
Still living with mortality: the longevity transfer market after one decade

Abstract of the Edinburgh Discussion

[Institute and Faculty of Actuaries, Sessional Research Event, Edinburgh, 29 January 2018]

This abstract relates to the following paper: Blake, D., Cairns, A. J. G., Dowd, K. and Kessler, A. R. Still living with mortality: the longevity risk transfer market after one decade. *British Actuarial Journal*, 23. doi: 10.1017/S1357321718000314.

The Chairman (Mr W. D. B. Anderson, F.I.A.): It is a pleasure to chair this evening's sessional meeting. This paper has been prepared by the Actuarial Research Centre and sets out their review of how the longevity risk transfer market has developed over the last 10 years.

The paper is a rich resource and provides a fantastic number of insights which are to be covered in the presentation.

I have worked in this market for the last 10 years. It is an exciting and innovative place to be. One forgets all the little milestones that have been passed over the last 10 years.

It is worth reminding ourselves that the UK leads the world in this area and that there are several countries now trying to mimic the UK in developing expertise in longevity risk trading.

It is also useful to consider the amount of pension-type risk, which includes longevity risk, that is likely to be transferred from defined benefit pension schemes to insurance companies.

There have been several commentators who have been talking about 2018 as being a record year for the volume of risk transfer with figures of around £30 billion. That is potentially 100% more than in 2017.

There has been discussion about participation from the capital markets. To date, the majority of the longevity risk transfer activity has taken place between insurance companies. The paper covers the activity in the capital markets, which is less obvious than the reinsurance activity.

The paper also highlights many of the barriers that have progressively been overtaken as the market develops.

Turning to our speakers, Professor Andrew Cairns, who works for Heriot-Watt University here in Edinburgh, and David Blake, who works for CASS in the City of London.

Professor A. J. G. Cairns, F.F.A. (general introduction): Before we start to talk about the paper, I will give an overview of where this paper fits into the wider project.

We have the Actuarial Research Centre of the actuarial profession in the UK. The ARC, as it is called, is funding some major commissioned research projects, of which this is one. It is also aiming to develop a wider international network of researchers, particularly with the aim of focusing on research which is a genuine collaboration between academics and practitioners and, with that in mind, focusing on problems which are of a very practical nature mostly motivated by industry.

Our particular research programme over the 4 years running up to 2020 also has significant funding from North America.

Our research covers four main themes. Tonight's paper is a small part of that but also covers both the modelling work that we are doing on a bigger scale and also longevity risk management.

That is an introduction to the project. I will now hand over to David Blake to introduce the paper and make some comments.

Professor D. Blake (introducing the paper): It was 12 years ago when we had the first paper, "Living with Mortality."

Today we are presenting our thoughts 12 years ago about how this market would develop. Then how it turned out, followed by what is going to have to happen in the future for the market to really grow given all the activities taking place in the UK. The UK is a relatively small market compared with the other countries that have the same problem of longevity risk, which is not being borne in the right place and has to be transferred to other players and investors who are more capable and interested in holding that risk in return for a risk premium.

Our agenda is as follows; first, a little background; then we look at the successful solutions that have taken place so far; then we look at the challenge between index and customised solutions and the issue of basis risk, which has been a real problem to get the capital market solutions off the ground. We want to analyse that and see what can be done; following on to the important role of mortality models. Andrew Cairns has been one of the leaders in looking at those issues including stochastic mortality models. Then to give a very brief coverage of what has happened since 2006. We will be able to do only a taster. Finally again looking into the future about what the future might look like.

Starting with the background, the longevity risk transfer market started in the UK in 2006. It is now a global market. At the time, the capital market solutions were expected to dominate.

Three years before Swiss Re had issued a mortality bond called Vita, and there were subsequent volumes. There is a tartan version of that but the volumes were actually quite low.

Then in late 2004, the European Investment Bank and BNP Paribas announced the issue of a longevity bond. That was actually based very closely on a paper that Billy Burrows and I wrote and published in the Journal of Risk and Insurance in 2001. We thought that this is it. The market is going to take off.

But there was no interest among pension plans to buy this bond. The bond was never launched. Since then it has been insurance-based solutions that dominated.

In retrospect, it was more sensible that the insurance companies stepped in and started to look and solve that problem. But are we going to argue that there is not enough capacity in the insurance industry to deal with all the global longevity risk that is out there.

So our two questions are: why was it that insurance solutions dominated, and then how long can this continue?

The other issue was that in 2004, the EIB longevity bond, prompted me to hold the first longevity risk conference in London at CASS business school. The most recent one was the 13th conference in Taiwan and we were encouraged to see Douglas Anderson talk at that conference. It is a conference in which academics and practitioners get together and try to learn from each other. The next one is going to be in Amsterdam in September.

First, we needed to quantify the size of the market as discussed in section 2 of our paper. The market is enormous if you take into account the public commitment, the state commitment and the private sector commitment estimates of the potential size of pension liabilities, estimated to be between \$60 trillion and \$80 trillion. The various holdings of this risk include the USA with the biggest, \$14 trillion; the UK with \$2.6 trillion; then Australia, Canada, Holland, and so on. These are huge amounts of quantifiable risks on corporate balance sheets. These countries are simply not able to manage this risk.

That is partly because actuaries told them it was a deterministic risk. It was a fixed known risk. As mortality tables were only being updated once every 10 years, and there were shocks to find it was not a fixed deterministic risk at all, it was a growing risk with a trend in it that was not anticipated.

But at the moment only the UK, the USA, Canada, and Holland have the conditions for a longevity risk transfer market to develop. These conditions include low-interest rates. We know that these low-interest rates have had the effect of telescoping the liabilities in pension plans.

