

Product options for enhanced retirement income

C. Donnelly* and J. Young

Abstract

A new stream of research proposes how people can increase their income in retirement by pooling their mortality risk. How one of these mortality risk-sharing rules could be implemented in practice, as part of a retirement income scheme, is considered. A potential advantage of the scheme is that a retiree's housing wealth can be monetised to provide an income stream. This would mean that retirees can continue living in their home, without needing to downsize. It may be most attractive to the millions of single pensioners, particularly those who are “asset-rich and cash-poor”. Other types of assets that could be included and how to mitigate selection risks are assessed. A way of smoothing the raw mortality credits in order to make the scheme more appealing to potential members is proposed. An illustrative premium calculation suggests that the cost of the smoothing is very small compared to the potential attractiveness of an enhanced, smoothed income.

Keywords

Pensions; Lifetime Savings; Mortality; Mutual Risk-Sharing

1. Introduction

The problems facing ageing individuals and societies as the baby boom generation reaches retirement are well known and well documented:

- The United Kingdom has a savings gap that has grown worse over time. Based on 2010 data, Aviva & Deloitte (2010) estimate that there is €379 billion a year difference between the pension provision that people retiring between 2011 and 2051 in the United Kingdom will need, and the pension amount that they can expect to receive. This is further shown by the UK's savings ratio which fell from over 10% in 1992 to <2% in 2008 before rising to around 7% in 2012 (Bunn *et al.*, 2012, chart 1).

Looking at pension scheme membership, the proportion of employees in workplace pension schemes fell from 55% in 1997 to 46% in 2012, with defined-benefit pension scheme membership falling from 46% to 28% in the same period (Office for National Statistics, 2013, page 11). The new mandatory automatic enrolment of UK employees into a workplace pension scheme may help improve pension savings, but it may be too little or too late for people who are approaching or in retirement.

- The Government will increasingly struggle to fund State pensions as the dependency ratio, that is, the number of people over the State Pension Age for every 1,000 people of working age, grows. The long-term average dependency ratio from the mid-1970s to 2006 has been around 300 people

*Correspondence to: Catherine Donnelly, Department of Actuarial Mathematics and Statistics, and the Maxwell Institute for Mathematical Sciences, Heriot-Watt University, Edinburgh EH14 4AS, UK. E-mail: C.Donnelly@hw.ac.uk

over the State Pension Age for every 1,000 people of working age (Office for National Statistics, 2012, page 4). With the increases to the State Retirement Age, the dependency ratio is projected to rise to 342 (Office for National Statistics, 2012, page 4). However, the increase in State Pension Age ignores the challenge of older people finding work or having alternative sources of income to assist them financially.

- With low interest rates, the cost of securing a lifelong retirement income is higher. Annuity rates¹ have fallen from 11% in 1994 to 7% in 2007 (Cannon & Tonks, 2009, figure 5.3). In January 2017, annuity rates were below 6%.² Further, Solvency II may have had an impact on annuity rates.³
- Longer life expectancies has further increased the cost of providing a pension income for life (Cannon & Tonks, 2009).
- The 2014 Budget gave greater freedom for people to not invest their pension savings in annuities. Product design has not yet caught up with this opportunity.
- Many pensioners are asset-rich and income-poor. As an indicative example, based on the English Longitudinal Study of Ageing (ELSA) 2002/2003 wave, there were over 1 million pensioner households⁴ who had incomes below £800 per month but lived in a house worth over £100,000.⁵

These problems have been exacerbated by a number of myths.

1.1. Myth 1: “My house is my pension”

1.1.1. Reality

People still need somewhere to live, which makes it hard to release the “value” in a house.

Over 63% Of single pensioners own their home (around 2.7 million people), with the figure increasing to 86% for pensioner couples (around 4 million couples) (Department for Work and Pensions, 2014, table 1.1; Office for National Statistics, 2014, figure 3.12). The median net property wealth of single pensioners is £165,000, and the comparable value for pensioner couples is over £210,000 (Office for National Statistics, 2014, figure 3.13).

Downsizing, that is, moving to a smaller house of lesser value, in order to generate savings or investment would be considered by 41% of over 55-year olds (Prudential, 2014).

Equity release schemes generate an upfront capital sum which could be used to provide a retirement income. In practice, it is used primarily to pay off debts, pay for home improvements and to fund a holiday (Adams & James, 2009, table 4). Of the people with housing wealth of < £500,000 who are

¹ For a 65-year-old male based on a single life level annuity with no guarantee period.

² Financial Times Annuity Rates Table, accessed on 11 January 2017. http://www.ft.com/personal-finance/annuity-table?ft_site=falcon&desktop=true

³ <http://www.actuarialpost.co.uk/article/tamsin-abbey-looks-at-the-solvency-ii-annuity-battle-4670.htm>

⁴ A pensioner household consists of either a single pensioner or a couple, of whom at least one person is a pensioner. A pensioner is someone who is older than their State Pension Age.

⁵ Figures from the ELSA data in Sodha (2005, table 4) extrapolated to the 6,850,000 pensioner households in Britain in 2002/03 (Balchin & Shah, 2004, page 8, paragraph 2). Figures are 2002/2003 values; based on the House Price Index for England and Wales, house prices have increased by ~140% from December 2002 to December 2013, and there were 8,750,000 pensioner households in 2012/2013 (Department for Work and Pensions, 2014, table 1.1).

in an equity release scheme, only 1% use it to provide a retirement income (Adams & James, 2009, table 4). In 2005, <1% of pensioner housing wealth was tied up in equity release schemes (Adams & James, 2009, page 21). However, there has recently been a growth in new products with greater flexibility (Equity Release Council, 2017, pages 4–5), which suggests an increasing level of housing wealth in equity release schemes.

1.1.2. Problem

Houses are illiquid assets for which it is hard to generate income while you still live in them. Barriers to downsizing are an emotional attachment to the individual's existing home, the cost of downsizing (e.g. stamp duty and removal expenses) and the inconvenience of moving home (Adams & James, 2009, page 21).

Equity release schemes tend to give a low upfront loan against the house that compounds rapidly. For example, in July 2017, the average equity release rate was 5.30% per annum (Equity Release Council, 2017, page 5) and, to provide some context, the average standard variable rate for a mortgage was around 2% per annum less.⁶

1.1.3. Solution

A scheme designed on the actuarially fair and transparent principles of mortality cross-subsidy can monetise housing wealth and help generate additional retirement income. The value of the house in the scheme may be more readily understood by the participants, as it would be expressed either as an absolute amount or as a percentage of the market value of the house. In contrast, the value of a compounding loan is more difficult for an average person to understand or calculate.⁷

1.2. Myth 2: "Annuities are a rip off"

1.2.1. Reality

The cost of annuities reflect the low interest rate environment: buying guaranteed future income is expensive.

