Actuarial Research Centre
Institute and Faculty of Actuaries

ARC Programme – Optimising Future Pension Plans

Heterogeneity among new entrants to a pooled annuity fund

by
Professor Catherine Donnelly

Preprint

March 2023

The Optimising Future Pensions research was commissioned by the Institute and Faculty of Actuaries' (IFoA) Actuarial Research Centre that has sponsored a number of large research collaborations with academia in recent years on topics such as longevity and morbidity risk, population ageing and big health. The ARC programme currently consists of two remaining academic programmes which conclude in 2023. The IFoA's research and thought leadership programme has broadened its approach to reflect current strategy and themes, and invests in industry-led, cutting-edge research through its thought leadership activity, policy reports, and member-led research initiatives; for instance, its work relating to the Great Risk Transfer and Poverty Premium.

Heterogeneity among new entrants to a pooled annuity fund

Catherine Donnelly*

https://risk-insight-lab.com/

Risk Insight Lab,
Department of Actuarial Mathematics and Statistics,
Heriot-Watt University, Edinburgh, Scotland EH14 4AS

March 27, 2023

Abstract

Pooled annuity funds pay an income for life to their members, by pooling directly the members' longevity risk together. It is well known that the more members in the fund, the less volatile are the income payments to the members. However, this is true only as long as the members are independent and identical copies of each other.

A group of heterogeneous members, who are of different ages and have different amounts of money in the fund, have a higher income volatility than a same-sized group of homogeneous members. A measure of heterogeneity is proposed and studied in conjunction with a measure of income volatility.

Despite heterogeneity increasing income volatility, the effects are in general small. The general rule that pooling more participants together is better, regardless of their characteristics, is still broadly supported. It is only in a fund with very heterogeneous participants that restrictions on the membership should be considered, such as capping the amount of money that an individual can bring to the fund.

Keywords: tontine; longevity risk; mortality; retirement; decumulation.

1 Introduction

Pooled annuity funds provide their participants with a retirement income for life. They do this by pooling directly the longevity risk of their participants. Just like in life annuities and defined benefit pension schemes, deaths among the shorter-lived subsiside payments to the longer-lived. A

*C.Donnelly@hw.ac.uk

key difference with the former structures is that there is no life insurance or sponsoring employer between or behind the participants. Instead, the participants are exposed to the volatility of deaths in the fund, which manifests itself through volatility in the income paid to the participants. The relationship between how the collective characteristics of the fund's participants affects the income volatility is the focus of this paper.

Structures in which longevity risk is pooled directly among participants have increasingly gained industry and academic interest in recent years. Various ways of doing longevity risk-sharing have been proposed. Some of these work for single-cohort pools in which every member is an independent and identical copy of each other (for example, Milevsky and Salisbury 2015; Stamos 2008) and some are intended for multi-cohort pools (for example, Piggott et al. 2005; Stamos 2008; Sabin 2010; Qiao and Sherris 2013; Donnelly et al. 2014; Milevsky and Salisbury 2016). There is heterogeneity among the participants in the latter funds since successive cohorts age and have different amounts of money in the fund, according to what has occurred.

Denuit and Robert (2021) present various fair linear risk-sharing rules, and a conditional mean risk-sharing rule and study their convergence. Weinert and Gründl (2021) derive a distribution to model the longevity credits paid from the pooled annuity fund, rather than modelling directly the mortality experience of the pool of participants.

Milevsky and Salisbury (2015) calculate the payout from a pooled annuity fund which maximizes the expected discounted value of lifetime consumption. The optimal payout to participants varies according to the utility-maximization problem considered, which can also be observed in Chen et al. (2020, 2021) who use the approach of Milevsky and Salisbury (2015).

Some authors calculate what proportion of the funds of those who have died should be received by each participant – whether explicitly (Stamos, 2008; Sabin, 2010; Donnelly et al., 2014) or implicitly (Piggott et al., 2005; Qiao and Sherris, 2013). Both of these explicit and implicit schemes then use an annuity value to calculate the income paid out. Here, an explicit rule is used to share out the funds of those who have died.

Due to a pooled annuity fund being akin to a life annuity without the implicit financial and mortality guarantees, a study of the demand for pooled annuity funds compared to life annuities is a natural one. It has been studied by authors such as Piggott et al. 2005; Valdez et al. 2006; Donnelly et al. 2013; Hanewald et al. 2013; Milevsky and Salisbury 2016; Chen et al. 2021. The results show that pooled annuity funds become increasingly preferred to life annuities as the loadings on life annuities increase. Additionally, the attractiveness of pooled annuity funds increases as the risk aversion of the retiree reduces. Retirees who are less risk averse are happier to bear the volatility of pooled annuity funds in exchange for their higher expected return.

The income paid from a pooled annuity fund is volatile whereas the income paid under a conventional life annuity contract is not. Some investors may prefer to switch their pensions savings to a life annuity at some future age, as they prefer the income stability of an annuity contract at higher ages. Chen et al. (2018) study this problem, considering an individual receiving income from a pooled annuity fund up to some fixed age, followed by income from a deferred life annuity. They determine the optimal age at which to switch between these contracts, with optimality determined using a CRRA utility function-based criterion.

Bequests are another important area of research, since many individuals value a death benefit which can be used by their dependents. This is studied by various authors in different settings, such as Bernhardt and Donnelly (2019); Dagpunar (2021); Zhou et al. (2021); Chen and Rach (2022).

Turning to the literature most relevant to the study of this paper, Bernhardt and Donnelly (2020) calculate how income stability is affected by the number of homogeneous members – who share the same characteristics – in a single cohort fund. They propose a measure of income stability which is applied here. They derive an analytical expression that can be used to determine for how many years is the income stable in a pooled annuity fund, for a given number of members. A numerical study of the same results, seen in Bernhardt and Donnelly (2021, Figure 2), shows that at least 2 000 participants are needed to eliminate most of the idiosyncratic longevity risk.

Bernhardt and Qu (2022) extend the study to a heterogeneous fund, in which some of the members are wealthier than the others. In all other aspects, members are identical and the setting is again a single cohort fund. They again derive an analytical expression that can be used to determine for how many years is the income stable. In the analytical expression, a measure of heterogeneity is observed. Their study then focuses on determining criteria for when a particular collection of members should stay in a large fund rather than form smaller but homogeneous funds to be beneficial.

