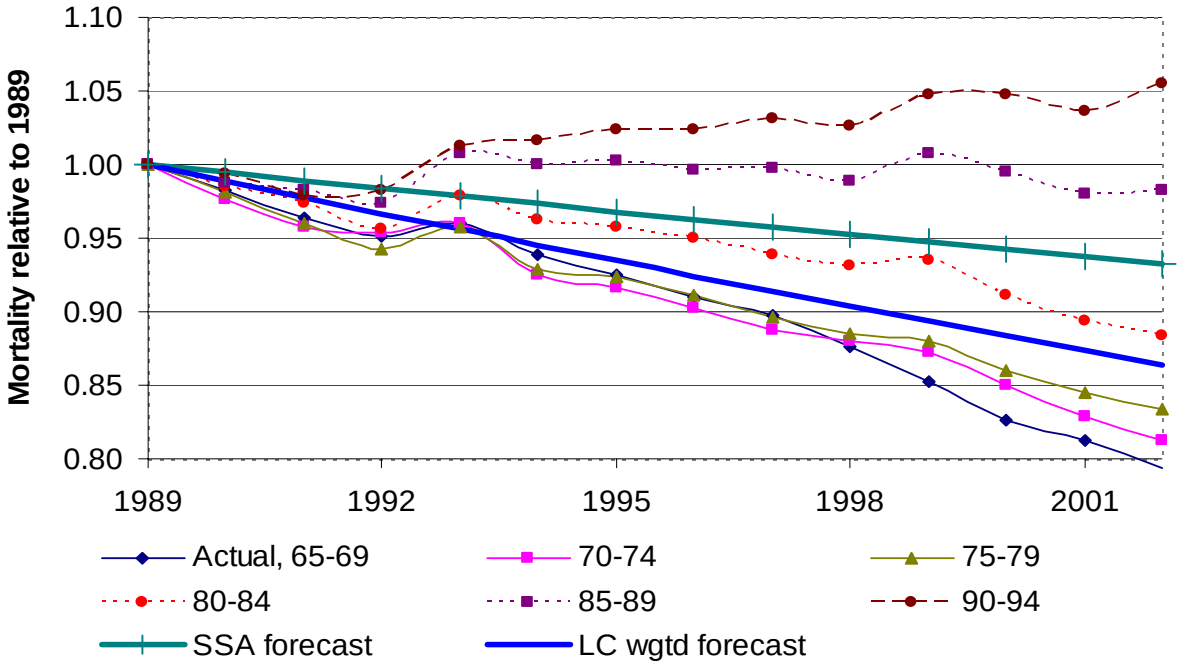
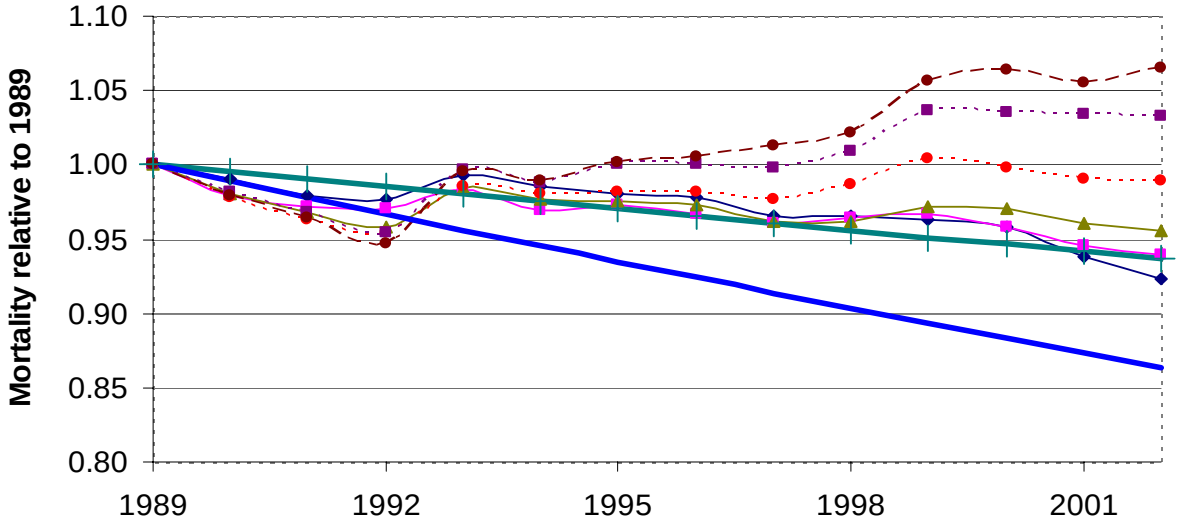
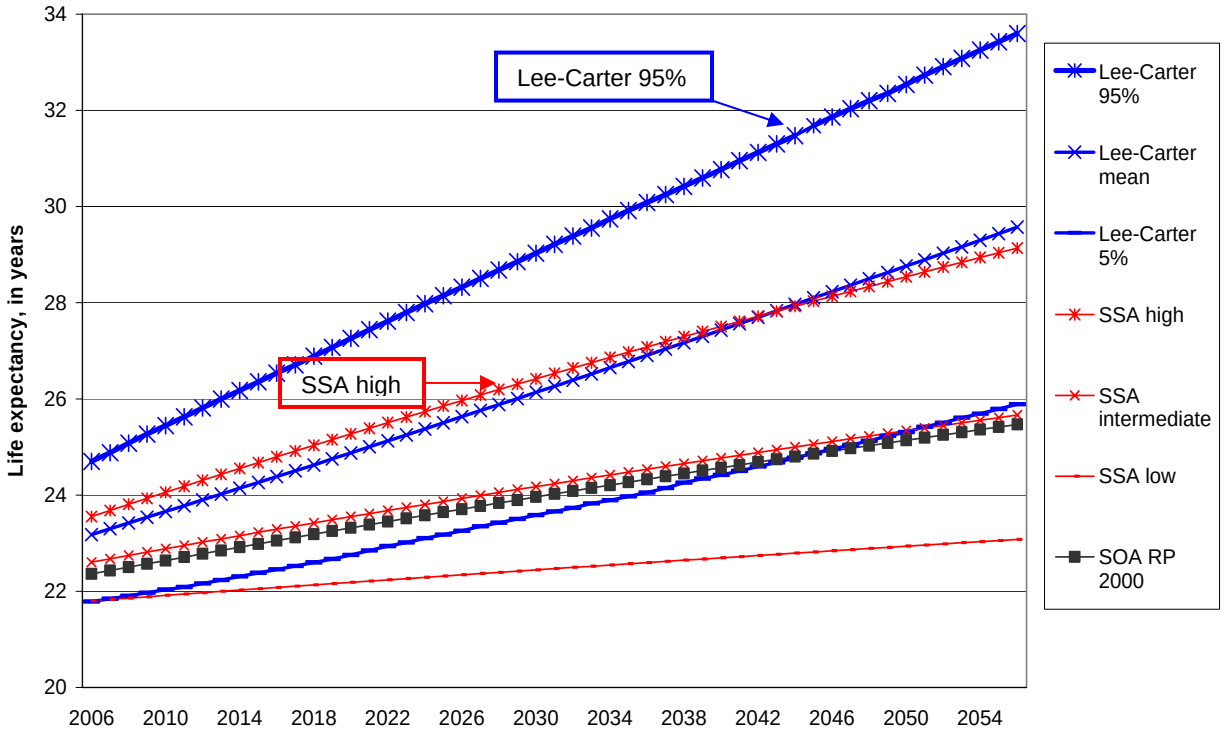


