Killing the Law of Large Numbers:
Mortality Risk Premiums and the Sharpe Ratio

M. A. Milevsky[1], S. D. Promislow and V. R. Young
York University, York University and University of Michigan

Version: 10 June 2006

---

[1] Milevsky is the contact author and can be reached via email at: milevsky@yorku.ca, or at Tel: 416-736-2100 x 66014, Fax: 416-736-5487. His mailing address is: 4700 Keele Street, Toronto, Ontario, Canada, M3J 1P3. This paper is a summary of the presentation given at the 2nd annual Longevity Risk Management conference (Bowles Symposium) in Chicago, April 24, 2006. The authors would like to acknowledge helpful comments and feedback from the participants at this event, as well as the conference organizer and JRI editor, Richard MacMinn.

Abstract

We provide an overview of how the classical law of large numbers breaks down when pricing life-contingent claims under stochastic as opposed to deterministic mortality (probability, hazard) rates. In a stylized situation we derive the limiting per-policy risk and show that it goes to a non-zero constant, which is in contrast to the classical situation when the underlying mortality decrements are known with certainty. We decompose the standard deviation per policy into systematic and non-systematic components, akin to the analysis of individual stock (equity) risk in a Markowitz portfolio framework. We then draw upon the financial analogy of the Sharpe Ratio to develop a premium pricing methodology under aggregate mortality risk. Our paper is presented in discrete time and is a companion to the continuous-time formulation developed in Milevsky, Promislow and Young (2005).

1 Introduction and Motivation

A basic textbook assumption when pricing insurance is that mortality risk is completely diversifiable and therefore not priced by capital markets in economic equilibrium. Under this traditional paradigm the law of large numbers (LLN) is invoked to argue that the standard deviation per policy (SDP) vanishes in the limit. Therefore, a large enough insurance company portfolio is sufficient to eliminate mortality risk from the pricing and valuation equations. And, although the literature and especially practitioners are well aware that this theorem only works when claims are independent random variables, it is commonly accepted that industrial portfolios of individual life insurance and/or pension annuities satisfy the requisite assumptions. In fact, many finance and economics textbooks use life insurance as the quintessential counterexample to illustrate the fundamental idea that diversifiable (a.k.a. idiosyncratic) risk is not rewarded.

In this paper we use elementary techniques to investigate and illustrate exactly what happens when this so-called independence assumption breaks down. However, we generate this dependency in a subtle way. We link individual mortality (hazard) rates to aggregate population uncertainty via an unknown parameter value which is shared by all insured lives. Thus, while the distribution of lives is independent *conditional* on knowing population mortality, they become dependent once the conditioning is removed. Our paper is motivated by the increasing concern that the evolution of aggregate mortality rates is unknown because of the possibility of fundamental medical breakthroughs (or perhaps natural disasters) that will impact large portfolios of policies in unpredictable ways.

We compute the limiting standard deviation per policy (SDP) under this *stochastic mortality* structure and we use the concept of a Sharpe Ratio to investigate the market price of this risk, since it is obviously no longer diversifiable. Boiled down to its essence we argue that the uncertainty regarding the evolution of the instantaneous force of mortality might induce a dependence than can not be diversified away by selling more claims. In fact, a larger portfolio of policies might actually increase as opposed to reduce the mortality exposure of the insurance company. In practice, this would induce a mortality risk premium which is priced by the market, but one whose magnitude is dependent on the extent to which opposing claims – for example life insurance against pension annuities – can be used as a natural hedge. In this paper we do not examine the hedging arguments and thus stop short of developing a full-fledged model of mortality risk compensation in equilibrium. Instead we focus our attention on the breakdown in the law of large numbers and how it can be translated to the world of financial Sharpe Ratios.

Insurance actuaries are well aware that underlying (population) mortality rates can change and therefore build contingency margins into quoted prices to account for this risk. What we are doing in this paper is providing a more rigorous way of calculating these margins. Indeed, the existence of such a mortality risk premium might breath new life into the "options" that are embedded within insurance and pension policies that have traditionally been viewed as being out-of-the-money and hence valueless. Our theoretical framework – which justifies charging for mortality risk – is consistent with a number of recent industry trends, namely (i) the emergence of longevity-linked bonds, (ii) capacity constraints in the immediate annuity market and (iii) the decline of defined benefit pension plans.

Other scholars have explored the implications of mortality risk premiums vis a vis the uncertainty in aggregate mortality rates. For example Lee Carter (1992) develop a time series model for aggregate mortality rates. Olivieri (2001) examines the implications of this uncertainty for the required reserves for insurance and annuity polices. In related work, Milevsky and Promislow (2001), Dahl (2004), Biffis (2005), Schrager (2006), Ballota and Haberman (2006) as well as Biffis and Millossovich (2006) use continuous-time modeling for the hazard rate to develop pricing formulae for options that are embedded within various insurance and annuity policies. Likewise, Denuit and Dhaene (2006) use comonotonicity methods to compute probabilities, conditional tail expectations and values-at-risk for portfolios of dependent insurance policies. Along the same lines, Cairns, Blake and Dowd (2006), Cox, Lin and Wang (2006) as well as Webb and Friedberg (2006) attempt to price mortality risk using the concept of Wang transform and/or the Consumption-based Capital Asset Pricing Model (CCAPM). And, while all of these are trying to achieve the same goal of modeling and pricing mortality risk in a non-deterministic world, we feel that our paper provides a much-needed pedagogical foundation to all of the above mentioned papers by using basic tools to illustrate the impact of stochastic hazard (mortality, probability) rates on the pricing of mortality-contingent claims.

