



Heriot-Watt University
Research Gateway

Exchanging uncertain mortality for a cost

Citation for published version:

Donnelly, C, Guillen, M & Nielsen, JP 2013, 'Exchanging uncertain mortality for a cost', *Insurance: Mathematics and Economics*, vol. 52, no. 1, pp. 65-76. <https://doi.org/10.1016/j.insmatheco.2012.11.001>

Digital Object Identifier (DOI):

[10.1016/j.insmatheco.2012.11.001](https://doi.org/10.1016/j.insmatheco.2012.11.001)

Link:

[Link to publication record in Heriot-Watt Research Portal](#)

Document Version:

Early version, also known as pre-print

Published In:

Insurance: Mathematics and Economics

General rights

Copyright for the publications made accessible via Heriot-Watt Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

Heriot-Watt University has made every reasonable effort to ensure that the content in Heriot-Watt Research Portal complies with UK legislation. If you believe that the public display of this file breaches copyright please contact open.access@hw.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Exchanging uncertain mortality for a cost

Catherine Donnelly*
Montserrat Guillén†
Jens Perch Nielsen‡

November 2, 2012

Abstract

We analyze a pooled annuity fund from a participant's perspective by comparing it to a mortality-linked fund, a type of variable payout life annuity, that gives a return linked to the force of mortality but subject to a cost. Fixing the instantaneous volatility of return on wealth, we find that the expected return on the pooled annuity fund is higher except when the costs are very low in the mortality-linked fund. Similar results are obtained when maximizing the expected lifetime utility of consumption, assuming a constant relative risk aversion utility function. In both settings, our results indicate that a participant may be willing to accept the mortality risk of the pooled annuity fund, even when only 100 individuals are pooling their mortality in the pooled annuity fund.

JEL classification: G22, G23.

Keywords: longevity risk; pensions; pooled annuity fund.

Subject Category and Insurance Branch Category: IM51, IB81.

**Corresponding author.* Department of Actuarial Mathematics and Statistics, and the Maxwell Institute for Mathematical Sciences, Heriot-Watt University, Edinburgh EH14 4AS, United Kingdom. Email: C.Donnelly@hw.ac.uk. Tel: +44 131 451 3251. Fax: +44 131 451 3249.

†Department of Econometrics, Riskcenter-IREA, University of Barcelona, Avinguda Diagonal 690, 08034 Barcelona, Spain. Email: mguillen@ub.edu

‡Cass Business School, City University, 106 Bunhill Row, London EC1Y 8TZ, United Kingdom. Email: Jens.Nielsen.1@city.ac.uk

1 Introduction

In a conventional life annuity, the annuitant pays the life insurance company a lump sum and in return receives regular payments until the time of death. The amount of the regular payments offered in exchange for the lump sum depends on expectations of future interest rates and mortality at the time that the annuity is sold. The insurance company is fully exposed to longevity risk, that is the risk of the annuitant living longer than anticipated. In some annuity contracts, such as variable annuities with a baseline guarantee, the annuitant can have limited control of the underlying investments. However, total investment freedom is not possible since the insurer has guaranteed a payment until death and thus the insurer must be able to choose an investment strategy in order to meet that guarantee.

We study a *pooled annuity fund* that does not transfer longevity risk to the insurer. This is very attractive because regulatory costs are not incurred for longevity risk, as it is not borne by the insurer. In principle, the pooled annuity fund allows the participants to have investment freedom; however, we assume that all participants follow the same investment strategy.

The pooled annuity fund has a very different structure to that of a conventional annuity. Participants in the fund agree to pool their mortality risk together. In contrast, buyers of a conventional annuity transfer their mortality risk to an insurance company. Additionally, the participants invest their wealth in a financial market according to their chosen investment strategy. As participants die, their wealth is shared among the surviving participants as a *mortality credit* (sometimes called a survivor credit). As an example, if there are 50 people in a pool and one of them dies, leaving wealth \$12 250, then each of the 49 survivors receives a mortality credit of $\$12\,250/49 = \250 . Stamos [2008] analyzes this particular pooled annuity fund.

The mortality credit is random in both magnitude and timing. Its magnitude is random since it depends on the wealth of the dying participants, which they had invested in the financial market. It is received at random times since the times of deaths are unknown. As the *pool size*, i.e. the number of participants in the pooled annuity fund, increases, deaths occur more often and the mortality credit should become less volatile. In the limit, as the pool size becomes infinite, deaths occur at a constant rate and the mortality credit is received continuously by the survivors. Note that it will still be random in magnitude as it depends on the magnitude of the wealth which is invested in the financial market.

Our aim is to determine how much an individual would pay to remove the mortality risk in a pooled annuity fund, under various criteria. If the actual amount charged by the insurer is higher, then individuals may benefit from joining a pooled annuity fund. A related analysis is Stamos [2008], who considers the welfare benefits of pooling mortality. He analyzes a pooled annuity fund and assumes, as we do in the sequel, that individuals are independent copies of one another. He calculates the gains in utility caused by participating in a pooled annuity fund, with either a finite or an infinite number of participants, compared to self-annuitization and a conventional annuity.

We distinguish the pooled annuity fund described above from other funds which pool mortality risk, such as *group self-annuitization products* (the latter are sometimes also called pooled annuity funds). In group self-annuitization products, an initial group of participants pool their wealth together and in re-

turn receive a regular annuity payment, calculated on a chosen mortality and investment basis. Periodically, the annuity payment is adjusted for actual mortality and investment experience. Thus there is some smoothing of mortality and investment experience over time. The fundamental analysis of their theoretical operation is by Piggott et al. [2005], and its extension by Qiao and Sherris [2012]. Valdez et al. [2006] consider the practical issues of demand and adverse selection in these funds. Another variation on the operation of a group self-annuitization product is by Richter and Weber [2011], who propose an adjustment to the insurer's reserves to allow for actual experience, and the annuity payment is re-calculated based on the adjusted reserves. While these products are very interesting and worthy of further study, we want to be able to distinguish easily between financial risk and mortality risk. Thus even if we will keep the spirit of these products, we will propose slight variations that will enable a different perspective on longevity risk.

The natural point of comparison with the pooled annuity fund is a fund in which individuals are exposed only to investment risk. We propose such a fund, which we call a *mortality-linked fund*¹. (A detailed description of the two funds is given in Section 3.) Members of the mortality-linked fund transfer their mortality risk to the seller of the fund. The seller pays a deterministic interest rate on each member's wealth, which is proportional to the latter's force of mortality. We call the interest rate the *mortality-linked interest rate*. In return for accepting the mortality risk, the seller charges suitable costs to the members.

The costs are pivotal to our analyses as they are used as a link between the two funds. They can be seen as how much an individual must pay to enjoy the full transfer of the mortality risk. We determine how high the costs have to be before an individual would favor the pooled annuity fund, in which they receive a random, but free of costs, mortality credit, to the mortality-linked fund, in which the deterministic mortality-linked interest rate is paid. The costs are driven by the mortality risk, itself a function of the pool size and force of mortality of the participants in the pooled annuity fund. The determination of the costs is done over two distinct time horizons, in which we allow for the investment and mortality risk in different ways. We make the simplifying assumption that participants in the fund are identical, independent copies of each other: they have the same mortality, wealth and risk preferences.

In the first analysis, we look at the impact of financial and mortality risk through a single lens: their combined effect on the surviving participants' instantaneous return on wealth. Risk is measured as the volatility of the instantaneous return on wealth, which allows us to view financial risk and mortality risk as interchangeable. For example, a member of the mortality-linked fund, who is not exposed to any mortality risk, can achieve the same amount of risk as someone in a pooled annuity fund, who may be heavily exposed to mortality risk. This is done by increasing the former's investment in a risky financial asset, so that they have the same volatility of the instantaneous return on wealth. We find that, even at a pool size of around 100 participants, the costs charged by the seller of the mortality-linked fund must be very low to be attractive to individuals; see Section 4.

In the second analysis in Section 5, we look at the effect of the two funds on an

¹Stamos [2008] says that such products can be thought of as ideal variable payout life annuities. Vittas [2011] refers more specifically to them as standard unit-linked annuities.

individual's utility of lifetime consumption. We assume that individuals follow the consumption and investment strategies that maximize their expected utility of lifetime consumption, assuming a constant relative risk aversion (CRRA) utility function. Then we find the costs at which the expected utilities are the same in both funds. Our results indicate that, even with this radically different time horizon and approach, the costs are similar to those calculated in the first analysis.

We assume throughout that individuals have no bequest motive². In both of the proposed funds, individuals lose their invested wealth upon death. Thus any wealth set aside to satisfy a bequest motive must be an *ex ante* decision; we assume that any such decision has already been made.

In summary, in this paper we compare a pooled annuity fund, in which participants pool their mortality, to another type of fund, which we call a mortality-linked fund, in which members transfer their mortality risk to the seller in exchange for paying some costs. We find that the participants in the pooled annuity fund can achieve

- a higher expected instantaneous return on wealth, for the same volatility of the instantaneous return on wealth, and
- a higher expected utility of lifetime consumption,

compared to members of the mortality-linked fund. Our findings that a pool sizes of 100 or more have the potential to give large benefits to the participants are consistent with the results of Stamos [2008].

