Institute for Research in Economic and Fiscal issues

IREF Working Paper Series

Projecting State Pensions

Kevin Dowd

IREF Working Paper No. 201304

June 2013

In English: en.irefeurope.org
In French: fr.irefeurope.org
In German: de.irefeurope.org

Institute for Research in Economic and Fiscal issues

PROJECTING STATE PENSIONS

By
Kevin Dowd[1]
July 29 2013

Abstract

This article sets out a modelling framework for the analysis of the long-term financial health of state pension systems. It also presents some illustrative results based on a simulation model applied to a hypothetical but not too unrealistic state pension scheme calibrated to the UK. Results indicate that this pension scheme is massively underfunded. These results cast major doubt on the solvency of all defined-benefit-type state pension schemes in western economies.

1. Introduction

There has been much concern in recent years about the solvency of state pension schemes, i.e., are these schemes sufficiently well-funded that we can be confident of their long-term ability to deliver on their commitments to future pensioners? However, existing projections tend to be somewhat crude and only give a partial picture of schemes' financial health. They are usually based on simplistic deterministic projections of investment risks supplemented by ad hoc assumptions,[2] and they ignore factors such as longevity risk – the risk of pensioners living too long - despite evidence that longevity risk is an important consideration in this context.[3]

What is needed is a complete, internally coherent and practically implementable modelling approach that can be calibrated to the specifics of any given economy and to any type of state pension system, and which enables us to provide quantifiable answers to important questions such as, e.g., their likely funding shortfalls, the contributions needed to ensure long-term solvency, and so on. Implicit in this is also the need to be able to compare the consequences of alternative policy choices. Only then can policymakers obtain a quantifiable[4]

---

[1] Partner, Cobden Partners, London, and Professor of Finance and Economics, Durham Business School, Old Elvet, Durham DH1 3HY, UK. Email:kevin.dowd@hotmail.co.uk. I thank David Blake, Enrico Colombatto and Kaetana Leontjeva for numerous helpful contributions.

[2] An idea of the state of the art is given by the work of the UK Pension Commission (Turner et alia, 2004, 2005): this gives an otherwise good analysis of the problem, but lacks a coherent modeling methodology. Such projections give no idea of the risks or uncertainties involved.

[3] For example, using the APC mortality model (see note 9 below) I estimate the life expectancy of a U.K. male currently aged 25 to be about 51.8 years if one ignores longevity risk and 60.9 if one takes account of it – a difference of just over 9 years. Though these estimates are calibrated for U.K. data, they are not likely to be very different from those obtainable using mortality data sets for other western European countries. See also, e.g., Dowd et alia (2010).

[4] One needs to use the term 'quantifiable' with considerable caution. Not all uncertainty is quantifiable, and it is by definition impossible to quantify uncertainty that is unquantifiable. For what it is worth, I subscribe to a Bayesian statistical school of thought caveated by the 'black

assessment of the issues they face and of the empirically realistic options available to them.[5]

The aim of this article is to outline such an approach – in effect, to offer a modelling template - and illustrate it with example outputs based on a model calibrated to the U.K. and based on a hypothetical U.K. state pension system. The template itself is extendable to any other country – data permitting – and to any form of state pension system, hypothetical or otherwise.

A good starting point is to consider that pension systems can be distinguished by the contributions made (e.g., by employees, employers and the state) and by the risk-sharing arrangements involved (e.g., who bears the investment risk, who bears the financial risk if pensioners live too long, etc.). Two polar opposite cases are: a Defined Benefit (DB) scheme, in which benefits are defined and risks are born by the pension provider, and a Defined Contribution (DC) scheme, in which only contributions are defined and the risks are born by the plan member.

For their part, state pension schemes have traditionally taken some version of the DB model, although there is no reason why a state pension system should necessarily take this form – and, indeed, state pensions in some countries are already moving to a DC model. In any case, we need a modelling framework that is general enough to handle any form of state pension, including hypothetical ones. Only then can we assess the range of alternative choices available.

