Immediate Annuity Pricing in the Presence of Unobserved Heterogeneity

Kim G. Balls*

Abstract

One of the acknowledged difficulties with pricing immediate annuities is that underwriting the annuitant's life is the exception rather than the rule. In the absence of underwriting, the price paid for a life-contingent annuity is the same for all sales at a given age. This exposes the market (insurance company and potential policyholder alike) to antiselection. The insurance company worries that only the healthiest people choose a life-contingent annuity and therefore adjust mortality accordingly. The potential policyholders worry that they are not being compensated for their relatively poor health and choose not to purchase what would otherwise be a very beneficial product.

This paper develops a model of underlying, unobserved health. Health is a state variable that follows a first-order Markov process. An individual reaches the state ''death'' either by accident from any health state or by progressively declining health state. Health state is one-dimensional, in the sense that health can either ''improve'' or ''deteriorate'' by moving farther from or close to the ''death'' state, respectively. The probability of death in a given year is a function of health state, not of age. Therefore, in this model a person is exactly as old as he or she feels.

I first demonstrate that a multistate, ageless Markov model can match the mortality patterns in the common annuity mortality tables. The model is extended to consider several types of mortality improvements: permanent through decreasing probability of deteriorating health, temporary through improved distribution of initial health state, and plateau through the effects of past health improvements.

I then construct an economic model of optimal policyholder behavior, assuming that the policyholder either knows his or her health state or has some limited information. The value of mortality risk transfer through purchasing a life-contingent annuity is estimated for each health state under various risk-aversion parameters. Given the economic model for optimal purchasing of annuities, the value of underwriting (limited information about policyholder health state) is demonstrated.

1. Introduction

One of the controlling realities of life insurance is that the policyholder generally has more accurate information about the state of his or her health than the insurance company. In the case of life insurance, the degree of asymmetry can be reduced through underwriting; almost all life insurance policies written are underwritten. Annuity policies are not generally underwritten. The typical exception is a structured settlement case involving an impaired life.

In the theoretical and academic community, one of the remaining puzzles of the insurance market is the thinness of the lifetime payout annuity market. One possible contributor is this asymmetry of information. If the policyholder knows information about his or her longevity that cannot underwritten and priced, the available annuities may be unattractive to those potential policyholders with the lowest

* Kim G. Balls, PhD, ASA, MAAA, is Vice President for North American Life Product Development at DFA Capital Management, 100 Manhattanville Rd., Purchase, NY 10577. kb@dfa.com.

survival expectations. Likewise, insurance companies pad annuity mortality tables expecting that only people with optimistic views of their life expectancy will purchase payout annuities.

This antiselection by policyholders, together with the insurance company response, could combine to eliminate the payout annuity market. The process operates as follows. Suppose population mortality suggests a population payout annuity rate that is too expensive to 40% of the population because of their private information about mortality. When this 40% is eliminated from the mortality pool, the life expectancy of potential policyholders improves, which increases the cost of the payout annuity to the buying public. This price increase, in turn, leads additional potential policyholders (always the least healthy or most likely to die) to avoid the market. The size of the market dwindles with each machination.

The available methods to reduce antiselection are to reduce the asymmetry of information (underwrite) or differentiate the market (offer a menu of products hoping to separate the heterogeneous population by product selection). If an insurance company tried to market to those with lower life expectancy by selecting those who initially refused an offer of an annuity, healthy (and wise) potential customers would refuse offers as well in expectation of a better offer.

This paper develops a model of annuity mortality that facilitates a more thorough evaluation of this issue. The annuitant has a health state. Aging occurs as the health state changes. Death occurs either as a terminal health state or as an accident. By knowing the current health state, the annuitant or the insurance company can improve its knowledge about the life expectancy of the annuitant.

The literature on the importance of adverse selection in the annuity market is in the same stage of development as the market itself. Poterba (2001) gives a good general overview of the issues related to adverse selection in the annuity market. Brunner and Pech (2000, 2002) develop a two-period general equilibrium model in which they can characterize the nature of the equilibria possible within a market with adverse selection. Finkelstein and Poterba (2004) analyze the policies of a large U.K. insurer and find evidence that different types of payout annuities have different mortality characteristics, which is evidence that knowledgeable insurers self-select the annuity that best matches their mortality risk. Finkelstein and McGarry (2003) find direct evidence that purchasers of long-term care (LTC) insurance (a health-contingent form of payout annuity) have private information about their probability of needing nursing home care. Mitchell and McCarthy (2002) use the difference between population mortality tables and the annuity mortality tables as evidence of self-selection in the annuity market.

The paper is organized as follows. First we estimate a Markov model that matches the most common annuitant mortality tables. We compare parameter estimates derived from various mortality tables. The striking result is that annuitant mortality can be modeled using very few parameters within a Markov model.

We then present an economic model of annuitant utility. The optimal consumption decision is derived for the case where no annuity market exists. Using the same utility framework, we derive the threshold price for an immediate annuity. Employing several competing utility models, we show the degree of market thinness created by the asymmetry of information.

2. MARKOV MODEL

Rather than focus on chronological age, the annuitant mortality model in this paper is concerned only with health status. As a gross approximation of actual human health, consider a model where an annuitant's mortality is described by a one-dimensional health status, S_t, where S_t is an integer from 0 to J. A person ''ages'' as health state decreases. $S_t = 0$ represents death. In any period the health status can either improve or deteriorate.

For the general population of annuitants aged t, let p_t be the $J + 1 \times 1$ vector of probabilities of annuitants with health status j, $j = 0, \ldots, J$. Let T represent the $J + 1$ by $J + 1$ transition matrix, with typical element $T_{ij} = \text{Prob}(S_{t+1} = j|S_t = i)$. Then $p_{t+1} = Tp_t$.

To restrict the number of parameters in the model, the transition matrix T is defined as follows:

- For $i > 1$, $T_{i,i-1} = p_d$, the probability of decreasing health state by 1
- For $J > i > 0$, $T_{i,i+1} = p_u$, the probability of improving health state by 1
- For $i > 1$, $T_{i,0} = p_a$, the probability of death by accident
- $T_{1,0} = p_d + p_a$
- $T_{i,i} = 1 - p_d - p_a - p_u$ (if $i < J$)
- $T_{i,j} = 0$, unless defined above.