2 Background

The need for prudent assumptions and allowances for costs result in the premium charged for a conventional annuity being higher than its expected value calculated on a "fair" net pricing basis, i.e. a basis which reflects current expectations of future interest rates and mortality, and ignores expenses and profit. This is not at all surprising: given the long-term, guaranteed nature of the annuity product with the consequent uncertainty about predicting far into the future, the insurer will choose a cautious pricing basis. For example, the insurance company may use a lower than expected discount rate to calculate the premium charged to the policyholder, and assume a more prudent mortality basis than that which is anticipated. Moreover, there are the transaction costs of buying assets, and risk management costs such as reserves, reinsurance premiums and hedging costs. Regulatory and administration costs must be included. And all this without even mentioning the insurance company's profit.

However, the question has been asked if annuities really do offer value for money to the policyholder. Part of the motivation for asking this question is the *annuity puzzle* observed in particular in the United States, the phenomenon where fewer individuals buy annuities than predicted by economic theory. Cannon and Tonks [2009] and Valdez et al. [2006] summarize the reasons for the

²Lockwood [2012] analyzes the impact of the bequest motive on the demand for both actuarially fair and actuarially unfair annuities. He finds that the bequest motive can have a large impact on the demand for the latter, even if it has little effect on the demand for the former.

annuity puzzle that have been proposed in the literature. One possible reason is a high level of expenses and profits included in the annuity premiums.

Unfortunately, the information required to determine if the annuity offers value for money is not available from companies, due to commercial sensitivity and the difficulty in disentangling life annuity costs from the costs in the general business mix of a life insurer. Nevertheless, several authors have attempted to measure the cost of life annuities. Typically, they do this by calculating the excess of the annuity premium charged to the individual over the expected value calculated on a fair net pricing basis. Then the excess amount is divided by the expected value calculated on the fair net pricing basis. The result is called the *load factor* and is a measure of the amount of expenses, profit and cost of the guarantee made to the annuitant. Clearly, the load factor is sensitive to the choice of the fair net pricing basis. In the UK, Cannon and Tonks [2009] report an average load factor of 10% from 1994 to 2007, with an increase to 12 – 14% in 2007, even after taking into account increases in life expectancies. In an earlier attempt looking at the US market in 1995, Mitchell et al. [1999] suggest a load factor in the range 10 – 15%.

For life insurance companies subject to Solvency II Directive, it is possible that load factors will increase. It is expected that the Solvency II Directive will result in insurers holding higher capital in reserve and facing additional administration requirements. They may also have to hold more government bonds and fewer corporate bonds, with a corresponding reduction in the discount rate used to value annuities. The business advisory firm Deloitte suggests that U.K. annuity rates may fall between 5% and 20%; see Deloitte [2012].

Unanticipated increases in lifespans may also increase load factors, since it suggests uncertainty about the correct choice of mortality model. In van de Ven and Weale [2009] it is suggested that, in the U.K. pensions buyout market, the capital charge levied by insurance companies due to uncertainty about future mortality rates is higher than the charge that even a highly risk averse individual would accept (in which “highly risk averse” means having a coefficient of relative risk aversion of 20).

Whether or not we believe that annuities offer value for money, the existence of the annuity puzzle suggests a problem with the attractiveness of conventional annuities. This is likely to be compounded further by the continuing uncertainty about the correct choice of mortality model (i.e. longevity risk) and, for insurers who are subject to the Solvency II Directive, increased regulatory costs. For these reasons, it is imperative to suggest and analyze alternative products that have the potential to meet at least some of the needs of retirees and which may turn out to be attractive to them.

3 Description of the funds

Here we describe the two funds and look at some practical issues surrounding them. The funds are

- a pooled annuity fund, in which participants pool their mortality risk, and
- a mortality-linked fund, subject to costs. Similar to a conventional annuity, members of the mortality-linked fund transfer their mortality risk to the seller of the fund and, in exchange, pay the costs.

We make the assumption that individuals are independent copies of one another: they are the same age, and have the same deterministic force of mortality and risk preferences at all times. At time 0, they all have the same constant wealth $w_0 > 0$.

Once an individual has joined one of the funds, they are not permitted to exit it for reasons other than death. We further assume that there is a fixed number L_0 of participants at time 0 in the pooled annuity fund. No new participants can join after time zero, i.e. the pooled annuity fund is a closed fund. The assumption could be relaxed, so that the pooled annuity fund is an open fund and new entrants are permitted after time zero. However, we have kept the closed fund assumption to avoid unnecessarily complicating our analysis.

The number of deaths which have occurred in the pooled annuity fund is modeled by a counting process N . The number of individuals alive at time t is denoted by

$$L_t = L_0 - N_t.$$

We let N_t be a Poisson process at rate $\lambda_t(L_{t-} - 1)$, in which λ_t is the deterministic force of mortality at time t of an individual. As we model the wealth of the surviving individuals, the rate of N_t is the force of mortality multiplied by $L_{t-} - 1$. In other words, we exclude the individual whose wealth we are modeling and who is alive at time t from the potential deaths over the short time interval (t_-, t) .

In both the pooled annuity fund and the mortality-linked fund, the wealth of each individual is invested in a frictionless financial market which consists of two traded assets: a risky asset and a risk-free asset. The risk-free asset has price B_t at time t with dynamics

$$dB_t = rB_t dt, \quad B_0 > 0 \text{ constant},$$

with constant risk-free rate of return r . The price process S of the risky asset is driven by a 1-dimensional standard Brownian motion \mathcal{Z} , so that at time t it has dynamics

$$dS_t = S_t(\mu dt + \sigma d\mathcal{Z}_t), \quad S_0 > 0 \text{ constant},$$

with constant $\mu > r$ and constant $\sigma > 0$.

The Brownian motion \mathcal{Z} and the Poisson process N are defined on the same complete probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and are independent processes. With $\mathcal{N}(\mathbb{P}) := \{A \in \mathcal{F} : \mathbb{P}(A) = 0\}$, the information available to individuals at time t is represented by the filtration

$$\mathcal{F}_t = \sigma\{(N_s, \mathcal{Z}_s), s \in [0, t]\} \vee \mathcal{N}(\mathbb{P}), \quad \forall t \geq 0.$$

In other words, at each time t , it is known how many individuals have died in the pooled annuity fund and the price of the risky asset at all times up to and including time t .

3.1 Pooled annuity fund

In the pooled annuity fund, the wealth of the dead participants is shared equally among the survivors as a mortality credit. If a participant dies in the time

interval $(t-, t)$, leaving wealth W_{t-} , then each survivor receives at time t an additional amount of wealth

$$\frac{1}{L_{t-} - 1} W_{t-}.$$

All participants have identical consumption and investment strategies. Then, as long as $L_{t-} > 1$, the wealth W of a survivor evolves as

$$\frac{dW_t}{W_{t-}} = (r + \pi_t(\mu - r) - c_t) dt + \sigma \pi_t dZ_t + \frac{1}{L_{t-} - 1} dN_t, \quad (3.1)$$

in which c_t is the rate of consumption by the individual and π_t is the proportion of wealth invested in the risky asset at time t . The remaining proportion of wealth $1 - \pi_t$ is invested in the risk-free asset. Note that, upon denoting the random time of death of an arbitrary participant by τ , the change in wealth at the time of death is $dW_\tau = -W_{\tau-}$, with the “lost” wealth $W_{\tau-}$ shared among the surviving participants.

When there is only one survivor left in the pool, i.e. $L_{t-} = 1$, the wealth dynamics are

$$\frac{dW_t}{W_{t-}} = (r + \pi_t(\mu - r) - c_t) dt + \sigma \pi_t dZ_t, \quad (3.2)$$

so that changes in the wealth are due solely to investment in the financial market. When the last survivor dies, their wealth is transferred to their estate.

The principal advantage of the pooled annuity fund is that the cost to the participants should be significantly lower than a conventional annuity. There are no capital reserving requirements, no need for reinsurance and no hedging costs because there are no guarantees. In particular, it is the participants who bear the longevity risk.

Despite the lower costs, there will still some expenses incurred. The fund must be set-up, advertised and potential participants educated as to the risks involved. Investment management fees must be paid. Deaths of participants must be notified to and verified by an administrative body, and the wealth released upon death must be allocated among the survivors. Our analysis ignores these expenses.

Another advantage of the pooled annuity fund is that, up to the time of their own death, participants always gain from pooling mortality while retaining investment freedom. They can consume as much as they want of their wealth. The main disadvantage compared to a conventional life annuity is that the pooled fund does not guarantee a minimum lifetime payment.

We assume that all pool participants are indistinguishable from each other. Then distributing the wealth of the dead participants equally among the survivors is fair, as all participants have the same wealth and the same probability of dying. In practice, we can only approximate such an ideal pool of participants; for example, by grouping participants into suitable age bands. To have similar amounts of wealth in the fund, participants’ wealth could be pooled and managed by a single investment manager so that their investment returns were the same. The manager would market the pooled annuity fund as one following a particular investment strategy. Individuals would choose to join the pooled annuity fund that best met their required investment objectives. There would likely be annual limits on how much can be consumed.

While these are not strict requirements for the operation of the pooled annuity fund in a realistic setting, the pooled annuity fund structure described above is not appropriate for participants with diverse characteristics. For example, a pool consisting only of 20-year-olds and 90-year-olds is likely to benefit the young much more than the old, assuming that the wealth of the dead participants is divided equally among all participants. The same holds true for participants with different risk preferences³. With the described pooled annuity fund structure, the only attempt that we know of to account for inhomogeneous participants is Sabin [2010]. Even then, his sharing rule has strong constraints on the inhomogeneity of the participants.

