

Colm Fagan: Presentation to Society of Actuaries in Ireland 7 February 2017: accompanying commentary

1) Title

Thank you, Roma.

A couple of comments before we get into the meat of the presentation:

First, the views expressed tonight are my own, not those of any of the organisations with which I'm associated.

Second, there's a lot of material to get through – 70 slides, to be precise. If I were to try to do justice to them all, we could be here all night. I don't want to inflict that on you. My plan is to finish presenting in 50 minutes or so and to leave plenty of time for discussion. That means I can only spend an average of 45 seconds on each slide. I will have to skim over some of the content, but the slides will be put onto the Society's website, so you can study them in detail later at your leisure.

2) Disclaimer

This is the Society's standard disclaimer.

3) Drawdown: where are we now?

This slide sets out the three main drawbacks of the current drawdown regime. I will deal with each in turn.

4) High charges

Trustees have an obligation to look after members' interests, but that obligation ceases on retirement. Once a member retires, they are on their own. The trustees and sponsoring employer wash their hands of them. A 2016 report of the Pensions Council concluded that charges on individual insurance-based arrangements are equivalent to a yield reduction of between 1.5% and 2% per annum.

5) Low investment returns

I was astonished when I read the first statistic shown here, that over 40% of insured ARF's are 100% in cash. That plays havoc with investment returns. Risk aversion is the main reason why so much is in cash. We will explore risk aversion in detail in later slides.

6) No security of income

For someone in drawdown, all that matters is their life expectancy, not the average. It's no good telling them that their life expectancy is 26.4 years or whatever. Their life expectancy could be anywhere between zero and 40 years. There is a risk that they could draw down too much or too little.

As far as annuities are concerned, the problem is not just that the money is tied up in low yielding bonds, but also that the undrawn funds are lost if the annuitant dies early.

7) How the new approach addresses current deficiencies

This slide summarises how the new approach addresses each of these drawbacks. It results in lower costs; higher investment returns and greater security of income, even into extreme old age. The risk of people outliving their savings is eliminated.

8) If Carlsberg did pensions ...

The proposed approach is so close to perfection that I decided to retitle the slide!

9) End result: potential income more than double that of an annuity

And this is the end result.

10) Prerequisites for proposed approach

In relation to the first prerequisite, the Pensions Authority is already moving towards allowing members to remain in the scheme post retirement. On the second prerequisite, the approach works best for large schemes: think Diageo, CRH, the banks, Intel. For the third one, some relatively minor tweaks of pensions and tax regulations may be required, but nothing insurmountable.

11) Key challenges

These are two key challenges. The third challenge, of lower costs, is achieved by allowing members to stay in the scheme post retirement.

12) Key challenges – maximise the expected net investment return.

The first challenge is to maximise the expected investment return at an acceptable level of risk. That means capturing the equity risk premium. When

I refer to the equity risk premium, by the way, I include other real assets, such as property, in the definition, not just equities.

13) The equity risk premium.

The Reserve Bank of New York completed a comprehensive analysis in 2015 and concluded that the equity risk premium was between 5% and 6% per annum.

14) ERP – a prospective assessment

KPMG Netherlands completes regular assessments of the prospective ERP. At end 2017, they estimated it at 5.5% per annum.

15) ERP – a prospective assessment (2)

A simplistic approach to estimating the prospective ERP comes up with a figure of 5.25% per annum. It's obtained by assuming a dividend or rental yield of 3.25% and real growth of 2% per annum. I should emphasise that it's precisely what it says - simplistic – but it's in the right ballpark.

Private equity and other illiquid investments, including property, can deliver a higher return. The approach I'm proposing allows significant investment in illiquid assets.

The bottom line is that it is reasonable to expect an equity risk premium in the region of 4% to 6% per annum over the long term.

16) Rich rewards for capturing the ERP

This slide shows the rewards to be reaped from capturing the equity risk premium. As noted at the bottom, returns can be further boosted by the lower costs resulting from allowing members to stay in the scheme post retirement.