Inflation, at least in the UK and Holland, is another cause of the increase in those liabilities. The frequent updating now by the actuarial profession of longevity projections; the introduction of market-consistent valuation methods; and increased accounting transparency of pension assets and liabilities on balance sheets are all affecting stock prices with increasing deficits. It is affecting insurers, shareholders and whether money will go into increasing investment for future growth in companies or to paying off or reducing, these pension deficits.

In addition, there are increased powers by the regulator to allow them to dictate how company profits are to be used.

Who are the stakeholders in this market? In section 3 Figure 1 of our paper we set out a chart that we have used for many years. We have potential buyers of longevity risk protection on the left and potential sellers of longevity risk protection on the right. The defined benefit plans exposed to longevity risk are going to be the buyers. Annuity providers are also exposed. They could be providers, but they could also write protections to gain exposure to longevity risk if they want to do so.

Life insurance companies are also potential sellers that can form a natural hedge in their book of business because US insurers, for example, have a lot of life book risk. People do not buy so many

annuities in the USA, so they are very happy to buy an annuity book risk, even if it comes from, say, the UK. There are reinsurers, life settlement funds, and so on.

Clearly, there is the potential for a market, and these players are playing in that market, but it is not a liquid market. It is not a capital market yet.

The first thing to talk about is the successful insurance-based solutions. These can be classified as customised indemnification solutions, because the insurer fully indemnifies the hedger against its specific risk exposure. They can also be thought of as “at the money” hedges, since the hedge provider is responsible for any increase in the liability above the current best estimate assumption on a pound for pound basis.

The classic solutions here are buy-outs and buy-ins. Then there are also quite a lot of longevity insurance contracts, or we might call those insurance-based longevity swaps.

The successful capital market solutions are very small in number compared with what we had anticipated in 2006. Swiss Re issued a longevity spread bond in 2010. The longevity swap was actually invented by an insurance company. That was the natural transaction that the capital market, the investment banks, JP Morgan and so on, tried to sell; the q-forward contract which I will explain.

Then more recently what has been successful is the tail-risk protection in the form of longevity bull call spreads. Andrew Cairns will discuss some examples of that later.

The key feature of these solutions is that they are indexed rather than customised solutions. That itself has caused problems because the trustees did not understand what index-based solutions were. They did not like this idea of basis risk. They wanted a solution that removed the problem entirely from their book of business. Even if potentially these are cheaper and more liquid, the trustees at the time, in 2006–2007, did not adopt this kind of solution as perhaps they did not understand what was on offer.

In Figure 4 of our paper, we show the classic, simplest forms of solution. Only one or two of these have been issued, the q-forward or mortality forward contract; a very simple swap structure between a pension plan and a hedge provider. It can be a fixed or floating transaction and is similar to an interest rate swap or a currency swap in its structure.

Further successful capital market solutions include the tail-risk protection, or what has been classified as a longevity bull call spread: all of these have been developed in Holland, but could apply to the UK. There have been five of them to date, two involved Aegon. One was Deutsche Bank and the other was Société Générale (Soc Gen). Two involved Delta Lloyd and RGA in 2014 and 2015. The most recent one was in December between NN Life and Hannover Re.

Prof Cairns (introducing the paper): The Soc Gen deal is quite similar in many respects to the other transactions in the Netherlands. It is perhaps less familiar to the people in the audience compared to some of the other deals that David Blake mentioned, like longevity swaps.

The Soc Gen deal highlights a particular difference compared to most of the activity in the UK in that these deals are focused on insurers themselves wanting to hedge their own exposure to longevity risk,

whereas many UK transactions involve pension plans. Later we cover what the differences might be between pension plans and insurers in terms of what type of longevity management they want.

Insurance companies are one of the key features of the Dutch deals. What they are trying to do through these contracts in particular is to manage their regulatory capital, but you can equally argue that it is about economic capital in a broader sense.

All of them have been indexed-based hedges rather than customised, using national population mortality. That is easier for the insurance companies because they have a more diversified portfolio of risks compared to a typical pension plan.

As David Blake mentioned, it is an option type of contract, a bull spread, but it is out of the money so there is only going to be something paid back to the insurer if the longevity goes up beyond what is anticipated by a certain level beyond the current expectations.

With the type of liabilities, these might be 50, or 60 years out into the future in terms of the cash flows. This is our challenge, a possible capital markets solution, because ultimately the capital market investors in these types of contracts do not want to have contracts that are tying them in for such a long-time period.

The particular example of Soc Gen was that it was for about 15 years' duration. Not all of the other contracts are for 15 years, with the most recent about 20 years. That is rather shorter than the 60 plus years for the entire liabilities. The design of the hedge, which is, given the type of contract, how you minimise the amount of population basis risk in order to maximise the amount of capital relief that you get. Thinking about minimising basis risk, there are three things put into the contract design that help with that. First, it is about specifying all of the individual cohorts that are going to be covered in both males and females in all their individual ages. Second, you can select the exposure vector. In particular, in some cohorts you might have much more in the way of financial commitments than other cohorts as well as thinking about their decline over time, or how you expect them to decline. That goes all the way out to cover all of the potential payments and not just up to 15 years.

Third, there is an "experience ratio matrix" that is included in the contract specification. The purpose of that is to adjust the national mortality rates to push them up or down so that they match more closely the hedger's own expected rates of mortality, which could be quite different because of the socio-economic mix of their business.

In section 5 Table 4, Exposure Vector, of our paper, we show an example of what the cohorts might look like. This is the first two of those points – the males and females at a variety of ages. This is all relative. If you look at 1,000 in the top left, the reference point, the amounts are relative for each cohort, and you can see how the numbers are expected to decline over time.