3.2 Mortality-linked fund

The second fund that we consider is a *mortality-linked fund*. The mortality-linked fund is related in a natural way to the pooled annuity fund, as we demonstrate in Subsection 3.3. Just as in the pooled annuity fund, members of the mortality-linked fund invest in the financial market, cannot exit the fund before their death and lose all their wealth upon death. However, there is no direct exposure to the pooling of mortality.

Instead, the wealth of each fund member increases at a *mortality-linked interest rate* that is proportional to their force of mortality. The force of mortality is deterministic and it could be calculated from a life table produced by an independent body, like a national statistical office. We make the simplifying assumption that the force of mortality is the same as the one used to model the pooled annuity fund; the implications of this assumption is discussed below.

The mortality-linked interest rate is in addition to the return due to investment in the financial market⁴. The underlying idea is that the seller, for example a life insurance company, has sold the mortality-linked fund to many people. The mortality-linked fund allows its members to invest in the financial market and also benefit from indirectly pooling mortality via the seller. However, unlike the pooled annuity fund participants, the members are not exposed directly to the volatility of the actual number of deaths occurring. Instead, the members receive the deterministic mortality-linked interest rate on their wealth. It is the seller who bears the mortality volatility and, accordingly, must be compensated.

Specifically, the mortality-linked interest rate is equal to $\lambda_t(1 - a_t)$, where λ_t is the member's force of mortality and a_t are the *costs* applying at time t . The costs are compensation to the seller for bearing the mortality volatility. It makes sense to express the costs as a percentage of the force of mortality since the volatility from deaths are proportional to λ_t .

Upon joining the mortality-linked fund, the wealth W^g of a surviving member evolves as

$$\frac{dW_t^g}{W_{t-}^g} = (r + \pi_t^g(\mu - r) - c_t^g + \lambda_t(1 - a_t)) dt + \sigma \pi_t^g dZ_t, \quad (3.3)$$

³It would be interesting to analyze a pooled annuity fund which allows true consumption and investment freedom, and the implications for the optimal consumption and investment strategies of the participants, for participants who differed only in their risk preferences. The analysis could allow for opportunistic behavior by some of the participants.

⁴When all of the wealth is invested in an asset which gives a deterministic return, the mortality-linked fund is similar to the actuarial note product of Yaari [1965].

in which c_t^g is the rate of consumption by the individual and π_t^g is the proportion of wealth invested in the risky asset at time t . The remaining proportion of wealth $1 - \pi_t^g$ is invested in the risk-free asset. As in the pooled annuity fund, all participants are assumed to follow the same consumption and investment strategies, so that c_t^g and π_t^g are identical for the surviving members at all times t .

There is only one source of volatility for a mortality-linked fund member: that from the financial market. In comparison, each participant in the pooled annuity fund experiences an additional source of uncertainty from the timing and number of deaths occurring in the pool.

Unlike the pooled annuity fund, there is a systematic mortality risk inherent in the mortality-linked fund. It is ultimately borne by the seller. To calculate the mortality-linked interest rate, the seller chooses a force of mortality that is intended to reflect the future mortality experience of the members. An incorrect choice of the force of mortality means that the seller is either charging too much or too little for the mortality risk. The risk could be mitigated by occasionally allowing the insurer the opportunity to update the force of mortality to reflect the emerging mortality experience, or by increasing the costs charged to the members. Having chosen a force of mortality, there is also a risk of adverse selection against the seller. It may be reduced in similar ways. We do not analyze the impact of systematic risk or adverse selection in this paper.

The seller decides the level of costs a_t . They may be increased above their fair value by the seller's requirement to satisfy regulations. It may also be the case that some sellers, for example large insurance companies, may be able to pool the mortality risk from the mortality-linked fund with other products they sell. Thus the costs charged may differ between sellers, even if they have identical beliefs about future mortality rates.

Excluding the costs a_t , the expenses of operating the mortality-linked fund should be higher than for the pooled annuity fund. The seller of the mortality-linked fund has guaranteed the mortality-linked interest rate, which implies a capital reserving requirement. Correspondingly, there may also be reinsurance and hedging costs. Other expenses should be the same as for the pooled annuity fund: for example, setting up the fund, managing the assets and monitoring deaths should incur a comparable amount of expenditure.

In practice, the seller of the mortality-linked fund would decide the investment strategy, as they have guaranteed to pay the mortality-linked interest rate on each member's wealth. Correspondingly, they would also place strict limits on the amount of wealth that could be consumed by the members. Unlike the pooled annuity fund, these restrictions are very likely to be required for the operation of the mortality-linked fund. However, in our analysis we assume simply that the members have no such constraints. We take the viewpoint that individuals can choose the mortality-linked fund that best meets their consumption and investment objectives.

3.3 Relationship between the two funds

Finally, we explore the important relationship between the pooled annuity fund and the mortality-linked fund. Consider a pooled annuity fund with an infinite number of participants: deaths occur at a constant rate equal to the force of mortality, without any variability. Under the assumption that there is no sys-

tematic mortality risk, in the infinite pooled annuity fund there is no idiosyncratic mortality risk, i.e. the mortality risk is fully diversified. Consequently, the dynamics of a survivor's wealth process W^∞ are obtained from (3.1) with $1/(L_{t-} - 1)dN_t$ replaced by $\lambda_t dt$, resulting in⁵

$$\frac{dW_t^\infty}{W_{t-}^\infty} = (r + \pi_t(\mu - r) - c_t + \lambda_t) dt + \sigma\pi_t dZ_t.$$

Compare the above dynamics with the wealth dynamics (3.3) for the mortality-linked fund. The only difference lies in the costs a_t charged by the seller, which act to reduce the expected return on wealth for members of the mortality-linked fund. Clearly, given the choice between the infinite pooled annuity fund and the mortality-linked fund with costs $a_t > 0$, a rational individual would choose the infinite pooled annuity fund.

In real life we cannot construct an infinite pooled annuity fund as there are only a finite number of people alive in the world. Even assuming that systematic risk does not exist, this means that mortality risk can not be completely eliminated. From this viewpoint, the costs a_t can be interpreted as the price that an individual must pay to access the equivalent of an infinite pooled annuity fund. The costs will depend on the alternative offered to the mortality-linked fund, namely the exact specifications of the pooled annuity fund. In our simple setting, the pooled annuity fund is specified by the number of participants (i.e. the pool size) and their characteristics: their force of mortality, consumption and investment strategies. For this reason, we set the costs of the mortality-linked fund to be a function of these variables of the pooled annuity fund.

4 Instantaneous approach

In this section we only care about the short-term dynamics, so we do not consider consumption and concentrate only on the instantaneous view.

We calculate the level of costs such that, for equal volatilities of return on the wealth, an individual has the same expected return on wealth from the pooled annuity fund as from the mortality-linked fund. We call the level of costs at which this occurs the *instantaneous breakeven costs*. If the actual costs charged by the mortality-linked fund are higher than the instantaneous breakeven costs, then the individual can obtain a higher expected return from the pooled annuity fund for the same amount of volatility of return. We find that the breakeven costs are approximately inversely proportional to the pooled annuity fund size.

We begin by calculating the expected value and volatility of the instantaneous return on wealth for a surviving individual in both funds. Then we determine the instantaneous breakeven costs, illustrated with some numerical results.

4.1 Pooled annuity fund

With zero consumption, the wealth dynamics (3.1) become

$$\frac{dW_t}{W_{t-}} = (r + \pi_t(\mu - r)) dt + \sigma\pi_t dZ_t + \frac{1}{L_{t-} - 1} dN_t. \quad (4.1)$$

⁵A more detailed explanation of the derivation of the wealth dynamics for an infinite pool can be found in Stamos [2008].

As N_t is a Poisson process at rate $\lambda_t(L_{t-} - 1)$, the instantaneous expected return is calculated from (4.1) as

$$\mathbb{E} \left(\frac{dW_t}{W_{t-}} \middle| \mathcal{F}_{t-} \right) = (r + \pi_t(\mu - r) + \lambda_t) dt. \quad (4.2)$$

Thus the instantaneous expected return from pooling mortality in the pooled annuity fund equals the force of mortality λ_t and is independent of the pool size L_{t-} . This also means that the expected return on wealth due to deaths increases as participants age, rather than being constant. The reason is that the force of mortality increases as the participants age, meaning that deaths occur more frequently in the fund. This may result in increasingly high levels of wealth as the participants age, although these may offset higher long-term care and healthcare costs incurred at higher ages. The remaining part of the instantaneous expected return is derived from investment in the financial market.

As deaths occur randomly in the pooled annuity fund, the variability of their occurrence feeds into the volatility on the return on wealth. Intuitively, the predictable volatility process v of the return on wealth should equal the conditional variance, i.e.

$$v_t^2 dt = \text{Var} \left(\frac{dW_t}{W_{t-}} \middle| \mathcal{F}_{t-} \right).$$

Mathematically, we define v through

$$v_t^2 dt = \frac{d\langle W, W \rangle_t}{(W_{t-})^2}, \quad (4.3)$$

in which $\langle \cdot, \cdot \rangle_t$ denotes the usual predictable bracket process (for example, see [Protter, 2005, page 125]). From (4.3) it is straightforward to show that, for $L_{t-} > 1$, the instantaneous volatility of wealth satisfies

$$v_t^2 dt = \left((\sigma\pi_t)^2 + \frac{\lambda_t}{L_{t-} - 1} \right) dt. \quad (4.4)$$

Now we can easily see the two sources of volatility of the return on wealth: one from the financial market, $(\sigma\pi_t)^2$, and the other arising from deaths occurring in the pool, $\lambda_t/(L_{t-} - 1)$. Thus, the greater the number of participants in the pooled annuity fund, the lower the volatility of the return on wealth. For an infinite pooled annuity fund, mortality risk is fully diversified and there is volatility only from the financial market.