17) But ERP rewards come at a high price

But there is no such thing as a free lunch. Not only can the index fall very suddenly, as it did in October 1987, when it fell by over 20% in two days, but market downturns can also be prolonged: the index remained below its August 2000 level for five years.

18) ERP rewards carry a high level of risk.

Linda Evangelista famously said that she wouldn't get out of bed for less than \$10,000. I see nods of agreement from the consultants in the audience. A

saver could equally claim that it wouldn't be worth their while putting money in the stock market if they could earn more by leaving it in the post office. That was the case between December 1999 and October 2013.

19). FTSE All-Share index: 1986 – 2017.

Stock market volatility is shown graphically on this slide. The index falls more frequently than one month in every three. Over the last 32 years, there were 12 monthly falls of more than 8%, including one of 26.5%.

20) Loss aversion makes matters worse.

Loss aversion is probably hard wired into us by evolution: he who fights and runs away will live to fight another day. As someone who claims to be an experienced investor, I can vouch for the fact that price falls, such as those experienced earlier this week, can be quite unnerving.

21) Hindsight bias militates against stocks.

This quotation from the Nobel Prize-winning behavioural psychologist Daniel Kahneman is particularly apposite for investment advisers and brokers. One can easily understand their reluctance to advise customers to put money into something where the odds are less than 2-1 against that they will lose money – possibly a significant amount – in the first month.

I have a personal recollection of Tony Taylor placing a full-page ad in a Sunday Newspaper on the day before Black Monday in 1987, advising clients to buy equities. The market fell by more than 20% over the next two days. Other advisers would have taken the lesson from that experience that they should be careful about selling the merits of equities to their clients.

The clear conclusion at this stage is that, while investment in real assets delivers superior long-term returns - of the order of 4% to 6% per annum over bonds - the short-term risks militate strongly against this form of investment.

How can we resolve the conundrum?

22) The solution to ERP conundrum?

And here is my solution- smoothing.

23) Key principles of smoothing formula.

The smoothing formula must be transparent. We don't want an actuarial black box. It must also be calculated objectively and be easy to apply.

It must strike a balance between damping the volatility of short-term changes in market values while remaining faithful to long-term trends. Two other important criteria are that the formula should remain unchanged over time and should minimise the risk of adverse selection.

24) Proposed smoothing formula.

I am proposing an exponential smoothing formula. It gives a weighting of just 1.5% to the current month's market value and a 98.5% weighting to last month's smoothed value, increased by one month's interest.

25) Formula applied to FTSE all share index 1986 – 2017.

This graph shows the results of applying the smoothing formula to the FTSE All-Share index between 1986 and 2017. Visually, it appears to have done a very good job of damping the volatility of the index.

26) Revisit monthly changes in index

In order to see how well it has damped short-term fluctuations in market values, we revisit the graph we showed earlier of monthly changes in the All-Share index.

27) Monthly changes (smoothed).

We now look at the graph of monthly changes in smoothed value. It contrasts sharply with the previous graph. There were only two months out of 384 when the smoothed value fell, and it fell by less than 0.1% in both occasions.

28) Smoothing: interim scorecard.

We now revisit the criteria for determining the quality of the smoothing formula.

It gets high marks for four of the five criteria, but what about the fifth, that it should avoid the risk of adverse selection?

29) Risk of adverse selection.

Whenever smoothed values are less than market values, there is a risk of adverse selection by people buying in on the cheap. Conversely, when the

smoothed value is greater than the market value, there is a risk that investors will leave the fund, pocketing more than the market value of their investments.

30) Adverse selection risk.

Someone joining in the late 1990s could have bought in at less than 70% of market value, while someone selling in the early 2000's or in 2008/2009 could have netted 150% of market value. Continuing members lose out in both cases. They foot the bill, so it's vitally important to take effective action to reduce the risk of adverse selection to the absolute minimum.

31) Minimising the risk of adverse selection.