The modelling template I suggest goes as follows. We start with a model of a fully-funded individual-member DC scheme. A number of such models already exist, of which one (and the first, I believe) is the PensionMetrics model designed by Blake et alia (2001). Such models give an indication of what pension outputs might be for any given set of inputs such as contribution rate, pension fund investment strategies and so on. We then extend this model by degrees:

- We extend it to incorporate guarantees that may require outside (i.e., additional) funding (typically from that universal pot of gold, the long-suffering taxpayer).
- We extend the model so that it handles not just individuals, but cohorts of similar individuals.
- Once we have models for cohorts, we can then model populations, which are, in essence, merely sums of cohorts.
- We can then apply the model and assess the implications of different choices available to us, i.e., we can do the policy analysis.

---

swan' view that most key uncertainties are fundamentally unquantifiable: after all, most financial risk models have proven utterly useless over the last few years, at least. Nonetheless, I believe it is possible to do stochastic scenario analysis that is useful up to a point: for example, in the pensions context, such analysis can give us at least an idea of the contributions needed other things being equal to make a pension scheme solvent. I would also go further and say that without some way of modeling quantifiable risks, however limited it and they might be, there is no way to tell what contributions should be.

[5] I do not, however, address the much more difficult questions of the politics involved in making good policy choices or, going deeper, whether existing political institutions are able to rise above short-term considerations and take the longer-term view. If the modelling task is difficult, the political one is truly Herculean (or perhaps Augean).

2. A Defined-Contribution Pension Model

As a starting point, consider a model in which the individual saves for their retirement, invests the contributions into a pension fund with a specified investment or asset-allocation strategy, and anticipates retiring at a particular age. When he/she retires, the accumulated pension fund is used to provide a retirement income for the rest of his/her life. The way in which the pension fund is converted to a retirement income is known as the decumulation strategy.[6]

To illustrate, suppose that a 25 year old male with a (fairly typical) starting salary of £24k individual contributes 5% of his salary income to his pension fund, that the fund is allocated 25% to equities and 75% to bonds based on a 'funding for an annuity'[7] decumulation strategy, that the anticipated retirement age is 65, and that the decumulation strategy involves the purchase of a real or inflation-protected life annuity, which gives a guaranteed inflation-protected retirement pension income for life.

I make a number of key assumptions about the economy itself, which are generally applicable to most western European economies. These assumptions include:

- An assumed underlying real salary growth rate of 2%,.
- An assumed equity premium of 3%.
- An assumed mean interest rate of 4%.
- An assumed mean inflation rate of 2%.

I then make other assumptions calibrated to the UK economy. In particular, I assume that the employer contributes 3% of the employee salary to the pension fund, and that the government provides an additional top-up contribution of 25% of the individual's contribution.[8]

In addition, since the decumulation strategy involves annuitisation, we have to make assumptions to calibrate the annuity pricing: I assume that the annuity involved is priced using the Age-Period-Cohort (APC) mortality model[9] and that this model is calibrated using LifeMetrics data for England and Wales.[10]

---

[6] For more on these issues, see e.g. Blake et alia (2001).

[7] A funding for an annuity strategy is one in which the fixed income component of the pension fund is designed to have a duration on retirement equal to the life expectancy on retirement. This strategy helps protect the member against interest-rate risk on retirement: should interest rates be low, then the annuity would be expensive, but the bond component of the pension fund would also be higher in value and so compensate for the expensive annuity.

[8] I gloss over ancillary assumptions, e.g., about the stock market process, the interest rate and inflation rate processes and the parameters involved.

[9] The APC model is a standard one in the mortality projection literature. Relative to other models, it has the attractions of being robust, allowing for a cohort or year-of-birth effect (which is regarded as a considerable advance in this literature) and of giving reliable forecasts for all adult ages (and not just, e.g., for old ages). For more on the APC model, see, e.g. Currie (2006) or Cairns et alia (2009).

[10] Data are taken from the LifeMetrics website: http://www.jpmorgan.com/pages/jpmorgan/investbk/solutions/lifemetrics/historictable

Thus, the model can be regarded as calibrated to the UK economy going forward from the time the member starts contributing to their pension fund until the end of their life.

However, several caveats should be kept in mind:

- The resulting outputs should *not* be interpreted as forecasts – the next 100 years or so are highly unpredictable! – but rather as projections that tell us what might happen if certain (I believe fairly reasonable, but maybe not: the modeller can always insert their preferred assumptions instead) assumptions hold true. We are therefore dealing with a form of stochastic scenario analysis rather than forecasts per se.
- To repeat what was said in the introduction, we are dealing with a form of quantifiable uncertainty or risk, which, by definition, is unable to anticipate or quantify the unknown unknowns and so forth.

