To Pensions Institute

DISCUSSION PAPER PI-0823

A Computationally Efficient Algorithm for Estimating the Distribution of Future Annuity Values under Interest-Rate and Longevity Risks

Kevin Dowd, David Blake, and Andrew J.G. Cairns

Published in the North American Actuarial Journal, Volume 15, Number 2, 2011

ISSN 1367-580X

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

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

A COMPUTATIONALLY EFFICIENT ALGORITHM FOR ESTIMATING THE DISTRIBUTION OF FUTURE ANNUITY VALUES UNDER INTEREST-RATE AND LONGEVITY RISKS

Kevin Dowd,* David Blake,[†] and Andrew J. G. Cairns[‡]

Abstract

This paper proposes a computationally efficient algorithm for quantifying the impact of interestrate risk and longevity risk on the distribution of annuity values in the distant future. The algorithm simulates the state variables out to the end of the horizon period and then uses a Taylor series approximation to compute approximate annuity values at the end of that period, thereby avoiding a computationally expensive "simulation-within-simulation" problem. Illustrative results suggest that annuity values are likely to rise considerably but are also quite uncertain. These findings have some unpleasant implications both for defined contribution pension plans and for defined benefit plan sponsors considering using annuities to hedge their exposure to these risks at some point in the future.

1. INTRODUCTION

This paper examines the impact of interest-rate risk and longevity risk on the distribution of annuity values in the distant future. At first sight, this might appear to be a rather arcane issue. Most people do not care much about current annuity values, so why would they care about possible annuity values in, say, 40 years' time? Nonetheless, we would suggest that the future values of annuities are much more important than they might first appear to be. One reason is suggested by the global increase in longevity: People are living longer than previously anticipated, and it is natural to ask how much this is likely to cost. To answer this question, we need some index of the cost of life expectancy, and a natural index is the cost of a life annuity, that is, the expected present value of an income steam of \$1 per annum until the buyer of the annuity dies. The cost of increased longevity is then the likely increase in future annuity values. Thus, an assessment of likely future annuity values is a key to evaluating the likely cost of higher longevity.

There is a second and often more personal reason why we should care about future annuity values. Consider the illustrative case of a male currently aged 25 who is starting a defined contribution (DC) pension plan and is planning to retire in, say, 40 years' time at the age of 65. He anticipates that when he reaches that age, he will convert his accumulated pension fund into a life annuity to hedge his own longevity risk and so avoid outliving his own financial resources. The value of his retirement income will depend not only on the value of his pension fund, but also on the price of annuities at that time. Other things being equal, this means that his retirement income prospects will be affected by the distribution of future annuity values: The greater the dispersion of that distribution, the more risky his retirement income will be. Hitherto, analyses of

^{*} Kevin Dowd, PhD, is Visiting Professor, Pensions Institute, Cass Business School, 106 Bunhill Row, London EC1Y 8TZ, UK, kevin.dowd@hotmail.co.uk.

[†] David Blake, PhD, is Professor of Pension Economics and Director of the Pensions Institute at Cass Business School, 106 Bunhill Row, London EC1Y 8TZ, UK, D.Blake@city.ac.uk.

[‡] Andrew J. G. Cairns, PhD, is Professor of Financial Mathematics, Actuarial Mathematics and Statistics, School of Mathematical and Computer Sciences, Heriot-Watt University, Edinburgh EH14 4AS, United Kingdom, a.j.g.cairns@ma.hw.ac.uk.

DC plans have tended to focus on the risks facing the member that arise from the uncertainty of the future value of the pension fund itself; that is, they have tended to focus on investment and contribution risks during the accumulation stage of the plan and, in so doing, have tended to overlook the impact of distribution stage risks, most notably interest rate and longevity risks, that can affect retirement income through their impact on annuity values. The superficially arcane issue of the distribution of future annuity values therefore turns out to be a key ingredient in determining the risks associated with DC pension plans. The same holds for defined benefit plan sponsors who consider using annuities to hedge their exposure to interest-rate and longevity risks at some point in the future.