Figure 2 - Actual versus forecasted mortality declines
Females, 1924 birth cohort



Looking forward in Figure 3, 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 2002 as a baseline, the SSA projects annual mortality reductions of 0.24, 0.47, and 0.70% under the three forecasts, jumping to 0.29, 0.67, and 1.16% in 2030.[16] The intermediate forecast (0.47%, then 0.67%) predicts that mortality will initially decline at the pace experienced between 1979 and 2002 and then almost at the pace of 0.72% experienced between 1900 and 2002.[17] The high forecast starts out at the higher rate of mortality decline and then accelerates. 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. Based on 1900-89 data, Lee-Carter forecasts a future age-sex weighted average mortality reduction of 1.12% a year.[18] This is almost the same as the SSA "high" forecast of 1.16% from 2030 on.

Figure 3 - Life expectancy at age 60 under various mortality assumptions, 2006-2058



[16] The SSA does not break down their projected declines by age within this group.
[17] 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.
[18] 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.

1.3 Comparing Lee Carter with Actuarial Forecasts

Next, we compare mortality declines projected by the Lee-Carter model with the projection scales that actuaries commonly use when forecasting mortality declines among annuitants and pension plan members. Projection Scale AA is the most recent one that the Society of Actuaries (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 RP-2000 life tables, the SOA considered whether Scale AA should be updated and concluded that it was consistent with recent trends.[19] Potentially, the population on which Scale AA was based differs from annuitants. Annuity purchasers in the U.S. are much wealthier than average (Dushi and Webb 2004). Rates of mortality declines may differ between annuitants and non-annuitants due to widening socio-economic mortality differentials, as documented by Willetts (1999) using U.K. data and Schalick et al (2000) using older U.S. data from 1967-1986. However, it is difficult to infer whether mortality differentials have continued to widen since 1986 in the U.S. If so, then forecasts based on population trends, like those derived from Lee-Carter, will underestimate reductions in annuitant mortality. Uncertainty as to the future course of socio-economic mortality differences will, in any case, contribute to the non-diversifiable mortality risk faced by insurance companies.[20]

Figure 3 shows that the SOA forecast, based on Scale AA from the RP-2000 life tables, lies extremely closely to the SSA intermediate forecast. Table 1 compares the mortality declines by age group projected by the Lee-Carter model and by the SOA in their Projection Scale AA. Scale AA forecasts a slower pace of overall mortality reductions than Lee-Carter does, which is unsurprising since the SOA data is based on a limited time period over which mortality declined relatively slowly. The age-sex weighted average for Scale AA is 0.71% for the whole population and 0.69% for the over 65s. From ages 45-79, the Scale AA and Lee-Carter forecasts are remarkably close. From age 80 upwards, Scale AA projects successively smaller percentage

[19] "The RP-2000 Mortality Tables" page 73. Projection Scale AA is also reported in this publication in Table 7-3.
[20] 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 discern no systematic differences among them that might relate to their presumed socio-economic status. However, as the report points out in Chapter 4, "the measurement of mortality improvement requires voluminous consistent data covering many years."

reductions in mortality, approaching zero at age 100 plus. In contrast, Lee-Carter projects somewhat larger reductions in mortality over age 80 than at ages 45-79.