4.2 Mortality-linked fund

With zero consumption, the wealth dynamics (3.3) become

$$\frac{dW_t^g}{W_{t-}^g} = (r + \pi_t^g(\mu - r) + \lambda_t(1 - a_t)) dt + \sigma\pi_t^g dZ_t. \quad (4.5)$$

The instantaneous expected return on wealth for a survivor in the mortality-linked fund is calculated from (4.5) as

$$\mathbb{E} \left(\frac{dW_t^g}{W_{t-}^g} \middle| \mathcal{F}_{t-} \right) = (r + \pi_t^g(\mu - r) + \lambda_t(1 - a_t)) dt. \quad (4.6)$$

Defining the instantaneous volatility process v^g similarly to (4.3), we find

$$(v_t^g)^2 dt = (\sigma\pi_t^g)^2 dt. \quad (4.7)$$

As the mortality credit $\lambda_t(1 - a_t)$ is deterministic, the only source of volatility of return on wealth comes from the financial market. This means that for identical proportions of wealth invested in the risky asset, i.e. $\pi_t^g = \pi_t$, the volatility in the pooled annuity fund v_t is higher than that of the mortality-linked fund v_t^g ; compare (4.4) and (4.7). However, the expected return on wealth in the pooled annuity fund is higher by the amount $\lambda_t a_t$; compare (4.2) and (4.6).

4.3 Instantaneous breakeven costs

In practice, the costs applied in the mortality-linked fund would be determined by the seller. However, we allow them to vary with the number of participants in the pooled annuity fund. This is how we link the two funds in our analysis. It allows us to calculate the amount that an individual would pay to remove the mortality risk of a pooled annuity fund of a given pool size.

Definition 4.1. The *instantaneous breakeven costs* applying at time t are the costs such that, for equal instantaneous volatilities of return on the wealth, a surviving individual has the same instantaneous expected return on wealth from the pooled annuity fund as from the mortality-linked fund at time t .

If the actual costs charged by the mortality-linked fund are higher than the instantaneous breakeven costs, then the individual can obtain a higher expected return from the pooled annuity fund for the same amount of volatility of return on wealth.

We expect the instantaneous breakeven costs to depend on the number of participants in the pooled annuity fund, the proportion invested in the risky asset and the force of mortality, as these three factors drive the expected return and volatility of the return on wealth. The exact relationship is made clear by the next proposition.

Proposition 4.2 (Instantaneous breakeven costs). *Suppose that a pooled annuity fund participant invests the proportion π_t of their wealth in the risky asset, with the volatility of the return on wealth denoted by v_t .*

Let the proportion of wealth π_t^g invested in the risky asset by a member of the mortality-linked fund be such that the volatilities of the return on wealth are equal, i.e. $v_t = v_t^g$.

Then $\pi_t^g = \pi^g(L_{t-}, \pi_t, \lambda_t)$ with

$$\pi^g(\ell, \pi, \lambda) = \begin{cases} \pi & \text{for } \ell = 1 \\ \left(\pi^2 + \frac{1}{\sigma^2} \frac{\lambda}{\ell-1}\right)^{1/2} & \text{for } \ell = 2, 3, \dots \end{cases} \quad (4.8)$$

and the instantaneous breakeven costs are $a_t^ = a^*(L_{t-}, \pi_t, \lambda_t)$ for*

$$a^*(\ell, \pi, \lambda) = \begin{cases} 1 & \text{for } \ell = 1 \\ \frac{(\mu-r)\pi}{\lambda} \left[\left(1 + \frac{1}{(\pi\sigma)^2} \frac{\lambda}{\ell-1}\right)^{1/2} - 1 \right] & \text{for } \ell = 2, 3, \dots \end{cases} \quad (4.9)$$

for all $\pi \geq 0$ and $\lambda > 0$.

Proof. See Appendix A.1. \square

Notice that as the pool size ℓ tends to infinity, the breakeven costs tend to zero, i.e. $\lim_{\ell \rightarrow \infty} a^*(\ell, \pi, \lambda) = 0$. As the pool size becomes larger, the volatility due to deaths declines and the finite pooled annuity fund becomes closer to the infinite pooled annuity fund discussed in Subsection 3.2. Correspondingly, the seller of the mortality-linked fund has to charge less in order to attract individuals to join the mortality-linked fund.

Next we develop a first-order approximation to the breakeven costs (4.9), to help us understand the main factors affecting the costs for finite pool sizes.

Proposition 4.3 (First-order approximation). *If $\ell - 1 > \lambda \left(\frac{1}{\pi\sigma}\right)^2$ then*

$$a^*(\ell, \pi, \lambda) \approx \frac{1}{2} \frac{(\mu - r)}{\sigma^2} \frac{1}{\pi} \frac{1}{\ell - 1}.$$

Proof. Apply a Taylor series expansion to (4.9) and ignore terms of order higher than one. \square

Since the volatility of the return on wealth due to deaths in the pool is inversely proportional to the size of the pool, it is not surprising that the first-order approximation suggests that the breakeven costs are too. We can also explain why the breakeven costs should be roughly inversely proportional to the proportion invested in the risky asset. The more risk-averse an investor, the less that is invested in the risky asset and the higher the breakeven costs. Thus the costs must be relatively high in the mortality-linked fund before a risk-averse investor chooses the pooled annuity fund.

4.4 Numerical illustration

In the numerical illustrations, we calculate the instantaneous breakeven costs under various scenarios. The results indicate that individuals will not pay much to transfer their mortality risk to the seller of the mortality-linked fund, compared to the alternative of the pooled annuity fund.

We make the following financial assumptions: $r = 0.02$, $\mu = 0.06$ and $\sigma = 0.18$. We consider different pool sizes, from $\ell = 10$ to $\ell = 10\,000$, different investment strategies, from $\pi = 10\%$ to $\pi = 75\%$, and various forces of mortality, from $\lambda = 0.005$, to $\lambda = 0.04$.

Table 1 shows $\pi^g(\ell, \pi, \lambda) - \pi$, i.e. the excess proportion of wealth invested in the risky asset for the mortality-linked fund over a pooled annuity fund with pool size ℓ and participants with force of mortality λ and investing the proportion π of their wealth in the risky asset, required to obtain the same volatility of return on wealth. The values are overall rather low, which indicates that the volatility in the pooled annuity fund due to mortality risk is quite low. For example, compared to a pooled annuity fund with 100 participants, an individual in the mortality-linked fund has to invest at most an additional 4.99% of their wealth in the risky asset to obtain the same volatility of return on wealth. For a pooled annuity fund with 1000 participants, this declines to 0.60% of their wealth. The results show that the excess percentage increases as the force of mortality λ increases and the pooled annuity fund size ℓ decreases.

Table 2 gives the instantaneous breakeven costs $a^*(\ell, \pi, \lambda)$ of Proposition 4.2. The breakeven costs appear approximately inversely proportional to the size ℓ

of the pooled annuity fund. They increase as the proportion of wealth invested in the risky asset increases. These observations demonstrate that the rule of thumb of Proposition 4.3 gives a reasonable indication of how the instantaneous breakeven costs behave.

Table 2 also shows in italics the breakeven costs expressed as a monetary cost rate per 100 units of wealth, calculated as $100(1 - e^{-\lambda a^*(\ell, \pi, \lambda)})$. Again, the values are surprising low. For example, compared to a pooled annuity fund of 100 participants who have a force of mortality of 0.04 and who invest 10% of their wealth in the risky asset, the costs charged by the seller of the mortality-linked fund would have to be less than 0.1994% of a member's wealth before an individual can obtain a higher expected return from joining the mortality-linked fund, for the same volatility of return on wealth as in the pooled annuity fund. Even though the mortality-linked fund member invests an additional 4.99% in the risky asset compared to the pooled annuity fund (see Table 1), the additional expected return is $(\pi^g - \pi)(\mu - r) = 0.2\%$ in this case, which explains why the monetary cost rate is very low. With a pool size of 1000, the costs must be even lower: less than 0.024% per annum of the member's wealth.

5 Lifetime approach

Here we analyze the pooled annuity fund in a much longer-term setting than in the previous section. We suppose that individuals consume their wealth over their lifetime and that they gain utility from the consumption.

First we calculate, for both funds, the consumption and investment strategies that maximize the expected utility of lifetime consumption. In our analysis, we use a CRRA utility function. Clearly, in the mortality-linked fund the optimal strategies depend on the costs charged. As in the previous approach, to link the two funds we let the costs at each time t be a function of the characteristics of the pooled annuity fund at time t .

Definition 5.1. Suppose the individuals in the pooled annuity fund and the mortality-linked fund follow their optimal consumption and investment strategies. Then the *lifetime breakeven costs* at time t are the costs at which a member of the mortality-linked fund has the same expected utility of lifetime consumption, starting from time t , as a participant in the pooled annuity fund.

It turns out, that for the chosen CRRA utility function, the lifetime breakeven costs depend only on the number of survivors in the pooled annuity fund. We find how

- (a) the lifetime breakeven costs compare to the instantaneous breakeven costs of Section 4, and
- (b) the amount of withdrawal varies for a pooled annuity fund participant compared to a mortality-linked fund member, when they both follow their optimal consumption and investment strategies and the lifetime breakeven costs are charged.