Qiao and Sherris (2013) study systematic longevity risk in a pooled annuity fund. They find that allowing a pooled annuity fund to remain open improves income stability. This illustrates again the general rule of pooled annuity funds: the more members, the better. While their charts show income falling gradually over time, this is likely due to the life annuity value used to calculate the income paid out, being discordant with the calculation of the longevity credit rather than a feature of pooled annuity funds.

Donnelly (2022) also studies systematic longevity risk, studying the question of how many members are needed to join an open fund each year, to have sufficient longevity pooling. She finds that if around 100 homogeneous members, most of the idiosyncratic longevity risk is eliminated. In a study of the distribution of income paid to each cohort, she finds that cohorts experience a similar outcome, in terms of the distribution of income, with the exception of the last cohorts. The last cohorts to join the fund have a similar experience up until the time when they are old. At that point, there are no new cohorts to join to continue to diversify idiosyncratic longevity risk. The last cohorts are, in effect, in a version of single cohort fund when they are old. These observations motivate a study of only single cohort funds in this paper.

The structure of the pooled annuity fund is set out in Section 2. The definitions of a measure of heterogeneity and a measure of income volatility are discussed in Section 3. Section 4 shows the numerical results and the conclusion is in Section 5.

2 The pooled annuity fund structure

2.1 Participants in the fund

Suppose there are N cohorts who join the fund, with $N \in \mathbb{N}$. Cohort m joins at integer time m when they are integer age $x > 0$ years old, for $m = 0, 1, \ldots, N-1$. The assumption that all members in a cohort are the same age is removed in one of the numerical simulations. However, it is notationally easier not to allow for this relaxation.

Each cohort is further sub-divided into two groups: a rich and poor group. The rich group members have the same mortality as the poor group members, but will bring more money to the fund. At integer time $n \geq 0$, there are $L^{(m,k)}(n)$ members of group $k \in \{\text{rich}, \text{poor}\}$ and cohort $m \in \{0, 1, \ldots, N\}$ in the fund. Noting that the time that cohort m joins the fund is time m, let $L^{(m,k)}(m) > 0$ for each group $k \in \{\text{rich}, \text{poor}\}$ and cohort $m \in \{0, 1, \ldots, N\}$.

The future lifetimes of the participants are independent random variables. The mortality of each survivor follows a known mortality table. As the focus is on establishing a relationship between a measure of heterogeneity in the fund and the stability of income payments, a simple mortality distribution is used. It is assumed that the mortality distribution faced by each survivor depends on their current age only.

2.2 Participants' fund value and income payments

Each new participant who belongs to cohort m and group k brings an amount $F^{(m,k)}(0_-) > 0$ units to the fund. Their money is deposited into an account that is, until the participant's death, ring-fenced for the participant.

The account value of each participant changes over time as investment returns and longevity credits are earned and as income is paid out. Investment returns and longevity credits are added to the account values at the end of each year, just before income payments are made. The account value of a surviving participant in cohort n and group $k \in \{\text{rich}, \text{poor}\}$ is modelled by a stochastic process, which has value $F^{(m,k)}(n)$ at time $n \geq m$, for $m = 0, 1, \ldots, N$.

Let ${}_kp_x$ be the probability that a life age x years survives for $k \geq 0$ years. Using a constant interest rate i effective per annum,

$$\ddot{a}_x = \sum_{k=0}^{\infty} (1+i)^{-k} {}_kp_x$$

represents the expected present value of a payment of 1 unit per annum, paid annually in arrears to a life who is currently age x years.

The amount of income paid out to a participant in cohort m and group k when they first join the fund at time m, is

$$C^{(m,k)}(m) = F^{(m,k)}(m_-)/\ddot{a}_x.$$

his leaves an amount $F^{(m,k)}(m) = F^{(m,k)}(m_-) - C^{(m,k)}(m)$ units in the participant's account, which is then invested over the next year. Investment returns are earned at the constant effective rate i per annum on the entire fund value. The calculation of the longevity credit, $M^{(m,k)}(n)$, which is the amount earned at time $n \geq m$ due to the sharing out of the account values of those who have died over $(n-1, n]$, is detailed in Section 2.4.

At integer time $n > m$, the income paid to a survivor of cohort m and group k, who is age $x + n - m$, is calculated as

$$C^{(m,k)}(n) = \frac{(1+i)\, F^{(m,k)}(n-1) + M^{(m,k)}(n)}{\ddot{a}_{x+n-m}}. \tag{1}$$

Then the fund value at time $n > m$ of each survivor in cohort m and group k is

$$F^{(m,k)}(n) = (1+i)\, F^{(m,k)}(n-1) + M^{(m,k)}(n) - C^{(m,k)}(n).$$

No longevity credit is paid to and no income is paid out at integer time n to a participant who dies over the time period $(n-1, n]$. Instead, the accumulated amount in their account, $(1+i)\, F^{(m,k)}(n-1)$, is shared out among the surviving participants at time n, as detailed in Section 2.4, as a longevity credit.

2.3 Motivation for the expression of the longevity credit

The next step is to allocate the funds of those who have died among the survivors, in the form of a payment called a longevity credit. The amount of longevity credit paid to a participant ℓ has the general form

$$\begin{aligned}&\text{Total account values released by deaths}\\ &\times \frac{\text{Account value of participant } \ell \times \text{Probability assigned to participant } \ell}{\sum_{\text{All surviving participants } k} \text{Account value of participant } k \times \text{Probability assigned to participant } k}.\end{aligned}$$

To assign a probability to participant ℓ, consider the implied expected longevity credit in the calculation of the expected present value of the life annuity. Suppose the life has an account value of amount $\ddot{a}_x$ units at age x. Assuming the life survives to age $x+1$, their account value will have accumulated to amount $(1+i)(\ddot{a}_x - 1)$ units at age $x+1$.

Since the life should have an amount

$$\ddot{a}_{x+1} = \sum_{k=0}^{\infty} (1+i)^{-k} {}_k p_{x+1}$$

at age $x+1$, to expect to continue to pay them 1 unit per annum annually in arrears, the implied longevity credit is

$$\ddot{a}_{x+1} - (1+i)(\ddot{a}_x - 1) = \frac{1 - {}_1p_x}{{}_1p_x}(1+i)(\ddot{a}_x - 1).$$

In standard international actuarial notation, $p_x := {}_1p_x$ and $q_x := 1 - p_x$. Thus the probability assigned to a participant ℓ at age $x+1$ years is q_x/p_x. This means that the probability used to calculate the longevity credit at age $x+1$ reflects the mortality risk assessment from age x to age $x+1$. More generally, if the longevity credit is done over the time period $(n - \frac{1}{m}, n]$ then the probability assigned to a participant ℓ when calculated at age $x+n$ is ${}_{\frac{1}{m}}q_{x+n-\frac{1}{m}} / {}_{\frac{1}{m}}p_{x+n-\frac{1}{m}}$.