Estimating a Markov mortality model to calibrate to any particular data set requires estimating the parameters p_u, p_d, and p_a, and the initial state conditions p_0. Define $q_{0|t}$ (probability of death in period t conditional on being alive in period 0) as $(p_{t0} - p_{t-1,0})$, where $p_{t,0}$ is the first element of the state probability vector p_t. For death data $\{d_t\}$, the likelihood of the data conditional on the model is

$$L(d|P) = \prod_t q_{0|t}^{d_t}, \text{ or } \ln(L) = \sum_t d_t \ln(q_{0|t}).$$

The statistical estimation exercise is to maximize $\ln(L)$ with respect to (p_d, p_a, p_u) and the initial state conditions p_0. The data used in this exercise are mortality tables, which are themselves aggregations of raw mortality data. Since no specific sample or sample size is available, statistical inference or testing is beyond the scope of this paper. However, a limited gauge of goodness of fit is possible. The nonparametric population estimates maximize $\ln(L)$ by choice of $\ln(q_{0|t})$, which delivers the population mean $q_{0|t} = d_t/n_0$, where n_0 is the lives at risk at the outset. For an assumed initial sample size n_0 it is possible to construct a synthetic data set $\{d_t\}$, which is consistent with the mortality table estimates $q_{0|t}$. The quality of the parameterized model is calculated based on the difference between the nonparametric log-likelihood and the parametric log-likelihood.

As a simple example, suppose a population is described by $n_0 = 100$ and $d_t = 10$ for $t = 1$ to 10. Then the mortality rates are 0.1, and the log-likelihood is 100*ln(0.1) = −230.259. If the parametric model has 10 states, then the model produces exactly the population estimates with $p_0 = 0.1$ for each state, $p_u = p_a = 0$, and $p_d = 1$. If the parametric model has only two positive health states, then

$$\begin{aligned}
p_{t2} &= p_{t-1,2} * (1 - p_d - p_a) \\
p_{t1} &= p_{t-1,2} * p_d + p_{t-1,1} * (1 - p_d - p_a) \text{ and} \\
p_{t0} &= p_{t-1,2} * p_a + p_{t-1,1} * (p_d + p_a) + p_{t-1,0}.
\end{aligned}$$

The maximum likelihood solution is $p_d = 0.327281$, $p_a = 0.096599$, and $p_{0,2} = 1$ (all start in the healthy state), with $\ln(L) = -248.523$. Twice the difference in log-likelihood functions (36.53) follows a X^2 distribution with seven degrees of freedom. The null hypothesis that the model fits the data is strongly rejected. The largest sample size for which the goodness of fit statistic is at the 5th percentile of its distribution is 38.

A similar exercise of maximum sample size that does not reject the model is possible for any mortality table and model. First, calculate the nonparametric log-likelihood of the underlying table conditional on one observation, $L_{NP} = \Sigma_t q_{0|t} \ln(q_{0|t})$. Then calculate the model log-likelihood as $L_M = \Sigma_t q_{0|t} \ln(q_{Mt})$, where q_{Mt} are the death probabilities in the model. The goodness-of-fit statistic for a population of size N would be $2N * (L_{NP} - L_M)$. Setting this equal to the threshold level of the χ^2 statistic (e.g., $\chi^2(0.05, 7) = 14.07$ in the previous case), delivers a threshold sample size of $N = \chi^2(\alpha, d.f.)/[2 * (L_{NP} - L_M)]$ $(= 14.07/[2 * (-2.30 - -2.49)] = 38.5$ in the previous example).

The above model is fit to U.S. Census population mortality tables for the years 1900–1990. Since the focus of this research is retirement, the model was calibrated to the ages 50–90. To obtain snapshots of specific birth cohorts, several mortality tables were interpolated to obtain cohort mortality tables. For example, the 1850 birth cohort uses the age-50 mortality from the 1900 census, the age-60 mortality from the 1910 census, and so forth. All the analysis in this paper assumes $J = 10$, which provides significant flexibility in mortality without becoming computationally cumbersome.

Table 1 shows maximum likelihood point estimates for each of the raw mortality tables and for birth cohort tables constructed for decennial birth cohorts from 1850 to 1900. In all cases p_u and p_a are estimated at zero and not reported. Health state initial probabilities are grouped, although health

Table 1
Mortality Estimates Based on Common Mortality Tables

Data Set (Mortality Table)	p_d	State 10	States 4–6	States 1–3	Model Fit
1996 US Annuity 2000 Basic, male	26.9%	81.7%	13.0%	5.3%	6,250
1996 US Annuity 2000 Basic, female	25.0	89.2	8.8	2.0	2,611
RP-2000 male combined healthy	29.0	85.3	12.0	2.7	1,704
RP-2000 male healthy annuitant	28.9	82.5	12.3	5.2	1,811
RP-2000 female combined healthy	26.5	85.7	12.0	2.3	2,963
RP-2000 female healthy annuitant	26.5	84.7	12.1	3.2	3,134
US (SSA AS 107) 1900, age nearest, male	35.2	54.3	30.0	15.7	5,567
US (SSA AS 107) 1910, age nearest, male	35.1	54.3	30.8	14.9	7,750
US (SSA AS 107) 1920, age nearest, male	34.7	58.9	27.9	13.2	6,233
US (SSA AS 107) 1930, age nearest, male	34.4	55.3	29.6	15.1	8,670
US (SSA AS 107) 1940, age nearest, male	34.4	57.2	28.3	14.5	7,084
US (SSA AS 107) 1950, age nearest, male	33.0	57.9	28.1	14.0	7,323
US (SSA AS 107) 1960, age nearest, male	32.7	57.7	28.7	13.6	7,027
US (SSA AS 107) 1970, age nearest, male	32.4	57.7	29.7	12.6	7,671
US (SSA AS 107) 1980, age nearest, male	31.4	62.2	27.4	10.4	7,219
US (SSA AS 107) 1990, age nearest, male	30.9	67.7	22.9	9.4	5,815
US (SSA AS 107) 1900, age nearest, female	34.4	57.6	26.8	15.6	5,607
US (SSA AS 107) 1910, age nearest, female	34.2	59.2	27.7	13.1	6,179
US (SSA AS 107) 1920, age nearest, female	34.0	61.1	26.0	12.9	6,002
US (SSA AS 107) 1930, age nearest, female	33.2	62.0	24.9	13.1	7,084
US (SSA AS 107) 1940, age nearest, female	33.0	67.6	21.5	10.9	4,675
US (SSA AS 107) 1950, age nearest, female	31.2	73.0	18.8	8.2	4,416
US (SSA AS 107) 1960, age nearest, female	30.5	77.2	16.5	6.3	2,647
US (SSA AS 107) 1970, age nearest, female	29.2	78.3	14.8	6.9	3,269
US (SSA AS 107) 1980, age nearest, female	27.8	79.1	15.2	5.7	3,055
US (SSA AS 107) 1990, age nearest, female	27.2	78.8	16.0	5.2	3,349
Male cohort 1850	34.5	54.6	29.1	16.3	8,491
Male cohort 1860	34.2	57.0	27.3	15.7	12,330
Male cohort 1870	33.1	52.4	33.1	14.5	12,184
Male cohort 1880	32.7	54.6	29.3	16.1	9,528
Male cohort 1890	32.1	54.1	31.6	14.3	14,309
Male cohort 1900	31.3	53.5	32.5	14.0	10,851
Female cohort 1850	33.4	56.9	28.6	14.5	8,248
Female cohort 1860	32.6	58.2	28.1	13.7	12,481
Female cohort 1870	31.2	58.6	27.1	14.3	7,275
Female cohort 1880	30.3	65.3	21.7	13.0	7,712
Female cohort 1890	28.8	68.7	20.2	11.1	8,758
Female cohort 1900	27.7	72.1	19.3	8.6	5,381

states 7–9 are estimated to be zero for all mortality tables. The model fit statistic represents the largest hypothetical sample size that would result in a goodness-of-fit statistic that does not reject the null hypothesis of model fit at the 5% significance level. (Higher numbers represent a better model fit.)