The cost of annuities reflects the trend of improving mortality. Companies have been forced to make more prudent (and expensive) assumptions (Wilson, 2007).

The cost of capital is high due to the difficulty in quantifying the future risks of longevity and investment returns.

1.2.2. Problem

Alternative products, like drawdown products, do not offer insurance for longevity. In comparison, annuities include a benefit from mortality cross-subsidy (in an annuity those who die earlier cover much of the insurance cost of those who live longer).

⁶ The average standard variable rate for a mortgage was 3.33% for July 2017, from Series CFMZ6IX published by the Bank of England.

⁷ In a study, highlighted in (Lusardi & Mitchell, 2007), of around 2,000 US respondents who were between ages 51 and 56, 78.5% were unable to calculate the value of a bank account after 2 years of accumulating interest at a constant rate.

Alternative products may offer much greater flexibility to the pensioner than annuities, to choose how and when to spend their funds. The downside is that funds may be reduced too fast and the pensioner may suffer unanticipated poverty.

Alternative products may also expose pensioners to investment risk at a stage when they are least able to recover financially from poor investment returns. Conventional annuities involve no investment risk for the retiree.

1.2.3. Solution

Create a fair and transparent product, as discussed here,

- That enables the mortality cross-subsidy to be credited to participants.
- In which risk and reward can be understood, and investment risk can be reduced while maintaining the expected reward, by replacing investment risk by mortality risk.

1.3. Myth 3: “Savings rates are unfair”

1.3.1. Reality

Savings rates reflect the current low interest rate environment which:

- Protects wider society against the greater risk of rising bad personal debts and credit contraction.
- Protects asset values for those approaching or in retirement with non-cash assets, albeit less than the rising cost of annuities.

1.3.2. Problem

Low savings rates:

- Harm those with no ability to increase wealth.
- Can encourage greater risk-taking in order to chase yield.

1.3.3. Solution

Develop flexible options that can augment savings income.

In summary, there is a current and growing problem in UK society as pensioners have tended to save too little leading up to retirement. Even those who have saved have seen the cost of retirement provision rising unexpectedly quickly. This creates a growing pressure as the baby boom generation starts to retire.

Between 2010 and 2051, the proportion of people aged 65 and over is projected to increase from 17% to 24% (Office for National Statistics, 2012, page 4). This demographic also creates a tremendous opportunity for product innovation. Thus far pension solutions have tended to focus on variants of traditional annuities and income drawdown-type products.

Solutions to generate additional income from other assets, such as housing, have generally focussed on lifetime mortgages, a type of equity release scheme. This paper proposes a whole new class of product solutions that utilise a simple insight on how annuities generate the mortality cross-subsidy that protects annuitants from outliving their assets. It nullifies the well-known issue of mortality drag

that affects traditional income drawdown products. Unlike conventional annuities, the product solutions do not require customers to buy a mortality and investment guarantee, but instead allow such guarantees to be added at the explicit decision of the customer. There is also the possibility that the product solutions can be designed to allow customers to withdraw before they die, rather than being forced into an irreversible decision as in conventional annuities.

2. Key Innovation: Mortality Risk-Sharing

The rule proposed by Donnelly *et al.* (2014) is a generalisation of the key mortality risk-sharing benefit of an annuity. The generalisation allows this benefit to be applied to arrangements of different assets, like houses.

With an annuity, investment risk and mortality risk are all transferred to the annuity provider. This can hide an underlying mechanism of what is actually happening: the annuitants who die early transfer the value of their fund to the annuitants who live longer. In effect, the surviving annuitants become the heirs of the dying annuitants. It is possible to design a scheme that focusses on this mechanism alone. This creates a wealth of options to help pensioners generate additional income.

2.1. Mortality Risk-Sharing Rule

Consider a scheme that:

- ring fences each customer's assets within the scheme; and
- estimates each customer's probability of dying over each future month given the customer's age and other risk factors.

This creates a unique exposure-to-risk for every person in the scheme, calculated as:

$$\begin{aligned} &\text{Customer's exposure-to-risk} \\ &= \text{Value of customer's assets} \times \text{Probability of customer's death over the next month} \end{aligned}$$

The exposure-to-risk value varies between customers within the same scheme, as customers bring different amounts of assets and have different probabilities of death over the same time period.

When anybody dies in the scheme, the value of the assets which they have put into the scheme are split up and paid out to everybody in the scheme in proportion to each person's individual exposure-to-risk. To ensure this is strictly actuarially fair, the estate of anybody who dies will also get a payment equal to that person's own share of their exposure-to-risk. This may seem like an odd idea, but it ensures the scheme still works for small numbers of people, or when there is someone who has a disproportionately large exposure-to-risk (in that case, the dead person's estate gets back a fair proportion of their assets).

The actuarially fair payment made at the end of each month to any customer or their estate is calculated using the rule:

$$\begin{aligned} &\text{Mortality credit paid to customer at end of the month} \\ &= \frac{\text{Customer's exposure-to-risk}}{\text{Total scheme exposure-to-risk}} \\ &\quad \times \text{Total value of scheme assets of those who died over the month} \end{aligned} \tag{1}$$

as detailed in Donnelly *et al.* (2014). The monthly mortality credit is always non-negative (being zero only when none of the scheme members die over a month). Effectively, it is an additional, non-negative return to the investment return on a member's assets. A member only makes a "loss" on the mortality risk-sharing when they die, since then they "lose" the value of their assets placed in the scheme which will outweigh the value of the mortality credit paid to their estate at the end of the month of their death.

The scheme can be overlaid onto existing investment products or assets; it extends rather than replaces the universe of investments. It benefits its members by allowing them to interchange investment risk with mortality risk, and gain a non-negative return as long as they are alive.

With enough diversification in the membership profile and sufficiently liquid assets, the scheme could be operated like a life annuity (Donnelly *et al.*, 2014). However, there are three key differences from the classical life annuity contract. First, the "premium" that is paid by the members, namely the loss of their own fund value, is paid upon death rather than upfront. Second, members can choose, up to some scheme-specific limit, how much income to withdraw from their ring-fenced assets. Third, the risks associated with mis-estimating the members' mortality probabilities are borne by the members, and not by the insurer (assuming that no mortality guarantee is purchased by the members).

The mortality risk-sharing rule diversifies the idiosyncratic risk that we do not know whether an individual member will live or die each month, even though we may know their probability of death. It works by pooling this idiosyncratic risk among the members of the scheme. Indeed, if there are sufficiently large numbers of members in the scheme then the frequency of deaths will occur in line with the probabilities of death over a month, even when members face different probabilities of death (Donnelly *et al.*, 2014).