Here, I set out the proposed rules to minimise the risk of adverse selection. The first is that people can only join on retirement. If they don't join then, they are not allowed to join later. People will not choose their retirement date purely to take advantage of smoothed values being less than market values. On the other hand, they are unlikely to decide not to join if smoothed values are greater than market values when they retire, as that would cause them to lose the benefits of smoothing for the entire duration of the retirement. An additional protection with new joiners is to drip feed the money into the fund over two years. As far as taking money out of the fund is concerned, retired members will be obliged to take regular withdrawals that follow a smooth progression. On death, similar phasing rules will apply to those imposed on entry.

32) How much could adverse selection cost (1)?

Despite the safeguards, there is likely to be some adverse selection. As mentioned earlier, the cost of adverse selection is met by continuing members, in the form of a reduction in the return credited to their accounts. I modelled the possible cost of adverse selection by assuming that the natural order is a steady monthly stream of new retirees, each with identical investments and each taking a regular income of 5% per annum.

In that situation, with no adverse selection, the smoothed price rose by an average of 8.5% per annum over the last 32 years. The fact that it's less than the 9.5% average return since 1985 is simply due to the fact that the 9.5% is an average for money left invested for 32 years, while the 8.5% is the average for regular monthly investments over the entire period.

33) How much could adverse selection cost (2)?

We now assume that new investment falls by 50% in any month in which the market value is less than 90% of the smoothed value. This would imply new investment falling by 50% from 2000 to 2004 and through all of 2009.

The result is that the smoothed return for continuing members would fall by just 0.1% per annum, from 8.5% to 8.4% on average.

In the extreme, if new money were to dry up completely whenever market values were less than 90% of smoothed values, the impact would be a reduction of just 0.2% per annum in the average return over the 32-year period. It is worth adding however, that this is an average over the entire period: the reduction would be more than 0.2% in the affected years, but not enough to cause any concern.

34) US similar to UK.

The graph of actual versus smoothed market values for the US is similar to that for the UK.

35) Risk of adverse selection for US stocks?

And the exposure to the risk of anti-selection is also very low in the US.

36) Looking good for smoothing formula ...

At this stage, everything is looking good. The formula works brilliantly for the US and the UK over the last 32 years. Even on the draconian assumption of no new money in any month in which the market value is less than 90% of the smoothed value, the reduction in yield is equivalent to less than 0.2% per annum over the entire period. It also looks good for the last 118 years in the UK and the last 92 years in the US.

37) But...

But there are problems ahead. The clue is in the flag in the bottom right. A big shark is hiding under that flag.

38) Japan: Topix index.

The Japanese market experienced a phenomenal boom in the late 1980s and then experienced a catastrophic collapse in the early 1990's. It took 25 years to recover. The smoothed index assuming constant new business (the red line) looks reasonably okay, but market values are less than smoothed values for

almost the entire period from 1991 to 2014. It is reasonable to assume that new money would have dried up during that period.

39) Japan: impact of adverse selection.

On the more realistic assumption of new money drying up after 1991, this graph shows the smoothed price falling year on year. The falling price is due to the increasing burden on continuing members of paying more than asset value to people leaving the fund.

40) Japan: unsustainable burden for continuing members.

This graph, which compares unsmoothed values assuming constant new money with unsmoothed values if new money dries up completely when market values are less than 90% of smoothed values shows Japan going into a death spiral.

41) Cumulative cost of adverse selection contrasting UK/US experience with Japan

This graph shows the contrast between the cost of adverse selection in Japan and the corresponding cost in the US and the UK. The cost of adverse selection is unsustainable in Japan whilst it is minimal in the UK and US.

42) Japan: aberration or precedent – the bubble?

In asking ourselves whether the Japanese experience was an aberration, which can be ignored, or a precedent, where we must allow for the possibility of a repetition, we look first at the boom and secondly at the bust and the hangover. This slide shows some statistics from the bubble period in the late 1980s. You can take your pick as to which is the most ridiculous. Some would give their vote to the fact that the Imperial Palace in Tokyo was worth as much as the entire state of California. Personally, I look at the more mundane statistic that Japanese bank shares were trading on a price earnings ratio of around 60 as late as end 1991, when the market had already experienced a severe fall.