Table 1 - Annual Mortality Reductions under Lee-Carter versus Projection Scale AA

Age	0	1-4	5-9	10-14	15-19	20-24	25-29	30-34
Percentage Reduction:								
Lee-Carter	3.31%	4.03%	3.35%	3.05%	1.73%	1.95%	2.18%	2.25%
Scale AA	2.00%	2.00%	2.00%	1.95%	1.77%	1.63%	0.81%	0.64%
Age	35-39	40-44	45-49	50-54	55-59	60-64	65-69	
Percentage Reduction:								
Lee-Carter	2.15%	1.93%	1.63%	1.40%	1.23%	1.08%	1.05%	
Scale AA	0.85%	1.20%	1.60%	1.72%	1.25%	1.07%	1.00%	
Age	70-74	75-79	80-84	85-89	90-94	95-99	100-104	65-104
Percentage Reduction:								
Lee-Carter	1.06%	1.18%	1.13%	1.13%	1.13%	1.13%	1.13%	1.12%
Scale AA	1.12%	1.03%	0.77%	0.52%	0.30%	0.16%	0.04%	0.69%

The disagreement above age 80 is particularly important. Mortality rates are much higher at these ages, so a given percentage reduction in mortality has 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.[21] Moreover, if the annuitant population continues to enjoy better than average mortality reductions, even the Lee-Carter forecasts may prove to be too conservative.

2. Quantifying the Mortality Risk Faced by Annuity Providers

In this section, we use simulations to quantify the aggregate mortality risk faced by an insurance company 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.[22] To focus on the effect of aggregate mortality risk, we impose a number of simplifications. 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 mortality and changes

[21] This distinction is less important for pension funds, since DB pensions, as deferred annuities, are "bought" by people of all ages.

[22] 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 assumptions about expenses, profits, current period mortality, and the cost of capital.

in mortality rates as predicted by the Lee-Carter model.[23] In reality, as we noted earlier, higher socio-economic classes may experience more rapid declines in mortality, which would amplify our findings. We also ignore adverse selection, which likely affects this market already (Finkelstein and Poterba 2004). We assume that the annuity price is set so that the provider breaks even at expected mortality rates and that the provider has zero administrative expenses. We further assume that the provider invests in a risk-free asset offering a real return of either 3 or 5% while selling a level real annuity, or else invests in a risk-free asset offering a 3% real return while selling a real annuity that increases at 3%.[24] We select values of a to fit the Social Security Administration's projections of 2006 period mortality for males and females separately.

There are three 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, so 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 (1), Appendix B in Lee-Carter shows that it is reasonable to ignore those sources of error when forecasting life expectancy, and also discusses the properties of the error terms more generally. Third, there may be gender differences in the rate of mortality declines, exposing insurers to additional risk unless they have a balanced portfolio of male and female lives.

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 -0.365 and standard deviation 0.069, the values associated with Lee-Carter's estimated k trend in (2). This gives us the baseline time trend 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. We fill in the value of k year by year using these draws together with the particular k trend and use equation (1) to calculate the associated annual mortality risk.

[23] 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.

[24] This last possibility is equivalent to earning a 0% real return and selling a level real annuity. A 3% real return is the common assumption in the academic literature, though the long-run risk-free rate as measured by TIPS is

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 birth cohort.[25] 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.[26]

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 requires 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 the same degree, 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 or 100%), single women, and single men aged 65, 70, 75, 80, and 85.[27] For an insurance company selling joint life and 50% survivor annuities with a real return of 3% to couples aged 65, a combination of shareholders' capital and premium loading equal to 2.69% of the premium is required to reduce the probability of loss or insolvency to 5%. Shareholders' capital or loading of 3.80% is required to reduce the probability to 1%.

Annuities sold to younger lives obviously have a longer duration than those sold to older lives. Joint life and 100% survivor annuities have a longer duration than joint life and 50% survivor annuities, which in turn have longer durations than female and finally male single life annuities. At the 3 and 5% interest rates, there is a negative relationship between risk and duration. Payments that fall due in the distant future carry more aggregate mortality risk, but are subject to greater time discounting, and the effects of time discounting dominate. At lower

---

currently well below 3%. Zero administrative expenses are common simplifying assumptions as well, and insurance companies are starting to offer real annuities.

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

[26] For the simulations, we chose as a starting point for *k* its value in 1989 of -11.06. Using a smaller value, consistent with more recent mortality data, would make the population appear slightly younger but would have little effect on the results.

[27] 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.

interest rates, the effects of time discounting are reduced, and at a 3% interest rate, duration has almost no effect on the risk of a 3% increasing annuity.