Notice that the trend in mortality is markedly different between males and females when estimated using this model. Based on 1900 census data, we see little difference between male and female mortality. By 1990 female mortality improvement is dramatic both with respect to the percentage of the age-50 population in the highest health state as well as with respect to the probability of declining health in a year. The common annuity tables show even further pads and improvements, consistent with the findings of Mitchell and McCarthy (2002).

Table 2 and Figures 1 and 2 show the health dynamics based on estimates using the Annuity 2000 Basic Male table. Notice that 82% of the male population is estimated to be in the top health state at age 50. By age 65, there is considerable heterogeneity of health status. The population of most unhealthy lives (state = 1) never gets large, but as a percentage of the remaining lives it becomes important starting after age 65. Recall that the population mortality rate in this model is the percentage of lives with health status 1 times the probability of declining health.

The notion that mortality is determined by health status, and not by age, is at a certain level a matter of symantics, since age can be described as one dimension of health status. Certainly the above model

Table 2
ML Estimates of Health Status as Percentage of Surviving Population A2000B-Male

Age	10	9	8	7	6	5	4	3	2	1
50	81.6%	0.0%	0.0%	0.0%	4.8%	8.2%	0.0%	2.3%	1.8%	1.2%
51	59.9	22.0	0.0	0.0	3.5	7.3	2.2	1.7	1.9	1.4
52	44.0	32.3	5.9	0.0	2.6	6.3	3.6	1.9	1.9	1.5
53	32.3	35.6	13.1	1.6	1.9	5.3	4.3	2.3	1.9	1.6
54	23.7	34.9	19.2	4.7	1.8	4.4	4.6	2.9	2.0	1.7
55	17.4	32.0	23.5	8.7	2.6	3.7	4.6	3.4	2.3	1.8
56	12.8	28.2	25.9	12.7	4.3	3.5	4.4	3.7	2.6	1.9
57	9.4	24.2	26.7	16.3	6.6	3.7	4.2	3.9	2.9	2.1
58	6.9	20.3	26.2	19.2	9.2	4.5	4.1	4.0	3.2	2.3
59	5.1	16.8	24.8	21.2	12.0	5.8	4.2	4.0	3.4	2.6
60	3.8	13.8	22.8	22.3	14.6	7.5	4.7	4.1	3.6	2.8
61	2.8	11.2	20.5	22.6	16.8	9.5	5.5	4.3	3.8	3.1
62	2.0	9.0	18.2	22.3	18.5	11.5	6.6	4.6	4.0	3.3
63	1.5	7.2	15.8	21.3	19.7	13.5	8.0	5.2	4.2	3.5
64	1.1	5.7	13.6	20.1	20.3	15.3	9.6	6.0	4.5	3.7
65	0.8	4.5	11.6	18.5	20.5	16.8	11.2	7.0	5.0	4.0
66	0.6	3.6	9.8	16.8	20.2	18.0	12.9	8.3	5.6	4.3
67	0.4	2.8	8.2	15.1	19.5	18.8	14.4	9.6	6.4	4.7
68	0.3	2.2	6.9	13.4	18.5	19.2	15.8	11.0	7.3	5.2
69	0.2	1.7	5.7	11.8	17.4	19.3	17.0	12.5	8.5	5.9
70	0.2	1.3	4.7	10.4	16.2	19.1	17.9	13.9	9.7	6.7
71	0.1	1.1	3.9	9.0	14.9	18.7	18.5	15.3	11.0	7.6
72	0.1	0.8	3.2	7.8	13.6	18.0	19.0	16.5	12.4	8.7
73	0.1	0.6	2.6	6.7	12.3	17.2	19.1	17.6	13.8	9.9
74	0.1	0.5	2.1	5.7	11.1	16.3	19.1	18.5	15.2	11.3
75	0.0	0.4	1.7	4.9	10.0	15.4	19.0	19.2	16.6	12.7
76	0.0	0.3	1.4	4.2	8.9	14.4	18.6	19.8	17.9	14.3
77	0.0	0.2	1.2	3.6	8.0	13.5	18.2	20.3	19.2	15.9
78	0.0	0.2	1.0	3.1	7.1	12.5	17.7	20.6	20.3	17.5
79	0.0	0.2	0.8	2.6	6.3	11.6	17.1	20.8	21.4	19.2
80	0.0	0.1	0.7	2.3	5.6	10.7	16.5	20.9	22.4	20.8
81	0.0	0.1	0.5	1.9	5.0	9.9	15.8	20.9	23.3	22.5
82	0.0	0.1	0.5	1.7	4.4	9.1	15.1	20.8	24.1	24.2
83	0.0	0.1	0.4	1.4	4.0	8.4	14.5	20.6	24.8	25.9
84	0.0	0.1	0.3	1.2	3.5	7.8	13.8	20.4	25.5	27.5
85	0.0	0.0	0.3	1.1	3.1	7.2	13.2	20.1	26.0	29.1
86	0.0	0.0	0.2	0.9	2.8	6.6	12.5	19.8	26.5	30.7
87	0.0	0.0	0.2	0.8	2.5	6.1	11.9	19.4	26.9	32.2
88	0.0	0.0	0.2	0.7	2.2	5.6	11.3	19.0	27.2	33.7
89	0.0	0.0	0.1	0.6	2.0	5.2	10.8	18.7	27.5	35.1
90	0.0	0.0	0.1	0.5	1.8	4.8	10.2	18.3	27.8	36.5

can be generalized to allow for state transitions of more than one state per period, the only problem being estimation. Survey data that ask respondents to rate their health condition may not correctly capture the health status data required here because the meaning of good health in the eyes of a respondent is relative to his or her cohort. (The meaning of a response of good health is different between a respondent of age 55 and a respondent of age 85.) Certainly objective underwriting techniques could eliminate this reporting bias.

3. ECONOMIC MODEL

The policyholder is assumed to make asset and consumption decisions with the intent of maximizing their expected lifetime utility. Utility in each period is possibly influenced by the health state of the policyholder. Expected lifetime utility is the net present value of the utility derived from consumption in each period of the policyholder's life. Assets not consumed are invested, earning a fixed rate of return.