Thus, if biostatisticians estimate that 50% of 65 year-olds in the year 2005 will live to see their 85th birthday, then a sufficiently homogenous group of 65 year-olds should be charged $0.5e^{-r20}$ for a longevity insurance (a.k.a. pure endowment) policy that pays \$1 dollar upon reaching and surviving to age 85, where r denotes the relevant market interest rate. The underlying assumption is the existence of a known population survival curve – often denoted by $({}_Tp_x)$ in the actuarial literature – which determines the fraction of the population living to any given time T. This is also a probability of survival for any given individual within the group. Thus, if the insurance company sells N of these policies and charges each policyholder $({}_{20}p_{65})e^{-rT}$, on average the company will have enough to pay the \$1 to the survivors. Using the language of portfolio theory, the idiosyncractic risk – a.k.a standard deviation per policy – will go to zero if they sell enough, and so the risk is not priced. In the language of derivative pricing, the Biometrical P measure and the Financial Q measure are one and the same so that $E_P[e^{-rT}] = E_Q[e^{-rT}]$. However, when $({}_Tp_x)$ itself is an estimate, there is always a risk that selling more will increase as opposed to reduce the insurance company's risk exposure. In this case the company might add a risk-charge or risk-premium to $({}_Tp_x)$ to account for this risk independently of any commissions, loadings, etc. In our paper we focus on the relationship between these probabilities in an attempt to flush-out the link with the non-diversifiable component of aggregate mortality risk.

In a related paper – which is a continuous time extension of the current work – by Milevsky, Promislow and Young (2005), the authors value mortality-contingent claims by assuming that the insurance company issuing the claims will be compensated for mortality risk via the Instantaneous Sharpe Ratio of a suitably defined hedging portfolio. In other words, they assume that the insurance company picks a target (profit) expected return to standard-deviation ratio α and then locates a hedging strategy as well as quoted price for any given mortality-contingent claim that will lead to this pre-determined α. This premium principle is at the heart of what we attempt, albeit in a simple one-period discrete framework.

The remainder of this paper is organized as follows. In the next Section 2 we review the basic framework for how the law of large numbers is invoked to kill mortality risk in the limit. In Section 3 we illustrate what happens when the underlying probability is unknown. Section 4 relates this discussion to Sharpe Ratios and provides a very brief overview of the extension to continuous-time, while Section 5 concludes the paper.

2 Longevity Insurance: Deterministic Hazard Rates

To make our basic point we assume that a generic insurance company sells a one-period longevity insurance (a.k.a. endowment) policy which pays \$2 if the annuitant survives to the end of the period, but pays nothing if the annuitant dies prior. We assume the insurance company issues N of these policies – at a market price of $(1+L)$ per policy – to N independent lives, each of whom has an identical probability p of surviving and probability $(1-p)$ of dying prior to the payout date. Each longevity insurance policy generates an end-of-period Bernoulli liability for the insurance company of the form:

$$\text{Longevity Insurance Payoff} := w_i = \begin{cases} \$2, & p \\ \$0, & 1-p \end{cases}. \tag{1}$$

The expected value $E[w_i] = 2p$ and the variance is $var[w_i] = p(2-2p)^2 + (1-p)(0-2p)^2 = 4p(1-p)$. In the simple case that $p = 0.5$ our model collapses to $E[w_i] = 1$ and $var[w_i] = 1$ as well as $SD[w_i] = 1$, which is the most basic type of mortality-contingent claim. Obviously in the event that $p = 0.5$ the fair (unloaded) actuarial premium is \$1, ignoring the time value of money, in which case the market price which we denoted by $(1+L)$ includes a loading of L. Figure #1 [**goes here**] is our basic payoff structure. We now let the random variable,

$$W_N = \sum_{i=1}^{N} w_i \tag{2}$$

denote the insurance company's aggregate liability at the end of the period. Naturally, we are abstracting from the multiperiod nature of the real-world problem where companies might issue these claims on an ongoing basis, and also might have other (insurance) liabilities on their books to partially hedge this exposure.

Either way, if we add-up these independent longevity insurance exposures, the issuing company faces an aggregate expected payout of $E[W_N] = N2p$, an aggregate variance of $var[W_N] = N4p(1-p)$ and an aggregate standard deviation of $SD[W_N] = 2\sqrt{Np(1-p)}$. This immediately leads us to the well-known results that the standard deviation per policy goes to zero, in the limit as $N \to \infty$. Stated technically:

$$\lim_{N\to\infty} \frac{1}{N} SD[W_N] = \lim_{N\to\infty} 2\frac{\sqrt{p(1-p)}}{\sqrt{N}} \to 0 \tag{3}$$

This is a special case or manifestation of the law of large numbers (LLN) which states that:

$$\lim_{N\to\infty} \frac{1}{N} W_N \to p. \tag{4}$$

If the insurance company sells enough of these longevity insurance policies their risk exposure (per policy) goes to zero. Yet another interpretation of this statement is that longevity risk is completely diversifiable and not compensated by markets in equilibrium.

We can go beyond the first two moments of the company's exposure and compute the entire distribution by recognizing that a sum of independent Bernoulli trials is Binomially distributed. The probability distribution for the aggregate payout can be computed via:

$$\Pr[\frac{W_N}{2} \leq k] = \sum_{i=0}^{k} B(i \mid N, p), \tag{5}$$

where $B(i \mid N, p)$ denotes the Binomial distribution with parameter N, p. More formally, $B(i \mid N, p)$ is defined using $\binom{N}{i} p^i (1-p)^{N-i}$. In fact, for large enough values of N we can use the Central Limit Theorem (CLT) to compute the probability the company's aggregate payout will be less than (or greater than) some prespecified level. As an example, assume a company starts their (longevity insurance) business venture with initial capital denoted by C and then issues or sells N of these mortality-contingent claims at a market price of $(1+L)$ per policy. In this case, the probability that the company will completely exhaust (wipe out) their initial capital C, is denoted by $\Pr[W_N \geq N(1+L)+C]$. Using the CLT the relevant quantity can be computed via:

$$\begin{aligned} \Pr[W_N \geq N(1+L)+C] &= \Pr\left[\frac{W_N - E[W_N]}{SD[W_N]} \geq \frac{N(1+L)+C-E[W_N]}{SD[W_N]}\right] \\ &\approx \Pr\left[\mathbf{Z} \geq \frac{N(1+L)+C-2Np}{\sqrt{4Np(1-p)}}\right] \end{aligned} \tag{6}$$

where $\mathbf{Z}$ denotes the standard normal random variable and the probability in question is the right-tail of the distribution. Using the same logic, L can be chosen to keep this probability under a certain tolerance or ruin level, etc. This is all standard.