Section 5 Table 5 shows the contract specification and the experience ratios. This is a very important part of the design of this contract.

The key point is near the bottom we have the population mortality. That is the reference population, the national population, q^{pop} . That will get scaled up or down using these experience ratios so that approximately that will be equal to the hedger's own mortality rate, the $q^{\text{hedger}}(t, x)$.

It will be a bit different because the actual outcome in terms of the uncertainty might be a bit different because of a lack of perfect correlation. But, as a starting point, this is a very important part in terms of optimising the hedge and getting a minimum amount of population basis risk.

After you have adjusted the mortality rates, you calculate the so-called "synthetic survivorship rates"; and with the experience ratios put in place, then the survivorship will decline at approximately the same rate as the hedger's own survivorship function for the individual cohorts.

Thinking further ahead, about the exposure period of 55 years versus the desire for a more limited term contract, all of the five Dutch deals include some form of a "commutation function" which is effectively trying to value at time 15 the remaining cash flows, a sort of value hedge rather than trying to do a perfect cash flow type of hedge.

This is achieved through a combination of factors. You are using the actual population mortality experience over the first 15 years of the contract and then the remaining cash flows, the remaining exposure, is calculated, so there is a valuation which is done by using a "re-calibrated" projection model. That is using the mortality experience that you have had over that first 15-year period.

That re-parameterisation is a part of thinking about how things would work in practice. In practice, you have more data and you recalibrate your model. For a contract itself the way that you do that has to be specified in the contract right at the outset rather than after 15 years and then there are some negotiations that could certainly fail.

At the outset you need to select the model and also the initial parameterisation. Then there is an objective process that is well documented for updating the parameters when you have the actual data. After that you carry out that re-parameterisation of that model. So you are following the prescribed procedures, including the model itself used for forecasting mortality, whatever that model might be.

Most important is incorporating that 15 years' worth of extra mortality data. Then ultimately you are using the forecast that you get at that time to calculate the present value of the remaining exposure.

The graph in section 5 Figure 5 gives a taster of that. These would be the mortality rates relative to some starting point. You have the historical mortality covering the first years up to, in this case, 2014. We then have the 15-year exposure period for the contract. In the context of carrying out risk management, there is more than one sample path but many possible outcomes of what could happen over the next 15 years.

Stochastic modelling can be adopted so that for each one of the outcomes shown there is then, after the recalibration of the model, a single projection based on the combination of the historical and the 15-year risk period. So instead of the lumpy bits in the middle, you have a smoother forecast. Each one of those longer term projections is reflecting what happens in that 15-year risk period.

Coming back to the contracts, it was an “out of the money” contract based on these future mortality outcomes. There are two strike prices – an attachment point and an exhaustion point. These are normally specified with reference to the distribution of what you simulate as being the “final index values.”

For example, you might choose the attachment point to be the 75% quantile and the exhaustion point to be the 95% quantile. You are covered in terms of the risks in that particular slice of the distribution.

Then the actual mortality experience feeds into the final index value commutation calculation, and includes the base table at that 15-year horizon as well as mortality improvement rates after that. It is base table plus improvement rates that give you the commutation calculation.

Given the resulting distribution, there is a process of negotiation between the hedger and the provider of the hedge in which they are attempting to optimise the balance between the hedger’s capital relief on the one hand and the price of the option.

In the example in our paper, where we are showing some of the graphical examples coming from the Michaelson and Mulholland paper, there is a best estimate of longevity liability, \$1 billion. This is based on not a 1-year horizon in terms of 99% confidence but a longer term horizon. There is \$1.2 billion of the quantile at the 15-year horizon to try hedge. At the moment, without a hedge in place, they would need \$200 million of economic capital.

After review, they put a hedge in place with an attachment point of \$1.1 billion and an exhaustion point of \$1.2 billion meaning a maximum payout of \$100 million. There is no payout at all if the liability turns out to be \$1.1 billion or less than that. In between \$1.1 billion the \$1.2 billion the payout from the hedge provider gradually increases and then it maxes out at \$100 million if the actual liability at that 15-year horizon turns out to be more than \$1.2 billion. This is set out graphically in section 5 Figure 6.

With the contract, in the example, you would get up to, but not necessarily, \$100 million. The reason why you might get less than that \$100 million would partly depend on the assessment for the impact of population basis risk, which cover later. A very important point, which some companies in the Netherlands discovered a bit too late, was that it was actually a good idea to discuss what you are doing with the regulators early and get approval for the contract, and approval in terms of the level of capital relief you can claim.

Moving on now more generally to indexed versus customised hedges. There is a great deal of debate about this, particularly at the conferences that David Blake organises every year. This is about whether standardised index-based hedges or customised hedges are better or worse for individual stakeholders.

Section 6 Table 6 of the paper, with some further discussion now, sets out the advantages and disadvantages of each.

The advantages of an index-based hedge, theoretically at least, the index-based hedges linked to a national mortality should be cheaper because you are not paying for the additional cost and maybe

an additional risk premium to hedge your pension plan's very specific idiosyncratic risks, whereas the national population does not have that so, in theory, it should be cheaper.

It is a well-understood index so it should be cheaper to set up, particularly if the market begins to develop. A shorter maturity has the advantage of lower counterparty credit exposure. That can be managed in a variety of ways.

The disadvantages of the index-based hedge are that it is not a perfect hedge for various reasons, including the much discussed basis risk, roll risk, and base table estimation risk. Although it is not mentioned, I would also argue that for one or other stakeholders for a customised hedge a base table estimation risk is also something that needs to be thought about.