This model assumes a single policyholder. For two policyholders making joint asset and consumption decisions based on their combined mortality, the health state model could be assumed to apply to the

Figure 1
Health Status as Percentage of Surviving Population

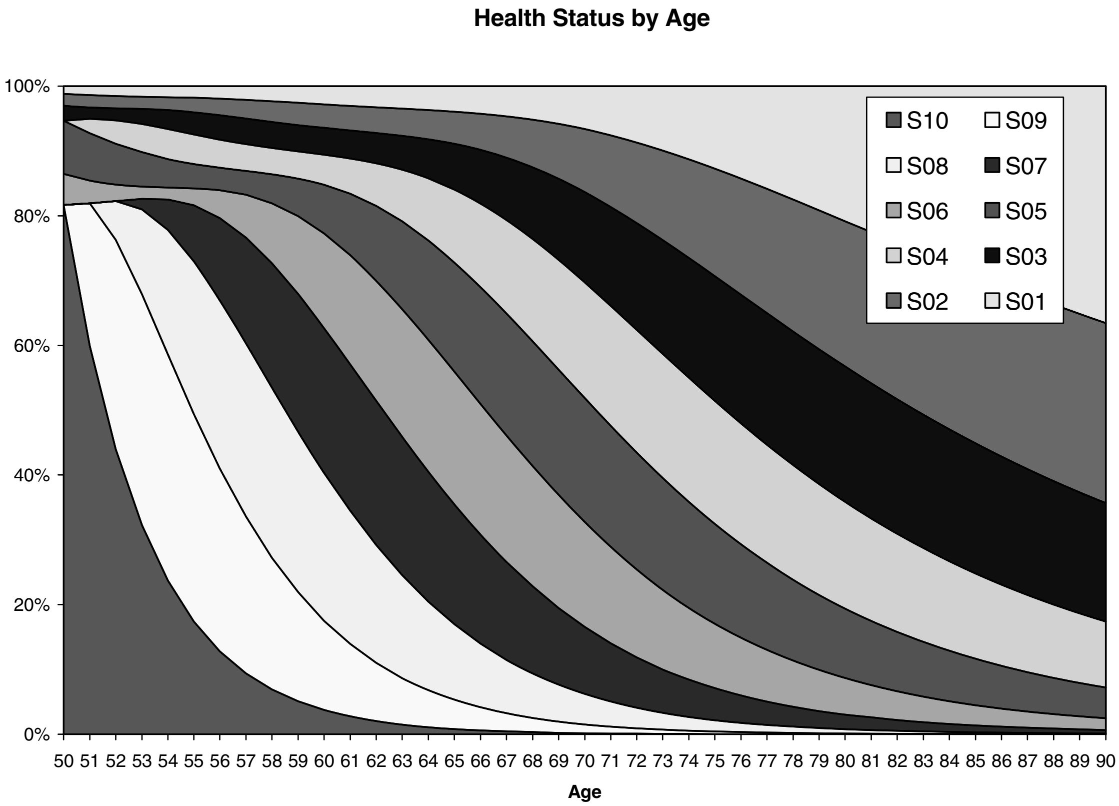

pair. Likewise, this model assumes no bequests, or no purpose for assets other than consumption. One could adjust the following model to handle bequests by either removing assets earmarked for bequests from the initial assets, or by assigning positive utility (and comparable utility function) to assets in the death state.

Mathematically, assume that per-period utility, conditional on health state j, is defined as

$$U_j(c_t) = K_c + (\alpha + 1)^{-1}(\theta_j c_t)^{(\alpha+1)}, \tag{3.1}$$

where c_t is consumption in period t, θ_t is a health-state consumption modifier, $\alpha \leq 0$ is the risk-aversion parameter, and K_c is a utility constant whose sole purpose is to make utility positive. Marginal utility with respect to consumption is $U_j'(c_t) = \theta_j c_t)^{\alpha}$.

The choice of utility function is based on mathematical simplicity and the fact that this particular class of utility functions is characterized by constant relative risk aversion. Absolute risk aversion is measured by $U'(c)/U''(c)$. For this class of utility functions, absolute risk aversion is linear in consumption. Relative risk aversion (risk aversion divided by consumption or wealth) is constant. For a discussion of alternative utility functions used in retirement wealth analysis, see Gerber and Shiu (2000). Brown (2003) uses linearly separable log utility, which is a special case of the power utility model used here for $\alpha = -1$.

The discounted net present expected value of total lifetime utility is a function of the health state and the current level of assets. Mathematically, this is the value of being in health state j with assets a_t:

$$V_j(a_t) = \max\{c_s, s \geq t\} \textstyle\sum_{s \geq t} \beta^{s-t} \mathrm{E}[U_j(c_t)], \tag{3.2}$$

subject to the budget constraints that $a_s \geq 0$, and

Figure 2
Health Status with Death as the Omitted State

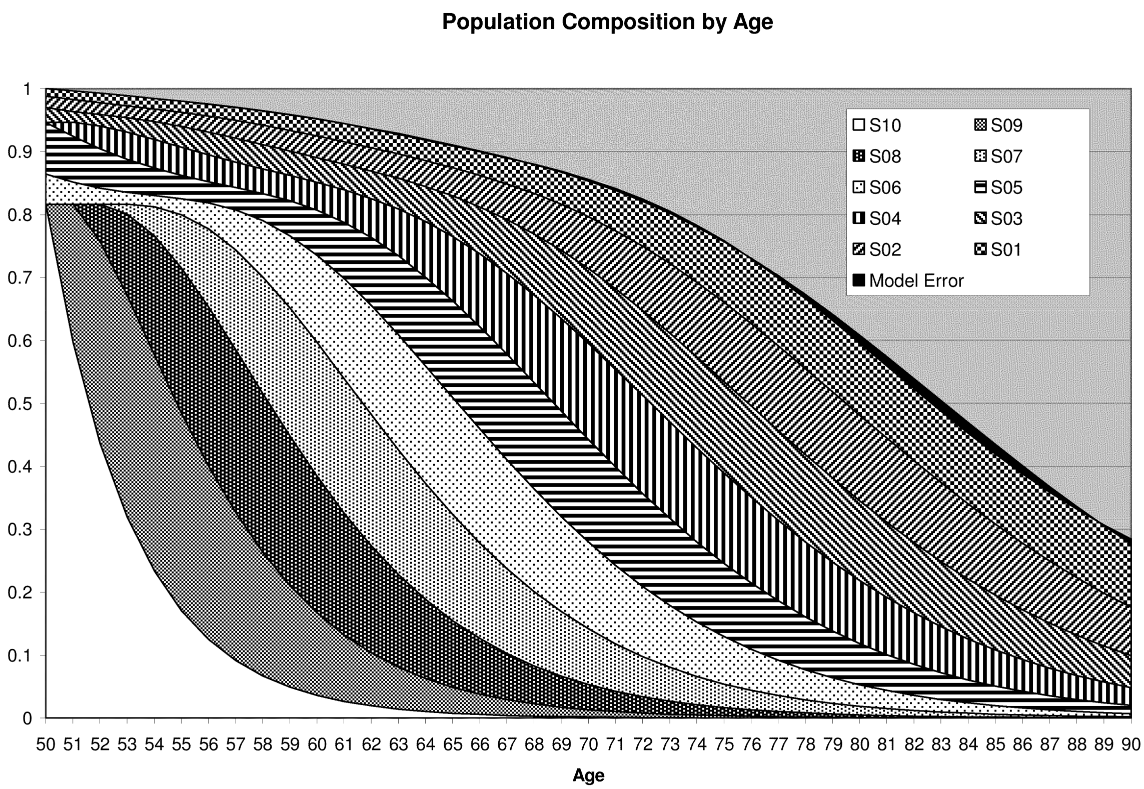

$$a_{t+1} = (a_t - c_t) * (1 + r), \tag{3.3}$$

where r is the investment rate of return. The expectation (3.2) is with respect to the health state j. Equation (3.2) can be written recursively, as a maximization with respect to only current consumption. This is the set of Bellman equations.