In practice, full diversification of the idiosyncratic risk may not be achieved due to having too few members in the scheme to adequately pool mortality risk. This means that the crude mortality risk-sharing rule given by equation (1) can result in lumpy mortality credits: in some years, nobody dies while in other years, many people die. This creates the opportunity for an insurance company to smooth the mortality payments in order to increase the attractiveness of the scheme.

There also exists the systematic risk (sometimes called longevity risk) that the incorrect probability of death has been chosen. The mortality risk-sharing rule does not protect the scheme members against this risk. For the scheme members to remove the systematic risk, they would have to buy a guarantee on the amount of mortality credit that they receive each month, based on the anticipated probabilities of death.

Unless a guarantee is bought by the members from an insurance company, the risk associated with mis-estimating the probabilities of death lies with the members. It is only with the sale of a guarantee, whether partial or full, that the insurer may bear some of this risk. Suggestions for ways of smoothing the mortality credits are discussed next, which all require the insurer to bear, either partially or fully, the systematic risk of mis-estimating the probabilities of death.

2.2. Mortality Credit Smoothing

We outline two possible methods to smoothing the mortality credit. In one method, the insurer bears all of the mortality risk. In the second method, the insurer bears only the downside mortality risk that fewer deaths occur than expected.

2.2.1. Insurer takes all of the risk for a fee

One method of doing the smoothing of the mortality credits is for the insurer to pay the expected value of the mortality credits in exchange for the (random) flow of funds arising from deaths in the scheme. Essentially, the insurer is selling the group of customers a sequence of instantaneous temporary life annuities, with the premium paid upon death rather than before death. It can also be thought of as a type of swap, with the group exchanging their “floating leg”, that is, the volatile mortality credits, for an insurer’s “fixed leg”, that is, a mortality credit based on each customer’s wealth and force of mortality.

For example, suppose the scheme operates for a fixed 3-year time span. Upon a customer joining the scheme,

- The customer specifies the assets that they bring to the scheme. These are the assets that are “at risk” if the customer dies and whose value is used to calculate the customer’s exposure-to-risk.
- The insurer allocates a time-varying force of mortality function to the customer.
- The insurer guarantees to pay a customer, who is alive at the start of month n , the amount:

Value of customer’s assets at the start of the n th month \times Probability of customer’s death over the n th month

at the end of month n , for $n = 1, 2, \dots, 36$.

- The insurer controls the investment strategy, in order to control the investment risk that they take on by guaranteeing the mortality credit based on the start-of-the-month asset values.
- In exchange, upon the death of any customer in the scheme during the 3 years of operation, their assets belong to the insurer.

At the end of the 3 years, surviving customers can choose whether to exit the scheme entirely, or to continue for another 3 years. At this point, the customer re-specifies the assets that are at risk.

This method leaves the insurer still having to manage investment and systematic mortality risk (the risk of choosing the wrong mortality model), albeit without the lifetime guarantee of a conventional life annuity.

Considering fairness to the customers, there is no onus on the insurer to choose their best estimate of the probabilities of death over each month. Indeed, the insurer may underestimate the probabilities of death in order to reduce the amount of payments that it makes to the customers. However, competition between insurers would mitigate the risk of a bad deal for customers.

2.2.2. Insurer takes the downside risk for a fee

For the individual, the advantage of sharing mortality risk is to gain a mortality credit payment. However, the timing and amounts of the mortality credits received may not coincide with the timing and amounts of the desired income of the customer, if they are using the scheme to provide an income in retirement. The scheme may be more attractive if a mortality credit payment is made regularly and is of a relatively stable amount.

One idea is that the insurer guarantees that the mortality credit is at least a fixed value. In other words, if the mortality credit is below the fixed value, then the insurer makes up the difference. The

attraction for the customers is that it gives them a guaranteed minimum income stream while allowing them to gain a higher mortality credit if more deaths occur than anticipated.

With this partial guarantee on the mortality credit, the scheme members share the systematic mortality risk with the insurer. However, they reduce the downside risk of idiosyncratic mortality risk, that fewer members die than anticipated so that the mortality credit paid is lower than expected. They keep the upside risk (as long as they survive) that more members die than anticipated so that the mortality credit paid is higher than expected.

For example, suppose there are 500 identical members in the scheme and each member has £100,000 allocated to the scheme at the start of the year. The insurer guarantees that the annual mortality credit payment to each member is at least £2,400 over the year. In return for this guarantee, each member pays the insurer £85.02.

Suppose that by the end of the year, eight members died over the year. Their total wealth in the scheme is £800,000. Splitting this evenly among the 500 members of the scheme results in an actual mortality credit – calculated using the mortality risk-sharing rule given by equation (1) – of £1,600 per member. Due to the guaranteed minimum of £2,400, the insurer must pay each member an additional £800. (The full assumptions underlying these values are detailed in Example 2.1.)

The guaranteed minimum mortality credit should provide an incentive for the insurer to try to choose the correct mortality probabilities for each member. With the minimum guarantee,

- If the mortality probabilities assigned by the insurer are too low, and the “true” mortality probabilities are higher, then:
 - more deaths occur than suggested by the assigned mortality probabilities; and
 - it is less likely that the insurer has to pay out on the guarantee, which means
 - the premium charged for the guarantee is too high (i.e. a financial gain for the insurer).
 - With market competition, another insurer can offer a lower premium to the group for the same guarantee.
- If the mortality probabilities assigned by the insurer are too high, and the “true” mortality probabilities are lower, then:
 - fewer deaths occur than suggested by the assigned mortality probabilities; and
 - it is more likely that the insurer has to pay out on the guarantee, which means
 - the premium charged for the guarantee is too low (i.e. a financial loss for the insurer).
 - With market competition, more groups move to the insurer and the insurer’s losses increase. In response, the insurer decreases its assigned mortality probabilities for the same guarantee.

Example 2.1. Suppose 1,000 people share their mortality risk in the scheme. All have independent and identically distributed future lifetimes. They each allocate £100,000 of their wealth to the scheme. Imagine that two people die over the month, releasing £200,000 in total. Each of the 1,000 people (i.e. including the estates of the two deceased members) receive £200 ($= \frac{200,000}{1,000}$). At the end of the month, there are 998 people left alive in the scheme. Assuming that none of these 998 people leave the scheme and that the scheme continues for another month, then any mortality credits payable due to deaths in the second month are paid to the 998 people who started the month in the scheme.

With no minimum mortality credit guaranteed by the insurer, the customer has no certainty in advance about the amount of mortality credit to be paid. If no-one dies, then there is no mortality credit to be paid. Deaths can be volatile from year to year. This may be unattractive to the scheme participants who wish to have a stable income.

We show how a minimum guarantee can be incorporated. To keep the example simple, we assume that there is no investment risk and investment returns are zero. Suppose that the insurance company believes that each of the 1,000 people has probability 0.003 of dying over the month and deaths occur independently.