We find that, for a given pool size, the lifetime breakeven costs are comparable in value to the instantaneous breakeven costs. Furthermore, when the

lifetime breakeven costs are charged, the individuals in both funds follow identical consumption and investment strategies. Our numerical results indicate that participants in the pooled annuity fund have the potential to withdraw a greater amount of money than the members of the mortality-linked fund, albeit with greater volatility in the withdrawn amount.

5.1 Value function

Individuals are assumed to have the same preferences as expressed by

$$U(\tau) = \int_0^\tau e^{-\delta s} u(C_s) ds,$$

in which δ denotes the time preference rate, τ denotes the stochastic time of death and $C_t = c_t W_t$ is the amount of consumption at time t . The utility function $u(C)$ is assumed to be of the standard CRRA type given below:

$$u(C) = \frac{C^{1-\gamma}}{1-\gamma}, \quad \gamma \neq 1, \quad \gamma > 0,$$

where γ is the level of relative risk aversion.

As we want to use the costs to compare the two funds, we allow them to depend on the number of survivors in the pooled annuity fund. Consequently, the value function of a individual at time t in either fund is of the form

$$V(w, \ell, t) = \sup_{(\pi_s, c_s)_{s=t}^\infty} \mathbb{E} \left(\int_t^\tau e^{-\delta(s-t)} u(C_s) ds \mid W_t = w, L_t = \ell, \tau > t \right).$$

By integrating over the random time of death, we obtain

$$V(w, \ell, t) = \sup_{(\pi_s, c_s)_{s=t}^\infty} \mathbb{E} \left[\int_t^\infty e^{-\int_t^s (\lambda_u + \delta) du} u(C_s) ds \mid W_t = w, L_t = \ell \right].$$

The difference in the value function for the two funds lies in the dynamics of the wealth process. We use $V(\cdot, \cdot, \cdot)$ to denote the value function for a participant in the pooled annuity fund, and $V^g(\cdot, \cdot, \cdot)$ to denote the value function for a member of the mortality-linked fund.

5.2 Derivation of the lifetime breakeven costs

To calculate the lifetime breakeven costs, we must first find the optimal consumption and investment strategies for individuals in the pooled annuity fund and the mortality-linked fund.

5.2.1 Optimal strategies for the pooled annuity fund

The wealth process of a surviving participant in the pooled annuity fund evolves with consumption as in (3.1). Since (3.1) is identical to that in Stamos [2008], we can use the solution in the latter; we summarize his results next. The optimal investment strategy is the constant proportion

$$\pi^* = \frac{\mu - r}{\gamma \sigma^2}. \quad (5.1)$$

The value function has the form

$$V(w, \ell, t) = f(\ell, t) \frac{w^{1-\gamma}}{1-\gamma}. \quad (5.2)$$

The optimal consumption rate c^* is then obtained by setting

$$c^*(\ell, t) = f(\ell, t)^{-\frac{1}{\gamma}}. \quad (5.3)$$

For $\ell > 1$, the function $f(\ell, t)$ satisfies

$$\frac{f_t(\ell, t)}{f(\ell, t)} + \gamma f(\ell, t)^{-\frac{1}{\gamma}} + (A - \lambda_t \ell) + \lambda_t(\ell - 1) \left(\frac{\ell}{\ell - 1} \right)^{1-\gamma} \frac{f(\ell - 1, t)}{f(\ell, t)} = 0, \quad (5.4)$$

with $f_t(\ell, t)$ denoting the partial derivative of $f(\ell, t)$ with respect to t , and the constant A defined by

$$A = (1 - \gamma) \left(r + \frac{1}{2\gamma} \left(\frac{\mu - r}{\sigma} \right)^2 \right) - \delta.$$

For $\ell = 1$, the function $f(1, t)$ satisfies

$$\frac{f_t(1, t)}{f(1, t)} + \gamma f(1, t)^{-\frac{1}{\gamma}} + (A - \lambda_t) = 0. \quad (5.5)$$

The boundary conditions are $\lim_{t \rightarrow \infty} f(\ell, t) = 0$ and $f(0, t) = 0$. As there is no analytical solution to (5.4), it must be solved numerically.

5.2.2 Optimal strategies for the mortality-linked fund

The wealth process of a surviving member of the mortality-linked fund, allowing for costs a_t , evolves with consumption as in (3.3). From the calculations shown in Appendix A.2, we find that the form of the value function is

$$V^g(w, \ell, t) = h(\ell, t) \frac{w^{1-\gamma}}{1-\gamma}, \quad (5.6)$$

and the optimal consumption rate c^{g^*} satisfies

$$c^{g^*}(\ell, t) = h(\ell, t)^{-1/\gamma}.$$

The optimal investment strategy is

$$\pi^{g^*} = \frac{\mu - r}{\gamma \sigma^2}. \quad (5.7)$$

In other words, a utility-maximizing individual invests the same constant proportion of their wealth in the risky asset, regardless of whether they join the pooled annuity fund or the mortality-linked fund; compare (5.1) and (5.7).

For $\ell > 1$, the function $h(\ell, t)$ satisfies

$$\frac{h_t(\ell, t)}{h(\ell, t)} + \gamma h(\ell, t)^{-\frac{1}{\gamma}} + A - \lambda_t \ell + \lambda_t(1 - \gamma)(1 - a(\ell, t)) + \frac{h(\ell - 1, t)}{h(\ell, t)} \lambda_t(\ell - 1) = 0. \quad (5.8)$$

For $\ell = 1$, i.e. when there is only one survivor left in the pooled annuity fund, we set the costs $a(1, t) = 1$. The assumption is that, for individuals following the same consumption and investment strategies, the costs are such that the wealth dynamics (3.3) in the mortality-linked fund are identical to those (3.2) of the pooled annuity fund. Then the function $h(1, t)$ satisfies (5.5).

The boundary conditions are $\lim_{t \rightarrow \infty} h(\ell, t) = 0$ and $h(0, t) = 0$. As for the pooled annuity fund, we must solve (5.8) numerically.

5.2.3 Lifetime breakeven costs

Having calculated the optimal consumption and investment strategies for both funds, we can now find the lifetime breakeven costs. We assume that individuals in the pooled annuity fund and the mortality-linked fund follow their optimal strategies, and calculate the corresponding expected utility of lifetime consumption. The lifetime breakeven costs at each time t are the costs at which the expected utilities are equal at that time t .

Proposition 5.2. *For each integer $\ell \geq 2$, the lifetime breakeven costs are*

$$a^*(\ell, t) = 1 - \frac{\ell - 1}{1 - \gamma} \frac{f(\ell - 1, t)}{f(\ell, t)} \left[\left(\frac{\ell}{\ell - 1} \right)^{1 - \gamma} - 1 \right], \quad (5.9)$$

for $f(\cdot, \cdot)$ satisfying (5.4). For $\ell = 1$, the lifetime breakeven costs are $a^*(1, t) = 1$.

Furthermore, when the lifetime breakeven costs apply, the optimal consumption rates for individuals in each fund are identical.

Proof. For $\ell = 1$, we assume that the costs, and hence the lifetime breakeven costs, are $a^*(1, t) = a(1, t) = 1$ for all $t \geq 0$.

Fix an integer $\ell \geq 2$. The lifetime breakeven costs are the costs $a^*(\ell, t)$ at which $V(w, \ell, t) = V^g(w, \ell, t)$. From (5.2) and (5.6), we see immediately that equating the value functions is equivalent to equating $h(\ell, t) = f(\ell, t)$, so that the rates of consumption are identical. Setting $h(\ell, t) = f(\ell, t)$ in (5.8) and then equating with (5.4), results in (5.9). \square

As there is no analytical solution for $f(\ell, t)$ we must solve for $a^*(\ell, t)$ numerically.

5.3 Numerical results

For the numerical results, we use the same assumptions in [Stamos, 2008, Section 4.1]. This means that, as before, for the financial market we set $r = 0.02$, $\mu = 0.06$ and $\sigma = 0.18$. The risk preferences of each individual is specified by the parameters $\delta = 0.04$ and $\gamma = 5$. The optimal investment in the risky asset for individuals in either fund is then

$$\pi^* = \pi^{g^*} = \frac{\mu - r}{\gamma \sigma^2} \approx 25\%. \quad (5.10)$$

We suppose that individuals are age 60 at time 0 and their force of mortality at time t is given by the Gompertz law of mortality

$$\lambda_t = \frac{1}{b} e^{(60+t-m)/b}, \quad \text{with } m = 86.85 \text{ and } b = 9.98, \quad \forall t \geq 0. \quad (5.11)$$

For example, in our notation, the force of mortality of an 80-year-old individual is λ_{20} . The parameter m is the modal age at death and b is the dispersion coefficient. The above law was fitted by Stamos [2008] to US female population mortality data.

5.3.1 Comparison of the breakeven costs

A natural question is how the lifetime breakeven costs compare to the instantaneous breakeven costs for different pool sizes. Figure 1 shows the lifetime breakeven costs and Figure 2 shows the instantaneous breakeven costs. For all calculated pool sizes, the two types of breakeven costs are of the same order of magnitude. They vary most across ages for a pool size of 10. However, for pool sizes of 100 or more, the costs are similar in value and quite stable, no matter which approach is used.