$$V_j(a_t) = \max\{c_t\}[U_j(c_t) + \beta\{p_s V_j(a_{t+1}) + p_u V_{j+1}(a_{t+1}) + p_d V_{j-1}(a_{t+1})\}], \tag{3.4}$$

where p_s is the probability of staying in the current health state, p_u is the probability of improving health state, and p_d is the probability of deteriorating health state.

Recall from equation (3.3) that a_{t+1} is a function of c_t, which means that the value equation can be solved by differentiating the right-hand side of equation (3.4) with respect to c_t. The difficulty is that although we have specified the functional form for the utility function $U_j(c)$, we have not specified $V_j(a)$. However, for the class of utility functions assumed, the functional form for $V_j(a)$ must be derived from $U_j(c)$. The appendix verifies that the consistent functional form is

$$V_j(a) = K_j + \delta_j U_j(a). \tag{3.5}$$

Furthermore, the appendix verifies that the optimal consumption is a set percentage of remaining assets, with the percentage varying by health state:

$$c_t^* = (1 - \gamma_j)a_t, \text{ and } a_{t+1}^* = \gamma_j a_t(1 + r). \tag{3.6}$$

The health-state constants δ_j and γ_j must match the utility parameters. Substituting equations (3.5) and (3.6) into equation (3.4) and differentiating the right-hand side of equation (3.4) with respect to c_t delivers the following level and first-order conditions, (3.7) and (3.8), respectively:

$$K_j + \theta_j^{\alpha+1}\delta_j(\alpha + 1)^{-1}a^{\alpha+1} = K_c + (\alpha + 1)^{-1}\theta_j^{\alpha+1} \times [((1 - \gamma_j)a)^{\alpha+1} + \beta\{K_{jp} + [(1 + r)\gamma_j a]^{\alpha+1} \times (p_s\delta_j + p_u(\theta_{j+1}/\theta_j)^{\alpha+1}\delta_{j+1} + p_d(\theta_{j-1}/\theta_j)^{\alpha+1}\delta_{j-1})\}, \text{ or} \tag{3.7a}$$

$$(K_c + \beta K_{jp} - K_j)(\alpha + 1)a^{-(\alpha+1)} = \delta_j - [(1 - \gamma_j)^{\alpha+1} + \beta[(1 + r)\gamma_j]^{\alpha+1} \times (p_s\delta_j + p_u(\theta_{j+1}/\theta_j)^{\alpha+1}\delta_{j+1} + p_d(\theta_{j-1}/\theta_j)^{\alpha+1}\delta_{j-1})] \tag{3.7b}$$

$$0 = (1 - \gamma_j)^{\alpha} - \beta[(1 + r)\gamma_j]^{\alpha} \times (p_s\delta_j + p_u(\theta_{j+1}/\theta_j)^{\alpha+1}\delta_{j+1} + p_d(\theta_{j-1}/\theta_j)^{\alpha+1}\delta_{j-1}), \text{ where}$$

$$K_{jp} = p_sK_j + p_uK_{j+1} + p_dK_{j-1}. \tag{3.8}$$

Setting first K_c and K_j's to zero, conditions (3.7) and (3.8) can solve for the optimal asset allocations, γ_j, and the value function multipliers, δ_j. Setting assets to zero, condition (3.7) can separately solve for utility and value function constants K_c and K_j, respectively.

Table 3 shows the solutions to this dynamic programming problem for varying values of the fundamental parameters of risk aversion (α). For all three examples, the utility function discount rate β is assumed to be 98%, and θ_j is assumed to be 1 for all j. Notice that in this self-funding mode, the optimal behavior is to consume a very small fraction of assets in the healthy states.

Note that the policyholder in this model cannot experience financial ruin. Unlike the self-funded retirement model of Milevsky and Robinson (2000), consumption is reoptimized each period as a function of current wealth. For each health state, the policyholder consumes a constant proportion of his

Table 3
Utility Model Parameters by State—Constant Utility

State	δ	γ	$\ddot{a}$	π	$1 - \gamma$	$\pi/\ddot{a}$
$\alpha = -0.75$, $\beta = 98.00\%$, $r = 5.00\%$						
10	25.95	95.96%	5.82%	5.54%	4.04%	95.21%
9	24.12	95.66	6.07	5.65	4.34	93.13
8	22.16	95.27	6.39	5.81	4.73	90.86
7	20.04	94.78	6.83	6.03	5.22	88.33
6	17.77	94.11	7.43	6.35	5.89	85.40
5	15.32	93.17	8.30	6.79	6.83	81.87
4	12.69	91.76	9.63	7.45	8.24	77.40
3	9.86	89.40	11.89	8.47	10.60	71.31
2	6.81	84.69	16.47	10.22	15.31	62.08
1	3.53	70.66	30.35	13.79	29.34	45.44
$\alpha = -1.50$, $\beta = 98.00\%$, $r = 5.00\%$						
10	111.60	95.55%	5.82%	5.41%	4.45%	92.96%
9	102.73	95.30	6.07	5.52	4.70	90.88
8	93.14	94.98	6.39	5.66	5.02	88.53
7	82.81	94.57	6.83	5.86	5.43	85.80
6	71.76	94.03	7.43	6.13	5.97	82.55
5	60.03	93.28	8.30	6.52	6.72	78.55
4	47.74	92.17	9.63	7.07	7.83	73.43
3	35.07	90.39	11.89	7.90	9.61	66.49
2	22.36	87.04	16.47	9.27	12.96	56.29
1	10.24	78.25	30.35	11.88	21.75	39.12
$\alpha = -2.25$, $\beta = 98.00\%$, $r = 5.00\%$						
10	1,022.12	95.31%	5.82%	5.29%	4.69%	90.99%
9	929.25	95.10	6.07	5.39	4.90	88.80
8	829.02	94.85	6.39	5.52	5.15	86.26
7	721.96	94.52	6.83	5.69	5.48	83.25
6	609.11	94.09	7.43	5.92	5.91	79.61
5	492.22	93.51	8.30	6.23	6.49	75.11
4	374.01	92.67	9.63	6.68	7.33	69.35
3	258.48	91.36	11.89	7.33	8.64	61.67
2	151.45	89.05	16.47	8.36	10.95	50.77
1	61.20	83.64	30.35	10.20	16.36	33.61

or her wealth in each time period. Although continuing good health may result in continually decreasing consumption, assets are never exhausted.