We model the number of deaths over the month among the 1,000 people by a binomially distributed random variable $N \sim \text{BIN}(1,000, 0.003)$. Based on this model, three people are expected to die over the month, with a s.d. 1.73 deaths. In terms of mortality credits, this means that members can expect to receive £300 at the end of the month, with a s.d. £173.

Suppose that the insurance company guarantees that the mortality credit is at least £250 at the end of the month. If only one person dies over the month, then the mortality credit (calculated by the mortality risk-sharing rule (1)) paid to each member is £100 per member. Due to the minimum guarantee, the insurer would pay an additional £150 to each of the 1,000 members, so that each member gets £250 in total.

If instead three people had died over the month, then the mortality credit (calculated by the mortality risk-sharing rule (1)) paid to each member is £300 per member. No top-up payment would be required by the insurer, as the mortality credit of £300 is above the minimum guarantee of £250.

The insurer needs to be financially compensated for providing the guarantees. A sample premium calculation is shown in section 4.

3. Scheme Design Choices

Here the design of the scheme is examined. The feasibility of including assets like bank deposits and property in the scheme is considered, as well as their impact on a scheme member's income.

3.1. Open or Closed Scheme

An open scheme would allow investment in any funds on a platform. A closed scheme would restrict scheme assets to those managed by the life office.

A closed scheme offers greater operational simplicity (assets are easily calculated and converted to cash on death) and the prospect of easier choices for customers. An open scheme could leverage existing platforms and the belief that investment fund choice is important and valuable to customers.

Note that the scheme needs at least two members to be non-trivial. With only one member, the member would "lose" the value of their assets on death, but would then be paid the full value of their assets at the end of the month of their death as a mortality credit. There is no pooling of mortality risk with only one member in the scheme.

3.2. Tontine Rules

The scheme would need to ensure data security so members do not know who else is in the scheme. This is to remove the tontine risk that makes tontines illegal. The risk would be mitigated further by a mortality smoothing method run by an insurance company, so that the mortality payment to survivors is less or not dependent upon the experienced mortality in a period.

3.3. Type of Assets Included in the Scheme

A simple scheme could be applied to income drawdown products, Individual Savings Accounts and share portfolios. For example, the assets in the drawdown product would become assets of the scheme. This would create a product that combines the mortality benefit of an annuity with a more flexible income drawdown product. This is a halfway house between an annuity and income drawdown product, with features and benefits of both.

The use of illiquid assets must be balanced against the purpose of the scheme. If it is to provide a regular, monthly income to the members, then enough assets must be liquid to allow income payments to be made to members. Alternatively, the manager of the scheme may provide short-term liquidity in order to make income payments. However, the cost of this will ultimately be borne by the scheme members through either an explicit fee charged by the scheme manager or a mortality credit that is lower than that generated by the expected selling price of the illiquid assets.

The design of the scheme may become more complex to develop and manage. However, it offers greater benefits to customers and society if additional types of assets can be placed into the mortality sharing scheme. Some other possible assets are discussed next.

3.3.1. Property

This might be the most interesting alternative asset class to include in a mortality risk-sharing scheme.

Some people have under-invested in retirement savings with the belief that rising house prices would create an asset that could help fund and pay for their retirement. The difficulty is that houses are illiquid and do not generate income (while you live in them). Alternatives like lifetime mortgages or downsizing that monetise property are expensive (Adams & James, 2009, box 1, page 13).

It would be neither sensible nor beneficial to put the whole value of a house into a mortality risk-sharing scheme because:

- Household members with no residual stake in the house on death may well ignore repairs and the house value will fall due to neglect. There is a difficulty and cost involved in tracking the decline in house value.
- People wish to leave value to their heirs and dependants.
- It is hard to value property and there is the risk that a property value is inflated in order to (unfairly) increase the mortality credit paid to an individual, or vice versa.

A first thought may be to allow anyone to join, with their home included in the scheme. However, it is unreasonable to have a situation in which the member dies so their house is sold, thereby depriving their surviving partner of a home. This means that the joint life, “last survivor” mortality of a

cohabiting couple should be used in the mortality credit calculation (1). After the first life dies, the individual mortality rate of the surviving life should be used.

It is straightforward to incorporate joint mortality: we replace the individual probabilities of death by the probability that both cohabitants die in the specified time-frame (e.g. the month or year over which the mortality credit is calculated). We assume here that the cohabitants die independently of each other. This is a simplification: in reality, it would be reasonable to expect that their deaths would be dependent events as they travel and live together, and due to the “broken heart” syndrome (Ji *et al.*, 2011; Spreeuw & Owadallya, 2013).

However, joint mortality probabilities are lower than individual mortality probabilities, as it is unlikely that the partners die in the same month. This means that mortality credits a couple would receive, while they both survive, are very low. It may be that some couples would not see the income as worthwhile, when weighed against the value of the house that they lose upon their joint death. For other couples, they may be happy to have a product that will automatically continue after the first death, and produce a higher income to the surviving partner.

Certainly, including housing wealth in the scheme is an option for someone who lives alone. There are millions of single pensioners who own their home. As stated above, there are around 2.7 million single pensioners who own their home (Department for Work and Pensions, 2014, table 1.1; Office for National Statistics, 2014, figure 3.12), with a median net property wealth of £165,000 (Office for National Statistics, 2014, figure 3.13).

Allocating 80% of the median house value to the scheme gives a value of £132,000. This could provide additional annual income of around £800 at age 65, £1,800 for 70-year olds, and £3,700 for 75-year olds.⁸

Including part of the value of a house in the scheme would require house prices to be estimated and trusts to be set up to hold title deeds against the house for the scheme. This is something that banks do everyday with traditional mortgage lending, where partnerships could be set up to benefit bank customers.

3.3.2. Bank deposits and savings

These assets often have low current returns, the average annual return for instant access deposits is currently 0.47%,⁹ while the best instant access annual rate has fallen from around 3% in 2012 to just 1.65% in February 2016.¹⁰ The best 1-year fixed-term deposit rates have fallen from 3.5% at the start of 2013 to just 2% in February 2016.¹¹

⁸ Based on the life table PMA92C20 produced by the Continuous Mortality Investigation Bureau, Institute and Faculty of Actuaries, UK. <http://www.actuaries.org.uk/research-and-resources/pages/92-series-mortality-tables-assured-lives-annuitants-and-pensioners>, and the figures rounded to the nearest \$100.

⁹ Source: Bank of England, *Monthly interest rate of UK monetary financial institutions (excluding Central Bank) sterling instant access deposits including unconditional bonuses from households (in percent) not seasonally adjusted*, Data series code IUMB6VJ, 31 January 2016.