This paper investigates this issue further and provides some illustrative results for a long-term (40-year) horizon. As is well known, the principal factors that determine annuity values are lifeexpectancy prospects (which depend on future mortality rates) and interest rates.¹ To model future annuity values, we therefore need a model of future mortality rates. We also need a model of future interest rates, since these will influence the future price of the bonds the annuity provider buys to make annuity payments and, in turn, the discount rate used to value the annuity itself. The models we use involve stochastic simulation, but any naive attempt to use stochastic simulation runs into a problem: If we simulate the state variables out to some future horizon period T, we then face the problem of how to obtain the future time-T annuity values, contingent on the time-T values of the state variables. The most obvious approach is to use stochastic simulation for this purpose, but we would then find ourselves running "simulations within simulations" that can be computationally very expensive. To operate within feasible real-time constraints, we need to find an alternative method to obtain the future annuity values contingent on the outcomes of the time-T state variables. We solve this problem using a Taylor series approximation: We simulate

¹ If interest rates are low at the point of retirement, the price of the annuity purchased will be high—equivalently the yield on the annuity will be low—since the annuity provider buys bonds to make the annuity payments and these will be expensive when interest rates are low.

the state variables out to T, and then use the Taylor series approximation to estimate the annuity values at T as functions of the values of the state variables.

This paper is organized as follows. Section 2 describes the stochastic mortality model used to model longevity risk, and Section 3 outlines the stochastic interest-rate model that we use.² Section 4 explains computational issues and proposes a "Taylor series approximation within simulation" approach that allows us to estimate future annuity values using simulated future values of the underlying state variables. Section 5 presents some illustrative results, and Section 6 concludes and elaborates on the implications of our findings for the riskiness of DC plan retirement incomes.

2. A STOCHASTIC MORTALITY MODEL

We model mortality stochastically using the Cairns-Blake-Dowd (CBD) model described in Cairns et al. (2006a) but with the parameterization defined in Cairns et al. (2009). Let q(t + 1, x) be the realized mortality rate in year t + 1 (that is, from time t to time t + 1) for individuals aged x at time t. We assume that q(t + 1, x) is governed by a two-factor Perks stochastic process:

$$q(t + 1, x) = \exp[\kappa_1(t + 1) + \kappa_2(t + 1)(x - \bar{x})] / \{1 + \exp[\kappa_1(t + 1) + \kappa_2(t + 1)(x - \bar{x})]\},$$
(1)

where $\kappa_1(t + 1)$ and $\kappa_2(t + 1)$ are themselves stochastic processes that are measurable at time t + 1 (see Perks 1932; Benjamin and Pollard 1993), and \bar{x} is a constant that is typically set to the mean of the range of ages used to calibrate the

² The present paper assumes that longevity and interest-rate risks are independent of each other. We believe that this is a reasonable assumption for most normal states of the world in the short run. We recognize that there will be some extreme states—such as those associated with a global war, a pandemic, or a huge meteor hitting Earth—where the assumption is not valid. In the long run, however, increasing longevity, if it is associated with an increase in the length of the working life, will increase the supply of labor relative to capital, which, in turn, will increase the return on capital (i.e., the rate of interest rate) relative to the wage rate.

model. Now let $\kappa(t + 1) = (\kappa_1(t + 1), \kappa_2(t + 1))'$ and assume that $\kappa(t + 1)$ is a random walk with drift

$$\kappa(t+1) = \kappa(t) + \mu + CZ(t+1), \quad (2)$$

where μ is a constant 2 × 1 vector of drift parameters, *C* is a constant 2 × 2 lower triangular Choleski square root matrix of the covariance matrix *V* (that is, $V = CC^T$), and Z(t + 1) is a 2 × 1 vector of independent standard normal variables. Cairns et al. (2006a, 2009) show that this model provides a good fit to U.K. Office for National Statistics (ONS) data for English and Welsh males over 1961–2004.

Now let S(t + 1, x) be the survivor index at time t + 1 of a cohort aged x in year 0: that is, S(t + 1, x) is the probability, measured retrospectively, that an individual aged x at time 0 survives to time t. For any given x, S(0, x) = 1 and S(t + 1, x) will decrease as t gets bigger and eventually approach 0 as t gets large. Given any path of q(t + 1, x) as obtained above, we then derive a corresponding path of S(t + 1, x) from the relationship between mortality rates and the survivor index:

$$S(t + 1, x) = (1 - q(t + 1, x))S(t, x).$$
 (3)

Note that these survivor rates are driven off the state variables $\kappa_1(t + 1)$ and $\kappa_2(t + 1)$. For our purposes, we wish to simulate sets of state variables out to a future time *T*, and then estimate the expectations of (3) conditional on surviving to the specified future date and conditional on the future values of the state variables $\kappa_1(T)$ and $\kappa_2(T)$ at that date.