The above analysis assumes that the policyholder cannot purchase payout annuities and must bear the longevity risk associated with retirement. Suppose instead that the policyholder can purchase an annuity that pays amounts that result in level utility for all remaining periods of the policyholder's life. The value of wealth is simply the utility value of the annuity purchased with the assets:

$$W_j(c) = U_j(c) + \beta\{p_s W_j(c) + p_u W_{j+1}(c) + p_d W_{j-1}(c)\}, \quad (3.9)$$

where $c = \pi a$ is the payout annuity. The purchase price of the annuity is calculated in the next section. In the estimated case where $p_u = 0$, $p_s + p_d = 1$, and $\beta = 1$, the value of an annuity is linear in life expectancy. In general, one can show that W_j is of the form

$$W_j(c) = \varepsilon_j U_j(c). \quad (3.10)$$

(From equation (3.9), each U_j can be expressed as a linear combination of W_j, W_{j+1}, and W_{j-1}. Inverting the set of linear equations $AW = U$, one obtains $W = A^{-1}U$.)

Knowing both V_j and W_j, one can solve for the value of π_j, which makes the policyholder indifferent between purchasing an annuity and self-insuring retirement: $\pi_j = (\delta_j/\varepsilon_j)^{1/(\alpha+1)}$. Table 3 also shows this as a percentage of the naive price of the annuity assuming the same interest rates and mortality. Not surprisingly, the higher the risk-aversion parameter, the greater the policyholder's premium for retirement risk protection.

Notice as well that the least healthy seem also to be the most willing to purchase annuities at "actuarially unfair" prices. This is precisely because they face the more immediate mortality risk. Although the healthy may face more aggregate uncertainty as to the age of death, that uncertainty is discounted in the future. A person with a health status of 10% and 32% probability of declining health in any one period is certain to survive for 10 years and has less than an 8% probability of death in 20 years.

The discussion to this point has tacitly assumed constant θ, or the state-dependent value of consumption. Suppose instead that the policyholder needs more consumption in lower health states to obtain the same utility. Hypothetically, $\theta_j = 1$ for $j > 2$, $\theta_2 = 0.8$, and $\theta_1 = 0.6$. In this case the optimal self-funding parameters change per equations (3.7) and (3.8). In addition, the constant annuity is no longer optimal. One can show that the optimal annuity has an extra payment of 25% for $j = 2$, and 67% for $j = 1$. Table 4 shows the equivalent parameters for $a = -1.5$ for this case. In this case, the annuity value $\ddot{a}$ is for the base payment without the extra payments for states 2 and 1.

This model can be extended to the case where there are preannuitized assets, such as government pensions and employer-based defined benefit pensions. In this model it is optimal for the policyholder to annuitize either all or none of his or her assets. As such, preannuitized assets become part of the total assets available to the policyholder, who makes consumption decisions based on the entire asset portfolio. The model fails only when optimal consumption cannot be satisfied by liquid (non-preannuitized) assets and market imperfections prevent borrowing against future annuity payments.

4. Competitive Model

Given a vector $\{\pi_i\}$ that is monotone in health state, the optimal aggregate policyholder response in this economic environment is that the annuity is purchased by that (healthy) group for whom π_i is less than the annuity price offered by the insurance company. In this section we explore an equilibrium outcome where the price offered by the insurance company attracts the same population used in setting that price. (The price correctly accounts for antiselection.)

Consider, for example, the case of the age-65 male. Mortality is given by parameter estimates from the Annuity 2000 Basic mortality table. Assume that the base policyholder utility parameters are $\alpha = -1.5$, $\beta = 0.98$, and $r = 5\%$. A summary of all relevant parameters and coefficients is given in Table 5. The rows $\ddot{a}|r$ and $\pi|\alpha$ come from an extension of Table 3 and represent derived parameters

Table 4
Utility Model Parameters by State—State-Dependent Utility

State	δ	γ	$\ddot{a}$	π	$1-\gamma$	$\pi/\ddot{a}$
$\alpha = -0.75$, $\beta = 98.00\%$, $r = 5.00\%$						
10	12.48	96.32%	5.60%	4.15%	3.68%	74.10%
9	11.64	95.96	5.78	4.04	4.04	69.78
8	10.74	95.51	6.02	3.92	4.49	65.02
7	9.78	94.91	6.34	3.78	5.09	59.68
6	8.75	94.10	6.75	3.62	5.90	53.55
5	7.64	92.93	7.32	3.40	7.07	46.37
4	6.42	91.11	8.14	3.07	8.89	37.78
3	5.09	87.90	9.37	2.57	12.10	27.38
2	3.79	82.11	11.43	5.26	17.89	46.03
1	2.21	63.69	18.21	15.26	36.31	83.78
$\alpha = -1.50$, $\beta = 98.00\%$, $r = 5.00\%$						
10	114.62	95.63%	5.60%	5.82%	4.37%	104.09%
9	106.02	95.39	5.78	6.00	4.61	103.71
8	96.69	95.10	6.02	6.23	4.90	103.46
7	86.63	94.73	6.34	6.55	5.27	103.41
6	75.83	94.24	6.75	7.00	5.76	103.66
5	64.34	93.58	7.32	7.65	6.42	104.41
4	52.23	92.62	8.14	8.63	7.38	106.05
3	39.67	91.14	9.37	10.27	8.86	109.55
2	24.08	87.66	11.43	10.77	12.34	94.21
1	10.24	78.25	18.21	11.88	21.75	65.21
$\alpha = -2.25$, $\beta = 98.00\%$, $r = 5.00\%$						
10	1,096.64	95.45%	5.60%	5.27%	4.55%	94.10%
9	1,010.31	95.28	5.78	5.35	4.72	92.41
8	916.46	95.07	6.02	5.45	4.93	90.48
7	815.29	94.81	6.34	5.59	5.19	88.26
6	707.40	94.47	6.75	5.78	5.53	85.64
5	593.90	94.03	7.32	6.04	5.97	82.51
4	476.66	93.42	8.14	6.40	6.58	78.65
3	358.53	92.53	9.37	6.91	7.47	73.77
2	184.42	89.97	11.43	8.05	10.03	70.38
1	61.20	83.64	18.21	10.20	16.36	56.02

that are state-dependent. The rows $\ddot{a}_c|r$ represent the competitive annuity price assuming that the purchasing population is of that state or higher. This annuity price is a weighted average of the annuity prices for the states greater than or equal to the marginal state, with the weights given by the conditional state probabilities.