¹⁰ Source: <http://www.thisismoney.co.uk>, *Best savings rates: General savings – easy access online and branch*, updated on 26 February 2016.

¹¹ Source: <http://www.thisismoney.co.uk>, *Best savings rates: Fixed-rate accounts*, updated on 26 February 2016.

Including these low yield savings in a mortality risk-sharing scheme would enhance the returns for savings customers. Its implementation would likely require the savings to be put in a trust, in order to ring fence them for the scheme in the event that the person dies. This would increase operational complexity and require banks and insurance companies to develop a joint solution.

Banks could benefit from enabling deposits and savings to be included in the scheme: a successful scheme could attract additional savings to a bank that offered this type of savings account. In addition, these savings may well have higher value to the bank in terms of behavioural life and customer stickiness. It would be important, in the interests of customer fairness, to ensure that these savings accounts mirrored normal savings accounts in terms of the interest (non-mortality-linked) rates offered on them.

Including savings such as these could increase annual yields by 0.5% for 65-year-olds, 1.3% for 70-year olds, and 2.7% for 75-year olds.¹² There would be an impact on inheritance tax as the ring-fenced savings would no longer form part of the person's estate. Savings and investments can be significant: 30% of households with one or more people over State Pension Age have savings of £20,000 or more (Department for Work and Pensions, 2012, table 4.9).

Selection risks are explored in section 3.6, when constraints on leaving the scheme and income withdrawal are discussed.

3.4. Impact on Investment Choices for Assets Included in a Scheme

With income drawdown, there is always a trade-off between the maximum income withdrawn, investment strategy, the expected funds remaining and the actual funds left due to volatile investment performance.

Customers in income drawdown have a choice on the degree of investment risk that they wish to take. Lower risk will create less volatility and less risk of fund values falling catastrophically. The flip side is that the lower expected return will limit the amount of withdrawals if the customer's fund is not to be exhausted. Riskier funds may let customers achieve higher average returns to try and recover the monies withdrawn, but at the cost of volatile returns and the risk in a falling market of rapid fund depletion. The problem can be exacerbated if higher fund charges and non-optimal asset allocation are included, or if customers are sold "high performance funds" that actually rely on survivor bias to generate good historic performance (Elton *et al.*, 2001; Choi *et al.*, 2010).

The product sold to enable participation in a mortality risk-sharing scheme can be structured to help customers understand the trade-off between risks and returns. For example, using suitable asset allocation strategies, customers could be shown a graduated choice of the estimated volatility of return on assets versus the expected returns. The customer could be shown the distribution of their expected fund values in 5, 10 or 20 years under different choices of income and amount of investment risk taken. This would help the customer to understand the risk–return trade-off.

¹² Based on the life table PMA92C20 produced by the Continuous Mortality Investigation Bureau, Institute and Faculty of Actuaries, UK. <http://www.actuaries.org.uk/research-and-resources/pages/92-series-mortality-tables-assured-lives-annuitants-and-pensioners>.

The effect of mortality cross-subsidy can be easily included to show its benefit of a higher expected income, or from the targeting of the same income but with a lower risk investment strategy.

The scheme design may also enable exploration of investment smoothing options. This may be easier with a mutual business model and reflect elements of with-profits smoothing.

3.5. Payment Frequency

Monthly payments of any mortality credit probably give the best trade-off between costs, risks and customer expectations. They allow customers to have a regular payment and the insurer to update the anticipated future distribution of deaths.

The main decision is when the mortality credits attributable to a particular death should be paid. The two main possibilities are in the month that the death is notified to the scheme, or when the assets of the deceased are liquidated.

If assets like property are included in the scheme, then there is less risk to the insurance company if the proceeds of the sale are distributed after the sale is completed. This reduces the risk from inheritance disputes, the property taking a long time to sell or the financial risk if close to 100% of the property value is allocated to the scheme.

However, waiting until property is sold slows down the mortality credits paid to members, and some members will die before they receive their mortality credit due from the sale of a property. To get by this, the manager of the scheme could pay the relevant mortality credit to members at the end of the month of death, and sell the property later. To diminish the risk involved in this provision of short-term liquidity, either an explicit fee could be charged or a reduced property value is used to calculate the mortality credits.

3.6. Reducing Selection Risks

An example of adverse selection arising in the scheme would be a member finding out that they are ill, and exiting the scheme to avoid any chance that they can “lose” their assets by dying. This can be mitigated by restricting when a customer can exit the scheme, and also restricting how much they can withdraw as income from their ring-fenced scheme assets.

- Exit at a certain age – for example, age 75. This reduces selection risk against the scheme as the exit date and event are not chosen by the customer.

The problem for the customer if the exit date is fixed rigidly is that the customer’s circumstances might change. For example, suppose the customer’s house is included in the scheme. The customer may need to go into care or sheltered accommodation, and no longer need the house as a home. They may wish to move to be close to family who can provide them with care, or they may need to downsize to a house that better meets their needs. This may require an additional exit clause for members who are admitted to a care home.

- Notice period before withdrawal. This will reduce but not eliminate selection risk – for example, if a 2-year or 5-year notice period is required before somebody leaves the scheme. In this case, the notice period means that somebody, who is withdrawing because they are aware that they have an increased mortality risk, needs to survive the notice period to take advantage of this knowledge. If they exit the scheme based on this information before they die, then the mortality experience in the scheme will be lighter than expected.

There is likely to be a scarcity of existing mortality data that can be used to model the customer's mortality with the notice period. Mortality data would be difficult to collect as no information would be collected on the customer after they had withdrawn from the scheme. Even if the mortality data existed, withdrawals may also be affected by the general economic environment, for example, if investment returns had been poor for a few years, the customer may withdraw their assets from the scheme in order to consume the assets, rather than for health reasons.

- Free withdrawal (say at end of year). While this generates selection risk, in theory the scheme can allow for this by using the appropriate mortality basis. The mortality experienced in the scheme will be select (i.e. a subset of all mortality) as it will only be when scheme members do not have enough advance notice to withdraw from the scheme before they die. This is true for all scheme members, so a lighter mortality based on these deaths alone should still provide an actuarially fair mortality benefit, albeit lower than a more restrictive scheme.

3.7. Costs and Charges

These will reflect the complexity of any scheme. There will be:

- product charges for income drawdown-type products;
- mortality sharing costs – to collect and pay out mortality sharing benefits;
- mortality smoothing costs – as set out in Section 4; and
- costs associated with other assets in the scheme. For example, allowing property as a scheme asset is likely to incur costs around its valuation and potentially ongoing monitoring costs of its valuation.
- Costs associated with allowing members to exit the scheme to, for example, move into a care home.