For the calculations in this paper, using q_x/p_x to calculate the longevity credits will lead to a distribution of the funds released by deaths that is consistent with the income calculation, which uses the expected present value of future payments.

2.4 Longevity credit calculation

Let $G(n)$ represent the total account value at integer time n of all the participants who died in the time interval $(n-1, n]$, i.e.

$$\begin{aligned} G(n) =& (1+i) \sum_{m=0}^{n-1} \left(L^{(m,\text{poor})}(n-1) - L^{(m,\text{poor})}(n) \right) F^{(m,\text{poor})}(n-1) \\ &+ (1+i) \sum_{m=0}^{n-1} \left(L^{(m,\text{rich})}(n-1) - L^{(m,\text{rich})}(n) \right) F^{(m,\text{rich})}(n-1). \end{aligned}$$

The account values of the recently deceased are just after investment returns have been credited at time n. The longevity credit added and income deducted at time n are excluded since the recently deceased benefit from neither.

Based on the risk-sharing rule outlined above, the longevity credit awarded at integer time n to a surviving member of cohort $m \in \{0, 1, \ldots, n-1\}$ and group $k \in \{\text{poor}, \text{rich}\}$ is

$$\begin{aligned} &M^{(m,k)}(n) = G(n) \\ &\times \frac{\frac{q_{x+n-m}}{p_{x+n-m}}(1+i)F^{(m,k)}(n-1)}{\sum_{m=0}^{n-1} \frac{q_{x+n-m}}{p_{x+n-m}} L^{(m,\text{poor})}(n)(1+i)F^{(m,\text{poor})}(n-1) + \sum_{m=0}^{n-1} \frac{q_{x+n-m}}{p_{x+n-m}} L^{(m,\text{rich})}(n)(1+i)F^{(m,\text{rich})}(n-1)}. \end{aligned}$$

The amount of longevity credit $M^{(m,k)}(n)$ increases as the amount of money released to the survivors by recent deaths increases. Additionally, a participant gets a greater proportion of the money released by deaths as

- Their chance of dying increases as q_{x+n-m} increases. This follows from the term $\frac{q_{x+n-m}}{p_{x+n-m}} = \frac{q_{x+n-m}}{1-q_{x+n-m}}$ increasing in value with q_{x+n-m}; and
- Their account value at time n, $(1+i)F^{(m,k)}(n-1)$, increases.

The more money a participant stands to lose when they die and the more likely they are to die, the higher the longevity credit they receive upon survival.

The amount of longevity credit is volatile, as it depends on who has died in the pool and when they die. The question of quantifying how the volatility relates to the profile of the participants – their mortality rates and account values – is studied next.

3 Relating volatility to heterogeneity

The income paid out in a pooled annuity fund fluctuates over time. There are two sources of the fluctuations: investment returns and mortality rates. Difference in the experience of these factors from what is assumed in annuity values like $\ddot{a}_x$, which is used to calculate the income, results in variations in the income paid out to participants. Here, investment returns are fixed so that experienced mortality rates are the only source of fluctuations.

Experienced mortality rates differing from their expected values results in volatility in the income paid out to participants. The differences manifest themselves directly via the longevity credit payments. The calculation of the expected longevity credit in the annuity value is shown in Section 2.3 and the actual longevity credit is exhibited in Section 2.4. The former assumes perfect pooling and a homogenous group of participants who all have the same account values and same chance of dying. The latter's value depends on who died over the last time period in the fund and the value of their accounts.

It would appear that the more participants in the pooled annuity fund, the better in terms of reducing income volatility, through a general appeal to the Law of Large Numbers. With more participants to pool their longevity risk together, the more likely it is that the expected value of longevity credits are received in practice.

However, the situation is more nuanced than it appears at first blush, as is shown in Bernhardt and Qu (2022). If all participants have future lifetimes which can be modelled as independent and identically distributed random variables, and they bring the same amount of wealth to a pooled annuity fund, then, indeed, the more participants the better (Bernhardt and Donnelly, 2020). With increasing numbers of participants, the longevity credits paid to the surviving participants will be closer to their limiting expected value. From a risk perspective, the random fluctuations element of longevity risk is reduced as the number of members in the pooled annuity fund increases.

However, adding members who are different to the other – homogeneous – participants changes that result. Adding someone who is different in some way increases instantly the random fluctuations element of longevity risk. For example, suppose a group of 500 identical members, each with an account value of 10 000 units, experience a 2% annual volatility in their income due to longevity risk. Adding 50 people, who each have an account value of 100 000 units, causes the income volatility to increase to over 3%. In fact, this new group of 550 people, of whom 10% have ten times the account value of the others, experiences the same volatility as a group of 185 identical people.

On the other hand, an open pooled annuity fund also results in an inhomogeneous membership. When there is a constant flow of new members, joining at the same age, then the results of Donnelly (2022) suggest that the level of income volatility is fairly constant across cohorts as long as the fund stays open. It is only the last cohorts to join who experience an increasing level of income volatility. The earlier-joining cohorts die off, leaving the last cohorts pooling longevity risk with an ever-decreasing number of members.

The main questions examined in this paper are as follows.

- How does heterogeneity among the participants affect the income paid out?
- What are the main drivers of the heterogeneity which affect the income paid out?
- Can a measure of heterogeneity predict future income volatility?

To study the questions, both a measure of heterogeneity and a measure of volatility of the income need to be defined. Begin by considering what has been studied in the academic literature on the measure of heterogeneity within a pooled annuity fund.

3.1 Measure of heterogeneity

A measure of heterogeneity calculates a value for how ‘spread out’ are the participants, in some sense. Donnelly et al. (2014) propose a measure that arises naturally from considering the volatility of account values which arises from longevity pooling. As their setting is in continuous time, an adjustment is required to their proposed measure to make it suitable for the discrete-time setting in this paper. Specifically, while a mortality rate is used in the continuous time setting, this is not appropriate in a discrete-time setting.