43) Japan: aberration or precedent – the bust and the long hangover?

The bust was almost as awesome as the bubble: the market fell by over 70% from peak to trough. A bigger problem however from the smoothing perspective was the long hangover, which was prolonged for decades by the Japanese government's failure to deal with zombie banks and zombie corporates. The zombies were kept on life support for far too long by an inappropriate actions or inaction by government and central bank. The swift actions by Western governments and central banks in the wake of the market collapse of 2008/2009 was informed in part by the lessons learned from the Japanese authorities' inept response to the collapse in their stock and property markets nearly two decades previously.

44) Could a smoothed DC fund suffer a Japan like fate?

The Japanese experience in the period 1995 to 2010 was unique. I contend that a fund invested in a range of real assets, including property, equities, infrastructure, with the investments spread across different geographies, would have a minimal risk of a similar fate. There is a lot of evidence to support that belief.

A paper published in December 2017, "The rate of return on everything, 1870 – 2015" is highly encouraging in this regard. My conclusion from studying that paper – albeit briefly - is that the smoothed returns on a fund invested in a combination of equities and housing across 16 developed economies would have sailed through the entire 145-year period from 1870 to 2015 without experiencing anything close to a Japan like death spiral.

45) Context for investment strategy.

The context for setting the investment strategy is ideal. Inflows and outflows are highly predictable; there is no risk of sudden unanticipated cash outflows, as is the case for banks, insurance companies and unit trusts, all of which are subject to the risk of sudden cash calls. Because of the low weighting given to current market value, it doesn't matter if accurate, up-to-date, market values aren't available. The result is that a sizeable proportion of the fund can be invested in illiquid assets. The final two background conditions set out on this slide mean that the investment managers can concentrate on delivering good long-term returns without having to worry too much about the short-term.

46) Proposed investment strategy - every investment should a real return of 4% per annum.

The first requirement of the investment strategy is that every single investment in the fund (without exception) should be expected to earn the risk-free return plus 4% per annum. One immediate consequence of this is that there would be no room for bonds in the portfolio. I could get my name in the newspapers for that!

47) Don't touch bonds.

As you can gather, I don't believe in the adage that you should invest your age in percentage terms in bonds.

48) Proposed investment strategy.

Other aspects of the proposed investment strategy don't need elaboration.

49) "It works in practice but will it work in theory?"

I hope I've proven by now that the proposed smoothing approach works in practice, but it will take more to convince the theoreticians.

50) "It works in practice but will it work in theory?"

Subject to the qualification that I only have monthly price changes for the UK and the US for 32 years, I'm confident that the smoothing approach would be able to withstand adverse selection for the last 117 years in the UK and the last 90 in the US - and that covers the great crash of 1929. As mentioned a few slides back, the December 2017 paper "The rate of return on everything, 1870 to 2015" indicates strongly that the smoothing approach would have worked for the entire period since 1870 for a portfolio invested in real assets across 16 developed economies. Theoreticians, many of whom are strong believers in the random walk, are unlikely to be convinced by that however. I believe they're wrong. For one, they ignore the impact of politics on markets at the extremes. They don't recognise the importance of what has variously been called the Greenspan put or the Bernanke put. I don't think we ever needed the Yellen put. Actions by governments and Central Banks have a very real impact in stressed conditions and cannot be ignored. I accept that markets must follow a random walk in the short term, to eliminate arbitrage opportunities, but there is strong evidence of mean reversion in the longer-term.

51) Key challenges

This marks the end of the first section of the presentation, which shows – I hope - that smoothing on the lines proposed enables us to capture the equity risk premium at an acceptable level of volatility.

52) Key challenges.

We now come to the second part of the presentation, which addresses the longevity risk.

53) Lifetime Income Fund.