3.8. Scheme Design and Customer Proposition

Here, we suggest a scheme with a minimum guaranteed mortality credit that can be operated by the insurer. We propose a scheme that:

- requires members to join for 10 years to reduce selection risk;
- allows both housing wealth and investment fund assets;
- allows house sales and purchases; and
- allows income withdrawal from investment fund assets that are in excess of the mortality credits.

The scheme design is shown in Table 1.

4. Premium Calculation for the Minimum Guaranteed Mortality Credit

We show how the premium for the minimum guaranteed mortality credit can be calculated, and show some sample calculations.

4.1. Premium Calculation for the Minimum Guarantee

As the insurer is guaranteeing a minimum amount of mortality credit, they require a premium in return for paying the chosen minimum mortality credit to a member. For simplicity, investment returns, expenses and taxes are ignored. We do the calculations annually, but any other time interval, for example, months, could also be used.

Table 1. Potential Scheme Design Features

Scheme Design Feature	Conditions	Motivation
Scheme entry event	Entry after age 60 and before age 65	Motivation for scheme is to provide post-retirement income
Scheme exit event	Death or age 75	Reduce anti-selection risk
Included assets	House and/or insurer's investment funds	Monetise housing wealth and/or allow additional return on investment funds
Asset valuation	(1) House. Fixed monetary partial value of house, maximum 60% value of house value. Re-valuations may be done periodically (2) Insurer's investment fund valuation. Market value at end of month	Avoid risk that owner neglects house Value of investment funds fluctuate with investment returns
Switches between assets	(1) House can be sold, and either a new house bought or insurer's funds purchased. However, value assigned to scheme must remain unchanged (2) Insurer's funds can be replaced with housing, subject to the house valuation conditions holding	Balance customer choice with anti-selection risk Balance customer choice with anti-selection risk
Withdrawal of assets	(1) No withdrawal of house value except indirectly through a switch to the insurer's funds (2) Withdraw a maximum of $X\%$ of insurer's fund value over a M -year period	Nature of housing as an asset Balance customer choice with anti-selection risk
Mortality credit	(1) Calculated monthly based on customer's total asset valuation (potentially averaged over the month) and monthly mortality probability (2) Paid at the end of each month (3) Asset value shared out at end of month that death is notified to the scheme, rather than waiting, e.g., for house to be sold	Ensure a regular payment to the customer
Minimum guaranteed mortality credit	Customer chooses amount	Ensure a minimum income to the customer

4.1.1. Homogeneous scheme

Consider first a homogenous scheme in which all $\ell \geq 2$ members are independent and identical copies of each other. Each member starts the year with wealth ℓw and has probability $q \in (0,1)$ of dying over the year. The number of deaths over the year among the ℓ people is modelled by a random variable $N \sim \text{BIN}(\ell, q)$. Based on this model, ℓq people are expected to die over the year, with a variance of $\ell q(1-q)$ deaths.

Applying the mortality risk-sharing rule (1), the mortality credit paid to each of the members at the end of the year is the random variable

$$\text{MC} = wN \times \frac{wq}{\ell wq} = wN \times \frac{1}{\ell} \quad (2)$$

If no member dies over the year then the realisation of the random variable N is zero and no mortality credit is paid. If one member dies over the year then the realisation of N is one, and each

Table 2. Annual Premium Rates Calculated Using the Expected Value Premium Principle With No Safety Loading for a Homogeneous Scheme

Guaranteed Minimum Mortality Credit Per Scheme Member		Annual Premium Per Scheme Member	
As an Amount	As a Percentage of £100,000	As an Amount	As a Percentage of £100,000
£250	0.25	£3.30	0.00330
£500	0.50	£21.60	0.02160
£750	0.75	£73.48	0.07348
£1,000	1.0	£174.59	0.17459

The values are calculated assuming that there are $\ell = 500$ people in the scheme, each with a constant wealth $w = 100,000$. Each member has a probability of death over the year of $q = 0.01$ and deaths occur independently. For this group, the expected mortality credit is £1,000, and the guaranteed minimum mortality credit is shown up to this value.

member receives a mortality credit of w/ℓ . In that case, each of the $\ell - 1$ survivors has a net wealth of $w + w/\ell$ at the end of the year, and the dead member has a net wealth of w/ℓ .

Suppose that a member of the scheme buys a guarantee from the insurance company. The guarantee is that the member will receive at least the amount $g \in [0, w]$ as a mortality credit at the end of the year. Any payments made by the insurance company are a “top-up” to the mortality credit defined by equation (2). Thus the insurer is liable to pay the amount

$$\max\{g - \text{MC}, 0\}$$

to the member at the end of the year. The larger the value of the guaranteed minimum g , the more the insurer is liable to pay out.

The insurer charges an annual premium $p(g) \geq 0$ to the member in exchange for the guaranteed minimum g . The premium can be calculated as per any standard actuarial premium principle. For example, using the classical expected value premium principle with no safety loading,

$$p(g) = \mathbb{E}(\max\{g - \text{MC}, 0\})$$

For schemes in which there are too few members to effectively diversify mortality risk, the insurer can add an explicit loading for diversification risk to the premium.

Table 2 illustrates the values of the premium based on the expected value premium principle, for a homogeneous scheme with $\ell = 500$ members at time 0, a probability of dying over the year $q = 0.01$ and each member starting with wealth $w = 100,000$. Deaths occur independently. For this homogeneous group, we calculate that the expected mortality credit paid to each member at the end of the year is £1,000 ($= 500 \times 0.01 \times 100,000/500$), with a s.d. £445 ($= \sqrt{500 \times 0.01 \times 0.99} \times 100,000/500$), before any minimum guarantee is applied. Each member can individually choose their desired guarantee amount g , and are charged the appropriate premium. The annual premiums for the guaranteed minimum mortality credit are low, ranging from £3.30 for a minimum of £250 per £100,000, to £174.59 for a minimum of £1,000 per £100,000.

Table 3 shows the same calculations except that the members have a twice-higher probability of dying of $q = 0.02$. With $q = 0.02$, the expected mortality credit paid to each member at the end of the year is £2,000 ($= 500 \times 0.02 \times 100,000/500$), with a s.d. £626 ($= \sqrt{500 \times 0.02 \times 0.98} \times 100,000/500$), before

Table 3. The Same Calculations and Assumptions as in Table 2 Except That Here Each Member has a Probability of Death Over the Year of $q=0.02$

Guaranteed Minimum Mortality Credit Per Scheme Member		Annual Premium Per Scheme Member	
As an Amount	As a Percentage of £100,000	As an Amount	As a Percentage of £100,000
£500	0.5	£0.36	0.00036
£1,000	1.0	£8.21	0.00821
£1,500	1.5	£68.51	0.06851
£2,000	2.0	£247.70	0.24770

For this group, the expected mortality credit is £2,000, and the guaranteed minimum mortality credit is shown up to this value.

any minimum guarantee is applied. Due to the higher mortality rate, the annual premiums are lower for the same amount of guaranteed minimum mortality credit as in Table 2. For example, the annual premium for a guaranteed minimum mortality credit of £1,000 per £100,000 is £8.21, a reduction of around 95% from the premium when $q=0.01$.