For customised hedges, you get an exact hedge. That is obviously a major driver for pension plans in particular. Once it is in place, you can relax and forget about at least that particular block of business.

However, theoretically, it may not always in practice be exact, be more expensive than a standardised hedge and potentially higher costs to set up. A lot of time is spent in terms of getting all the data, et cetera, and analysing that. Also the customised hedge has lower liquidity relative to the index hedges, which might or might not be a problem for the hedger but the receiver of that risk might have more of an issue with that.

Credit risk is present and is less attractive to capital markets investors. It is possible to visualise various hedges in a variety of ways, for example by carrying out calculations first without a hedge in place and determine how much risk you are exposed to. Then you put the hedge in place, and by way of an example from our paper from about 10 years ago.

In section 6 Figure 8 is "Hedge of effectiveness of q -forwards" we summarise one process. If you have an index-based hedge in place using q -forwards, our results show an index hedge, so it is not perfect risk reduction, but nevertheless a calculated 80% risk reduction. It is not so much the detail of the risk calculation, it is more the processes that you need to be thinking about in order to assess what type of hedge you want to go for.

Basis risk, which is the residual risk associated with imperfect hedging, is not a perfect alignment between movements of the underlying exposure and hedging instrument. A much talked about part of basis risk is the population basis risk, that is, the national mortality versus pension plan mortality. Are they perfectly correlated or is it somewhat less? That varies a lot, depending upon how diversified or how concentrated your portfolio is in a particular socio-economic group, for example.

You also have base table risk. That can affect stakeholders in different ways.

Turning to structural risk: this is about the design of the hedging instrument, and in particular if you have something like the bull spread in the Soc Gen contract. It is not a perfect hedge compared to a swap. You have to think about that and what the consequences are.

Following on to the restatement risk: this covers the impact of errors or restatements of national mortality tables. Then idiosyncratic risk in your population is that which is being hedged. Again, that is another factor to think about in terms of basis risk.

Our next area is mortality models: there is quite a wide variety that are in use. They go from the very simple through to the three types of model that are more sophisticated in their approach. We have the extrapolative or time series models, where we are looking at the historical data, and then forecast on the basis of statistical analysis of that data.

Then we have other types of approaches to the explanatory or causal models where we are looking at things like socio-economic status, lifestyle factors, and so on. We are trying to use a bit more information about the individuals in the population to inform us better about future forecasts as well as the base table.

Last, we have process-based models that are looking at the more medical underpinnings of that mortality. So, the extrapolated models have plenty of advantages. In some ways, they are more objective than other approaches.

Depending on which models you are using, changes in the last few years in terms of the slowing down of mortality improvements have been problematic. How well do our models cope with that? Some extrapolative models are doing well. Some of the older models, particularly ones based on the random walk time series model, do not do so well.

Then there are issues about whether we now have a temporary blip in mortality rates, or a change in long-term trends and how long that will take to feed through.

Then in the other two groups we have, for example, the RMS longevity risk model. It is always a little bit difficult to find out what the real detail is in that model. The key point there is that it is more trying to model or subdivide the improvements in mortality down into a variety of medical type or related causes. That is what they call “vitagion categories.” One category is the lifestyle changes, including smoking, health environment, medical interventions; and then the more interesting ones, looking far into the future, are things like regenerative medicine, which you could separate from the last point as it is more about the repair of the body, whereas the retarding of ageing is more about simply slowing down the ageing process.

There is plenty of research that is going on in both of these last two areas. Although there has been little in the way of historical, medical mortality improvements for these last points, it is plausible in the future that some of these research areas will bear fruit and they will have an impact on actuarial liabilities.

Then more on the causal side of things, socio-economics. We have the Club Vita data, which is a fantastic resource, where we have collection of a lot of data that include a lot of socio-economic information about individuals. You can follow these individuals to see how mortality rates vary according to different socio-economic class.

There are also other datasets which you can piece together from public sources, both in the UK and other countries, including Denmark, in particular, which is one that I have spent a lot of time looking at.

The point of this is that through these causal factors you can, for different subgroups, do a more focused projection and try to look at all of the different groups at the same time and forecast them jointly into the future so that there is a consistency between them.

The Club Vita data, for example, looked at smoking as a factor, in addition. That is more commonly associated with lower socio-economic classes. On the other hand, their findings are that the malignant melanoma, skin cancer, is more commonly associated with the higher groups. They are going off in the sun in the summer and not taking their sun cream with them.

Other approaches are taking insights from medical science, improved data but also cause of death. That is a part mentioned in the middle, but there are other studies also looking at cause of death in order to better inform not just what has happened in the past but also what is maybe happening in the future.

Applications of mortality models: one application is to think about how you determine the longevity risk premium. You start off with your best estimate forecast of mortality. If you were doing a swap contract of some sort or some other deal, then as the hedger you are going to end up paying more than the best estimate. That is the risk premium.

The stochastic models, and other models, can be used first of all to assess how much risk there is for the different time horizons or the different cohorts. Then, on the basis of the amount of risk, the risk premium should be proportionate, or reflect these different levels of risk.

Another application is to look at using stochastic models to think about estimating regulatory capital relief. They are already used for assessing regulatory capital, particularly in the UK. If you put a hedge in place as an insurer, you have these discussions with the regulators and you want to know how much capital relief you can get as a result of that hedge.

The experiences that we have are particularly in the Netherlands rather than in the UK but there are good lessons for us in the UK to learn and we may see some transactions of this type with UK insurers. The Dutch national financial regulator looks at each transaction on a case-by-case basis.

There was a ruling in 2016 that led to some not great outcomes for a couple of insurance companies. That led to a temporary halt in the development of that market while people took stock, companies took stock, of what was being said. Then, based on some other research, they shifted the attachment and exhaustion points, a little bit away from the 99.5% quantile.