Define a measure of heterogeneity at time n for each member of cohort $g \in \{0, 1, \ldots, \min\{n, N-1\}\}$ as

$$H^{(g)}(n) := \frac{q_{x+n-g}}{p_{x+n-g}} \frac{\sum_{m=0}^{\min\{n,N-1\}} \sum_{k\in\{\text{poor},\text{rich}\}} L^{(m,k)}(n) \left(F^{(m,k)}(n)\right)^2 \frac{q_{x+n-m}}{p_{x+n-m}}}{\left(\sum_{m=0}^{\min\{n,N-1\}} \sum_{k\in\{\text{poor},\text{rich}\}} L^{(m,k)}(n)\, F^{(m,k)}(n)\, \frac{q_{x+n-m}}{p_{x+n-m}}\right)^2} \tag{2}$$

at time $n = 0, 1, 2, \ldots$. The higher the value of $H(n)$, the more heterogeneous is the membership and the higher the volatility in the longevity credits. It quantifies the spread of wealth among the participants, adjusted by their chance of dying. The measure is derived from the volatility of a participant’s account value that arises from longevity risk-sharing. Donnelly et al. (2014) simplify the expression but do not investigate the usefulness of the measure.

The expression $H(n)$ proposed here as a measure of heterogeneity has deliberate analogies with the form of the longevity credit. In particular, the ratio $\frac{q_{x+n-m}}{p_{x+n-m}}$, which replaces the mortality rate used in the original version in Donnelly et al. (2014) .

Versions of the heterogeneity measure emerges naturally in the study of Bernhardt and Donnelly (2020) and Bernhardt and Qu (2022), through the study of a measure of income volatility. It is extremely simplified in the setting of Bernhardt and Donnelly (2020), since their setting is a single cohort fund in which all participants have the same initial account value and future distribution of deaths. The study in Bernhardt and Qu (2022) is discussed in more detail below.

In the numerical study in Section 4, only the value of $H^{(g)}(0)$ is considered. Examining the evolution of $\{H^{(g)}(n) : n = 0, 1, 2, \ldots\}$ in, for example, a single cohort fund shows that it is extremely small for most of the surviving cohort’s lifetime. It begins increasing to very high levels when there are few members left alive. At the same time, income volatility climbs up too.

However, as shown in Bernhardt and Qu (2022), the value of $H^{(g)}(0)$ can be used to predict the future income volatility, at least in a single cohort setting. The study here looks at more membership profiles than in Bernhardt and Qu (2022), to see if the result continues to hold.

Example 3.1. *To understand the measure $H(n)$ better, consider a single cohort fund ($N = 1$) in which there all members are independent and identical copies of each other. Initially, there are $\ell(0) > 0$ participants in the fund. At time $n = 1, 2, \ldots$, there are $\ell(n)$ surviving participants at time n, who each have account value $F(n)$. The heterogeneity measure at time $n = 0, 1, 2, \ldots$,*

$$H(n) = \frac{1}{\ell(n)}.$$

Since all participants bring the same account value, there is no explicit dependence on the account values in $H(n)$. Similarly for the chance of dying.

First consider the behaviour of $H(0)$ for funds with different single cohorts. The higher the number of participants $\ell(0)$, the lower the heterogeneity measure. With more people with whom to pool longevity risk, the stream of longevity credits is more stable. The stream of longevity credits are more volatile due to volatility in the number of deaths.

Next consider the behaviour of $H(n)$ for a given single cohort fund over time. The value of $H(n)$ increases in the single cohort fund as time n increases. As time goes on, fewer of the original cohort survive and there are fewer people in the fund to pool their longevity risk together.

3.2 A measure of income volatility

Next turn to defining a measure of income volatility. It is assumed here that participants want their income to be close to their initial income. Reductions in income below the initial income are undesirable. It is assumed that participants are not concerned by increases in their income above the initial income.

The number of years for which the income stays above a specified lower bound can be used to assess this objective. The lower bound is a specified fraction of the initial income paid out to a particular group of participants. The income must stay above the lower bound in a specified number of future scenarios. This definition is proposed in Bernhardt and Donnelly (2020) and studied further in Bernhardt and Qu (2022). In both these papers, a mathematical relationship is established between a measure of heterogeneity among the participants and the measure of income volatility.

Suppose the number of years for which the income paid to participants stays above α times the initial income, where $\alpha \in (0, 1)$, is tracked. For example, $\alpha = 0.95$ means that the income stays above 95% of the initial income. This could be done by projecting the income in future scenarios. From the projections, the highest number of years for which the income is stable can be calculated for each future scenario. For example, in one scenario, the income may stay above 95% of the initial income for 18 years, before falling below it in hte 19th year. In another, the income may stay above 95% of the initial income for 23 years, before falling below it in the 24th year.

The maximum number of years for which the income stays above the specified lower bound across 100β% of the future scenarios, for $\beta \in (0, 1]$, can be calculated from these projections. For example, suppose $\beta = 0.9$. In each of the 100 future scenarios, the highest number of years for which the income is above 95% of the initial income has been calculated. In 90 of these scenarios, the income stays above the lower bound for 21 years or more. In the remaining 10 scenarios, the income stays above the lower bound for less than 21 years. In this simple example, the income is *stable* for a maximum of 20 years, for $\alpha = 0.95$ and $\beta = 0.9$.

Generalising the definition in Bernhardt and Donnelly (2020) to the multi-cohort setting, the income paid to cohort m and group $k \in \{\text{poor}, \text{rich}\}$ is stable until integer time $t^{(m,k)}$ if

$$\mathbb{P}\left[C^{(m,k)}(n) \geq \alpha C^{(m,k)}(m), \text{ for } n = m, m+1, \ldots, t^{(m,k)}\right] \geq \beta.$$

The amount $C^{(m,k)}(m)$ is the initial amount of income paid to participants of cohort m and group k. The event $\left[C^{(m,k)}(n) \geq \alpha C^{(m,k)}(m), \text{ for } n = m, m+1, \ldots, t\right]$ is the set of future states of the world in which the income paid to these participants stays above α times the initial income for $t - m$ years. For example, $\alpha = 0.95$ means that in these future states of the world, the income paid to those participants never falls below 95% of the initial income in the first $t - m$ years of payments to them. The entire expression means that the event should hold in at least the fraction β of all future states of the world.

For a single cohort fund in which all participants have the same future distribution of deaths, the income is stable for the same number of years, regardless of their individual account values. Since they have the same mortality distribution, their account values remain locked in the same ratios over time. However, once participants have different mortality distributions, their income is stable for a different number of years. This may be occur in the setting of either an open fund or a single cohort fund.

Having introduced the mathematical notation to study heterogeneity and income stability, the results of a numerical study are presented next. It is seen that 'the more, the merrier'. In general, higher numbers of participants outweigh the disadvantages of heterogeneity.