3. A STOCHASTIC INTEREST-RATE MODEL

We also need a model of the interest-rate process, and the simplest model that meets our requirements is the Cox-Ingersoll-Ross (CIR) model (1985). This model postulates that the instantaneous spot interest rate r obeys the following continuous-time process:

$$dr(t) = \alpha(\bar{r} - r(t))dt + \sigma \sqrt{r(t)}dW(t), \quad (4)$$

where α indicates the strength of the meanreversion process governing r, \bar{r} is the long-term mean instantaneous spot interest rate, σ is the interest-rate volatility, and dW(t) is a standard geometric Brownian motion. This model is attractive because it allows for interest rates to be mean-reverting but does not allow them to become negative.³ Another attractive feature is that it gives us a straightforward formula for the spotrate term structure based on the current instantaneous spot rate: If R(t, T) is the time-*t* spot rate for the fixed maturity date *T* (that is, T - t years to maturity), then

$$R(t, T) = -(T - t)^{-1} \ln P(t, T), \qquad (5)$$

where P(t, T) is the time-*t* price of a zero-coupon bond with maturity *T* and where

$$P(t, T) = A(t, T) \exp[-B(t, T)r(T)],$$

$$A(t, T) = (2\gamma \exp[(\alpha + \gamma)(T - t)/2] /{(\gamma + \alpha)(\exp[\gamma(T - t)] - 1)} + 2\gamma})^{2\bar{r}\alpha/\sigma^2},$$

$$B(t, T) = 2(\exp[\gamma(T - t)] - 1) /{(\alpha + \alpha)(\exp[\gamma(T - t)] - 1)} + 2\gamma)$$

$$\gamma \{(\gamma + \alpha)(\exp[\gamma(I - t)] - 1) + 2\gamma\},\$$
$$\gamma = \sqrt{\alpha^2 + 2\sigma^2}.$$

From a computational perspective, the CIR model is appealing because the exact distribution of the instantaneous spot interest rate under the CIR model is known. To be precise, if r(T) follows a CIR process, then $(4\alpha r(T))/{\sigma^2(1 - \alpha r(T))}$ $\exp[-\alpha T]$ has a noncentral chi-squared distribution with $4\alpha \bar{r}/\sigma^2$ degrees of freedom and a noncentrality parameter equal to $(4\alpha r(0))/$ $\{\sigma^2(1 - \exp[-\alpha T])\}$ (Cairns 2004, Theorem 4.8 [c]). This means that we can simulate values of r(T) directly from their exact distribution using the CIR parameters and the current instantaneous spot rate r(0) as inputs. Once we have these terminal instantaneous spot-rate values, we can then use equation (4) to infer the spot-rate term structures contingent on each of these values.

A typical distribution for r(T) is given in Figure 1 and Table 1, based on assumed values of $\alpha = 0.20$, $\sigma = 0.10$, and $\bar{r} = 0.04$. Figure 1 shows the probability density function (pdf), and Table 1 gives some of its key parameters. We can see that the pdf has a strong positive skew and a long right-hand tail. What is perhaps most striking about this pdf is the extent of its dispersion—for example, 30% of simulated values are under 0.02, 30% are above 0.049, and we get a

³ The model and its properties are discussed in Cairns (2004, p. 67).

Figure 1 Probability Density Function for Future CIR Instantaneous Spot Interest Rate at a 40-Year Horizon



Note: Figures shows a noncentral chi-squared distribution based on eq. (4) calibrated to $\alpha = 0.20$, $\sigma = 0.10$, $\bar{r} = 0.04$, r(0) = 0.04, and T = 40.

Table 1

Properties of the Probability Density Function of the Cox-Ingersoll-Ross Instantaneous Spot Interest Rate in 40 Years' Time

Parameter	Parameter Value		
Mean	0.040		
Standard deviation	0.032		
Skewness	1.581		
Kurtosis	6.742		
10th percentile	0.008		
20th percentile	0.014		
30th percentile	0.020		
40th percentile	0.026		
50th percentile	0.032		
60th percentile	0.040		
70th percentile	0.049		
80th percentile	0.061		
90th percentile	0.082		

Note: The instantaneous spot interest rate is assumed to be governed by a CIR process (given by eqs. [4]–[5] in the text) with parameters $\alpha = 0.20$, $\alpha = 0.10$, and $\bar{r} = 0.04$.

small number of very high values in the long right-hand tail—and this is the case even though the spot-rate process is mean reverting.⁴

These findings confirm that interest-rate risk will have a considerable impact on the distribution of future annuity values.