Table #1 [**placed here**] provides some numerical examples for the (very small) probability of deviating from the expected payout $2Np$, when $N = 10,000$ policies and $p = 0.5$ per policy. As one would expect, the probability of exhausting (even) \$500 of capital is quite small. This is a direct manifestation of the law of large numbers and is the theoretical argument used to justify the inexistence of a mortality (or longevity) risk premium.

3 Longevity Insurance: Stochastic Hazard Rate

With a review of the classic results behind us, in this section we examine what happens when the probability parameter p is unknown. This is also equivalent to not knowing the hazard rate or instantaneous force of mortality underlying the probabilities. Stated differently, what if you have an estimate of p versus a certainty for p? In this case Figure #2 [**placed here**] provides the relevant diagram in which p is replaced by a random variable denoted by $\widetilde{p}$. The underlying payoff function is now:

$$\text{Longevity Insurance Payoff} := \ w_i^* = \begin{cases} \$2, & \widetilde{p} \\ \$0, & 1-\widetilde{p} \end{cases}, \tag{7}$$

where the asterisk on top of the w_i^* reminds the reader that the parameter $\widetilde{p}$ itself has its own (symmetric) distribution denoted by:

$$\widetilde{p} = \begin{cases} p+\pi, & \text{Probability} = 1/2 \\ p-\pi, & \text{Probability} = 1/2 \end{cases}. \tag{8}$$

Obviously we must impose a restriction on the newly defined uncertainty parameter π, namely that $\pi \leq 1-p$ and $p > \pi$. Also, by definition $E[\widetilde{p}] = 0.5(p+\pi) + 0.5(p-\pi) = p$. In general one could envision a situation in which $\widetilde{p}$ can take on a multitude of values – and that is at the heart of the continuous time model developed in Milevsky, Promislow and Young (2005) – but in this paper we will restricted ourselves to the symmetric and binomial uncertainty case.

For the subsequent numerical examples we will assume that $\pi = 0.1$, $p = 0.5$ so that $\widetilde{p}$ takes on values of either 0.6 or 0.4. The intuitive interpretation for this would be that while the expected value $E[\widetilde{p}] = 0.5$ of the survival probability is 0.5, it is equally likely to take on a value of 0.4 (a decline in mortality) or 0.6 (an improvement in mortality). Later we also examine the impact of increasing the uncertainty (spread) π, while keeping the same mean $E[\widetilde{p}]$.

The point of this section – and to some extent, the paper – is to explore the aggregate and the per-policy risk dynamics when $E[\widetilde{p}] = 0.5$ versus the deterministic case in which $p = 0.5$. We now define the total (aggregate) exposure of the insurance company by the notation,

$$W_N^* = \sum_{i=1}^{N} w_i^*, \tag{9}$$

with the immediate implication that $E[W_N^*] = N2p$, which is identical to the traditional (deterministic) case. However, the key difference between W_N and W_N^* lies in the term for the variance. We are no longer entitled to add-up the individual variance terms due to the implicit dependence created by the common $\widetilde{p}$ factor. We use the variance decomposition relationship:

$$\begin{aligned} var[W_N^*] &= E[var[W_N^*|\widetilde{p}]] + var[E[W_N^*|\widetilde{p}]] \\ &= 4N(p - p^2 - \pi^2) + 4N^2\pi^2 \\ &= 4Np(1-p) + 4N\pi^2(N-1) \end{aligned} \tag{10}$$

which collapses to the familiar and intuitive $4Np(1-p)$ when $\pi = 0$, i.e. we are back to the deterministic mortality world. Note also that when $N = 1$, the variance of the payout is the same $4Np(1-p)$ it would be under the deterministic case, which means that an individual policy isn't any riskier under a stochastic $\widetilde{p}$ versus a deterministic $p = E[\widetilde{p}]$. It is the portfolio aggregation that creates the extra risk.

For example, if the company sells $N = 10,000$ longevity insurance policies under the parameter set $\pi = 0.1$, so that $\widetilde{p}$ takes on a value of either 0.6 or 0.4 with equal probability, the variance of the aggregate payout becomes $0.96N + 0.04N^2$ according to equation (10). Notice the N^2 term, and herein lies the crucial breakdown in the law of large numbers. Notice that the variance grows non-linearly in N, which means that that standard deviation per policy $\sqrt{0.96N + 0.04N^2}/N$, will converge to a constant 0.2 instead of zero. No matter how many insurance policies the company issues, the uncertainty (risk) per policy will never be less than \$0.20 per \$1.00 of expected payoff. The risk never goes away. This is akin to the much analyzed market risk in modern portfolio theory. Diversification can only reduce risk up to a certain point. In general, we have that

$$\lim_{N\to\infty} \frac{1}{N} SD[W_N^*] = \lim_{N\to\infty} \frac{2\sqrt{N(p - p^2 - \pi^2 + N\pi^2)}}{N} \to 2\pi \neq 0 \tag{11}$$

The standard deviation per policy (SDP) converges to the total spread (2π) in the probability (hazard) rate. Naturally, the only case in which this quantity would vanish is when $\pi = 0$,which would indicate there is no uncertainty in the estimate for the probability of survival.