4 Numerical study

The numerical study begins with a single cohort fund, first examining the results for a homogeneous membership before moving onto one in which members have different account values. These settings have been studied in Bernhardt and Donnelly (2020) and Bernhardt and Qu (2022). However, unlike those papers, the focus here is on the value of the heterogeneity measure and how it can be used to determine the number of years for which the income is stable.

The closing of a multi-cohort fund is also studied. While the fund remains open, income volatility remains constant over successive cohorts to join the fund. New entrants keep the income volatility low as is studied in Donnelly (2022). Once the fund closes to new members, the level of heterogeneity in the fund will increase over time. As shown in Donnelly (2022), the last members to join the fund bear the highest levels of income volatility. When the last cohorts join, there is – in general – sufficient pooling of longevity risk in the fund. However, as they reach old age, the earlier generations have died off. The last cohorts are left pooling own longevity risk with mostly themselves. The amount and frequency of longevity credits increases for the longest-live of the last cohorts. Hence the volatility of their account values and income increases, too.

For this reason, the study of a multi-cohort fund is restricted to the time at which the fund closes to new members. The closure situation can analysed using a single cohort fund, in which participants in a single cohort are of different ages and have different account values. The ages, numbers and account values are chosen to mimic their expected values in a multi-cohort fund which had just closed to new members.

4.1 Single cohort fund with homogeneous participants

Consider a single cohort fund in which all participants have the same account value and are all age x. All participants have the same distribution of deaths.

The focus is on how the measure of heterogeneity at time 0 relates to the maximal number of years for which the income is stable. From Example 3.1, the measure of heterogeneity in a single cohort fund in which all participants are independent and identical copies of each other is

$$H(0) = \frac{1}{\ell(0)}.$$

A relationship between $H(0)$ and the number of years T for which the income is stable is proved in Bernhardt and Qu (2022). Recall that α is the fraction of the initial income above which the income should stay, to remain stable. The value β determines the proportion of future scenarios in which the income is stable for a given number of years.

Bernhardt and Qu (2022) show that a close approximation T to the maximal number of years for which the income is stable satisfies

$$
{}_Tq_x = \left(1 + H(0)\left(\frac{1-\alpha}{\alpha}\right)^2 \left(\Phi^{-1}\left(\frac{1-\beta}{2}\right)\right)^2\right)^{-1}, \tag{3}
$$

in which ${}_Tq_x$ is the chance of someone age x years dying within the next T years and Φ^{-1} is the inverse of the standard normal distribution function. The approximation holds for any single cohort fund in which all members are the same age and have the same distribution of deaths.

The idea is to evaluate the right-hand side of the expression, which can be done once the values of α and β are chosen and using the composition of the membership at time 0 to calculate $H(0)$. The value obtained represents the proportion of the membership who obtain a stable income for T years, given that the members join the fund at age x. Additionally, the value obtained is independent of the chosen mortality distribution of the participants. Then, having chosen a suitable mortality distribution, the value of T can be calculated by finding the T which gives the same value of ${}_Tq_x$ as the right-hand side of the expression. Through a numerical study, Bernhardt and Qu (2022) show that the expression results in a very close approximation to T.

The value of the approximation (3) is two-fold. First, it avoids having to do numerical simulations of the future income in order to determine T. Second, it shows that it is the measure of heterogeneity $H(0)$ that determines T. Two different membership profiles may look different in terms of the numbers of members and account values. But if they have the same value of $H(0)$ then they will have the same number of years for which the income is stable, as long as the participants have the same distribution of deaths and there is only one cohort in the fund. Bernhardt and Qu (2022) do not study directly the value of $H(0)$ compared to the number of years for which the income is stable. Rather, it is used to determine when it is beneficial for a heterogeneous group to form one single fund or else smaller, more homogeneous, funds.

To study the value of $H(0)$ compared to the number of years for which the income is stable, choose

$$
x = 65, \qquad \beta = 0.9, \qquad \text{and the life table is S1PMA.}
$$

Table 4.1.1 shows the maximal number of years for which the income is stable, under three different lower bounds, when the approximation (3) is applied in the single cohort setting with all members being independent and identical copies of each other. For example, when there are initially 100 members in the fund, all age 65, the income does not fall below 90% ($\alpha = 0.9$) of the initial income for at least 15.9 years, in 90% of future scenarios.

The values in Table 4.1.1 are the benchmark values, for a single cohort fund in which all members are the same age. How do they change when heterogeneity is introduced?

4.2 Single cohort fund with two groups

Consider another fund in which there is again a single cohort, but two groups within that cohort. At time 0, there are $\ell^{\text{poor}}(0)$ poor participants, who each have account value $F^{\text{poor}}(0)$, and $\ell^{\text{rich}}(0)$ rich participants, who each have account value $F^{\text{rich}}(0)$.

Since all participants are assumed to have the same distribution of deaths and join at the same age, the measure of heterogeneity (2) at time 0 reduces to

$$
H(0) := \frac{\ell^{\text{poor}}(0)\,\left(F^{\text{poor}}(0)\right)^2 + \ell^{\text{rich}}(0)\,\left(F^{\text{rich}}(0)\right)^2}{\left(\ell^{\text{poor}}(0)\,F^{\text{poor}}(0) + \ell^{\text{rich}}(0)\,F^{\text{rich}}(0)\right)^2}
$$

$\ell(0)$	$H(0)$	T for $\alpha = 0.8$	T for $\alpha = 0.9$	T for $\alpha = 0.95$
100	0.01000	23.2	15.9	7.7
200	0.00500	26.1	18.1	10.4
500	0.00200	30.7	23.2	15.3
1 000	0.00100	32.6	26.1	19.4
2 000	0.00050	34.7	29.5	23.8
5 000	0.00020	36.6	32.6	27.7
10 000	0.00010	37.2	34.8	29.1

Table 4.1.1: Maximal number of years for which the income is stable, in a single cohort fund in which all participants have the same distribution of deaths, are the same age and have the same account value. The level of certainty is $\beta = 0.9$.

First fix the total number of participants who join the single cohort fund, so that $\ell^{\text{poor}}(0) + \ell^{\text{rich}}(0) = 1\,000$. The values of $H(0)$ are calculated in this two-group setting, to see how they relate to the one-group setting. This particular membership profile is studied in Bernhardt and Qu (2022) but they have a different goal; to see if it is better for one of the groups to not pool their longevity risk with the other group.