Assume that the insurance company has no information about policyholder health and that the policyholder has perfect information of health state. Further, assume that the insurance company can

Table 5
State Parameters for Age-65 Male Example Constant Utility

Parameter	State									
	10	9	8	7	6	5	4	3	2	1
Probability	0.82%	4.52%	11.63%	18.51%	20.46%	16.85%	11.24%	7.05%	4.96%	3.96%
$\ddot{a}\|r = 5\%$	5.82	6.07	6.39	6.83	7.43	8.30	9.63	11.89	16.47	30.35
$\ddot{a}\|r = 4\%$	5.13	5.39	5.74	6.19	6.81	7.69	9.03	11.30	15.87	29.68
$\ddot{a}_c\|r = 5\%$	5.82	6.03	6.27	6.55	6.85	7.14	7.39	7.62	7.83	8.07
$\ddot{a}_c\|r = 4\%$	5.13	5.35	5.61	5.90	6.20	6.49	6.74	6.96	7.17	7.39
$\pi\|\alpha = -2.25$	5.29	5.39	5.52	5.69	5.92	6.23	6.68	7.33	8.36	10.20
$\pi\|\alpha = -1.50$	5.41	5.52	5.66	5.86	6.13	6.52	7.07	7.90	9.27	11.88
$\pi\|\alpha = -0.75$	5.65	5.76	5.92	6.15	6.49	6.98	7.71	8.86	10.90	15.26

competitively price an annuity using the same interest rate facing the policyholder. Using the above population initial conditions, the insurance company can derive an annuity mortality table, which not coincidentally matches the Annuity 2000 Basic mortality table. Using an interest rate of 5%, the payout annuity factor is 8.07%. (This is the last entry in the $\ddot{a}_c|r = 5\%$ row of Table 3).

This payout factor is sufficiently low to make purchasing an annuity unprofitable for anyone with health states 1 or 2. Eliminating this portion of the population delivers a payout annuity factor of 7.62%. This in turn eliminates health state 3 from the buying population. The annuity factor for mortality conditional on initial health state > 3 at age 65 is 7.39%, which is optimal to purchase for any policyholder with health state > 3.

In this example the effect of self-selection is to decrease the payout factor from 8.07% to 7.39% and shrink the market by making the annuity unattractive to the least healthy 16% of the age-65 population. However, the market still exists for the remaining healthy population. In general, the purchasing states are those where $\pi|\alpha \leq \ddot{a}_c|r$.

Suppose instead that the insurance company bases its prices on an interest rate which is lower than that used by the policyholder. Even if the interest rate earned by the insurance company was 4% (a full 1% lower), the annuity reaches equilibrium for health state > 5, with annuity payout factor of 6.20% (down from the total population calculation of 7.39%), with the payout reduction eliminating 44% of the market).

This result is similar to other findings that are advantages to delaying annuitization, particularly when market returns available to the policyholder are superior to those available in the form of an annuity (see, e.g., Brown and Poterba 2000; Mitchell et al. 1999; and Milevsky 1998, 2000). However, the effect here is heterogeneous, depending also on the expected longevity of the policyholder.

Alternatively, suppose that the risk-aversion parameter is −0.75. To give some context, with this risk-aversion parameter the policyholder in health state 1 with no access to the annuity market optimally consumes over 29% of his or her assets in any period even though the probability of death is only 32.5%. Brown (2001) uses answers to simple risk-aversion questions to estimate that only about 10% of the population has risk-aversion parameters greater than −1. Using other techniques, Eisenhauer and Halek (2001) estimate a population mean risk aversion parameter of −0.89 with an intraquartile range of (−1.83, −0.54). In this environment, with the insurance company facing the same interest rate as the policyholder, annuities are not supported for the bottom four health states, but the equilibrium holds for 73% of the population.

Only if you combine a risk-aversion parameter of −0.75 with an interest rate differential of greater than 0.50% does the market collapse due to adverse selection. With a company interest rate of 4%, the mortality table conditional on initial health state 10 delivers an annuity payout factor of 5.13%. Health-state-10 policyholders with risk-aversion parameter $\alpha = -0.75$ prefer annuities with payout factors greater than 5.65%.

Based on Table 5, the generalization of this example to the general case is relatively straightforward. The row ($\ddot{a}_c|r$) is the inverse of the weighted-average value of the annuity given the assumed purchasing population. The weighted-average calculation depends only on the underlying state probabilities and the $\ddot{a}$ calculation. These state probabilities are age-dependent, but the critical annuity price values π_i are not. Table 6 shows the effect of self-selection on the annuity pricing model as a percentage of simply using the underlying Annuity 2000 table. Notice that until the model accepts all health states, the effect is roughly a 7–10% reduction in annuity values.

5. POLICY CONSIDERATIONS

It has been suggested that capital market considerations make it optimal for policyholders to buy a "term annuity" and invest the difference, at least at the retirement age of 65. The key assumption in this conclusion is that direct access to capital markets provides a superior return to the typical fixed return earned on a payout annuity. Individual knowledge of mortality increases the richness of this decision.

Table 6
Effect of Self-Selection on Annuity Factors

Age	Min. State	Market Pct.	Annuity Factor	Pct. of Table
50	8	81.7%	5.82%	93.17%
51	7	81.9	5.88	93.05
52	7	82.2	5.95	92.96
53	7	82.6	6.03	92.90
54	7	82.5	6.10	92.77
55	6	84.2	6.21	92.97
56	6	83.9	6.29	92.77
57	6	83.2	6.37	92.47
58	5	86.4	6.53	93.12
59	5	85.7	6.62	92.83
60	5	84.8	6.72	92.45
61	5	83.4	6.81	91.97
62	5	81.5	6.90	91.34
63	4	87.1	7.16	92.90
64	4	85.8	7.28	92.33
65	4	84.0	7.39	91.60
66	4	81.9	7.50	90.70
67	3	88.9	7.91	93.23
68	3	87.5	8.06	92.44
69	3	85.7	8.21	91.49
70	3	83.6	8.36	90.40
71	3	81.4	8.50	89.18
72	3	78.9	8.63	87.86
73	2	90.1	9.44	93.15
74	2	88.7	9.65	92.30
75	2	87.3	9.85	91.40
76	2	85.7	10.05	90.46
77	2	84.1	10.25	89.50
78	2	82.5	10.43	88.52
79	1	100.0	12.13	100.00
80	1	100.0	12.47	100.00

The major finding of this paper is that the annuity puzzle may be explained by a heterogeneous population with relatively minor risk aversion, facing a product offering that is slightly uncompetitive. The adverse selection in this model is admittedly stark, with the consumption-smoothing advantages of annuities competing strongly with the adverse selection issues. Based on the utility functions used in this paper, the percentage of the market restricted due to selection issues merits further empirical investigation.