In the language of portfolio theory, one can think of 2π as the *market risk* and the difference between 2π and the SDP as the idiosyncractic or diversifiable risk. We can define this component as:

$$\text{Idiosyncractic Mortality Risk} = 2\left(\sqrt{\frac{p-p^2-\pi^2}{N}+\pi^2}-\pi\right) \tag{12}$$

Table #2 [**placed here**] is a numerical comparison of the standard deviation per policy as a function of the number of policies sold N, for the two different hazard rate assumptions. In the first case the mortality parameter p is known with certainty to be 0.5, in the second column the mortality parameter is unknown, but is expected to be 0.5. Notice that although the SDP values are identical for $N = 1$, the rate at which the SDP value decays is much higher for the deterministic case vs. the stochastic case. In fact, in the limit we are left with an extra (non-diversifiable) \$0.20 per policy under the stochastic case. When $N = 100$ policies the SDP can be decomposed into two parts. The first portion of $2\pi = \$0.20$ is systematic risk and the second portion of \$0.023 is diversifiable or idiosyncractic. The two add-up to the total SDP of \$0.223. Figure #3 [**placed here**] provides a graphical illustration of the relationship between the SDP and the number of policies issued under deterministic versus stochastic mortality or hazard rates. Notice that as N gets large the SDP curve converges to zero (i.e. no market risk) under the deterministic p rate, but plateaus at 2π under the stochastic $\widetilde{p}$ mortality case. One can readily see the connection between the 2π level and the so-called market risk in modern portfolio theory.

We can go beyond the SDP and compute the actual distribution as follows of W_N^* as follows. Notice that conditional on a given value of $\widetilde{p}$, the distribution of the aggregate exposure is Binomial. Thus, since $\widetilde{p}$ can take on two distinct values we are left with a mixture of Binomial distributions.

$$\Pr[W_N^*/2 \leq k] = \frac{1}{2}\sum_{i=0}^{k} B(i \mid N, p+\pi) + \frac{1}{2}\sum_{i=0}^{k} B(i \mid N, p-\pi), \tag{13}$$

where the coefficients are defined as before.

With an actual distribution in hand, Table #3 [**placed here**] provides additional estimates of how risk gets diversified under deterministic vs. stochastic probability values when the company sells $N = 100$ longevity insurance policies, under the same parameters as above, namely when $E[\widetilde{p}] = 0.5$. Note that the probability of exhausting \$2 of capital per \$100 of reserves is 0.382 under deterministic rates and a much higher 0.483 under stochastic rates. Likewise, the probability of exhausting \$30 of capital per \$100 of reserves is a mere 0.001 under deterministic rates versus 0.065 under stochastic rates. To create a link with the language of financial economics we label the first column the diversifiable risk case and the second column the non-diversifiable case. Table #3 also provides a rough indication of how much more (solvency) capital is required to support a given book of business when the underlying mortality rates are themselves unknown. For example, if capital is fixed so that the company will remain solvent with 94% probability, then one will require \$30 of capital in a stochastic mortality world versus \$15 in a deterministic mortality world. Of course, these are highly idealized examples, but the main idea should be clear. For any given level of confidence (or CTE, or ruin probability) the company's required capital is higher when the mortality probability is estimated as opposed to known with certainty.

4 Pricing via the Sharpe Ratio

Although there are many ways in which the Sharpe Ratio is presented in the finance literature, for the most part the definition is as follows. Let X denote the return from some risky investment asset and let R denote the risk-free return. The expected value and standard deviation are denoted by $E[X]$ and $SD[X]$ respectively. The Sharpe Ratio of this investment (gamble) is defined and denoted by:

$$\alpha = \frac{E[X] - R}{SD[X]}, \tag{14}$$

with the obvious understanding that $SD[X] > 0$. And, when the risk-free rate R is assumed to be zero, then Sharpe Ratio is simply the ratio of risk to expected return.

For example, the historical return from a broadly diversified portfolio of stocks as proxied by the Standard & Poors 500 index, has been in the vicinity of 11% (nominal) with a standard deviation of 20%. Over the same period of time the risk-free return has averaged 6%

(nominal), both according to Ibbotson Associates estimates. We would therefore state that the Sharpe Ratio from stocks or large-cap equity as an asset class is $\alpha = (0.11 - 0.06)/(0.20) = 0.25$. The Sharpe Ratio is one (very popular) and widely used method for adjusting realized returns for its level of risk. In general. Sharpe Ratios greater than a value of 1 are quite rare and hard to achieve.

Within the context of a financial valuation of longevity insurance, one could envision the following premium-pricing principle. *The loading L will be set so that the Sharpe Ratio α is consistent with other asset classes in the economy.* From the perspective of the insurance company, we can define the Sharpe Ratio in aggregate to be:

$$\text{Sharpe Ratio} := \frac{N(1+L) - E[W_N^*]}{SD[W_N^*]} = \frac{1 + L - 2p}{\frac{1}{N} SD[W_N^*]}. \tag{15}$$

Note that we can use an SDP (i.e. the denominator in the Sharpe Ratio) under $N \to \infty$ or under a finite value of N, depending on the context of the problem. We now provide some examples to illustrate how this would be applied. Suppose the issuing company sells $N < \infty$ policies for which they want a Sharpe Ratio compensation of 0.25 units, which is akin to the numbers mentioned above for equity markets. The total premium collected would be $N(1+L)$ and the expected payout $E[W^*]$ will be N, when $E[p] = 0.5$. For our values of $p + \pi = 0.6$ and $p - \pi = 0.4$ the payout will have a standard deviation of $\sqrt{0.96N + 0.04N^2}$. This gives a Sharpe Ratio of $NL/\sqrt{0.96N + 0.04N^2}$. We can then solve for L in terms of N and imply the required loading.