Table 2- Potential Losses Arising From Aggregate Mortality Risk

	Loss probability							
	Single men		Single women		Married Couples with survivor benefit			
					50%		100%	
	5%	1%	5%	1%	5%	1%	5%	1%
				3% interest rate				
65	3.89%	5.45%	3.67%	5.24%	3.74%	5.39%	3.91%	5.53%
70	4.08%	5.87%	3.74%	5.38%	3.55%	5.05%	3.69%	5.17%
75	3.96%	5.97%	3.87%	5.45%	3.90%	5.62%	4.10%	5.97%
80	3.91%	5.61%	3.81%	5.57%	3.77%	5.51%	4.12%	5.97%
85	3.70%	5.25%	3.64%	5.28%	3.67%	5.24%	4.01%	5.79%
				5% interest rate				
65	3.11%	4.46%	2.82%	3.95%	2.94%	4.29%	2.87%	4.08%
70	3.30%	4.80%	3.00%	4.27%	2.79%	3.87%	3.12%	4.54%
75	3.43%	4.94%	3.15%	4.61%	3.26%	4.70%	3.34%	4.71%
80	3.35%	4.82%	3.23%	4.74%	3.33%	4.71%	3.56%	5.25%
85	3.28%	4.87%	3.26%	4.72%	3.34%	4.73%	3.59%	5.13%
				3% interest rate and 3% increasing annuity				
65	5.67%	8.26%	5.64%	8.15%	5.77%	8.11%	5.99%	8.87%
70	5.52%	7.98%	5.40%	7.86%	4.45%	6.53%	4.36%	6.43%
75	5.28%	7.77%	5.14%	7.47%	5.22%	7.44%	5.65%	8.16%
80	4.88%	7.11%	4.90%	7.22%	4.94%	7.18%	5.34%	7.97%
85	4.88%	7.11%	4.90%	7.22%	4.94%	7.18%	5.06%	7.48%

Note: Computed for the 1924 birth cohort, assuming population mortality. Each cell shows the amount 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.

Both inflation-linked and increasing annuities are backloaded relative to level annuities, paying a greater proportion of their benefits at older ages. Both are therefore relatively more attractive to long-lived households but expose the insurance company to more 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 and attribute this to adverse selection, since increasing annuities are particularly attractive to low-mortality types. The higher prices may also reflect the greater amount of aggregate mortality risk to which sellers of increasing annuities are exposed.

The calculations in this section are similar in spirit to those in Dowd, Cairns, and Blake (2005) of value-at-risk and expected shortfall for U.K. annuity providers.[28] They estimate a 90% value-at-risk for long-dated coupon-paying mortality bonds of between 4 and 5% of the premium, depending on whether they allow for parameter uncertainty. Their results are of the same order of magnitude as ours, though somewhat higher. Given considerable differences in mortality patterns between the countries, it is valuable to undertake such calculations specifically for the U.S., and we build on the analysis in the rest of this paper.[29]

Lastly, we consider the possibility of a sudden medical breakthrough, although even major innovations are likely to have incremental effects on mortality. Olshansky, Carnes, and Cassel (1990) calculate that eliminating *all* cancer and ischemic heart disease, while an unrealistic prospect, would raise life expectancy at age 50 by 6.57 years for females and 7.83 years for males; and eliminating cancer, all circulatory diseases, and diabetes would raise it by 15.30 and 15.02 years. The latter gains correspond to a 75% reduction in mortality from all causes, which has a less than 0.01% probability of occurring over a 30-year horizon under the Lee-Carter model. If such a shock occurred instantaneously, it 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 they would raise prices to new purchasers while mortality declined gradually and payments to existing annuitants increased by lesser amounts.

To sum up, we find that the combination of shareholders' capital and/or premium loadings that would be required to reduce the probability of insolvency to 5% or to 1% is considerable. It lies in the range of 2.7-4.7% for a joint life and 50% survivor annuity with a real return of 3% sold to couples between the ages of 65 and 85. It should be noted that applying mark-ups such as these might deter high-expected mortality individuals from annuitizing, exacerbating adverse selection.

[28] They use the two-factor mortality model developed in Cairns, Blake, and Dowd (2006). 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 premia for mortality bonds of various maturities and initial cohort ages, assuming alternative weightings.

[29] British mortality is lower below the age of 65, but increasingly higher above that age (Deaton and Paxson 2004), and mortality improvements in the U.S. lead changes in Britain by about four years on average.

3. THE PRICING OF AGGREGATE MORTALITY RISK BY ANNUITY PROVIDERS

When we calculated the required premium loading for insurers to offset mortality risk, we assumed that insurers price annuities in accordance with Lee-Carter. To the best of our knowledge, no published research has worked out the financial implications of the differences between Lee-Carter predictions and those of actuarial life tables. We find that, by the Lee-Carter benchmark, insurance companies would 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 Scale AA appears to understate expected mortality improvements.

3.1 Methodology

It would be difficult to use pricing data from insurance companies to determine whether aggregate mortality risk is correctly priced. Annuity prices reflect not only an insurer's views about aggregate mortality risk but also unobservable information about expenses, asset returns, and the expected mortality of its particular annuitants. Given these difficulties, we focus instead on the up-to-date recommendations made by the SOA in their Projection Scale AA.