4. COMPUTATIONAL ISSUES

If we combine these models, we have three random state variables: the $\kappa_1(.)$ and $\kappa_2(.)$ state variables from the mortality model, and the instantaneous spot interest rate r(.) from the interest-rate model. Suppose that the current time is time 0 and we wish to simulate these state variables out to some future period *T*. Suppose, then, that we take *J* simulation paths of each

⁴ It is also worth noting that the CIR process is widely regarded by interest-rate practitioners as *under*estimating the true distribution of instantaneous spot rates at long horizons. To the extent that this is the case, then our later results will understate the impact of interest-rate risk on the distribution of future annuity values and therefore underestimate the riskiness of DC pension plans.

state variable out to period *T*, and let $[\kappa_1^i(T), \kappa_2^i(T), r^i(T)]$ be the *j*th set of simulated state variables for period *T*. The fair value of an annuity at time *T*, a(T), then depends on the values of these state variables, so we will, from time to time, use the extended notation $a(T) \equiv a(T; \kappa_1(T), \kappa_2(T), r(T))$ to reflect the fair value's dependence on the state variables. Similarly, in the case of interest rates, we have $R(T, T + i) \equiv R(T, T + i; r(T))$. The fair value of the annuity, conditional on the simulated state variables under simulation path *j*, is thus

$$a^{j}(T) = a(T; \kappa_{1}^{j}(T), \kappa_{2}^{j}(T), r^{j}(T))$$

$$= (1 + \phi) \sum_{i=1}^{50} \exp(-iR(T, T + i; r^{j}(T)))$$

$$E[S(T + i, x)/S(T, x)|$$

$$\kappa_{1}^{j}(T), \kappa_{2}^{j}(T)], \qquad (6)$$

where ϕ is the cost loading factor built into the annuity value and R(T, T + i; r(T)) is the spot interest rate prevailing over the period from *T* to T + i given r(T). Equation (6) is a sum of the product of future expected survivor rates from age 66 to age 115 (conditional on surviving to age 65) and the price of a zero-coupon bond maturing with unit value at the same age, and we assume that no one lives beyond age 115. The term $E[S(T + i, x)/S(T, x)|\kappa_1^i(T), \kappa_2^i(T)]$ is to be interpreted as the expected probability that an individual currently aged x will survive to year T + i conditional on their surviving to T and conditional on the mortality state parameters at T.

However, we cannot compute equation (6) directly because there is no simple formula for $E[S(T + i, x)/S(T, x)|\kappa_1(T), \kappa_2(T)]$ in terms of the mortality state variables $\kappa_1^j(T)$ and $\kappa_2^j(T)$. The most obvious solution to this problem would be to use stochastic simulation to estimate E[S(T + $i, x)/S(T, x)[\kappa_1(T), \kappa_2(T)],$ for each set of $\kappa_1^j(T)$ and $\kappa_{2}^{j}(T)$. For example, we might use M simulation trial paths to generate $j = 1, \ldots, M$ pairs of simulated time-T state variables $\kappa_1^j(T)$ and $\kappa_2^j(T)$. From each such pair, we might then generate k = 1, ..., M sets of paths for the post-T survivor rates $S^{jk}(T + i, x)/S^{jk}(T, x)$ and obtain $E[S^{j}(T + i, x)/S^{j}(T, x)|\kappa_{1}^{j}(T), \kappa_{1}^{j}(T)]$ as their average. In principle, this solution would work, but it involves our taking M^2 simulation paths,

and this would be computationally extremely expensive.⁵

A simple way to reduce this computational burden is to use a Taylor series approximation, as suggested by Cairns (2011). Let $\kappa = [\kappa_1(T), \kappa_2(T)]'$ and let $\hat{\kappa} = E[\kappa]$ be the expectation of the mortality state variables at *T*. Define $f(i, x, \kappa)$ $= \Phi^{-1}(E[S(T + i, x)/S(T, x) | \kappa = (\kappa_1(T), \kappa_2(T))])$ as the probit transformation of $E[S(T + i, x)/S(T, x) | \kappa = (\kappa_1(T), \kappa_2(T))]$, where $\Phi(.)$ is the standard normal distribution function.⁶ We then take the following second-order Taylor series expansion of $f(i, x, \kappa)$ around $\hat{\kappa}$:

$$f(i, x, \kappa) \approx \Delta_0(i, x) + \Delta_1(i, x)'(\kappa - \hat{\kappa}) + \frac{1}{2} (\kappa - \hat{\kappa})' \Delta_2(i, x)(\kappa - \hat{\kappa}), \quad (7)$$

where $\Delta_0(i, x)$ is a scalar function of *i* and *x*, $\Delta_1(i, x)$ is a 2 × 1 vector of first derivatives, and $\Delta_2(i, x)$ is a 2 × 2 matrix of second derivatives. For any given *i* and *x*, these Δ terms are parameters that are easily computed by Monte Carlo simulation.⁷ Once we have these, simulated time*i* expected survivor rates out to *i* = 50 years can be recovered from

$$E[S(T + i, x)/S(T, x) | \kappa_1^i(T), \kappa_2^j(T)]$$

$$\approx \Phi(\Delta_0(i, x) + \Delta_1(i, x)'(\kappa^j - \hat{\kappa}) + \frac{1}{2}(\kappa^j - \hat{\kappa})'\Delta_2(i, x)(\kappa^j - \hat{\kappa})). \quad (8)$$

Each simulated $E[S(T + i, x)/S(T, x)|\kappa_1^i(T), \kappa_2^j(T)]$ can then be plugged into equation (6) to

⁵ For example, with M = 10,000 simulation paths in each stage for each of the mortality state variables, this would require 200 million simulation paths for the mortality state variables alone; combined with all the other calculations required, this implies a computational burden that is not for practical purposes feasible under real-time constraints.

⁶ The purpose of the probit transformation is purely to extend the domain of the function from [0 to 1] to the full real line and hence eliminate potential boundary problems.

⁷ For more on the implementation of this approach, see Cairns (2011) who carries out numerical tests on the accuracy of the approximation. He reports results indicating that the quadratic approximation is accurate even at long horizons. Approximate values can be either above or below true values. However, a very small bias can creep in if $\hat{\kappa}$ is not calibrated to the relevant time horizon, i.e., $\hat{\kappa} = E[\kappa(T)]$. If we focused on a single simulation time horizon, T, however, then the errors above can be reduced by calibrating $\hat{\kappa} = E[\kappa(T)]$ to that specific time horizon.

give us the corresponding simulated future annuity value we are seeking.

5. THE DISTRIBUTION OF FUTURE ANNUITY VALUES UNDER INTEREST-RATE AND LONGEVITY RISKS: SOME ILLUSTRATIVE RESULTS

We now provide an example based on a deferred annuity with a starting age of 65 and purchased at the age of 25, a current instantaneous spot interest rate equal to 4%, and 10,000 simulation trials. The annuity makes level payments of \$1 for each year the annuitant survives, and the annuity value is assumed to incorporate a loading factor of 10%.⁸ Results are also presented for two versions of the mortality model: a version that assumes that the parameters of the mortality model are estimated with certainty (the PC case), and a version that takes account of uncertainty in those parameters (the PU case).⁹

We examine three different cases. In the first case, we allow for longevity risk but not interestrate risk: We model future longevity improvements using simulations from our mortality model, but we take the future instantaneous spot interest rate to be equal to its current value of 4%. In the second case, we assume that there are no changes in future longevity, but we allow the instantaneous spot interest rate to be stochastic and use our interest-rate model to simulate its value at T = 40. In the third case, we allow both interest rates and longevity to evolve stochastically over the period to T = 40. In all cases, the future annuity values are obtained by taking the T = 40 present value of later cash flows discounted at the relevant spot interest rate, where these later rates are obtained from the calibrated CIR interest rate model. Our results are presented in Table 2 and Figures 2–4.

Table 2 shows that the current fair value of an annuity for a 65-year-old male is 13.050 if we take the parameters of the mortality model to be certain, and 13.141 if we allow for uncertainty in the estimates of those parameters. These values provide benchmarks against which we can assess the prospective annuity prices that our current 25-year-old might face when he reaches 65.