A property of the Sharpe Ratio is that L decreases with greater N as we would intuitively expect. For example, the loading $L = 0.092$, which is a 9.2% mark-up when $N = 10$ policies are sold, but $L = 0.061$ when $N = 50$ and $L = 0.051$ when $N = 500$. However, L will never become zero as it would if $\pi = 0$, but rather it will approach a limit of 0.05 as N approaches ∞.

Once again, we are shying away from declaring that the market will compensate the insurance company for taking on the risk of only $N = 10$ or 50 or 500 policies. That is an economic equilibrium statement that would also depend on the extent to which the uncertainty in $\widetilde{p}$ can be hedged by selling other insurance claims with opposite aggregate exposures. We can provide an upper bound on the equilibrium Sharpe Ratio for mortality risk by arguing that it should not exceed 2π, since any risk beyond that can be diversified away by selling more policies.

Note, once again, that L in our framework denotes the loading intended to account for mortality risk and would be above and beyond other loadings to account for expenses, profits, etc. Another observation is that in the presence of greater uncertainly about our mortality parameter $\widetilde{p}$, the required loading will increase, as one would expect. The Sharpe Ratio provides the analytical link. Suppose for example that $\pi = 0.2$ instead of $\pi = 0.10$ and thus $\widetilde{p}$ now takes a value of either 0.7 or 0.3 with equal probability. The standard deviation of the aggregate payout will become $\sqrt{0.84N + 0.16N^2}$. And, as $N \to \infty$ the SDP will converge to $2\pi = \$0.4$ and the required loading will increase to $L = 0.10$ units, in order to maintain the same 0.25 Sharpe Ratio.

Finally, we briefly discuss the extension of this framework to continuous time. More extensive results are provided in Milevsky, Promislow and Young (2005). Here we only provide the highlights. Recall that under a continuous law of mortality, the survival probability which we labeled p in equation (1), becomes:

$$
{}_tp_x = e^{-\int_0^t \lambda(x+s)ds}, \tag{16}
$$

where $\lambda(x+s)$ now denotes the instantaneous force of mortality (IFM) at age $x+s$. In this case, the random payout from selling a longevity insurance policy to someone aged x, that expires at time T, would have an expected value of $w = 2({}_tp_x)$. The uncertainty in the survival probability can be mapped into uncertainty regarding the instantaneous force of mortality λ. Thus, in the event that the IFM itself is random, one can think of a diffusion hazard rate denoted by λ_t which would satisfy a stochastic differential equation (SDE). It can be viewed as the extension of the concept underlying π in equation (8). Figure #4 [**placed here**] illustrates some sample paths for stochastic hazard rates in continuous time.

We refer the interested reader to Milevsky, Promislow and Young (2005) for the precise details, but here is a high-level summary of those results. They prove that when the instantaneous force of mortality (IFM) is assumed constant, as in the traditional actuarial insurance models, the value of the mortality-contingent claim that is priced via the instantaneous Sharpe Ratio collapses to the discounted expected payoff using the Biometric probability measure as the number of policies $N \to \infty$. In other words, if mortality risk is diversifiable the insurance company can not charge any more regardless of their profit target α. However, they also prove that if the hazard rate for the IFM is non-zero, the financial value of the mortality-contingent claim is greater than the above mentioned discounted expected payoff, even as $N \to \infty$. Furthermore, their valuation operator is subadditive and satisfies a number of other appealing properties.

5 Conclusion

There have been a number of research papers that have focused attention on the problem of longevity risk, from a variety of different perspectives. For example, in this volume of the *Journal of Risk and Insurance* the papers by Cairns, Blake and Dowd (2006) as well as Cox, Lin and Wang (2006) examine the pricing of longevity-linked bonds using various specifications of the mortality risk premium.

In this brief paper we provide a pedagogical overview of how the classical law of large numbers breaks down when pricing life-contingent claims under stochastic, as opposed to deterministic, mortality (hazard) rates. In other words, we go back to basics and explain how the relationship between the uncertainty in our estimates translates directly into a non-zero standard deviation per policy. In a stylized situation, we derive the limiting per-policy risk and show that it goes to a non-zero constant, which is in contrast to the classical situation when the underlying mortality rates are known with certainty. We then draw upon the financial analogy of the Sharpe Ratio to develop a premium pricing methodology for the non-diversifiable component of aggregate mortality risk. Our main qualitative insight is that similar to the financial economic approach to analyzing stock market risk, the uncertainty embedded within mortality-contingent claims can be decomposed into idiosyncractic (diversifiable) and non-diversifiable components. And, while the equilibrium magnitude of these two components is largely an empirical question – and depends on the extent to which part of the uncertainty which can be hedged by selling opposing claims – our simple analysis provides yet additional proof for the existence of a mortality risk premium.

References

[1] Ballotta, L. and S. Haberman (2003), Valuation of guaranteed annuity conversion options, *Insurance: Mathematics and Economics*, Vol. 33, pg. 87-108.

[2] Blake, D. and W. Burrows (2001), Survivor bonds: Helping to hedge mortality risk, *Journal of Risk and Insurance*, Vol. 68, pg. 339-348

[3] Biffis, E. and P. Millossovich (2004), The fair value of guaranteed annuity options, *Working Paper, Universita degli Studi di Trieste*

[4] P.P. Boyle and M. Hardy (2003), Guaranteed annuity options, *ASTIN Bulletin*, Vol. 33, pg. 125-152.