Suppose that some of the 1 000 initial participants each bring 100 units to the fund; this is the poor group. The remainder of the 1 000 participants form the rich group. They each bring twice the amount of money as the poor group's members, each bringing 200 units. All participants are the same age when they join and have the same distribution of deaths. Table 4.2.1 shows the maximal number of years for which the income is stable, when the certainty level is 90%.

Broadly, the heterogeneity measure is fairly similar across the considered groups. Consequently the maximal number of years for which the income is stable are also fairly close in value. For example, the income is above 90% of the initial income for over 26 years in 90% of future scenarios, for all the listed groups in Table 4.2.1. The maximal number of years for which the income is stable, increases only slightly as the number of poor members increase. Effectively, even though the rich members are twice as rich as the poor members, the level of heterogeneity increases only slightly.

The value of $H(0)$ can be easily calculated at the time the single cohort joins the fund. Yet it can be used to determine the maximal number of years for which the income is stable. It is seen from Table 4.2.1 that, since there are only very small changes in the value of $H(0)$, the maximal number of years is very similar across the different memberships.

As the relative wealth of the rich members increases, for a fixed number of rich members, the heterogeneity measure $H(0)$ increases (Figure 4.2.1a). While this means that income becomes more volatile, the decrease in the number of years for which the income is stable is not as dramatic as the lines in Figure 4.2.1a may suggest. For example, consider a fund in which 400 of the 1 000 members are in the rich group, but the rich group members have the same account value as the poor members. In that case, the income stays above 90% of the initial income in 90% of scenarios for 26.1 years (since $H(0) = 0.001$ for this membership, the number of years can be read off Table 4.1.1). If the members of the rich group each had 12 times the account value of each poor members – justifying their label of being rich – then the heterogeneity measure doubles to $H(0) = 0.002$. However, the number of years for which the income is stable falls to 23.2 years (again, found by reading off Table 4.1.1 at the row for which $H(0) = 0.002$), a less dramatic change.

Number of poor $\ell^{\text{poor}}(0)$	Number of rich $\ell^{\text{rich}}(0)$	$H(0)$	T for $\alpha = 0.8$	T for $\alpha = 0.9$	T for $\alpha = 0.95$
0	1 000	0.00100	32.6	26.1	19.4
250	750	0.00106	32.7	26.4	19.7
500	500	0.00111	32.9	26.6	20.0
750	250	0.00112	32.9	26.6	20.0
1 000	0	0.00100	32.6	26.1	19.4

Table 4.2.1: Maximal number of years for which the income is stable, in a single cohort fund in which there is a poor group and a rich group. Each member of the rich group has twice the account value of each member of the poor group. All participants have the same distribution of deaths and are the same age. The level of certainty is $\beta = 0.9$.

The value $H(0) = 0.002$ is the same as for a homogeneous fund in which there are 500 identical members (see Table 4.1.1). In this example, the 600 poor members would have a lower heterogeneity value $H(0)$ if they formed their own pooled fund, excluding the 200 rich participants. Bernhardt and Qu (2022) focus their analysis on these type of cases. Adding heterogeneous members may, if there are not enough heterogeneous members added, increases the measure of heterogeneity $H(0)$. They develop a criterion to decide when it may be in the interests of a group not to pool their longevity risk with others.

The heterogeneity measure $H(0)$ appears to fall as the number of rich members increases in the fund in which the total number of members is fixed, as demonstrated by the relative position of the lines in Figure 4.2.1a). However, this is not entirely true as Figure 4.2.1b shows.

Figure 4.2.1b, a similar chart to Bernhardt and Qu (2022, Figure 3), shows that the heterogeneity measure $H(0)$ peaks when 10%-20% of the fund membership is rich. It is observed that the higher the account value of the rich group relative to the poor group,

- the greater the heterogeneity value $H(0)$ and hence the shorter the time for which the income is stable, and
- the smaller the number of rich members needed to maximize the heterogeneity.

Figure 4.2.1b also suggests that, to have significant changes in the value of the heterogeneity measure, requires extreme heterogeneity among the membership. Additionally, even if the heterogeneity measure doubles from $H(0) = 0.001$ to $H(0) = 0.002$, the number of years for which the income is stable in 90% of scenarios reduces by 3 to 4 years, depending on the value of α.

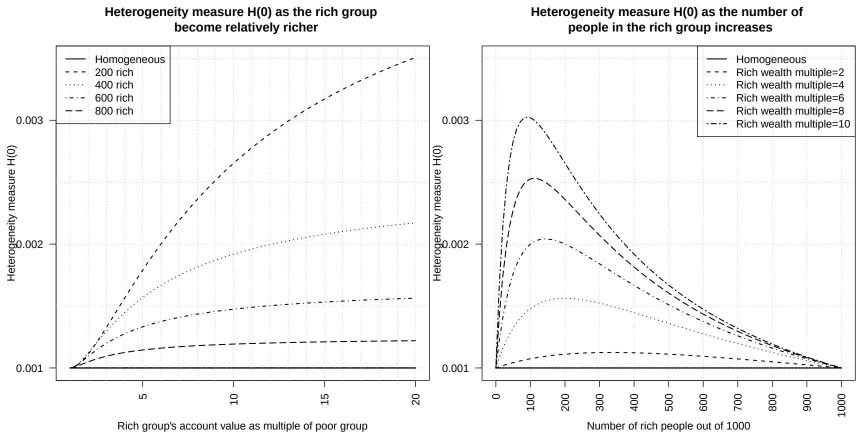

(a) Increasing relative wealth. (b) Increasing number of rich members.

Figure 4.2.1: Heterogeneity measure $H(0)$ in a single cohort fund in which there are a total of 1 000 members. These 1 000 members are divided into a poor group and a rich group. Figure 4.2.1a shows how $H(0)$ varies as the relative wealth of the rich group increases, for a given number of rich people in the fund. Figure 4.2.1b shows how $H(0)$ varies as the number of members in the rich group increases (and hence as the number of poor members decreases), for a given relative wealth multiple.

4.3 Closure of a pooled annuity fund

Now consider the situation of an open pooled fund that has just closed to new entrants. Suppose that, up to now, a new cohort consisting of a constant number of homogeneous participants had joined each year. Until the fund closed, it would have had a constant level of heterogeneity across its cohorts, more or less. Interpreting the results of Donnelly (2022), the first cohorts to join the fund would have higher levels of heterogeneity as there are few people in the fund. However, as the fund become increasingly mature, with new members joining every year, the level of heterogeneity would decrease.