5.3.2 Withdrawal amount rate

Here we consider how the amount consumed varies for a pooled annuity fund participant compared to a mortality-linked fund member. We assume that mortality-linked fund charges the lifetime breakeven costs and each individual follows their optimal consumption and investment strategies, as calculated in Subsection 5.2.

When the lifetime breakeven costs are charged, we saw in Subsection 5.2 that the optimal investment and the consumption strategies are identical for individuals in both funds. This means that

- the mortality-linked fund member has a lower expected return on wealth since they pay the breakeven costs. The mortality-linked fund member has an expected (and deterministic) mortality credit at the rate $\lambda_t(1 - a_t^*)$ whereas the pooled annuity fund participant's expected mortality credit is at the rate λ_t . But,
- the pooled annuity fund participant has a higher volatility of return on wealth since they are exposed to volatility in the mortality credit. For both funds, the volatility due to the investment market is the same.

To illustrate how this impacts on the amount consumed, we simulate the withdrawal amount rates $(1 - e^{-c_t})W_t$ and $(1 - e^{-c_t^g})W_t^g$ for participants in the pooled annuity fund and members of the mortality-linked fund, respectively, over a 50-year period. All individuals follow their optimal consumption and investment strategies. Note that, since $c(\ell, t) = c^g(\ell, t)$, this is really illustrating the differences in the wealth evolution of individuals in the two funds.

The mortality-linked fund is compared to a pooled annuity fund with different initial pool sizes ($L_0 = 100$, $L_0 = 500$ and $L_0 = 1000$). All individuals in both funds are age 60 and have initial wealth $w_0 = 1$ unit at time 0. Time was discretized into one month steps, $\Delta t = 1/12$, and the pool size was kept as integer steps over the discretization grid. In total, 100 000 samples of the withdrawal amount rates were produced.

Figure 3 indicates that individuals may benefit from joining the pooled annuity fund. It shows that the expected withdrawal amount is lower for the mortality-linked fund than for the pooled annuity fund for a given initial pool

size, except for an initial pool size of $L_0 = 100$ at ages above 105 years (at time 45 years). For the chosen parameters, the withdrawal amount is almost identical for ages up to 85 (at time 25 years), and they do not diverge significantly until age 95 (at time 35 years). Under the chosen mortality law, we expect about half of the initial pool to have survived to age 85, around 10% to survive to age 95, and less than 5% to survive to age 100. Therefore, for an initial pool size $L_0 = 100$, there are only about 10 survivors at age 95 on average and fewer than 5 are expected to survive to age 100.

Given that the individuals in both funds follow identical optimal consumption and investment strategies, it is not surprising that the expected withdrawal amount is higher for the participants of the pooled annuity fund: they do not have to pay any costs. Notice that the expected amount peaks at higher ages as the initial pool size increases. This can be attributed to having enough participants still alive at higher ages in the pooled annuity fund so that the mortality credit is received with a high probability, whereas in the mortality-linked fund the mortality-linked interest rate includes the lifetime breakeven costs.

The mean and various percentiles of the withdrawal amount rates for an initial group size of 100 is shown in Figure 4. For an initial group size of 1000, see Figure 5. For Figure 5, the greater volatility in the pooled annuity fund is clear but it is skewed favourably for the surviving individual.

6 Discussion

We have studied the stability of income streams from a pooled annuity fund. It is obvious that the size of the pool is critical so that statistical regularity produces a stable series of payments. There is a substantial list of literature examining pooled returns, for instance, surpluses in life insurance and performance of life settlement funds⁶, but only a few authors have studied annuities.

Our main result is about the equivalence between a pooled annuity fund and a mortality-linked fund, the latter charging some costs to exchange an uncertain mortality return for a certain, regular return. In order to compare the two, we fix the volatility of return on wealth and then find the level of costs such that the two funds have exactly equal expected returns. The level of costs decreases as the size of the pool increases. An identical conclusion is obtained using a lifetime approach.

These results are promising as they indicate that a mechanism different to the mortality-linked fund, a mechanism which does not have to adapt to uncertainty about mortality and that is equivalent in terms of volatility of wealth and expected return, may be attractive to an individual, regardless of the time horizon of the individual. This suggests that the study of pooled annuity funds is very relevant.

We have considered only homogeneous pools of individuals. However, in real-life, individuals are highly dissimilar, having different wealth, risk preferences and mortality rates. Heterogeneous pools have been studied in the context of group self-annuitization products by Piggott et al. [2005] and Qiao and Sherris

⁶Bohnert and Gatzert [2011] examine surplus appropriation schemes in participating life insurance contracts and the impact on the insurers shortfall risk and the net present value from an insured's perspective. Braun et al. [2012] indicate that caution is advised in order not to overestimate the performance of life settlement funds.

[2012]. Heterogeneous pooled annuity funds have been explored by Sabin [2010] but, unfortunately, the spreading rule that he proposes does not work for any heterogeneous pool. We aim to explore how this difficulty could be overcome for a pooled annuity fund in a later work.

Acknowledgements

The authors thank an anonymous referee whose comments greatly improved the paper. Support for this research from the Spanish Ministry of Science/FEDER grant ECO2010-21787-C03-01 and ICREA Academia is gratefully acknowledged.

References

- T. Björk. *Arbitrage Theory in Continuous Time*. Oxford University Press, Oxford, 3rd edition, 2009.
- A. Bohnert and N. Gatzert. Analyzing surplus appropriation schemes in participating life insurance from the insurer's and the policyholder's perspective. *Insurance: Mathematics and Economics*, 50:64–78, 2011.
- A. Braun, N. Gatzert, and H. Schmeiser. Performance and risks of open-end life settlement funds. *Journal of Risk and Insurance*, 79(1):193–230, 2012.
- E. Cannon and I. Tonks. Money's worth of pension annuities. *Department for Work and Pensions, UK Government*, 2009. Research Report No 563.
- Deloitte. Solvency II rules could make retirement more expensive, June 2012. Media release developed out of a survey conducted by the Economist Intelligence Unit. http://www.deloitte.com/view/en_GB/uk/62ef05d570ff7310VgnVCM2000001b56f00aRCRD.htm.
- L.M. Lockwood. Bequest motives and the annuity puzzle. *Review of Economic Dynamics*, 15(2):226 – 243, 2012.
- O.S. Mitchell, J.M. Poterba, M.J. Warshawsky, and J.R. Brown. New evidence on the money's worth of individual annuities. *The American Economic Review*, 89(5):1299–1318, 1999.
- J. Piggott, E.A. Valdez, and B. Detzel. The simple analytics of a pooled fund. *Journal of Risk and Insurance*, 72(3):497–520, 2005.
- P.E. Protter. *Stochastic Integration and Differential Equations*. Springer, Germany, 2nd edition, 2005.
- C. Qiao and M. Sherris. Managing systematic mortality risk with group self-pooling and annuitization schemes. To appear in the *Journal of Risk and Insurance*, 2012.
- A. Richter and F. Weber. Mortality-indexed annuities: Managing longevity risk via product design. *North American Actuarial Journal*, 15(2):212–236, 2011.
- M.J. Sabin. Fair tontine annuity. <http://ssrn.com/abstract=1579932>, March 2010.

- M.Z. Stamos. Optimal consumption and portfolio choice for pooled annuity funds. *Insurance: Mathematics and Economics*, 43(1):56–68, 2008.
- E.A. Valdez, J. Piggott, and L. Wang. Demand and adverse selection in a pooled annuity fund. *Insurance: Mathematics and Economics*, 39(2):251–266, 2006.
- J. van de Ven and M. Weale. Mortality risk and pricing of annuities. National Institute Discussion Paper Number 322, July 2009.
- D. Vittas. The mechanics and regulation of variable payout annuities, August 2011. The World Bank, Policy Research Working Paper 5762.
- M.E. Yaari. Uncertain lifetime, life insurance and the theory of the consumer. *Review of Economic Studies*, 32(2):137–150, 1965.

A Technical results

A.1 Proof of Proposition 4.2

Proof of Proposition 4.2. First consider the case that $L_{t-} = 1$. The relevant wealth dynamics are (3.2) for the pooled annuity fund and (3.3) for the mortality-linked fund, with the consumption set to zero, i.e. $c_t = c_t^g = 0$. To equate the volatilities of return on wealth, set $\pi_t = \pi_t^g$. Then equating the expected returns on wealth gives that the instantaneous breakeven costs are $a_t^* = 1$.

Next consider the case that $L_{t-} > 1$. The instantaneous volatility of return on wealth of the pooled annuity fund participant v_t satisfies (4.4). For a mortality-linked fund member who invests a proportion π_t^g of their wealth in the risky asset, the instantaneous volatility of return on wealth v_t^g satisfies (4.7). To have the same volatility of return on wealth, i.e. $v_t^g = v_t$, the proportion π_t^g must satisfy

$$(\sigma\pi_t^g)^2 = (\sigma\pi_t)^2 + \frac{\lambda_t}{L_{t-} - 1}.$$

Rearranging we find $\pi_t^g = \pi^g(L_{t-}, \pi_t, \lambda_t)$ with

$$\pi^g(\ell, \pi, \lambda) = \left(\pi^2 + \frac{1}{\sigma^2} \frac{\lambda}{\ell - 1} \right)^{1/2}, \quad (\text{A.1})$$

for all $\pi \geq 0$, $\lambda > 0$ and integer $\ell \geq 2$, and we have now shown that (4.8) holds.