[5] Cairns, A.J.G., D. Blake and K. Dowd (2004), A two-factor model for stochastic mortality with parameter uncertainty, *Journal of Risk and Insurance*, Vol. 73(4), pg.

[6] Cox, S.H., Y. Lin and S. Wang (2006), Multivariate exponential tilting and pricing implications for mortality securitization, *Journal of Risk and Insurance,* Vol. 73(4), pg.

[7] Dahl, M. (2004), Stochastic mortality in life insurance: Market reserves and mortality-linked insurance contracts, *Insurance: Mathematics and Economics,* Vol. 35, pg. 113-136

[8] Denuit, M. and J. Dhaene (2006), Comonotonic bounds on the survival probabilities in the Lee-Carter model for mortality projection, *working paper*

[9] DiLorenzo, E. and M. Sibillo (2003), Longevity risk: Measurement and Application Perspectives, *Working Paper, Universita degli Studi di Napoli*

[10] Lee, R.D. and L.R. Carter (1992), Modeling and forecasting U.S. mortality, *Journal of the American Statistical Association*, Vol. 87(419), pg. 659-671

[11] Milevsky, M.A. and S.D. Promislow (2001), Mortality derivatives and the option to annuitize, *Insurance: Mathematics and Economics*, Vol 29, pg. 299-318

[12] Milevsky, M.A., S.D. Promislow and V. Young (2005), Financial Valuation of Mortality Risk via the Instantaneous Sharpe Ratio, *working paper, under review.*

[13] Olivieri, A. (2001), Uncertainty in mortality projections: An actuarial perspective, *Insurance: Mathematics and Economics,* Vol. 29, pg. 231-245.

[14] Promislow, S.D. and V.R. Young (2004), Indifference pricing via the probability of ruin, *Working Paper, University of Michigan*

[15] Smith, A., I. Moran and D. Walczak (2003), Why can financial firms charge for diversifiable risk?, *Working Paper, Deloitte Touche Tohmatsu*

[16] Schrager (2006), Affine stochastic mortality, *Insurance: Mathematics and Economics*, 2006

[17] Soininen, P. (1995), Stochastic variation of interest and mortality, *Proceedings of the 5th AFIR International Colloquium*, pg. 871-904.

[18] Webb, T. and L. Friedberg (2006), Life is cheap: Using mortality bonds to hedge aggregate mortality risk, working paper, *Centre for Retirement Research, Boston College*

Longevity Insurance Payoff: Deterministic Hazard Parameters

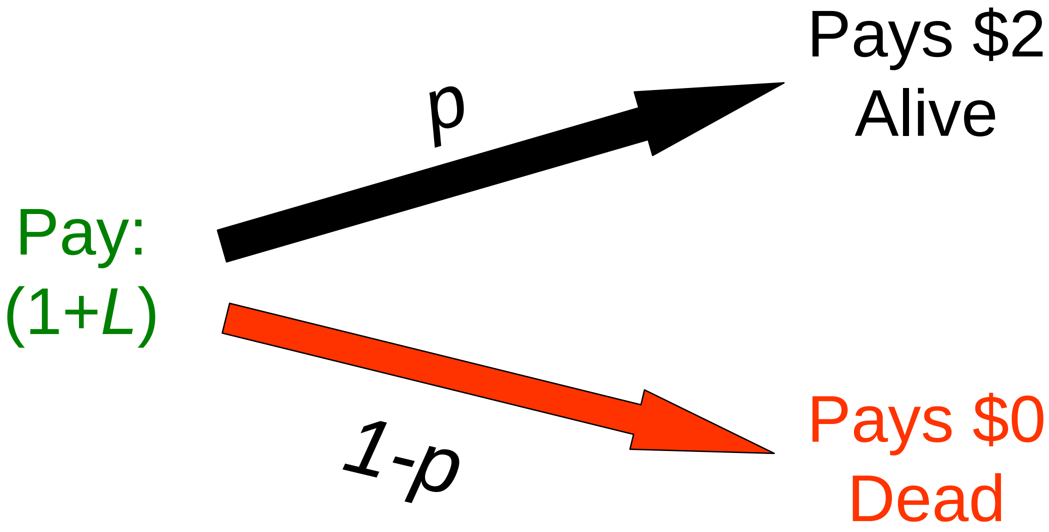

Pay:
(1+*L*)

p

Pays \$2
Alive

$1-p$

Pays \$0
Dead

Figure #1

Longevity Insurance Payoff: Stochastic Hazard Parameters

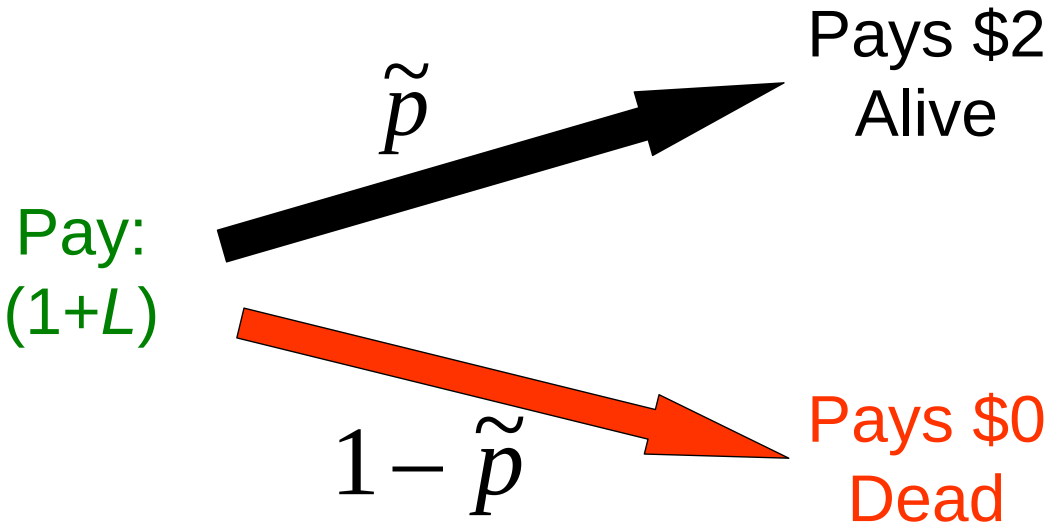

Pay: (1+*L*)