This is where I unveil my pièce de resistance, what I call the “Lifetime Income Fund”. As the name implies, the purpose is to ensure a lifetime income for new retirees who decide to join the club. There is no compulsion to join. I’m assuming a membership fee of 1% per annum. In practice, the fee will depend on the gender and occupation mix of the membership, but 1% is a reasonable estimate for someone joining at age 65.

54) Benefits of joining the LIF

The way the LIF works is that the retiree’s investment is divided into 100 identical sub-accounts. The member cashes one of the sub-accounts every quarter, or 4 every year. Thus, after 100 quarters (25 years), they will have cashed all their sub-accounts. Nothing will be left. At that point, when they’ve reached age 90 (i.e. 25 years after age 65), the LIF takes over and continues to pay a sub-account each quarter for the rest of their life. If they die before age 90, i.e. before they have cleaned out their account, any outstanding balance is paid to the estate.

55) LIF example.

The best way to illustrate how the LIF works is through an example.

We assume an average smoothed return of 5.5% per annum, which reduces to 4.5% after deducting the 1% per annum contribution to the LIF. The amount withdrawn is 1% of the original investment each quarter, plus interest.

56) LIF example

For every €1,000 invested, 1%, or €10 (plus interest) is withdrawn each quarter, €40 plus interest each year. In this example, the amount withdrawn in year one is €41.14 and the balance remaining at the end of the year is slightly more than €1000 because the interest credited was greater than the amount withdrawn. Another 4% is withdrawn in each subsequent year until we come to year 25. Looking at the first set of circled figures, we see that €117.08 was left in the account at the end of year 24. This had increased to €120.41 by the end of year 25 and was then withdrawn, leaving nothing in the account. Then, at the end of year 26, the Lifetime Income Fund pays out another 4% of the original investment plus interest, which equates to €125.92 and so on for the remainder of the member's life.

57) Focus on extremes of probability distributions

An interesting discovery from behavioural finance is that normal people (not actuaries, of course), when faced with a range of possible outcomes, tend to focus on the extremes and ignore the mid-points of the distribution. In this case, the pay-offs are good at the extremes. If someone dies early, they have contributed very little to the LIF, and the full balance remaining on their account is paid to the estate. On the other hand, if they live to 100, they receive a regular 4% from the LIF every year from age 90. The main losers are the people in the middle, particularly those who die between ages 85 and 90, who have contributed 1% every year to the LIF but get nothing back.

58) LIF: the trustees' perspective

This slide shows the year by year evolution of the finances of the LIF. The table assumes 10,000 new retirees join at age 65, each making an investment of 100 and contributing 1% to the LIF. The LIF builds up year by year until year 25, when it reaches a maximum of 402,548, at which points pay-outs commence. Just over 50% of the original cohort (5,096) are alive at that time. The pay-out per survivor in year 26 is 12.02 for every 100 contributed, which means that 4% of the amount invested has increased from 4 at the start to 12.02 by year 25.

The balance in the LIF reduces each year thereafter because of continuing payments to survivors. By year 37, when there are 853 survivors from the original cohort of 10,000, each aged 101, there is still a nice positive balance of 89,255 in the fund.

59) LIF: the small print

The small print is reasonably straightforward. I suspect that a high proportion of the LIF members will choose not to take their 4% entitlement each year, that they will leave something aside for the rainy day. Any such savings will boost the solvency of the LIF, because the 1% LIF contribution will still be charged on those amounts. It's also worth noting that someone can opt out of the LIF at any time, but they will lose any contributions paid to date.

60) Comparison annuity v drawdown

This slide compares an annuity (on the LHS) with drawdown plus longevity protection (on the RHS). The current market rate for a life annuity for a male aged 65 is €4,000 per €100,000 invested. Total payments, assuming survival to 100, amount to €140,000. The corresponding payments under the drawdown alternative, where the member has opted for the LIF, amount to €340,000, or well more than double, almost 2.5 times, the amounts payable under the annuity. Furthermore, on early death, there is a lump sum payment under the drawdown option. If the annuitant dies after age 70, nothing is payable under the annuity option.