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 specifically 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.[30] 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 underestimates of mortality improvements in Scale AA.

We calculate the expected present value (EPV) of payments from a $1-a-year annuity under various assumptions about mortality declines. We take the ratio of the EPV computed using forecasts from the Lee-Carter model to the EPV computed using Projection Scale AA.

[30] http://www.actuary.org/pdf/life/varwg_0305.pdf

This reveals the extent to which using Scale AA would contribute to underpricing of annuities. The ratio will exceed 100% if and only if the Lee-Carter forecasts are more optimistic about future mortality than Projection Scale AA. As before, we apply a 3% real interest rate and zero administrative costs. We assume that the Annuity 2000 mortality table from the SOA correctly describes current mortality of people buying the annuity.[31]

3.2 Analysis of Potential Underpricing

Table 3 reports our results. EPVs are in all cases greater than 100% of the premium paid, which indicates potential underpricing according to Lee-Carter, compared to Scale AA. The extent of the underpricing if annuity providers were using Scale AA ranges from 1.6 to 3.7%. Underpricing is greater at older ages because Lee-Carter assumes little age-related variation in the rate of mortality declines at older ages, whereas Scale AA assumes a sharp deceleration. Underpricing is, at all ages, greater for women than men, reflecting the fact that Scale AA is based on recent data, during which time mortality improvements for men exceeded those for women, whereas the Lee-Carter model only allows us to calculate unisex improvement factors. As we mentioned earlier, though, Preston and Wang (2005) predict a further narrowing of sex mortality differentials, in which case the lesser underpricing of male annuities may be illusory. To sum up, the results show that, if the Lee-Carter model indeed provides an unbiased forecast of mortality improvements, then insurers using Projection Scale AA would underprice annuities by amounts that could have a substantial effect on insurance company solvency.

Table 3 - Percentage Underpricing Under Lee-Carter versus
Single meningle womed couple with survivor I

			50%	100%
Age 65	1.64%	2.93%	2.31%	3.01%
70	2.06%	3.04%	2.57%	3.23%
75	2.52%	3.16%	2.86%	3.45%
80	2.84%	3.27%	3.07%	3.62%
85	3.02%	3.34%	3.18%	3.68%

Note: Ratio of expected present values computed, in the numerator, when mortality improvements are forecasted Lee-Carter and, in the denominator, when mortality improvements follow Projection Scale AA.

[31] The Annuity 2000 life tables can be downloaded from www.soa.org. We chose it over others, like those published by the Social Security Administration, because our focus is on the individual annuity market. As this applies to both EPV calculations, our conclusions hold irrespective of the current period life table that the insurance company may use, although the magnitude of the difference between the EPVs would be slightly affected.

Mitchell, Poterba, Warshawsky, and Brown (1999) undertake a related calculation to determine the value of annuities to potential annuitants, rather than the pricing of annuities by insurance companies. Their calculations are based on period mortality tables projected using SSA forecasts of mortality improvements, and those projections are underestimated according to Lee-Carter. Assuming, again, that the Lee-Carter projections are unbiased, our results imply that Mitchell et al somewhat overstate the actuarial unfairness of annuities. The low rate of voluntary annuitization may therefore be even more puzzling than it previously appeared.[32]

4. THE PRICING AGGREGATE MORTALITY RISK IN FINANCIAL MARKETS

Given our evidence of substantial aggregate mortality risk, this section discusses how insurers or other annuity providers might reduce their exposure. We describe mortality-contingent bonds and calculate the returns that investors would have experienced on one such bond, had it been available over the period 1959-1999. To price the 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 show how the use of mortality-contingent bonds would influence annuity prices.

4.1 Mortality-Contingent Bonds

As of 2006, only two mortality-contingent bonds had been issued and another one proposed. In December 2003 Swiss Re issued the first of two short term bonds, 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. They issued a second similar bond in 2005. In November 2004, the European Investment Bank (EIB) announced a £540 million bond, although it was not issued. It was to have a life of 25 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 U.K. male

[32] 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 who report 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.

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 reduce their exposure to aggregate mortality risk. Since we focus on annuities, we will make calculations to price this latter bond.

A non-diversifiable risk may require only a small premium on financial markets if that risk is both small in relation to the market and is uncorrelated with market returns. Cairns, Blake, and Dowd (2005) report that the EIB bond, based on its prospectus, was to be issued at an expected yield that was some 20 basis points lower than those at which similar non-mortality related EIB bonds traded – in other words, the bond price incorporated a mortality risk discount, not a premium.[33] 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, according to the CAPM and CCAPM, be determined by its covariance with other sources of risk – a point which drives our discussion of asset pricing.