We are interested mainly in the accumulated pension fund. Based on these assumptions, my DC pension model gives us the following fan chart projections for the pension fund value over the course of a male individual's working life.

Figure 1

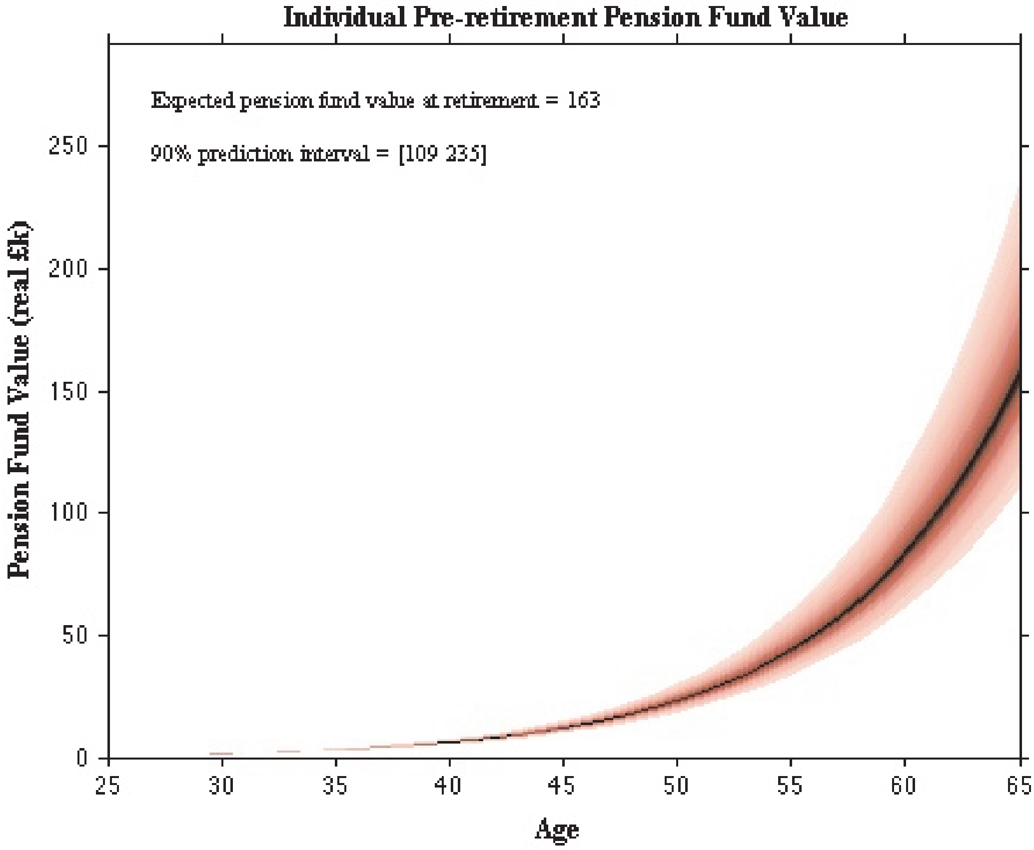

This shows how the pension fund value is projected to evolve over time given the assumptions made. This chart is actually a fan chart, in which the darkness of the projection corresponds to its likelihood: the more likely the outcome, the darker the shading. The expected PF value at retirement turns out to be £163k, and the bounds of the 90% prediction interval – given by the gap between the lowest coloured point and the highest – are [£109k, £235k]. The interpretation is that we can be 90% confident – assuming the model and calibration are correct, a very big if admittedly – that the pension fund at retirement will have a value between £109k and £235k, and is expected to be about £163k.

The next question is what this means for retirement income. Given our assumptions, the answer is provided in Figure 2.