$\tilde{p}$

Pays $2
Alive

$1-\tilde{p}$

Pays $0
Dead

Figure #2

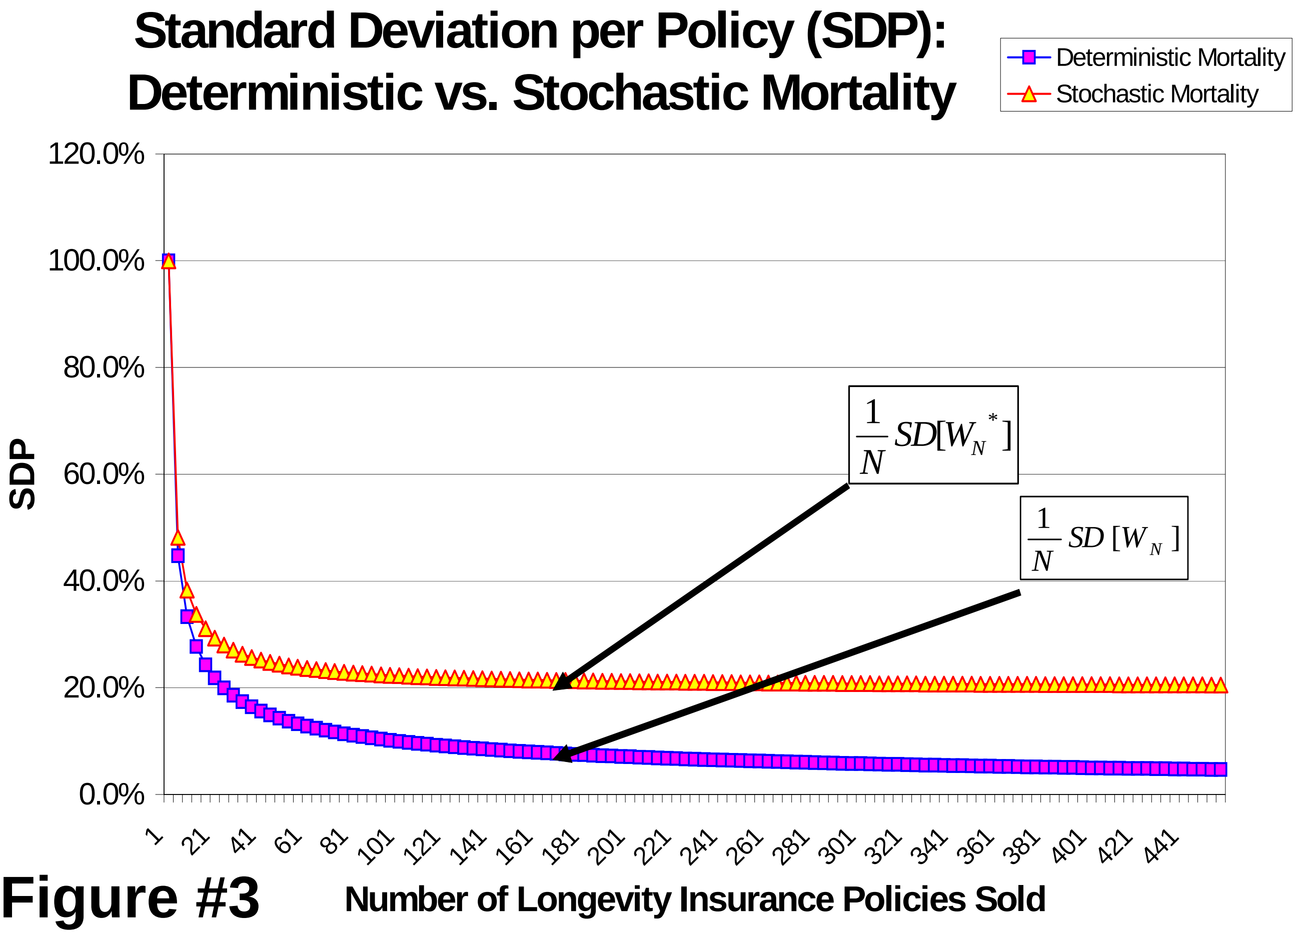
Standard Deviation per Policy (SDP): Deterministic vs. Stochastic Mortality
Deterministic Mortality
Stochastic Mortality
120.0%
100.0%
80.0%
60.0%
40.0%
20.0%
0.0%
SDP
$\frac{1}{N}SD[W_N^*]$
$\frac{1}{N}SD[W_N]$
1
21
41
61
81
101
121
141
161
181
201
221
241
261
281
301
321
341
361
381
401
421
441
Number of Longevity Insurance Policies Sold