61) Where's the catch?

There is no such thing as a free lunch, so where's the catch? The main risk is that the assumed investment return of 5.5% pa will not be achieved. That's unlikely, as the implied ERP of 3.75% is less than expert consensus.

The other major problem with the drawdown option is that money cannot be withdrawn at will. As the slide states, it's a pension, not a piggybank. The constraints on withdrawals are essential to reduce the risk of adverse selection. Nevertheless, we still must consider what would happen in the unlikely event of a Japan-like spiral. My proposal is that, if adverse selection cuts the smoothed return in any 12-month period by 1% compared to what it would have been without adverse selection, then the trustees must change the smoothing formula to give greater weight to current market values.

62) Capital implications?

The key argument is that no guarantees are being offered by either the trustees or the sponsoring employer – all the costs fall on the members - so there are no capital implications.

I am suggesting that the employer invest seed capital in the fund at the outset. That will have liquidity, but no capital, implications for the sponsoring company.

Given that the fund is at its most vulnerable in the early years, it would be a good idea for the employer to promise at the outset that the smoothed return will not go negative for the first three years.

63) Possible cost of promising a minimum return in early years

I have put together an extreme hypothetical example, to ascertain what a promise on these lines might cost the employer.

Assume that a fund was established on 31 October 2007, just prior to the crash of 2008/09, etc. (as per text of slide).

64) Results for October 2007 start date

Despite the steep falls from the stock market high of October 2007, the smoothed return in this hypothetical example would have been negative in just one month – February 2009, and then only marginally. (Continue reading slide).

65) Next frontiers?

Three other areas come to mind immediately where the proposed approach might be of interest. I will look at each in turn.

66) Individuals and small group schemes

The biggest problem with extending the approach to post retirement arrangements for individuals and small group schemes is the risk of adverse selection by self-employed and proprietors of small businesses in deciding when to “retire”. That same option doesn’t exist for employees of large companies.

The conclusion is that insurance-based arrangements could be used to extend the approach eventually to individuals and small groups, but the adverse selection risk would have to be addressed.

67) Why not extend to pre-retirement?

When I started on this project, the aim was to use smoothing to allow a higher proportion of the default fund under group DC arrangements to be invested in equities. I concluded that it wouldn’t work because of the variety of options

there are for scheme members, trustees and sponsoring employers to select against the fund.

A certain large insurance company in Ireland tried something on these lines a number of years ago but had to throw in the towel when the crisis of 2008/09 hit.

As noted at the bottom of the slide, it will seem incongruous to be encouraging active service members who have chosen the “lifestyle” option to shift to “safer” assets as retirement date approaches and then investing 100% of their money in so-called “risky” assets post retirement.

68) Suitable for auto-enrolment?

Read the slide.

69) Next steps

I have made some bold assertions this evening, in particular, that if we were to back-test the proposed approach over the last 145 years for a hypothetical portfolio invested in a broad mix of real assets across a number of geographies, it would have delivered stable smoothed returns, without experiencing anything close to a Japan-style death spiral.

I stress that I haven't done this detailed back-testing myself. I have based my conclusions on high-level results from the paper “The rate of return on everything, 1871 – 2015”, which was published in December last, just as I was finalising my conclusions for this paper. The paper helped to confirm my belief in the approach I'm proposing. I would like an academic researcher to take up the challenge of checking the veracity or otherwise of my assertion.

Assuming a favourable conclusion from that academic study, the next step is to find a champion, ideally the sponsoring employer and trustees of a large DC scheme, who are committed to delivering better outcomes for their members. I don't see this as an act of altruism. The first company to implement this approach for its DC scheme will be a very desirable place to work and will succeed in attracting and retaining high quality employees.

70) Can we do it? Yes, we can!

I believe that future generations of retirees will thank us if we in the actuarial profession become advocates for the approach being proposed here this evening. It will take courage on our part to challenge conventional wisdom,

but I believe that we as a profession are up to the challenge. Can we do it? All together: Yes, we can!