In particular, if you look at the deal done in November or December last year, the NN Life deal, with Hannover Re, it addressed a lot of the problems but it included early discussion with the regulator and they managed to get capital relief relating to the size of the layer that was being covered.

The last application is fair comparison of the longevity risk management options.

A good approach to risk management is not just to focus on one way of hedging your risk. You might want to look at two or three different options then weigh one up against the other. That might include indexed versus customised hedges. The hedge itself could be partial or possibly full risk reduction.

When you are comparing the results, you want to think what the metrics are. What measure are you going to use to compare the different options?

Then the comparison is going to depend on a number of things: the risk premium that the receiver is wanting or demanding for our particular deal; the risk premium versus how much risk reduction is there and; what your risk appetite is.

If you do not have risk appetite well specified in an objective way, it becomes rather more difficult to compare the different options. These are different for pensions plans and insurers, in terms of what your risk appetite would be, but also in terms of the mix, and what your portfolio looks like. Pension plan portfolios tend to be more concentrated; insurance portfolios perhaps more diversified.

Then population basis risk is another factor and the size of your portfolio.

The key point is that one size does not fit all, so the best solution for one company or one pension plan might not be the same as for another one.

Prof Blake (introducing the paper): I am very briefly going to go over developments in the market since 2006.

Figure 16 in our paper shows how strongly that market has been growing and the dominance of the transactions in the UK; over \$100 billion of buy-ins and buy-outs and \$88 billion of swaps. The USA, which is a much bigger market, is catching up quite rapidly. Canada as well.

Innovation has been a continuing feature of this market. Buy-ins and buy-outs with deferred premiums; phased de-risking using the sequence of partial buy-ins with an “umbrella” structure so that there is only one set of contracts to negotiate. For accelerated buy-ins; the insurer provides a loan to the plan equal to the deficit. A partial buy-in can take place immediately.

Forward start buy-ins are a standard buy-in with a start date delayed to reflect the level of planning available but with the ability to bring forward the start date for an additional fee. Automated bulk plan transfers were introduced in 2017 to reduce risks. Top slice buy-ins are to target the highest value liabilities.

Named life swaps focus on named member; tranching by age to reduce costs on the basis that future buy-out is cheaper in the case of a buy-in for pensioners if they happen to be up to age 70 compared with the buy-in for over age 70. So working out what tranches of the pension plan membership you want to deal with.

Longevity swaps used to be only available to very large plans, over £500 million, but now the smaller plans are able to access the swaps.

Novation is the ability to transfer a hedge from one provider to another – the beginning of the introduction of some liquidity in the market. Then longevity swap to buy-in conversions. Phoenix Life did this in December 2016. Insuring away the extreme tail of liabilities in a closed plan.

Increasing optionality in contracts improves flexibility to meet the real needs of different types of plans, for example, allowing for the option to convert from RPI to CPI – unlike British Telecom, who were not able to do that.

Combining longevity management solutions such as interest rate and inflation swaps, ETV and PIE exercises, and bulk annuities in a buy-out.

“Buy-out aware” investment portfolios are used to reduce the buy-out price volatility and close the funding shortfall. Improved arrangements for handling data errors are used to reduce pre-deal requirements and post-deal uncertainty.

Arrangements to handle the deferred members were one of the early problems that could not be handled, only pensioners in pensions in payment, but 45% of the DB plans are deferred members, so including deferreds is improving insurer appetite to assume additional risk and costs.

Then a look into the future. We looked into the future last time and we got it wrong.

A number of potential solutions were suggested in 2006 in our Living with Mortality paper. We thought longevity bonds would take off. We had a whole range of ideas for this, zero-coupon bonds, deferred bonds, principal-at-risk longevity bonds, and longevity spread bonds. Only one bond, the Kortis bond, has been issued.

Other potential solutions included mortality, longevity and annuity futures, mortality options, longevity caplets and floorlets, and mortality swaptions. These were all solutions available in the capital markets. They were not original for us. They had been around a long time in the capital markets. We thought let us design versions of that for longevity. So far that really has not yet been successful.

Prof Cairns: A potential way forward is building on some of the things that we have seen, such as thinking about the risk management chain because, particularly for pension plans in the UK, the population basis risk is one of the key issues that persuades them that longevity swaps are the way to go or buy-ins and buy-outs, and so on, because you are getting that pool risk reduction. But that can be part of a longer chain because their risks are often going to reinsurers or an insurance company; and then, further down the chain you could have something like a catastrophe type of bond. The Kortis bond is an example of that. Then there are also elements of the Dutch transaction that I mentioned earlier.

In Figure 17 of our paper we summarise the process: at the top we have the initial hedgers, such as pension plans. These are hedgers that really do not want to be bothered with the population basis risk. They just want to get longevity risk off their balance sheets, so the chain starts to step 1. They go for customised swaps, for example, and they go to one particular reinsurer.

The key point about the role of the reinsurer as part of this chain is that they do lots of these swaps. Through careful choice of who they deal with, they can much more easily get a diversified set of liabilities, for example, they can have a blue-collar pension plan offset by a white-collar pension plan, and get a rather better mix.

By doing that, they can end up with a portfolio that has a much higher correlation relative to the national population. Then that allows them to do something that would – and this is the middle layer of the chain – be our longevity linked derivative payoffs. That bull spread contract, for example. That would be interacting with something like a special purpose vehicle rather than, say, another

reinsurer, because the objective is that the special purpose vehicle then acts as a conduit through to the wider capital markets.