Figure 2

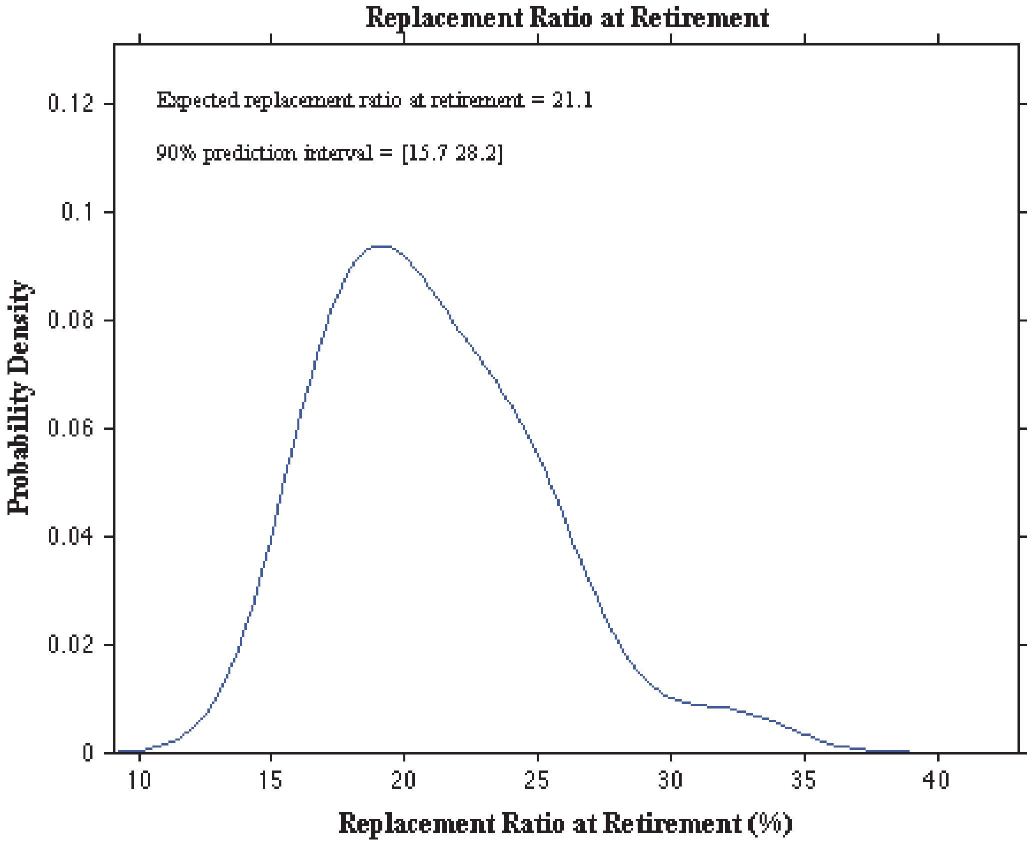

This Figure shows the projected probability density function for the projected replacement ratio – the ratio of retirement pension income to final salary – at retirement. The replacement ratio is expected to be 21.1% and has a 90% prediction interval of [15.7%, 28.2%].[11]

[11] One should keep in mind that the projected replacement ratio depends not just on the assumptions made about the economy, but also upon key decisions made, of which the most important are those relating to the contribution rate, the asset allocation strategy and the decumulation strategy. So, for example, a higher contribution rate will lead to a higher projected replacement ratio, a more risky asset allocation strategy will lead to a higher expected replacement ratio, but will typically lead the replacement ratio to be more dispersed. As for the decumulation strategy, a typical nominal (or non-indexed) annuitisation strategy would produce a higher projected replacement ratio at retirement (because the non-indexed annuity is cheaper) but will see the replacement ratio gradually eroded thereafter by inflation. An alternative decumulation strategy, much favoured in the U.S., is drawdown, in which the individual keeps their pension fund invested on the stock market, say, but draws down a pension income from it. Drawdown often leads to a riskier pension income in retirement and, of course, has the drawback

We might note, therefore, that whilst a retirement pension fund of £163k might seem a lot, it does not translate into a high retirement pension income because the projected replacement ratio is way below typical guaranteed replacement ratios such as, e.g., 67%.

Recall too that this pension scheme is a DC scheme and hence is self-funding, i.e., there is no outside party to top up the pension if it turns out to be low or to take from it if it turns out to be high. The member simply saves for their pension and gets what they get, low or high.

A state pension scheme may be set up as self-funding in this regard. However, most are not, so our next stage is to extend the model to incorporate a guaranteed pension outcome regardless of whether it has been fully funded or not. This, in turn, takes us to the central issue of the pension fund shortfall.