For the selection of the minimum guarantee for the homogeneous group, this could be a fixed value offered by the insurer to all members or the members may be offered a menu of minimum guaranteed mortality credits, and can choose individually which one they want, if any.

4.1.2. Heterogeneous scheme

Next consider a heterogeneous scheme consisting of two homogeneous groups, Group A and Group B. All the members in Group A are identical copies of each other, and similarly for the members of Group B. All members have independent future lifetime distributions.

For $X \in \{A, B\}$, in Group X there are $\ell^X \geq 1$ members, each of whom starts the year with wealth ℓw^X and has probability $q^X \in (0,1)$ of dying over the year. The number of deaths over the year among the ℓ^X people in Group X is modelled by a random variable $N^X \sim \text{BIN}(\ell^X, q^X)$.

Applying the mortality risk-sharing rule (1), the mortality credit paid to each member of Group A at the end of the year is the random variable:

$$MC^A : = (w^A N^A + w^B N^B) \times \frac{w^A q^A}{\ell^A w^A q^A + \ell^B w^B q^B} \tag{3}$$

with $\mathbb{E}(MC^A) = w^A q^A$. By symmetry, the mortality credit paid to each member of Group B at the end of the year is the random variable

$$MC^B : = (w^A N^A + w^B N^B) \times \frac{w^B q^B}{\ell^A w^A q^A + \ell^B w^B q^B} \tag{4}$$

with $\mathbb{E}(MC^B) = w^B q^B$.

Suppose that a member of Group A buys a guarantee from the insurance company that the member will receive at least $g \in [0, \frac{w^A q^A (w^A \ell^A + w^B \ell^B)}{\ell^A w^A q^A + \ell^B w^B q^B}]$ as a mortality credit at the end of the year. Thus the insurer is liable to pay the amount

$$\max\{g - MC^A, 0\}$$

Table 4. Annual Premium Rates Calculated Using the Expected Value Premium Principle With No Safety Loading for a Heterogeneous Scheme

Guaranteed Minimum Mortality Credit for Group A Member		Annual Premium Per Group A Member	
As an amount	As a percentage of £100,000	As an amount	As a percentage of £100,000
£500	0.5	£0.53	0.00053
£1,000	1.0	£14.15	0.01415
£1,500	1.5	£91.05	0.09105
£2,000	2.0	£289.29	0.28929
Guaranteed Minimum Mortality Credit for Group B Member		Annual Premium Per Group B Member	
As an Amount	As a Percentage of £350,000	As an Amount	As a Percentage of £350,000
£1,750	0.5	£5.30	0.00151
£3,500	1.0	£142.18	0.04062
£5,250	1.5	£759.39	0.21697
£7,000	2.0	£1971.93	0.56341

There are two groups in the scheme, Group A and Group B. Group A has $\ell^A = 450$ members, each starting the year with wealth $w^A = 100,000$ and probability $q^A = 0.02$ of dying over the year. Group B has $\ell^B = 50$ members, each starting with wealth $w^B = 350,000$ and the lower probability $q^B = 0.015$ of dying over the year. Deaths occur independently.

to the member of Group A at the end of the year. As before, the annual premium $p^A(g) > 0$ that the insurer charges to each member of Group A for the guaranteed minimum g can be calculated as per any standard actuarial premium principle. Using the expected value premium principle with no safety loading, $p^A(g) = \mathbb{E}(\max\{g - MC^A, 0\})$. Similarly, the annual premium charged to a member of Group B for a guaranteed minimum g can be calculated using the expected value premium principle with no safety loading as $p^B(g) = \mathbb{E}(\max\{g - MC^B, 0\})$.

Table 4 illustrates the values of the premium based on the expected value premium principle, for a heterogeneous scheme. Group A has $\ell^A = 450$ members, each starting with wealth $w^A = 100,000$ and probability $q^A = 0.02$ of death over the year. Group B has $\ell^B = 50$ members, each starting with wealth $w^B = 350,000$ and lower probability $q^B = 0.015$ of death over the year. Thus the Group B members, who are relatively wealthy, are assumed to have a lower probability of death than the Group A members. Deaths occur independently, both within and across the groups.

Suppose each member wishes to have the value of the mortality credit that they expect to receive as their guaranteed minimum mortality credit. The expected mortality credit is different for Group A members compared to Group B members. For Group A members, this means choosing 2.0% of their wealth as the guaranteed minimum mortality credit for a premium of 0.28929% of their wealth. In monetary amounts, the Group A members choose a guaranteed minimum mortality credit of £2,000 ($= 2.0\% \times 100,000$) for a premium of £289.29. For Group B members, it means choosing 1.5% as the guaranteed minimum mortality credit for a premium of 0.21697% of their wealth. In monetary amounts, the Group B members choose a guaranteed minimum mortality credit of £5,250 ($= 1.5\% \times 350,000$) for a premium of £759.39.

Table 5. The Same Calculations and Assumptions as in Table 4 Except that the Two Group's Membership has Doubled: Group A has $\ell^A = 900$ Members and Group B has $\ell^B = 100$ Members

Guaranteed Minimum Mortality Credit for Group A Member		Annual Premium Per Group A Member	
As an Amount	As a Percentage of £100,000	As an Amount	As a Percentage of £100,000
£500	0.5	£0.06	0.00001
£1,000	1.0	£1.81	0.00181
£1,500	1.5	£36.99	0.03699
£2,000	2.0	£204.91	0.20491

Guaranteed Minimum Mortality Credit for Group B Member		Annual Premium Per Group B Member	
As an Amount	As a Percentage of £350,000	As an Amount	As a Percentage of £350,000
£1,750	0.5	£0.17	0.00000
£3,500	1.0	£42.71	0.01220
£5,250	1.5	£537.89	0.15368
£7,000	2.0	£1829.56	0.52273

Next we look at the impact of heterogeneity within the scheme on the premiums. Broadly, we see below that to reduce the impact of heterogeneity on premiums, the number of members in Groups A and B should be increased. We compare the premiums for the Group A members in Table 4 to those in Table 3 (the latter table shows premiums for a homogeneous scheme). For Table 4, the Group A members each have a probability of death of 0.02 and wealth £100,000 and are in a heterogeneous scheme in which the Group B members have a lower probability of dying. The total number of members in the scheme is 500. The premiums in Table 3 are for a homogeneous scheme which effectively consists of 500 Group A members, and no Group B members.