The newly retired policyholder finds for the first time in his or her life that good health is expensive, at least in the sense that it lengthens the expected cost of providing for retirement. This effect exacerbates the negative correlation between health status and retirement age (see, e.g., Dwyer 2001). Healthy workers may optimally delay retirement to avoid drawing down assets that they perceive they will need to fund a long retirement period.

Based on population estimates, the annuitant population at the typical retirement age of 65 is much more homogeneous than the population at 75 or 85. If the annuity market can generate a product mix rich enough to differentiate (separate) between healthy and unhealthy annuitants, the problem of self-selection is ameliorated.

The brute-force solution to mortality differentiation would be to introduce underwriting. Alternatively, any annuity feature which can distinguish between health states might be useful. For example, a payout annuity could be bundled with LTC insurance. A payout annuity with direct inflation protection or increasing future payments would be viewed as more attractive to those with longer life expectancies.

Finally, proper accounting for the adverse selection in payout annuities suggests using a select-ultimate table for payout annuity pricing. Select-ultimate tables more accurately correct for selection effects than do loading factors.

APPENDIX

Here we will see that the value function (3.5) and optimal consumption function (3.6) solve the dynamic programming problem faced by the policyholder.

Lemma (C)

The consumption function (3.6) is optimal given value function (3.5).

The key result in (3.6) is not that a proportion of current assets are consumed each period, but that the proportion is independent of the level of assets a_t. Substituting equations (3.1) and (3.5) into equation (3.4) yields the optimization problem

$$V_j(a_t) = \text{Max}\{c_t\}[(\alpha + 1)^{-1}(\theta_j c_t)^{(\alpha+1)} + \beta(\alpha + 1)^{-1}[(a_t - c_t)(1 + r)]^{\alpha+1} \times \{p_s\delta_j + p_u\delta_{j+1} + p_d\delta_{j-1}\}] + K_j^*, \quad \text{(A.1)}$$

Differentiating the right-hand side of equation (A.1) gives the first-order condition

$$0 = \theta_j^{\alpha+1}c_t^{\alpha} - \beta(1 + r)^{\alpha+1} \times \{p_s\delta_j + p_u\delta_{j+1} + p_d\delta_{j-1}\}(a_t - c_t)^{\alpha}, \quad \text{(A.2)}$$

or equivalently

$$c_t = [1 + (\theta_j^{\alpha+1}/[\beta(1 + r)^{\alpha+1}\{p_s\delta_j + p_u\delta_{j+1} + p_d\delta_{j-1}\}])^{1/\alpha}]^{-1}a_t = (1 - \gamma_j)a_t. \text{ QED} \quad \text{(A.3)}$$

Lemma (V)

The value function (3.5) is consistent with the utility function (3.1) and the optimization problem (3.4).

Substituting equations (3.5) and (3.6) into equation (3.4) yields

$$\begin{aligned} V_j(a_t) &= (\alpha + 1)^{-1}(\theta_j(1 - \gamma_j)a_t)^{(\alpha+1)} + \beta(\alpha + 1)^{-1}[\gamma_j a_t(1 + r)]^{\alpha+1} \times \{p_s\delta_j + p_u\delta_{j+1} + p_d\delta_{j-1}\} + K_j^* \\ &= (\alpha + 1)^{-1}[(\theta_j(1 - \gamma_j))^{(\alpha+1)} + \beta[\gamma_j(1 + r)]^{\alpha+1} \times \{p_s\delta_j + p_u\delta_{j+1} + p_d\delta_{j-1}\}]a_t^{\alpha+1} + K_j^*. \text{ QED} \end{aligned} \quad \text{(A.4)}$$

ACKNOWLEDGMENTS

I would like to thank without implication William Cummings, Moshe Milevsky, Hal Pedersen, Brandon Welte, Sergey Zuev, and two anonymous referees for comments on earlier drafts of this work.

REFERENCES

BROWN, JEFFREY R. 2001. Private Pensions, Mortality Risk, and the Decision to Annuitize. *Journal of Public Economics* 82(1): 29–62.

———. 2003. Redistribution and Insurance: Mandatory Annuitization with Mortality Heterogeneity. *Journal of Risk and Insurance* 70(1): 17–41.

BROWN, JEFFREY R., AND JAMES M. POTERBA. 2000. Joint Life Annuities and Annuity Demand by Married Couples. *Journal of Risk and Insurance* 67(4): 527–53.

BRUNNER, JOHANN K., AND SUSAN PECH. 2000. Adverse Selection in the Annuity Market When Payoffs Vary over the Time of Retirement. Johannes Kepler University of Linz, Department of Economics Working Paper 0030.

———. 2002. Adverse Selection in the Annuity Market with Sequential and Simultaneous Insurance Demand. Johannes Kepler University of Linz, Department of Economics Working Paper 0204.

DWYER, DEBRA S. 2001. Planning for Retirement: Accuracy of Expected Retirement Dates and the Role of Health Shocks. Center for Retirement Research Working Paper 2001-08.

EISENHAUER, JOSEPH G., AND MARTIN HALEK. 2001. Demography of Risk Aversion. *Journal of Risk and Insurance* 68(1): 1–24.

FINKELSTEIN, AMY, AND KATHLEEN MCGARRY. 2003. Private Information and Its Effect on Market Equilibrium: New Evidence from Long-Term Care Insurance. NBER Working Paper 9957.

FINKELSTEIN, AMY, AND JAMES M. POTERBA. 2004. Adverse Selection in Insurance Markets: Evidence from the U.K. Annuity Market. *Journal of Political Economy,* forthcoming.

GERBER, HANS U., AND ELIAS S. W. SHIU. 2000. Investing for Retirement: Optimal Capital Growth and Dynamic Asset Allocation. *North American Actuarial Journal* 4(2): 42–58.

MILEVSKY, MOSHE A. 1998. Optimal Asset Allocation towards the End of the Life Cycle: To Annuitize or not to Annuitize? *Journal of Risk and Insurance* 65(3): 401–26.

———. 2000. Optimal Annuitization Policies: Analysis of the Options. *North American Actuarial Journal* 5(1): 57–69.

MILEVSKY, MOSHE A., AND CHRIS ROBINSON. 2000. Self-annuitization and Ruin in Retirement. *North American Actuarial Journal* 4(4): 112–24.

MITCHELL, OLIVIA S., AND DAVID MCCARTHY. 2002. Estimating International Adverse Selection in Annuities. *North American Actuarial Journal* 6(4): 38–54.

MITCHELL, OLIVIA S., JAMES M. POTERBA, MARK J. WARSHAWSKY, AND JEFFREY R. BROWN. 1999. New Evidence on the Money s Worth of Individual Annuities. *American Economic Review* 89(5): 1299–1318.

POTERBA, JAMES M. 2001. Annuity Markets and Retirement Security. Center for Retirement Research Working Paper 2001-10.

Discussions on this paper can be submitted until April 1, 2007. The author reserves 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.