4.2 Historical Returns

Our goal is to calculate the returns that investors would have experienced had the EIB bond been available in the U.S. market in past years. We assume that investors would have believed that the Lee-Carter model correctly described mortality improvements and would have used those beliefs to price the bond. 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

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

or $b_x e_{t+1} + e_{x,t+1}$. 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 on. We compare this with the change predicted by Lee-Carter and recover, for each year, a vector of combined mortality shocks, $b_x e_{t+1} + e_{x,t+1}$. The two components of the shock have different implications for the pricing of the bond. We are interested in e_{t+1}, which affects all ages equally, will affect the expected value of k in all future years, and will thus have a

[33] Based on government mortality projections. If the market believed that these projections were conservative, then the expected yield would be higher.

substantial effect on the value of the bond. The $e_{x,t+1}$ term varies with age, only affects current-period mortality, and will thus have a negligible effect on the value of the bond.

We take several steps to investigate this issue and other features of the mortality shocks. We begin by confirming that the Lee-Carter specification fits the data well.[34] We also expect the mortality shocks to have a mean close to zero, and the sample mean of e_{t+1} of 0.034 is not significantly different from zero. An important assumption is that the e_{t+1} shocks to k 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 combined shocks do 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. This indicates that 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. Importantly, it also suggests that the age-specific $e_{x,t+1}$ term is small relative to the common e_{t+1} term. We proceed by treating the mortality shocks that we estimate as revealing the annual value of e_{t+1}.

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 changes to payments that are now expected in all future years. We express the mortality shock for each of the years 1960-1999 in terms of the innovation to k and 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 forecasts from the Lee-Carter model. The effect on the income stream depends on the duration of the bond and the age of the reference population. 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. To illustrate, when the reference population is aged 65, the resulting mean of this percentage is 0.04%, very close to zero as expected, and the standard deviation 0.64. The standard deviation increases somewhat with the age of the reference population, but it is very

[34] The Lee-Carter model was estimated over 1933-1987, which does not precisely coincide with the period we examine here.

small – the general consensus is that the standard deviation of stock returns is about 17 to 20% by comparison, or roughly thirty times greater.[35]

This initial result suggests that, if investors used the Lee-Carter model to forecast mortality and price mortality bonds and if the Lee-Carter model provided unbiased estimates, then a bond structured like the EIB bond would have been 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 eliminated similarly through diversification. The magnitude of the risk premium that an investor requires to hold a risky asset like a mortality bond equals its beta, defined as

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

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

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

the risk-free rate plus beta multiplied by the excess of the expected market return over the risk-free rate. Insurance companies selling annuities are, of course, owned either directly or indirectly by individual investors. The company's cost of capital and the price it must charge for annuities will therefore reflect the return those investors require, which in turn depends on the company's beta.

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 investments is relatively more valuable at that time compared to

[35] 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.

others. 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 unexpectedly high and the marginal utility of consumption low. Conversely, investors place a high value on assets that offer high returns when consumption is unexpectedly 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

$$E_t(R_{b,t+1}) = R_{f,t+1} + Cov_t\left\{\frac{-U'(C_{t+1}), R_{b,t+1}}{E_t(U'(C_{t+1}))}\right\}, \qquad (6)$$

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:[36]

$$LogE_t\left[(1+R_{b,t+1})/(1+R_{f,t+1})\right] = \gamma\sigma_{b,c}\ , \qquad (7)$$

where γ is the coefficient of relative 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.

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 a beta on EIB-type mortality bonds, varying with the age of the reference population, had such bonds been available over the period 1959-1999.[37] For example, the age 65 result is a beta of 0.005 with a 95%

[36] The assumptions, which are standard (cf. Mehra 2003), are that equity and stock returns are jointly log-normally distributed and that consumption growth and risky asset returns are both i.i.d.

[37] We recognize that the S&P500 is a commonly used proxy for the undetermined "market portfolio" (Roll 1977). The S&P 500 is a capitalization-weighted index of large capitalization stocks that is strongly positively correlated with other domestic stock returns. 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.

confidence interval of [-0.005, 0.015].[38] The point estimates suggest little or no correlation with the market.[39] 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 2½ basis points (since 2½ = 0.005*500) above the risk-free rate when the reference population is aged 65, and premiums of similar magnitudes at other ages, with a 95% confidence interval of [-2½, 7½] basis points. In contrast to the results for mortality bonds, the fact that equity offers the highest returns when consumption growth is high should lead investors to require a modest risk premium for holding equities.

Pricing Using the CCAPM

We hypothesize that the correlation between mortality shocks and contemporaneous aggregate consumption shocks is negative, leading to a negative covariance between the growth in per-capita consumption and returns on an EIB-type bond. Mortality shocks have their greatest absolute effects at older ages, since that is when most deaths occur.[40] Suppose an unexpectedly small number of older people die in a particular year – then total national income is unlikely to rise appreciably in response because labor supply is largely fixed, as the extra numbers of mostly older people are unlikely to work, and the capital stock is fixed in the short-term as well. In consequence, per-capita income will decline because the relatively constant level of output must be shared among a greater-than-expected number of individuals.[41] Moreover, to the extent that the mortality shock is global, one country could not borrow from another to offset the consumption shock. A final point is that reductions in mortality may result from higher medical spending, leaving less money available for general consumption.

---

[38] The beta increases with age, so the required rate of return on the mortality bond and the resulting cost of insuring against aggregate mortality risk also rise with age. We do not account for this in our calculations, though, because the effect on the cost of obtaining such insurance through mortality bonds is never substantial.