The first two columns in the table give the main features of the distribution of future annuity values 40 years hence in the presence of longevity risk but no interest-rate risk. They show that future annuity values have a mean of 16.208 if we take the mortality parameters as certain (the PC case) and a mean of 16.044 if we allow for them to be uncertain (the PU case). These are, respectively, 24.2% and 22.1% higher than the values of comparable annuities for 65-year-olds bought now. Clearly, future annuity values are expected to rise because the model projects further longevity improvements in the future. But what is less clear is the impact on future annuity values of longevity risk, that is, the uncertainty attached to future longevity projections. Columns 1 and 2 show that future annuity values have an 80% confidence interval equal to [15.356 16.958] for the PC case and a somewhat wider 80% confidence interval of $[14.449\ 17.422]$ in the PU case.¹⁰ (The corresponding standard deviations are 0.666 and 1.253.) These results indicate that longevity risk has a considerable impact on both the mean future annuity value and the dispersion of future annuity values, and that the degree of dispersion increases as we take account of parameter uncertainty.

Columns 3 and 4 give the comparable results for the case where we allow for interest-rate risk but not longevity risk. The mean future annuity values are now 13.135 (in the PC case) and 13.227 (in the PU case), which are much closer

⁸ This is approximately equal to the average loading factor found by Finkelstein and Poterba (2002) for U.K. insurers selling level annuities.

⁹ The parameter-certain case uses the estimated values of the μ and C parameters, whereas the parameter-uncertain case makes use of simulated values of these parameters drawn from the appropriate distributions. More details of the model and the simulation procedures (including the method used to allow for parameter uncertainty) are given in Cairns et al. (2006a).

The simulations reported in this paper did not allow for any difference between the real-world and risk-neutral probability and pricing measures. We did not allow for this difference for two reasons. First, we do not have hard empirical evidence on what the relevant market price of risk in the mortality model would be, and we need this to specify the risk-neutral probability measure; and, second, some illustrative results presented in Cairns et al. (2006a) suggest that the prices of the longevity bond examined in their paper are fairly insensitive to the specific market prices of risk they assumed.

¹⁰ The 80% confidence interval comes from comparing the difference between the 10th and 90th percentiles.

Table 2

Current Annuity Values and the Probability Density Function of Annuity Values in 40 Years' Time¹

Inputs							
Age at retirement Years to retirement Current year Current instantaneous spot interest rate Loading factor in annuity value Number of simulation trials			65 40 2007 0.04 0.10 10,000				
Results for the Current Annuity Value							
Current annuity value (PC) ² Current annuity value (PU) ³			13.050 13.141				
Results for Annuity Values at $T = 40$							
Parameters of Annuity Value Distribution	With Longevity Risk ⁴ but No Interest Rate Risk ⁵		With Interest Rate Risk ⁶ but No Longevity Risk ⁷		With Longevity Risk ⁴ and Interest Rate Risk ⁶		
Mean Std Dev Skewness Kurtosis 10th perc. 20th perc. 30th perc. 40th perc. 50th perc. 60th perc. 80th perc. 90th perc.	PC 16.208 0.666 -0.952 4.484 15.356 15.735 15.968 16.146 16.298 16.431 16.581 16.581 16.739 16.957	PU 16.044 1.253 -1.108 5.883 14.449 15.171 15.620 15.940 16.227 16.476 16.731 17.018 17.422	PC 13.135 1.234 -1.121 4.391 11.428 12.203 12.693 13.082 13.398 13.685 13.949 14.208 14.474	PU 13.227 1.245 -1.121 4.389 11.504 12.286 12.781 13.174 13.493 13.782 14.049 14.310 14.579	PC 16.321 1.744 -0.8278 3.746 13.952 14.938 15.654 16.158 16.591 16.990 17.392 17.824 18.323	PU 16.155 2.037 -0.571 3.294 13.362 14.482 15.243 15.871 16.384 16.879 17.367 17.903 18.580	

¹ In all cases, discounting is carried out at the spot interest rate for the relevant maturity as determined by the calibrated CIR interest-rate model.

² PC means that the simulations assume the parameters of the stochastic mortality model to be certain.

³ PU means that the simulations allow for uncertainty in the parameters of the Cairns-Blake-Dowd (CBD) stochastic mortality model using the method outlined in Cairns et al. (2006a).

⁴ Longevity risk is modeled using simulated values of $\kappa_1(40)$ and $\kappa_2(40)$ obtained using the CBD model and taking account of interim stochastic mortality improvement in the period since 2002. These parameter values are based on estimates of the mortality of English and Welsh males aged 65 over the period 1982–2002. See eqs. (1)–(2).

⁵ The instantaneous spot interest rate at T = 40 is assumed to be equal to 0.04.

⁶ The instantaneous spot interest rate and term structure of interest rates are assumed to be governed by a CIR process (given by eq. [6]) with parameters $\alpha = 0.20$, $\sigma = 0.10$, and $\bar{r} = 0.04$.