On the one hand, you have a bull spread option for the reinsurer but the capital markets prefer a bond type of format. Catastrophe bonds have done quite well in the non-life insurance market. That is one example of where thinking about the bigger chain allows us more broadly to think that there is a place for index-based transactions as well as customised transactions.

Prof Blake (continuing introduction): The next idea, reinsurance sidecars, comes from one of our other co-authors, Amy Kessler, from Prudential Financial. That was announced at the Chicago conference we had, Longevity 12.

Figure 18 of our paper sets out the structure. This is a financial structure to allow external investors to take on the risk and benefit from the return of specific books of insurance or reinsurance business.

It is typically set up by reinsurers looking to either partner with another source of capital or an entity to enable reinsurers to accept capital from third-party investors. This is a way of slowly bringing in third-party investors who really do not understand what the risks are, but as long as the original reinsurer stays engaged and co-shares the downside risk, then there is some assurance to these new types of investors, the sovereign wealth funds, and so on, that they are not being “sold a lemon.”

This is an interesting new idea. Only just a few weeks ago a new sidecar structure called Langhorne was announced in the USA. RGA was one of the reinsurers. They now have other types of investor in a sidecar structure co-sharing the risk.

This is an exciting development in which the reinsurers are actually bringing in capital markets investors in what now seems to be a much more natural way than we actually thought at the time 10 years ago.

To conclude, longevity risk inherent in the world’s aggregate retirement obligations far exceeds the amount of risk capital that the global insurance industry can realistically bring to bear.

At the end of 2016, there was just under \$600 billion worth of reinsurance capital, to deal with the \$60 trillion–\$80 trillion pension liabilities that are potentially out there.

Vast sums of additional risk capital must be introduced to manage longevity risk. This can only come from the global capital markets.

There four major challenges.

The first is currently long-term investors prefer bonds but most of the successful solutions for hedging longevity risk to date have been in the form of longevity swaps. Short-term mortality bonds have been a success not long-term bonds so far. However, the Swiss Re strategy of gradual iteration from a successful innovation is a good way forward. Swiss Re have the mortality bond, and their development of a Kortis longevity spread bond was a modest adaptation of that Vita Bond.

More creative approaches to the risk management chain are needed to meet the differing requirements and risk appetites of the two sides of the market. The two key prizes are a much bigger

investor base and much greater market liquidity. That is the big benefit of capital markets, you start to get liquidity.

The second challenge is there needs to be a common agreement between the counterparties on the mortality model to use for the design and pricing of longevity linked solutions. That was the lesson from Aegon's deal with Soc Gen which went ahead in 2013. Both sides agreed to use the same mortality model, which happened to be the RMS model. Even if the mortality model produces the wrong forecasts – which it is bound to do – as long as those forecasts are not systematically biased, one way or another, then it can become a potential candidate for use in this model. That is another possible breakthrough here.

The third challenge is that there are a number of operational issues that need to be dealt with: these are basis risk, credit risk, collateral and liquidity, which are all discussed in the paper. That requires the market participants to work out the optimal trade-offs between basis risk and liquidity on the one hand and between credit risk and collateral on the other. The regulator needs to be involved right from the beginning in the negotiations over the deals. The regulator must be willing to grant fair levels of regulatory capital relief for these index-based solutions. If they do not do that, then the capital markets will not start.

We want a fair game here for a level playing field. So the capital relief has to be compatible with current solvency capital requirements but consistent with the levels of capital relief already made available to the fully customised hedges.

The fourth and final challenge is again a regulatory issue. The regulatory response to the global financial crisis had a very big effect in slowing down the establishment of the longevity risk capital market.

The Dodd-Frank Act, and Basel III, effectively meant that the investment banks had to leave that market. JP Morgan left the market. It was only investment banks like Goldman Sachs and Deutsche which happened in the UK to have insurance subsidiaries, Rothesay and Abbey Life, that enabled those banks to remain in the market. But for different reasons even those two banks have now left the market. We have to deal with long-dated illiquid credit exposure associated with longevity transactions because that increases the capital requirements. There are still many regulatory issues. Thankfully, the Dodd-Frank Act is being wound up, so that helps the US investment banks possibly re-enter the market.

These four challenges will need to be addressed for the next stage of the development of this market. But innovation has been an important feature of the market since 2006, and there is every reason to believe that this will continue as the different players in the industry seek to reduce costs, optimise capital and manage risks.

The Chairman (opening the discussion): In my firm the pension fund side find this area a little challenging, particularly operating life insurance. We are more familiar with the terminology and the ways of managing risk.

It certainly helps me to understand, when I am engaging with people from a pensions background, and to realise that in the risk transfer market that is now very active in the UK, virtually all the risk

that is being taken on by the bulk annuity companies in the buy-in format, which is by far and away the most dominant type of transfer, is automatically longevity reinsured at the moment.

It may not be apparent to all that this market is going on, and is very, very vibrant in what effectively has become back-to-back reinsurance trading on the back of every buy-in deal that takes place. So the buy-in price in the UK, for example, becomes a kind of add-on. The investment risk is layered on top of the longevity price that has been provided by longevity reinsurers, typically.

Mr S. D. Baxter, F.I.A: I would like to make a couple of observations on some of the challenges that the authors raise. First, one of the major challenges raised is the operational issues to be dealt with. Within that, I suggest that there is one milestone and one gremlin that we need to be particularly alert to.

The milestone is getting the PRA over the hurdle of giving regulatory relief for transactions without a material haircut. It still to this day remains untested, and unproven, whether there will be a haircut applied to an insurers suggested risk reduction, and if so the quantum of it.

That said, as soon as a transaction successfully attracts regulatory capital relief, this may well open the floodgates. It is a major concern that many in the industry have; the extent to which the PRA will be accepting of the capital benefits of such transactions.