The position of the fund at the point of closure can be represented by single cohort fund with an appropriately chosen membership. It is assumed here that the single cohort consists of 30 sub-groups. The members of each sub-group have the same age and account value. Each sub-group can be identified by the age of its constituents, which ranges from age 65 to age 94. The account value of each member is commensurate with their age, being calculated as $100\ddot{a}_x$, if their age is x years at time 0. There are 100 members age 65, $\lfloor 100p_{65} \rfloor$ members age 66, $\lfloor 100{}_2p_{65} \rfloor$ members age 67, and so on. In total, there are 1823 members in this single cohort fund at time 0 when the life table S1PMA is used.

For such a fund, the measure of heterogeneity given in equation (2) is adapted slightly. With $K = 30$ subgroups, let $x(k)$ represent the age of members of group k. In group k at time 0, there are $\ell^{(k)}(0)$ members, who each join with account value $F^{(k)}(0)$. Then the measure of heterogeneity applying to each member of group k at time 0 is

$$H^{(k)}(0) := \frac{q_{x(k)}}{p_{x(k)}} \frac{\sum_{m=0}^{K} \ell^{(m)}(0) \, \left(F^{(m)}(0)\right)^2 \frac{q_{x(m)}}{p_{x(m)}}}{\left(\sum_{m=1}^{K} \ell^{(m)}(0) \, F^{(m)}(0) \, \frac{q_{x(m)}}{p_{x(m)}}\right)^2}.$$

Note from the expression for the measure of heterogeneity that $H^{(k)}(0) \times \frac{p_{x(k)}}{q_{x(k)}}$ is constant for all

Group k	Age	Number in group	Account value	$H^{(k)}(0)$	T for $\alpha = 0.8$	T for $\alpha = 0.9$	T for $\alpha = 0.95$
1	65	100	2946	0.000149	31	26	21
2	66	98	2767	0.000166	30	26	21
3	67	97	2597	0.000186	29	25	20
4	68	97	2435	0.000186	29	24	19
5	69	97	2281	0.000186	28	24	19
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
29	93	11	401	0.003862	6	5	2
30	94	8	375	0.004312	5	4	2

Table 4.3.1: Maximal number of years for which the income is stable, in a single cohort fund in which participants have different ages and account values. However, all participants have the same distribution of deaths at each age. Only a selection of the membership profile is shown, as it extends down to age 94. The level of certainty is $\beta = 0.9$.

k. This means that the value $H^{(k)}(0) \times \frac{p_{x(k)}}{q_{x(k)}}$ is a measure of heterogeneity across a particular membership.

The benchmark membership profile

The question is: does the measure of heterogeneity correspond to the maximal number of years for which the income is stable? The results of Bernhardt and Qu (2022) cannot be applied here. In fact, using wrongly their approximation results in a large over-estimation of the maximal number of years for which the income is stable.

Instead, the maximal number of years must be calculated using a simulation. This was done using the statistical software package *R*. As the simulation was done on an annual basis, the maximal number of years calculated from the simulations are integers.

The maximal number of years for this particular fund, in which all members of the same age have the same account value, is shown in Table 4.3.1. It is observed that, as the heterogeneity measure increases, the maximal number of years decreases. The lower maximal number of years is due to the shorter future lifetimes of the older members. It is also calculated that $H^{(k)}(0) \times \frac{p_{x(k)}}{q_{x(k)}} = 0.013065$ for this particular membership.

This membership profile is the benchmark one. A benchmark is require since there is no approximation on which to calculate the maximal number of years for which the income is stable.

Heterogeneity within an age group

Now consider adding in further heterogeneity. Suppose that the age groups are further divided into rich and poor, with the rich subgroup four times as wealthy as the corresponding poor subgroup. Additionally, there are three times as many rich people in each subgroup as the number of poor people. In total, there are 1 802 members in the studied fund, comparable to the first fund considered which had only homogeneous age groups.

The results obtained earlier suggest that adding in wealth heterogeneity should reduce the maximal number of years for which the income is stable by a couple of years. In particular, the study of a single cohort fund in which there is only a rich group and a poor group, and all members are the same age, shows that the measure of heterogeneity increases when rich and poor are mixed

Group k	Age	Number of rich in group	Number of poor in group	$H^{(k)}(0)$	T for $\alpha = 0.8$	T for $\alpha = 0.9$	T for $\alpha = 0.95$
1	65	75	25	0.000175	30/29	26/25	21/21
2	66	73	24	0.000195	29/28	25/25	20/20
3	67	72	24	0.000219	29/27	24/24	19/19
4	68	72	24	0.000245	28/27	24/24	19/19
5	69	69	23	0.000275	27/25	23/23	18/18
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
29	93	8	2	0.004552	5/2	4/2	2/1
30	94	6	2	0.005082	4/2	3/2	2/1

Table 4.3.2: Maximal number of years for which the income is stable, in a single cohort fund in which participants have different ages and account values. Each age group is divided into a rich and poor subgroup, with the rich members having four times the account value of the poor. However, all participants have the same distribution of deaths. The maximal number of years for which the income is stable is shown as the values for rich/poor, for each value of α. Only a selection of the membership profile is shown, as it extends down to age 94. The level of certainty is $\beta = 0.9$.

(Figure 4.2.1). However, there was only a modest decrease in the maximal number of years for which the income is stable, of around 0.5 years.

In this fund of whom about three-quarters of the membership have four times the account value of the remaining one-quarter, the value of $H^{(k)}(0) \times \frac{p_{x(k)}}{q_{x(k)}} = 0.015398$, for all k. This is a higher value than for the fund in which all members of the same age have the same amount of wealth (0.013065). Since this fund with many rich members is more heterogeneous, it would be expected that the maximal number of years for which the income is stable is lower. The numerical simulations verify this (Figure 4.3.2).

Allowing the number of rich members in each age subgroup to decline, similar results are obtained. For example, when half the members are rich and half are poor, the fund becomes even more heterogeneous. This mirrors the results in the single cohort funds studied, in which all members had the same age (for example, see Figure 4.2.1).

In this fund in which half the members are rich and half are poor, $H^{(k)}(0) \times \frac{p_{x(k)}}{q_{x(k)}} = 0.0180486$, for all k). As this is a higher measure of heterogeneity compared to the previous two funds, it is expected that the maximal times for which the income is stable are lower. Again, the numerical simulations verify this. For example, the maximal times are mostly up to one year lower than the values in Table 4.3.2.

Decreasing the proportion of rich members to be one quarter of the membership, it is calculated that $H^{(k)}(0) \times \frac{p_{x(k)}}{q_{x(k)}} = 0.020720$, for all k. The numerical simulations again verify that the maximal times for which the income is stable decline with this higher level of heterogeneity, being around 1 year less than the values in Table 4.3.2.