To have the same volatility of return on wealth, the mortality-linked fund member must invest a greater fraction of wealth in the risky asset than the pooled annuity fund participant. This is to compensate for the mortality-linked fund member not being exposed to any mortality risk. From (4.8), we see that the mortality-linked fund member will increase their holdings in the risky asset π^g as

- λ increases (i.e. deaths occur more frequently in the pooled annuity fund),
- σ decreases (i.e. the volatility of the risky asset decreases), and
- ℓ decreases (i.e. the wealth released by each death is spread among fewer survivors in the pooled annuity fund).

Having equated the instantaneous volatility at time t , we determine the instantaneous breakeven costs a_t^* by equating the instantaneous expected returns in (4.2) and (4.6), i.e.

$$r + \pi_t^g(\mu - r) + \lambda_t(1 - a_t^*) = r + \pi_t(\mu - r) + \lambda_t.$$

Solving we find

$$a_t^* = \frac{\mu - r}{\lambda_t} (\pi_t^g - \pi_t). \quad (\text{A.2})$$

From (A.2), we see that the instantaneous breakeven costs increase as

- $\mu - r$ increases (i.e. the expected return from the risky asset increases),
- λ decreases (i.e. the expected return from deaths decreases), and
- $\pi_t^g - \pi_t$ increases (i.e. a greater fraction of wealth is invested in the risky asset by the mortality-linked fund member).

Substituting (A.1) into (A.2) we find that (4.9) holds. The assumptions that $\mu > r$ and $\pi \geq 0$ ensure that $a_t^* \geq 0$. \square

A.2 Optimal strategies for the mortality-linked fund

For notational ease, we drop the “g” superscript except to denote the optimal strategies, and define

$$\theta = \frac{\mu - r}{\sigma}.$$

For the mortality-linked fund, we wish to find the optimal investment strategy π^{g^*} and optimal consumption strategy c^{g^*} that solves:

$$V(w, \ell, t) = \sup_{(\pi_s, c_s)_{s=t}^{\infty}} \mathbb{E} \left[\int_t^{\infty} e^{-\int_t^s (\lambda_u + \delta) du} u(c_s \pi_s) ds \mid W_t = w, L_t = \ell \right],$$

for the wealth dynamics given by

$$\frac{dW_t}{W_{t-}} = (r + \pi_t(\mu - r) + \lambda_t(1 - a_t) - c_t) dt + \sigma \pi_t dZ_t. \quad (\text{A.3})$$

To apply dynamic programming theory, we assume that the consumption process and costs are of the form $c(L_{t-}, t_-)$ and $a(L_{t-}, t_-)$, respectively. We begin by deriving the relevant Hamilton-Jacobi-Bellman equation, following the technique in [Björk, 2009, Chapter 19].

Assume that from time t to time $t + h$, the participant follows arbitrary consumption and investment strategies (c_s, π_s) . After time $t + h$, the participant follows the optimal consumption and investment strategies. Then denoting by $\mathbb{E}^{w, \ell, t}$ the expectation conditional on $W_t = w$ and $L_t = \ell$,

$$\begin{aligned} V(w, \ell, t) \geq & \mathbb{E}^{w, \ell, t} \left[\int_t^{t+h} e^{-\int_t^s (\lambda_u + \delta) du} u(c_s W_s) ds \right] \\ & + \mathbb{E}^{w, \ell, t} \left[e^{-\int_t^{t+h} (\lambda_u + \delta) du} V(W_{t+h}, L_{t+h}, t + h) \right]. \end{aligned} \quad (\text{A.4})$$

Define the operator

$$\begin{aligned} \mathcal{L}_t^{c,\pi} h(w, \ell, t) &= h_t(w, \ell, t) + wh_w(w, \ell, t) [r + \pi\theta\sigma + \lambda_t(1 - a(\ell, t)) - c(\ell, t)] \\ &\quad + \frac{1}{2}w^2 h_{ww}(w, \ell, t)\sigma^2\pi^2 + [h(w, \ell - 1, t) - h(w, \ell, t)]\lambda_t(\ell - 1). \end{aligned} \quad (\text{A.5})$$

Assuming sufficient differentiability so that we can apply Ito's lemma to the product $e^{-\int_t^{t+h}(\lambda_u+\delta)du}V(W_{t+h}, L_{t+h}, t+h)$, we use the wealth dynamics in (A.3) to find

$$\begin{aligned} &e^{-\int_t^{t+h}(\lambda_u+\delta)du}V(W_{t+h}, L_{t+h}, t+h) \\ &= V(W_t, L_t, t) + \int_t^{t+h} \mathcal{L}_s^{c,\pi} e^{-\int_t^s(\lambda_u+\delta)du}V(W_{s-}, L_{s-}, s-)ds \\ &\quad + \int_t^{t+h} e^{-\int_t^s(\lambda_u+\delta)du}W_sV_w(W_{s-}, L_{s-}, s-)\sigma\pi_s d\mathcal{Z}_s \\ &\quad + \int_t^{t+h} e^{-\int_t^s(\lambda_u+\delta)du}(V(s) - V(s-))\lambda_s(L_{s-} - 1)d\mathcal{M}_s, \end{aligned} \quad (\text{A.6})$$

in which $\mathcal{M}_t = N_t - \int_0^t \lambda_s(L_{s-} - 1)ds$ is the Poisson martingale associated with the Poisson process N . Substituting (A.6) into (A.4), the expectation of the stochastic integrals vanish, leaving

$$\begin{aligned} V(w, \ell, t) &\geq \mathbb{E}^{w,\ell,t} \left[\int_t^{t+h} e^{-\int_t^s(\lambda_u+\delta)du} u(c_s W_s) ds \right] \\ &\quad + \mathbb{E}^{w,\ell,t} \left[V(W_t, L_t, t) + \int_t^{t+h} \mathcal{L}_s^{c,\pi} e^{-\int_t^s(\lambda_u+\delta)du} V(W_{s-}, L_{s-}, s-) ds \right]. \end{aligned}$$

Hence

$$\begin{aligned} 0 &\geq \mathbb{E}^{w,\ell,t} \left[\int_t^{t+h} e^{-\int_t^s(\lambda_u+\delta)du} u(c_s W_s) ds \right] \\ &\quad + \mathbb{E}^{w,\ell,t} \left[\int_t^{t+h} \mathcal{L}_s^{c,\pi} e^{-\int_t^s(\lambda_u+\delta)du} V(W_{s-}, L_{s-}, s-) ds \right]. \end{aligned} \quad (\text{A.7})$$

Divide by $h > 0$, let h go to zero and assume enough regularity so that we can take the limit within the expectation to obtain

$$\begin{aligned} 0 &\geq u(c_t w) - (\delta + \lambda_t)V(w, \ell, t) \\ &\quad + V_t(w, \ell, t) + wV_w(w, \ell, t)(r + \pi\theta\sigma + \lambda_t(1 - a(\ell, t)) - c(\ell, t)) \\ &\quad + \frac{1}{2}w^2 V_{ww}(w, \ell, t)\sigma^2\pi^2 + [V(w, \ell - 1, t) - V(w, \ell, t)]\lambda_t(\ell - 1). \end{aligned}$$

If we choose the optimal consumption and investment strategies, then we obtain equality:

$$\begin{aligned} &(\delta + \lambda_t)V(w, \ell, t) \\ &= \sup_{(c,\pi)} \left\{ u(cw) + V_t(w, \ell, t) + wV_w(w, \ell, t)(r + \pi\theta\sigma + \lambda_t(1 - a(\ell, t)) - c) \right. \\ &\quad \left. + \frac{1}{2}w^2 V_{ww}(w, \ell, t)\sigma^2\pi^2 + [V(w, \ell - 1, t) - V(w, \ell, t)]\lambda_t(\ell - 1) \right\}. \end{aligned} \quad (\text{A.8})$$

Guessing the form of the value function,

$$V(w, \ell, t) = h(\ell, t) \frac{w^{1-\gamma}}{1-\gamma}, \quad (\text{A.9})$$

with boundary conditions $\lim_{t \rightarrow \infty} h(\ell, t) = 0$ and $h(0, t) = 0$, the optimal investment strategy is

$$\pi^{g^*} = \frac{\theta}{\gamma\sigma}$$

and the optimal consumption rate is given by

$$c^{g^*}(\ell, t) = h(\ell, t)^{-\frac{1}{\gamma}}. \quad (\text{A.10})$$

Dropping the supremum in (A.8) and setting $\pi := \pi^{g^*}$, $c := c^{g^*}(\ell, t)$ and using the value function (A.9) results in

$$\frac{h_t(\ell, t)}{h(\ell, t)} + \gamma h(\ell, t)^{-\frac{1}{\gamma}} + A - \lambda_t \ell + \lambda_t (1 - \gamma)(1 - a(\ell, t)) + \frac{h(\ell - 1, t)}{h(\ell, t)} \lambda_t (\ell - 1) = 0,$$

which is precisely (5.8).

B Figures and tables

Table 1: $\pi^g(\ell, \pi, \lambda) - \pi$, i.e. the excess percentage of wealth invested in the risky asset for the mortality-linked fund over the pooled annuity fund, required to obtain the same instantaneous volatility of return on wealth. Different forces of mortality (λ), proportion of wealth invested in the risky asset (π) and pool sizes (ℓ) are shown.