[39] 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.

[40] This is not true in percentage terms – the value of b in the Lee-Carter model declines with age, and it is b that governs the percentage impact on mortality of a shock to k.

[41] 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.

We indeed find that per-capita consumption and the returns on the pseudo-EIB bond that we simulated above co-vary negatively.[42] Consumption growth is a relatively smooth series with a mean of 2.25% and a standard deviation of 1.20%. The correlation between consumption growth and mortality bond returns is -0.1958, when the reference population is aged 65, and is significantly different from zero irrespective of their age. 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 from the CCAPM that investors should be willing to accept a risk discount for holding mortality bonds, and given the statistical significance of our estimate, we can add that investors should not pay a risk premium.[43] In contrast, the literature has found a large positive correlation between stock returns and consumption growth of about 0.5. Though the correlation between mortality bond returns and consumption growth is some 40% smaller in absolute value (when the reference population is aged 65), 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 close to zero at -0.0013%. Applying the CCAPM, the risk discount is only two basis points when the coefficient of risk aversion equals ten.

Sensitivity

At this point we will examine the sensitivity of our analysis to some issues related to the nature of mortality shocks that we discussed earlier. We have done additional calculations to verify that market prices are unlikely to be influenced by expectations about transitory mortality shocks. One-off shocks would have little effect on the value of the EIB bond, since its duration is long. Since most or all one-off shocks are negative, decreasing but not increasing longevity

[42] We follow convention by focusing on the consumption of non-durables and services, using data for 1959-1999 from the National Income and Product Accounts. The correct variable, although more difficult to measure, would be the consumption of stockholders who own insurance companies and firms that offer DB pensions. Consumption of stockholders is more highly correlated with the stock market (Mankiw and Zeldes 1991), which would deepen the risk discount on an asset that negatively covaries with the market.

[43] Here we assume 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.

temporarily, we do not consider large positive transitory shocks.[44] 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 market price of the EIB-type bond that we have analyzed by only 0.28%.[45]

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 hypothesized market price of the EIB-type 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 Lee-Carter. The shock reduces the value of the bond by 0.61 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.[46]

Table 4 - Impact on Bond Price of One Standard Deviation Mortality Shock

	Interest Rate	
Age of Cohort	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.

[44] 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.

[45] Throughout this analysis, we assumes that investors used the Lee-Carter model to price the bond. Here, we also assume that they correctly understood that mortality experienced a one-period shock, after which it reverted to trend. It should be noted that Beelders and Colarossi's (2004) use of extreme value theory, with its focus on the extreme tail of the probability distribution, to price the Swiss Re bond is less relevant for the much longer-horizon EIB bond.

[46] We have not undertaken similar calculations of the impact of cures for cancer and ischemic heart disease, which we discussed earlier, because we do not know the time interval over which such cures might take effect.

Discussion

Both the CAPM and the CCAPM suggest that the Swiss Re bond should have commanded a much smaller risk discount than the proposed 20 basis points. The design of the bond may have contributed to the size of the risk discount. 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 passing it on 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 notes 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. Financial markets may be better positioned than smaller and constrained insurance markets to absorb these non-diversifiable risks, so the prospect of using reinsurance instead 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, for example by making bondholders residual claimants after survivor benefits are met.

It is also possible that the yield discount, as perceived by market participants, may be less than the 20 basis points referred to by Cairns, Blake, and Dowd (2005). The calculated yield discount is based on mortality estimates produced by the Actuary's Department of the British Government. If the market believed that future mortality will be lower, then the discount, as perceived by investors, would be correspondingly reduced.[47] We calculate that the 20 basis point discount on the EIB bond would be eliminated if investors believed that mortality declines were underestimated by 0.45% a year. This in turn represents about two-thirds of the gap between the intermediate SSA forecast (0.47% per year through 2030) and our Lee-Carter forecasts (1.12% per year) for the U.S.

To sum up, it is by no means clear what yield discount investors believed the EIB bond was offering. Differences in forecasts of aggregate mortality provide one reason why a bigger discount might have been justified for the EIB bond than what was predicted using either CAPM

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

or CCAPM. We have also noted a reason why the discount that EIB required might have reflected mispricing, due to the limitations of insurance markets as opposed to general financial markets in absorbing this risk.

4.5 The Impact of Holding Mortality-Contingent Bonds on Annuity Prices

Our last goal is to show the implications of our results for the cost of hedging aggregate mortality risk. We recompute 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 that is either predicted by the CCAPM or was required by the EIB bond.

In the previous subsection, 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 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 risk discount for bonds on each age until expiration. We find that it remains extremely small at all ages.[48]

We then recompute 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 under the CCAPM. This reveals the markup that the insurance company must charge to compensate for the lower discount rates. Given the very small mortality risk discounts, the cost of an annuity to men and women of all ages increases only very slightly – by less than 1% if the coefficient of risk aversion is 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 Lee-Carter and the CCAPM, the effect on annuity prices would be extremely small.