However, despite the increase in heterogeneity, the decline is not large in terms of the maximal number of years for which the income is stable. The impact of mixing rich members and poor members does not appear significant.

5 Conclusion

Members of a pooled annuity fund will differ in many ways, from their ages to their account values. For the manager of a pooled annuity fund, the question arises as to: do the differences matter? A measure of income volatility is used to study these differences. The number of years for which the income stays above some specified fraction of the initial income is used as the measure.

It seems to matter less on what each members brings to the fund in terms of account values. It is only when there is a small subgroup of extremely wealthy members that there is a large impact on income stability for all members. This can be dealt with by, for example, not allowing them to join or capping the amount of money they can bring to the fund. However, this should only be used in the most extreme cases of heterogeneity among the membership.

With moderate levels of heterogeneity, as would be the case in most funds, there is a modest reduction in the number of years for which the income is stable. It is likely that, with the incorporate of random investment returns, the reduction would contribute a relatively small part to income volatility.

Heterogeneity among the membership of a fund does not have a very large impact on the income stability. It does have some impact, but this should be small for typical membership profiles.

References

Bernhardt, T. and Donnelly, C. (2019). Modern tontine with bequest: Innovation in pooled annuity products. *Insurance: Mathematics and Economics*, 86:168–188.

Bernhardt, T. and Donnelly, C. (2020). Quantifying the trade-off between income stability and the number of members in a pooled annuity fund. *ASTIN Bulletin*, 51:101–130.

Bernhardt, T. and Donnelly, C. (2021). Staying the course: how pooled annuity funds are proving an attractive alternative. *The Actuary*. Published online on Wednesday 5th May 2021, https://www.theactuary.com/2021/04/30/staying-course-how-pooled-annuity-funds-are-proving-attractive-alternative.

Bernhardt, T. and Qu, G. (2022). Wealth heterogeneity in a closed pooled annuity fund. https://doi.org/10.48550/arXiv.2110.13467. Accessed 19 March 2023.

Chen, A., Guillen, M., and Rach, M. (2021). Fees in tontines. *Insurance: Mathematics and Economics*, 100:89–106.

Chen, A., Hieber, P., and Klein, J. (2018). Tonuity: a novel individual-oriented retirement plan. *ASTIN Bulletin*, 49(1):5–30.

Chen, A. and Rach, M. (2022). Bequest-embedded annuities and tontines. *Asia-Pacific Journal of Risk and Insurance*, 16(1):1–46.

Chen, A., Rach, M., and Sehner, T. (2020). On the optimal combination of annuities and tontines. *ASTIN Bulletin*, 50(1):95–129.

Dagpunar, J. (2021). Closed-form solutions for an explicit modern ideal tontine with bequest motive. *Insurance: Mathematics and Economics*, 100:261–273.

Denuit, M. and Robert, C. (2021). From risk sharing to pure premium for a large number of heterogeneous losses. *Insurance: Mathematics and Economics*, 96:116–126.

Donnelly, C. (2022). Practical analysis of a pooled annuity fund with integrated bequest. Technical report, Heriot-Watt University. Research commissioned by Institute and Faculty of Actuaries' Actuarial Research Centre. `https://www.actuaries.org.uk/documents/practical-analysis-pooled-annuity-fund-integrated-bequest`.

Donnelly, C., Guillén, M., and Nielsen, J. (2013). Exchanging uncertain mortality for a cost. *Insurance: Mathematics and Economics*, 52(1):65–76.

Donnelly, C., Guillén, M., and Nielsen, J. (2014). Bringing cost transparency to the life annuity market. *Insurance: Mathematics and Economics*, 56:14–27.

Hanewald, K., Piggott, J., and Sherris, M. (2013). Individual post-retirement longevity risk management under systematic mortality risk. *Insurance: Mathematics and Economics*, 52(1):87–97.

Milevsky, M. and Salisbury, T. (2015). Optimal retirement income tontines. *Insurance: Mathematics and Economics*, 64:91–105.

Milevsky, M. A. and Salisbury, T. S. (2016). Equitable retirement income tontines: Mixing cohorts without discriminating. *ASTIN Bulletin*, 46(3):571–604.

Piggott, J., Valdez, E., and Detzel, B. (2005). The simple analytics of a pooled annuity fund. *The Journal of Risk and Insurance*, 72(3):497–520.

Qiao, C. and Sherris, M. (2013). Managing systematic mortality risk with group self-pooling and annuitization schemes. *Journal of Risk and Insurance*, 80(4):949–974.

Sabin, M. (2010). Fair tontine annuity. Available at SSRN. `http://ssrn.com/abstract=1579932`.

Stamos, M. Z. (2008). Optimal consumption and portfolio choice for pooled annuity funds. *Insurance: Mathematics and Economics*, 43:56–68.

Valdez, E. A., Piggott, J., and Wang, L. (2006). Demand and adverse selection in a pooled annuity fund. *Insurance: Mathematics and Economics*, 39(2):251–266.

Weinert, J. and Gründl, H. (2021). The modern tontine. *European Actuarial Journal*, 11(2):49–86.

Zhou, Y., Sherris, S., Ziveyi, J., and Xu, M. (2021). An innovative design of flexible, bequest-enhanced life annuity with natural hedging. *Scandinavian Actuarial Journal*, 0(0):1–22.

Actuarial Research Centre
Institute and Faculty of Actuaries

Beijing
14F China World Office 1 · 1 Jianwai Avenue · Beijing · China 100004
Tel: +86 (10) 6535 0248

Edinburgh
Level 2 · Exchange Crescent · 7 Conference Square · Edinburgh · EH3 8RA
Tel: +44 (0) 131 240 1300 · Fax: +44 (0) 131 240 1313

Hong Kong
1803 Tower One · Lippo Centre · 89 Queensway · Hong Kong
Tel: +852 2147 9418

London (registered office)
7th Floor · Holborn Gate · 326-330 High Holborn · London · WC1V 7PP
Tel: +44 (0) 20 7632 2100 · Fax: +44 (0) 20 7632 2111

Oxford
1st Floor · Park Central · 40/41 Park End Street · Oxford · OX1 1JD
Tel: +44 (0) 1865 268 200 · Fax: +44 (0) 1865 268 211

Singapore
163 Tras Street · #07-05 Lian Huat Building · Singapore 079024
Tel: +65 6906 0889

www.actuaries.org.uk