⁷ Mortality rates are assumed to be unchanged from their current levels.

to the current annuity values. Thus, allowing for interest-rate risk on its own has a much smaller impact on expected future annuity values—in fact, it leads the expected value to rise by only 0.7% in each case—than allowing for longevity risk on its own. This follows because of the mean reversion embodied in the CIR model (4). The 80% confidence intervals are now [11.428 14.474] and [11.504 14.579] for the PC and PU cases, respectively. (The corresponding standard deviations are about 1.2 in both cases.) So although interest-rate risk on its own has a negligible effect on the mean future annuity value, it has an impact on the dispersion of future annuity values broadly comparable to that of longevity risk (in the PU case, although not in the PC case).

Finally, columns 5 and 6 show the results when we allow for both longevity and interest-rate risk. The mean future annuity values now rise to 16.321 and 16.155: The impact of both longevity risk and interest-rate risk is to increase expected future annuity values by 25.1% in the PC case and 22.9% in the PU case, relative to current values. The 80% confidence intervals are now [13.952 18.323] and [13.362 18.580], respectively, for the PC and PU cases, and the corresponding standard deviations are 1.744 and 2.037.

Figure 2 shows the histogram of simulated future annuity values if we allow for longevity risk

Figure 2 Histogram of Simulated Future Annuity Values under Longevity Risk but No Interest-Rate Risk



Note: As per notes 2, 3, 4, and 5 to Table 2.

but not interest-rate risk, Figure 3 shows the same histogram if we allow for interest-rate risk but not longevity risk, and Figure 4 shows the histogram if we allow for both these risks simultaneously. Figure 4 indicates that the distribution of future annuity values has a fairly strong negative skew when we allow for both risks. Figure 2 reveals that the distribution has a very small negative skew if only longevity risk is considered. Figure 2 confirms that the cause of the negative skew in Figure 4 is predominantly caused by the negative skew in the distribution of future annuity values due to interest rate risk, and this is the direct counterpart of the strong positive skew in the distribution of interest rates shown in Figure 1.

6. CONCLUSIONS AND IMPLICATIONS FOR DC PENSION PLAN MEMBERS

This paper proposes a simple computationally efficient algorithm for estimating the distribution

of future annuity values in the presence of both longevity and interest-rate risk. The algorithm is based on a second-order Taylor series expansion of the probit transformation of the expected values of a survivor index in future years, conditional on an individual surviving to the previous year and conditional on the state parameters governing the stochastic mortality model. The algorithm allows us to avoid the computational problem of simulation-within-simulation and hence increase computational efficiency by a factor of 10,000 in the case where we wish to estimate the distribution of future annuity prices using 10,000 simulations.

Illustrative results suggest that the combined effect of longevity and interest-rate risks is to considerably widen the dispersion of future annuity values, in comparison with the cases in which each is treated separately. The mean future annuity value is also considerably higher than the current annuity value, but, as we have seen, this is principally due to projected future improve-



Figure 3 Histogram of Simulated Future Annuity Values under Interest-Rate Risk but No Longevity Risk

ments in longevity rather than to any effects of interest rate risk. Interest-rate risk is largely responsible for giving the distribution of future annuity values a strong negative skew.

It is helpful if we end by elaborating a little on the implications of these findings for those in the early stages of a DC pension plan, such as our illustrative 25-year-old male plan member. Let us suppose that the value of the accumulated pension fund is given by F. For an individual aged 65 and retiring now, we see from Table 2 that his annual retirement income would be F/13.141 if we use the parameter-uncertain annuity valuation. However, our current 25-year-old is likely to be much less fortunate in terms of his retirement income. Since he faces an expected future annuity value of 16.155 (if we use with the parameteruncertain valuation and allow for both longevity and interest-rate risk), his expected retirement income is only F/16.155, or 18.7% lower, other things being equal. This reduction in expected retirement income is due primarily to projected longevity improvements over the course of his working lifetime.¹¹

Worse still, his pension also becomes more risky. One cause of this increase in risk is the dispersion in the distribution of future annuity values. Of particular concern is the positive tail of this distribution. If we examine the quantiles of this tail, there is a 20% probability of a value in excess of 17.903, a 10% probability of a value in excess of 18.580, and so forth. Translated into their retirement-income equivalents, there is a 20% probability that our 25-year-old will have a retirement income that is at least 26.6% (i.e., 13.141/17.903 - 1) lower than that received by a male retiring now, and there is a 10% probability that he will receive an income that is at least

¹¹ Perhaps the most obvious response to this reduction would be for him to anticipate working longer: After all, if he is anticipating living longer, it might be reasonable for him to be prepared to work longer. The alternative to this, of course, is contribute around nearly 20% more to his pension plan than his 40-year-older compatriot did.