The premiums for the Group A members, in the heterogeneous scheme, are higher for the same amount of minimum mortality credit. For example, for a minimum mortality credit of £1,000, a Group A member pays £14.15, but would pay only £8.21 in the homogeneous scheme. The higher premium in the heterogeneous scheme is due to the heterogeneity caused by the Group B members, who have a higher wealth and lower probability of death to the Group A members.

However, the effects of heterogeneity diminish as the number of members in the heterogeneous scheme increase. For example, if we double the members in the heterogeneous scheme, so that Group A has $\ell^A = 900$ members and Group B has $\ell^B = 100$ members, while keeping the probabilities of death and wealth values of each group member unchanged, then the effect of the heterogeneity diminishes. For a minimum mortality credit of £1,000 a Group A member in this larger heterogeneous scheme would only pay £1.81 (Table 5). While this premium is lower than in the homogeneous scheme, there are twice as many members in the enlarged heterogeneous scheme as in the homogeneous scheme.

It is straightforward to extend the premium calculations outlined in this section to encompass multiple homogeneous groups in the scheme.

The overall message is that members can receive a minimum, guaranteed income from mortality risk-sharing at a low annual cost. The figures shown are meant only to be illustrative, and do not include any safety loading. However, they show that heterogeneity within the scheme, in terms of number of members, their probabilities of death and wealth, should be taken into account when setting the premium for each scheme member. Heterogeneity increases the premiums, but it can be reduced by increasing the number of members within each homogeneous sub-group, where a sub-group in the scheme is homogeneous if the members have the same wealth and probability of death.

5. Summary

The potential to increase retirement income by pooling mortality risk is attractive in the current economic environment. We have considered how the mortality risk-sharing rule proposed in Donnelly *et al.* (2014) could be applied in a scheme, to enhance the income received by retirees. While standard assets could be permitted in the scheme, such as investment funds, property can also be included. Adding bank deposits and savings would require a partnership between the manager of the scheme and the banks.

The scheme offers the ability to monetise housing wealth without requiring people to sell their house to release its value. In the proposed scheme, the retiree knows how much of their housing wealth will be lost to their estate upon their death. The housing value at risk upon death may be more easily understood by an average person than that in an equity release scheme, since there is no interest calculated and thus people are not required to do a compound interest calculation. However, the scheme does not provide the same benefits as an equity release scheme. In an equity release scheme, housing wealth value is released upfront as cash to the customer. In return, the customer must eventually pay back this value with interest to the equity release company, generally through the sale of their house up death.

In the proposed scheme, the housing wealth is used to calculate a mortality credit; the higher the housing wealth value and the older the member, the larger is the mortality credit. The customer does not have a large cash lump-sum generated by joining the scheme, but rather an income derived from mortality credits. The customer only has to sell their house if they die while a member of the scheme. If they exit the scheme alive, they do not have any debt to the scheme and thus do not have to sell their house.

The risk of mortality anti-selection effects is high with the scheme, and this is the biggest challenge to be managed in a scheme that implements the proposed mortality risk. To mitigate these, restrictions on the ages of members at entry and exit are proposed. The restrictions aim to balance the ability of members to withdraw assets from the scheme as required, with the need to have reasonably accurate predictions of mortality, to ensure fairness between members.

As deaths occur sporadically, guaranteeing a minimum income stream is proposed. In return for a minimum income arising from mortality pooling in the scheme, the members would pay a suitable premium to an insurer. This should make the scheme more attractive to potential members as they gain the security of a minimum, regular payment.

Acknowledgements

The authors thank the reviewers whose comments greatly improved the paper. This paper forms part of the research project “Minimizing longevity and investment risk while optimizing future pension

plans” and Catherine Donnelly thanks the Institute and Faculty of Actuaries for their sponsorship of the project.

References

- Adams, J. & James, S. (2009). Retirement income and assets: how can housing support retirement?. Pensions Policy Institute Discussion Paper, London.
- Aviva & Deloitte (2010). Europe’s pension gap, available at <http://www.aviva.com/europe-pensions-gap/> (accessed 20 January 2017).
- Balchin, S. & Shah, D. (2004). The pensioners’ incomes series 2002/3. Pensions Analysts Division, National Statistics.
- Bunn, P., Le Roux, J., Johnson, R. & McLeay, M. (2012). Influences on household spending: evidence from the 2012 NMG Consulting survey. Quarterly Bulletin Household Sector Articles, Bank of England.
- Cannon, E. & Tonks, I. (2009). Money’s worth of pension annuities. Department for Work and Pensions, UK Government. Research Report No 563, London.
- Choi, J., Laibson, D. & Madrian, B. (2010). Why does the law of one price fail? An experiment on index mutual funds. *Review of Financial Studies*, 23(4), 1405–1432.
- Department for Work and Pensions (2012). Family resources survey, United Kingdom 2010/11.
- Department for Work and Pensions (2014). The pensioners’ incomes series 2012/3. National Statistics.
- Donnelly, C., Guillen, M. & Nielsen, J. (2014). Bringing cost transparency to the life annuity market. *Insurance: Mathematics and Economics*, 56, 14–27.
- Elton, E., Gruber, M. & Blake, C. (2001). A first look at the accuracy of CRSP mutual fund database and a comparison of the CRSP and Morningstar mutual fund databases. *Journal of Finance*, 56(6), 2415–2430.
- Equity Release Council (2017). Equity release market report, Autumn 2017.
- Ji, M., Hardy, M. & Li, J.S.-H. (2011). Markovian approaches to joint-life mortality. *North American Actuarial Journal*, 15(3), 357–376.
- Lusardi, A. & Mitchell, O. (2007). Financial literacy and retirement preparedness: evidence and implications for financial education. *Business Economics*, 42, 35–44.
- Office for National Statistics (2012). Pension trends – Chapter 2: Population change.
- Office for National Statistics (2013). Pension trends – Chapter 7: Private pension scheme membership.
- Office for National Statistics (2014). Pension trends – Chapter 3: Property wealth, wealth in Great Britain 2010–12.
- Prudential (2014). Over two million homeowners over 55 bank on downsizing, available at <http://uk.prweb.com/releases/2014/05/prweb11869634.htm> (accessed 20 January 2017).
- Sodha, S. (2005). Housing-rich, income-poor: the potential of housing wealth in old age. A paper for housing across the lifecycle, Institute for Public Policy Research, London.
- Spreeuw, J. & Owadallya, I. (2013). Investigating the broken-heart effect: a model for short-term dependence between the remaining lifetimes of joint lives. *Annals of Actuarial Science*, 7(2), 236–257.
- Wilson, S. (2007). Annuitant longevity improvements. UK “Dear Chief Executive Officer” letter. Financial Services Authority.