Any comments that the authors have in this regard, and things we can do to move the regulatory dialogue forward, will be much appreciated.

The operational gremlin that is not often talked about is what happens after we develop a liquid and traded index market. This makes longevity “mark to market” that does not necessarily sit very well in some of the current regulatory regimes for insurance company balance sheets where this could introduce additional market volatility to the liability side of the capital balance sheet.

This another important, but surmountable, operational hurdle. I hope the authors have some comments they can offer.

Looking forward, though, I feel excited by the market. It strikes me that the key attractiveness challenge you refer to in the paper is one of liquidity versus basis risk. How to ensure a liquid market that means we can get capital markets interested; versus removing the basis risk, which is the primary concern in achieving capital relief sought within insurance companies.

To date, it has been suggested that this is best achieved by population-based indices using age groups split by gender. I wonder whether there might be an alternative balance to be struck. Two thirds of the liabilities that will sit on any balance sheet are typically concentrated on just one third of the population in terms of socio-economic background. Perhaps there is some simple stratification of the UK population that would better close the basis risk?

In this regard, we should be aware of the modelling versus intuition paradigm. However, much our stochastic models can demonstrate a theoretical capital relief, there will always be an intuitive concern of: “Are my policyholders quite the same as the national population? Do I dare take that risk? Will the PRA, or other relevant regulatory body, view the index-based contract as providing a good hedge?”

There must be scope to bring some additional dimensions beyond age and gender to the index market to alleviate this concern. I would again value your views.

Finally, I remain very optimistic. If we can make longevity index transactions work, it opens up the floodgates for international capital to flow in to the longevity markets. You showed most of the risk is in the UK and the USA. There is a massive amount of capital available in the Asian markets, keen to access the UK and US longevity markets, but currently there is no simple way of doing that.

Index solutions and bonds may be just what is needed. I look forward to a very exciting next 10 years.

Prof Blake (responding): That was a very good set of questions. As we mentioned, and you have specified, and I have heard you speak often on the issue of the regulator and the PRA, it seems that they need some education in the same way as trustees need some education. We have not found a successful way of educating the regulator in this way.

I am just wondering what the best way will be to get the regulator to understand that what we are asking for are very realistic ways of introducing a new way of hedging this risk. They have to understand, all sides have to understand, that there needs to be capital there. It should be close to economic capital. There needs to be protection for the consumers.

This kind of aversion that the regulator seems to have to index-based solutions, even in Holland (which you would have thought would be much more open to this, given that it is a much closer, smaller country in which these people talk to each other all the time), the regulator had real problems with dealing with this.

The question that I would ask back is: what can we do about educating the regulator in this? There is a real hurdle and barrier until we do that.

The next thing mentioned is the issue of mark to market. If we have distorted markets, as we appear to have since the global financial crisis, in the form of quantitative easing, the distortion to the interest rate market, then what does mark to market mean? That has to be addressed as well.

Another interesting point is about the concentration of the risk in the top socio-economic group and focusing on that. How best to do that? It is interesting to see that you are telling us that there is a lot of capital in the Asian economies that could be applied to this. It will be interesting to see if there is a way that we can help to bring the parties together at conferences, and so on, actually to engage them better.

How otherwise will we get these investors to understand the issues, and that they are not being “sold a lemon” and there are ways in which they can co-share the risks, say, through sidecars?

You have raised a lot of interesting operational issues. We cannot answer them but we should find ways of jointly trying to do so.

Prof Cairns (responding): I will add one or two points on the regulatory side of things. I did some work with the Dutch regulators and I realise that although they are somehow operating under the

same Solvency II type regime, the way that they approach their calculation of regulatory capital is a little bit different and, in my mind, a bit behind how it is being done in the UK.

In the UK, for insurers, they are properly doing 1 year ahead, 99.5% VaR, whereas in the Netherlands it is a little bit hybrid. Nevertheless, the general mechanism that they are using to work out the regulatory capital relief I think would be equally applicable in the UK but with small modifications which should not be philosophical. It is just a case of changing time horizons and things of that sort.

There are lessons to be learnt. But as suggested, it needs an insurance company to go to the regulator and say, “We are thinking about this. Can we talk and discuss what can be done?”

When we spoke to the regulator and had a session just over a year ago, so far as I recall they said, “Well, nobody has come to us with this type of suggestion.” That might have changed.

It does need the insurance companies to start thinking about that.

Mr D. O. Forfar, F.F.A: I was struck to learn recently that the expectation of life (at birth) in America is falling because of opiates and various other reasons. We have all got used to longevity increasing every year, particularly in the UK, but whereas the rate of improvement has slowed down in the UK, it is actually falling in America!

I am wondering what the authors think the future might hold? In the future, our best predictions are based on anticipating improvements in medicine, gene therapy, etc. But if these improvements do not materialise as predicted, or even if there were, in future, actual falls in the UK expectation of life (as has happened in America), it would be startling but also very interesting.

Prof Cairns (responding): I will reply briefly to that. Perhaps it is a little bit beyond the scope of our paper but, as part of our wider research programme, this is very much a topic that is high on our list of things to be looking at. We are considering a lot of different countries, trying to see what are the similarities and differences to the USA, Canada, Denmark, and the UK. They all have their own idiosyncratic characteristics.

It is when you go down into the socio-economic groups – the Club Vita analysis will also bear that out, I suspect – you can see quite big differences between different groups.

In much of the slowdown in the USA, it is the lower working class groups that are really causing that slowdown, whereas other wealthier groups are perhaps carrying on at a faster rate of improvement and not slowing down so much.