Note: As per notes 2, 3, and 4 to Table 2.

29.3% (i.e., 13.141/18.580 - 1) lower than his older counterpart, other things remaining equal, in particular the size of the retirement fund.¹² These amount to a gloomy prognosis for a riskaverse pension plan member. Still, the good news (such as it is) is that although the news is always bad, it is not always quite so bad: If we look at the other, more fortunate, tail of the distribution, there is also a 10% probability that he will get a retirement income that is almost as good as that received by the older pensioner now. He might just get lucky.

Finally, our analysis has identified an important need for pension plans and their members to hedge the interest-rate and longevity risks in their future annuity purchases. They could do this using annuity futures or annuity futures options as suggested in Blake et al. (2006). Unfortunately, these contracts do not yet exist, in either capital markets or over-the-counter form. But we believe this is only a matter of time. The framework outlined in Cairns et al. (2006b) shows how these contracts can be priced, while the computationally efficient algorithm outlined above would allow the contract pricing to be done in real time.

7. ACKNOWLEDGMENT

The authors thank Richard MacMinn and Tony Webb for helpful comments on an earlier draft. The usual caveat applies.

¹² Our results about the riskiness of DC pensions are, however, open to two offsetting sources of bias. On the one hand, as mentioned in note 4, the distribution of future instantaneous spot interest rates is likely to be underestimated by the choice of a CIR interest-rate process, and a more empirically plausible interest rate process would lead the distribution of future annuity prices to become even more dispersed. On the other hand, we have implicitly assumed that the pension fund value, *F*, is fixed. In practice, it is much more likely that *F* would be stochastic and positively correlated with long-term interest rates, particularly if the assets in the pension fund were dominated by bonds. In this case, higher interest rates would be likely to produce both a high value of *F* and higher annuity prices, and DC pension outcomes would be more stable and less dispersed than we have suggested. We thank Tony Webb for this latter point.

REFERENCES

- BENJAMIN, B., AND J. H. POLLARD. 1993. The Analysis of Mortality and Other Actuarial Statistics. 3rd ed. London: Institute of Actuaries.
- BLAKE, D., A. J. G. CAIRNS, AND K. DOWD. 2006. Living with Mortality: Longevity Bonds and Other Mortality-Linked Securities. *British Actuarial Journal* 12: 153–197.
- CAIRNS, A. J. G. 2004. *Interest Rate Models: An Introduction*. Princeton: Princeton University Press.
- CAIRNS, A. J. G. 2011. Modelling and Management of Longevity Risk: Approximations to Survivor Functions and Dynamic Hedging, Pensions Institute Discussion Paper PI-1105.
- CAIRNS, A. J. G., D. BLAKE, AND K. DOWD. 2006a. A Two-Factor Model for Stochastic Mortality with Parameter Uncertainty: Theory and Calibration. *Journal of Risk and Insurance* 73: 687–718.
- CAIRNS, A. J. G., D. BLAKE, AND K. DOWD. 2006b. Pricing Death: Frameworks for the Valuation and Securitization of Mortality Risk. ASTIN Bulletin 36: 79–120.

- CAIRNS, A. J. G., D. BLAKE, K. DOWD, G. D. COUGHLAN, D. EPSTEIN, A. ONG, AND I. BALEVICH. 2009. A Quantitative Comparison of Stochastic Mortality Models Using Data from England and Wales and the United States. North American Actuarial Journal 13(1): 1–35.
- COX, J. C., J. E. INGERSOLL, AND S. A. ROSS. 1985. A Theory of the Term Structure of Interest Rates. *Econometrica* 53: 385–407.
- FINKELSTEIN, A., AND J. POTERBA. 2002. Selection Effects in the United Kingdom Individual Annuities Market. *Economic Journal* 112: 28–50.
- PERKS, W. 1932. On Some Experiments in the Graduation of Mortality Statistics. *Journal of the Institute of Actuaries* 63: 12–57.

Discussions on this paper can be submitted until October 1, 2011. The authors reserve the right to reply to any discussion. Please see the Submission Guidelines for Authors on the inside back cover for instructions on the submission of discussions.