Even if annuity providers were to pay the higher mortality risk discount of 20 basis points proposed on the EIB bond, the effect on the price of a joint life and 100% survivor annuity at age 65 would be only 1.84%, assuming a 3% real interest rate. At age 85, the effect is only 0.75%. These are considerably smaller amounts than the markups (or required reserves) of

[48] At a 3% interest rate, the risk discount doubles from age 65 to 90, but the additional risk premium at such advanced ages has little effect on the price of an annuity.

around 4% that are required, as we reported in Table 2, to reduce the probability of making a loss to just 5%.

5. CONCLUSIONS

The market for voluntary annuities in the U.S. is extremely small. In the near future, demand may well rise as unannuitized defined contribution pension plans displace annuitized defined benefit plans and if Social Security is similarly transformed using private accounts. 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.

Our results show that annuity business exposes providers of voluntary annuities to substantial aggregate mortality risk. While we use Lee-Carter as a benchmark to gauge the magnitude and implications of aggregate mortality risk, the divergence across credible forecasts by actuarial experts suggests that true confidence intervals on mortality predictions may be even wider. Moreover, many of the conclusions about aggregate mortality risk apply as well to employers that offer defined benefit pensions, as analyzed in Dushi, Friedberg, and Webb (2006), and to the U.S. government through Social Security and the Pension Benefit Guaranty Corporation.

Insurance companies could manage such risk by charging substantial markups or by maintaining substantial capital against mortality shocks. Alternatively, annuity providers of any type might shed that risk through the use of mortality-contingent bonds. Our calculations based on the CAPM and the CCAPM indicate that this might be accomplished at what is, for all practical purposes, zero cost to annuity providers. In the alternative, a workable scheme for transferring aggregate mortality risk to annuitants 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. Transferring such risk to financial markets instead would be more efficient as long as some investors differ in their tolerance for risk or would gain from taking the opposite position in aggregate mortality risk.

It is relevant to consider how 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.

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*, O. Mitchell, Z. Bodie, P. Hammond, and S. Zeldes, eds. Philadelphia: University of Pennsylvania Press, 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. 2006. "A Two-Factor Model for Stochastic Mortality with Parameter Uncertainty." *Journal of Risk and Insurance* forthcoming.

Cairns, Andrew, David Blake, and Kevin Dowd. 2005. "Pricing the Risk on Longevity Bonds." *Life and Pensions* October: 41-44.

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." In *Perspectives on the Economics of Aging*. D. Wise, ed. Chicago: University of Chicago Press, 247-280.

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, Leora Friedberg, and Anthony Webb. 2006. "The Impact of Aggregate Mortality Risk on Defined Benefit Pension Plans." Manuscript, Boston College Center for Retirement Research.

Dushi, Irena and Anthony Webb 2006. "Rethinking the Sources of Adverse Selection in the Annuity Market." In *Competitive Failures in Insurance Markets: Theory and Policy Implications*. P-A. Chiappori and C. Gollier, eds. Cambridge, MA: MIT Press, 185-212.

Finkelstein, Amy and James Poterba. 2004. "Adverse Selection in Insurance Markets: Policyholder Evidence from the U.K. Annuity Market." *Journal of Political Economy* 112(1): 183-208.

Gong, Guan and Anthony Webb. 2006. "Mortality heterogeneity and the distributional consequences of mandatory annuitization." Center for Retirement Research at Boston College Working Paper.2006-11.

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. 2005. "Securitization of Mortality Risk in Life Annuities." *Journal of Risk and Insurance* 72 (2): 227-252.

Mankiw, N.Gregory, and Stephen Zeldes. 1991. "The Consumption of Stockholders and Non-Stockholders." *Journal of Financial Economics* 29(1): 97-112.

Mehra, Rajnish. 2003. "The Equity Premium: Why is It a Puzzle?" *Financial Analysts Journal* January-February, 54-69.

Mehra, Rajnish, and Edward Prescott. 2003. "The Equity Premium in Retrospect." In *Handbook of the Economics of Finance,* G. Constantinides, M. Harris, and R. Stulz, eds. JAI Press.

Milevsky, Moshe and David Promislow. 2001. "Mortality Derivatives and the Option to Annuitize." *Insurance: Mathematics and Economics* 29: 299-318

Milevsky, Moshe and David Promislow. 2005. "Can Life Insurance be Used to Hedge Payout Annuities? Part 1: Modeling Longevity Risk " York University Individual Finance and Insurance Decisions Centre Working Paper

Milevsky, Moshe, David Promislow, and Virginia Young. 2005. "Financial Valuation of Mortality Risk Through the Instantaneous Sharpe Ratio." York University Individual Finance and Insurance Decisions Centre Working Paper

Milevsky, Moshe and David Promislow. 2003 "Can Life Insurance be Used to Hedge Payout Annuities? Part 1: Modeling Longevity Risk." York University 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, VA: 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. 2005. "The Simple Analytics of a Pooled Annuity Fund." *Journal of Risk and Insurance* 72(3): 497-520.

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

Roll, Richard. 1977. "A Critique of the Asset Pricing Theory's Tests: Part 1: On Past and Potential Testability of the Theory." *Journal of Financial Economics* 4, 129-176.

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.