Prof Blake (responding): I will add that you are just reflecting the debate among the demographers. On the one hand, in the USA, some say it is definitely slowing down, and for the reasons that you said. On the other hand, there are those who are saying that there is no upper limit to life. Aubrey de Grey was saying that the first person to live to 1,000 is currently alive. That is longevity risk.

The Chairman: It is very interesting that economists look at something called the Gini coefficient as a measure of income inequality in different countries. I think of a Gini coefficient as a leading indicator of longevity inequality. One follows the other.

We have to be very, very circumspect about interpreting these national population numbers that are very popular in newspapers, because it is just a single number that they can digest.

It is a pretty good example of the law of averages, that one single national population number.

Mr R. J. Mellows, F.I.A.: Thinking about the sources of basis risk and the things that worry trustees and also the PRA. On the one hand, you have the choice of model point, and issues around the basis risk between the population versus the portfolio. You can chip away at that as we have discussed. We have talked about this in your talks today. Steven Baxter talks about socio-economic sub-model points, so they can be incorporated if you like.

The thing that we have not really talked about but which must be a significant source of this basis risk, is this unwillingness of the capital markets to push their time period out beyond 15 years.

This leads to a messiness of the commutation and trying to agree in advance what basis you might use in 15 years' time to capture the remaining risk beyond that point.

How could this be quantified? You have the basis risk from the model points and the basis risk from the commutation. Is there a sense which is the bigger one, the most disliked by regulators? Do you have views on the possible hedges?

Given that the 15-year end-point feels significant, is there any wriggle room in that? Can you push the capital markets to take 20–25–30 years? That will start to become more appetising to those wanting to hedge the risk.

Prof Cairns (responding): That is a very good question for us to be thinking about, and for regulators to be thinking about.

First of all, in terms of the time horizon, you would actually be quite surprised about how good the 15-year maturity is at actually managing your exposure. That is viewing the exposure to longevity risk as of now at time zero. In other words, what you would get at 15 years is very highly correlated with the ultimate run-off.

If you are looking at a run-off of the entire risk, and looking at present values today, then you get a very high correlation.

Then the totally different approach is the 1-year approach with the regulators for Solvency II. That is where we need to have that dialogue with the regulator early on in terms of what we propose and how the contract will work. There is no doubt about what the model will be.

On the other hand, you also have to agree with the regulator that, for our own valuation purposes, this is the model that we will use in 1 year's time. If they agree that, or they agree an approach, then that is all you need, I would say, to get a decision on how much capital relief you should get. I am more optimistic about that. It is not a perfect solution, but if we can get that agreement, and the way

that Solvency II calculations get done, there is a load of approximations in terms of what is going on, being able to agree that our model will be used in a year's time is one of those, perhaps.

Mr Mellows: That is a good point. The solution is an incremental approach. However, we should recognise that the people who want to lay off the risk have the risk out for 70 years so 15 years remains quite short.

The initial group of investors that was approached were cat bond type of investors who were used to maturities of 2 or 3 years. So you have to bring them along incrementally. 15 years is a long time for hedge funds and private investors. But if you can try to do this incrementally, and it is not a catastrophe and the investors get roughly the return they expected for the risk premium expected, then you can incrementally move it out.

It is going to be a slow process. The investment banks at the very beginning realised that this was going to be a very, very slow process. Unfortunately, they were driven out of the market. They were planning to start with the cat bond investors, then bring in the sovereign wealth funds incrementally.

But at each stage you have to have success. The market has to be a success. There is enough liquidity for people to move out if they want to move out and you incrementally try to push those maturity periods. There is no alternative to doing this incrementally.

The Chairman: If we were to reconvene in 10 years' time, maybe you could put your necks on the line and in a few words describe what you think the world will look like in 10 years' time.

Prof Blake (responding): If the insurance companies support something, it should happen. The sidecar proposition looks as though that will work. Academics can stand up and say, "This is a great idea. Do it." But it was the insurers over the past 10 years who made this market work. It was not the investment banks. The investment banks were too far away from the trustees. It was the insurers who understood the trustees much better and understood their needs, understood they needed customised solutions, they needed 100% indemnification, and they got a really successful market going.

It was naive of us to think that the capital markets would be able to do this from the start. It was the insurers that did it. But as we said, there is not enough shareholder capacity out there for the insurers to go on, to grow this market much further, much faster, without additional investors. That still does not mean the investment banks will walk back in; it is the insurers who will control this because they will be the ones that do the deals with the sovereign wealth funds, and so on and so forth.

They are actually going to play the role that the investment banks would otherwise play, but they are going to do it at their own pace and for their own needs. The sidecar approach is potentially very interesting on that point.

Yes, there will be capital market investors, but it will not be the investment banks that are successful here. It will be the insurers and the reinsurers that do this.

Prof Cairns (responding): For me it is really about that risk management chain. The sidecars are one way to relieve some of the capacity constraints with the reinsurers. There are other ways through that risk management chain that will begin to develop through cat bonds or other sorts of formats. That may well be alongside the sidecar sorts of arrangements.

Ultimately, the demand is there and it is going to push things in a way that will result in some interesting solutions. They will gradually develop. They are pushing out the boundaries from 10, 15, 20 years perhaps in terms of how these deals work.

The other thing that I would like to see, and what is desirable, would be things like the socio-economic indices. Steven Blake mentioned that they offer very interesting prospects to make hedges of a particular type more attractive than they would be, helping or reassuring people that there does not need to be so much in the way of population basis risk.

The Chairman (closing): Thank you very much, Andrew Cairns and Steven Blake. We have it on record. In 10 years' time, we can replay today. Let us hope we are all here to see the event.

I thank everybody, particularly the authors, for their contributions, coming along and their questions.