Figure #3

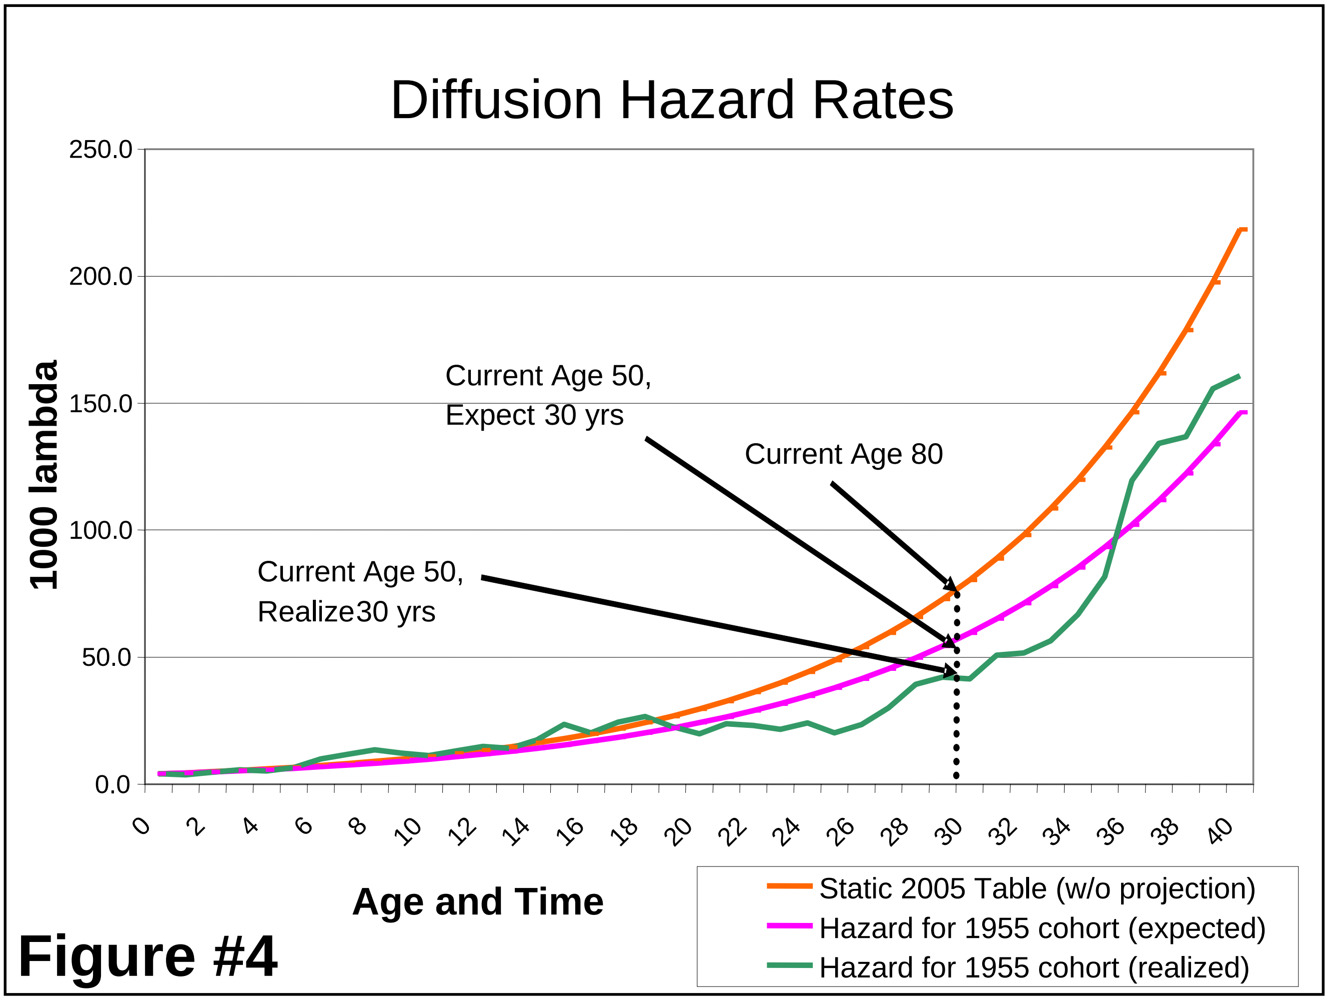
Diffusion Hazard Rates
1000 lambda
250.0
200.0
150.0
100.0
50.0
0.0
0
2
4
6
8
10
12
14
16
18
20
22
24
26
28
30
32
34
36
38
40
Age and Time
Current Age 50,
Expect 30 yrs
Current Age 80
Current Age 50,
Realize30 yrs
Static 2005 Table (w/o projection)
Hazard for 1955 cohort (expected)
Hazard for 1955 cohort (realized)

Figure #4

Numerical Example of Classic Model
Sell 10,000 Longevity Insurance Policies

Total Payout Outside the Range of:	Event Probability:
($5,000 -- $15,000)	0.000%
($7,500 -- $12,500)	0.000%
($9,000 -- $11,000)	0.000%
($9,500 -- $10,500)	0.000%
($9,700 -- $10,300)	0.542%

Note: Standard deviation of payout is only $100, since p = 0.5.

Table #1

How Fast Does Idiosyncratic Risk Decline?
Standard Deviation per Policy (SDP)

Number of Policies Sold: **N**	**Deterministic Hazard Parameter** **p = 0.5**	**Stochastic Hazard Parameter** **E[p] = 0.5**
1	$1.000	$1.000
2	$0.707	$0.721
5	$0.447	$0.482
100	$0.100	$0.223
1000	$0.032	$0.202
10,000	$0.010	$0.200
Infinity	$0.000	$0.200

p = 0.6 or 0.4 with even odds

Table #2

Company Sold 100 Longevity Policies: Deterministic vs. Stochastic Hazard Rates

Total Payout Larger Than: K	$\Pr[W_{100} > K]$	$\Pr[W_{100}^{*} > K]$
	Diversifiable	**Non-Diversifiable**
$102	0.382	0.483
$110	0.136	0.411
$120	0.018	0.231
$130	0.001	0.065

Note: Survivor gets $2

E[p] = (0.6)(0.5)+(0.4)(0.5) = p = 0.5

Table #3