		$\ell = 10$	$\ell = 100$	$\ell = 1\,000$	$\ell = 10\,000$
$\lambda = 0.005$	$\pi = 10\%$	6.48 %	0.75 %	0.08 %	0.01 %
	$\pi = 25\%$	3.22 %	0.31 %	0.03 %	0.00 %
	$\pi = 50\%$	1.69 %	0.16 %	0.02 %	0.00 %
	$\pi = 75\%$	1.13 %	0.10 %	0.01 %	0.00 %
$\lambda = 0.01$	$\pi = 10\%$	11.05 %	1.45 %	0.15 %	0.00 %
	$\pi = 25\%$	6.11 %	0.62 %	0.06 %	0.01 %
	$\pi = 50\%$	3.32 %	0.31 %	0.03 %	0.00 %
	$\pi = 75\%$	2.25 %	0.21 %	0.02 %	0.00 %
$\lambda = 0.02$	$\pi = 10\%$	18.03 %	2.74 %	0.30 %	0.03 %
	$\pi = 25\%$	11.21 %	1.22 %	0.12 %	0.01 %
	$\pi = 50\%$	6.44 %	0.62 %	0.06 %	0.01 %
	$\pi = 75\%$	4.44 %	0.41 %	0.04 %	0.00 %
$\lambda = 0.04$	$\pi = 10\%$	28.36 %	4.99 %	0.60 %	0.06 %
	$\pi = 25\%$	19.68 %	2.38 %	0.25 %	0.02 %
	$\pi = 50\%$	12.22 %	1.23 %	0.12 %	0.01 %
	$\pi = 75\%$	8.65 %	0.83 %	0.08 %	0.01 %

Table 2: Instantaneous breakeven costs $a^*(\ell, \pi, \lambda)$ expressed as a percentage. Underneath each number, in italics, is $100(1 - e^{-\lambda a^*(\ell, \pi, \lambda)})$, the instantaneous breakeven costs expressed as a monetary cost rate per 100 units of wealth. Different forces of mortality (λ), proportion of wealth invested in the risky asset (π) and pool sizes (ℓ) are shown.

		$\ell = 10$	$\ell = 100$	$\ell = 1\,000$	$\ell = 10\,000$
$\lambda = 0.005$	$\pi = 10\%$	51.81 % <i>0.2587</i>	6.01 % <i>0.0300</i>	0.62 % <i>0.0031</i>	0.06 % <i>0.0003</i>
	$\pi = 25\%$	25.77 % <i>0.1288</i>	2.48 % <i>0.0124</i>	0.25 % <i>0.0012</i>	0.02 % <i>0.0001</i>
	$\pi = 50\%$	13.49 % <i>0.0674</i>	1.25 % <i>0.0062</i>	0.12 % <i>0.0006</i>	0.01 % <i>0.0001</i>
	$\pi = 75\%$	9.08 % <i>0.0454</i>	0.83 % <i>0.0042</i>	0.08 % <i>0.0004</i>	0.01 % <i>0.00004</i>
	$\pi = 10\%$	44.18 % <i>0.4409</i>	5.81 % <i>0.0581</i>	0.61 % <i>0.0061</i>	0.06 % <i>0.0006</i>
	$\pi = 25\%$	24.45 % <i>0.2442</i>	2.46 % <i>0.0246</i>	0.25 % <i>0.0025</i>	0.02 % <i>0.0002</i>
	$\pi = 50\%$	13.28 % <i>0.1327</i>	1.24 % <i>0.0124</i>	0.12 % <i>0.0012</i>	0.01 % <i>0.0001</i>
	$\pi = 75\%$	9.01 % <i>0.0901</i>	0.83 % <i>0.0083</i>	0.08 % <i>0.0008</i>	0.01 % <i>0.0001</i>
$\lambda = 0.02$	$\pi = 10\%$	36.07 % <i>0.7187</i>	5.48 % <i>0.1096</i>	0.61 % <i>0.0122</i>	0.06 % <i>0.0012</i>
	$\pi = 25\%$	22.41 % <i>0.4472</i>	2.43 % <i>0.0487</i>	0.25 % <i>0.0049</i>	0.02 % <i>0.0005</i>
	$\pi = 50\%$	12.89 % <i>0.2574</i>	1.24 % <i>0.0248</i>	0.12 % <i>0.0025</i>	0.01 % <i>0.0002</i>
	$\pi = 75\%$	8.88 % <i>0.1775</i>	0.83 % <i>0.0166</i>	0.08 % <i>0.0016</i>	0.01 % <i>0.0002</i>
	$\pi = 10\%$	28.36 % <i>1.1281</i>	4.99 % <i>0.1994</i>	0.60 % <i>0.0240</i>	0.06 % <i>0.0025</i>
	$\pi = 25\%$	19.68 % <i>0.7843</i>	2.38 % <i>0.0952</i>	0.25 % <i>0.0098</i>	0.02 % <i>0.0010</i>
	$\pi = 50\%$	12.22 % <i>0.4877</i>	1.23 % <i>0.0493</i>	0.12 % <i>0.0049</i>	0.01 % <i>0.0005</i>
	$\pi = 75\%$	8.65 % <i>0.3453</i>	0.83 % <i>0.0033</i>	0.08 % <i>0.0033</i>	0.01 % <i>0.0003</i>

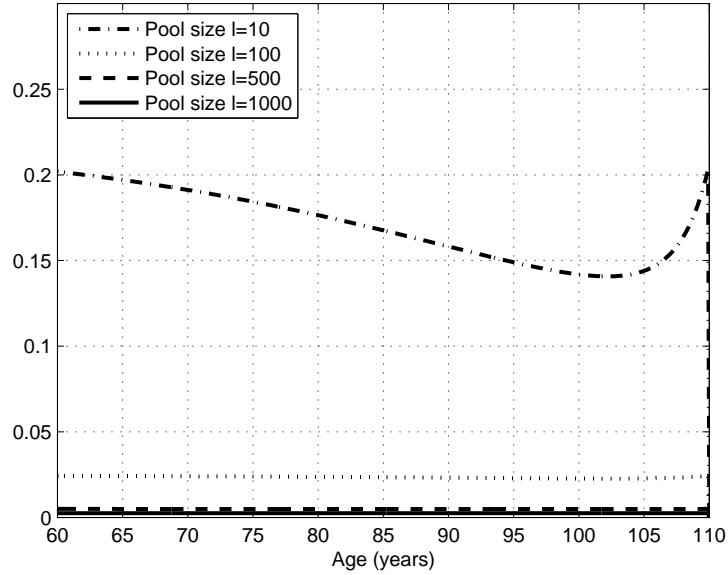


Figure 1: Lifetime breakeven costs, calculated from Proposition 5.2, for various pool sizes. The proportion invested in the risky asset is given by (5.10) and the force of mortality at each age is calculated using (5.11).

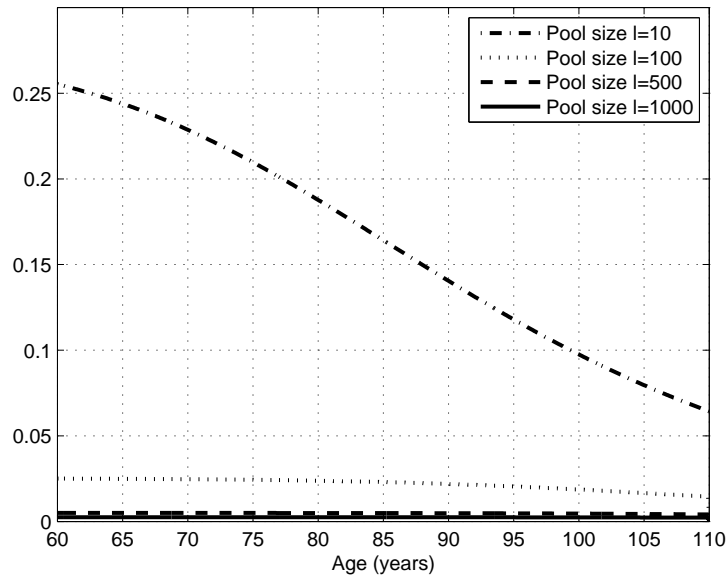


Figure 2: Instantaneous breakeven costs, calculated from (4.9), for various pool sizes. The proportion invested in the risky asset is given by (5.10) and the force of mortality at each age is calculated using (5.11).




Figure 3: Mean withdrawal amount rate when the lifetime breakeven costs are charged. The rates are calculated for the pooled annuity fund (“Pool”) and mortality-linked fund (“M-L”), for various initial pool sizes L_0 . All participants are age 60 and have initial wealth $w_0 = 1$ unit at time 0.




Figure 4: Mean withdrawal amount rate and percentiles when the lifetime breakeven costs are charged. They are calculated for the pooled annuity fund (“Pool”) and mortality-linked fund (“M-L”), for initial pool size $L_0 = 100$. All participants are age 60 and have initial wealth $w_0 = 1$ unit at time 0. For each fund, the central line is the mean, the next two lines above and below the mean are the 25% and 75% percentiles, and the outermost two lines are the 5% and 95% percentiles.




Figure 5: Mean withdrawal amount rate and percentiles when the lifetime breakeven costs are charged. They are calculated for the pooled annuity fund (“Pool”) and mortality-linked fund (“M-L”), for initial pool size $L_0 = 1000$. All participants are age 60 and have initial wealth $w_0 = 1$ unit at time 0. For each fund, the central line is the mean, the next two lines above and below the mean are the 25% and 75% percentiles, and the outermost two lines are the 5% and 95% percentiles.