3. Extending the Model to Incorporate Guarantees

If the actual contribution rate happens to match the required contribution rate needed to achieve the guaranteed replacement ratio, then the scheme is fully funded and there is no net shortfall, positive or negative. Hence:

- If the actual contribution rate falls short of the required contribution rate, then the scheme is underfunded.
- If the actual contribution rate exceeds the required contribution rate, then the scheme is overfunded.

Taking account of a potential funding shortfall or surplus, and focusing on the pension fund shortfall as the key outcome of interest, we now assume a hypothetical guaranteed replacement ratio of 67% inflation-protected. Recall that this is a fairly typical guaranteed replacement ratio and is in line with what most people these days regard as reasonable to expect.

The projected outcomes are given in Figure 3.

---

that the individual can run out of funds before they die. By contrast, annuitisation prevents such an outcome.

Figure 3

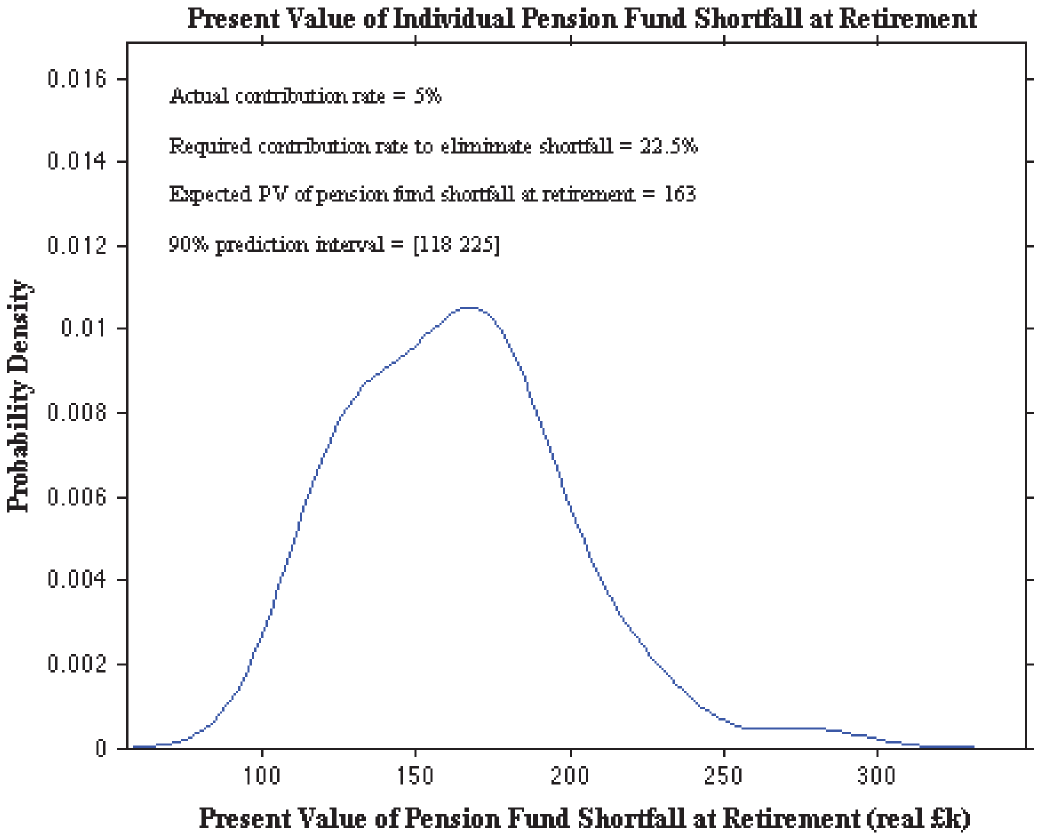

In this case, the individual contributes 5% of his salary to his contribution rate and gets a guaranteed inflation-protected replacement ratio of 67%.

That this guarantee is very expensive can be inferred from any of the following indicators:

- For the same contributions he now gets an inflation-protected replacement ratio of 67% instead of the expected inflation-protected replacement ratio he had earlier of only 21.1%.
- The expected pension fund shortfall at retirement in present-value terms is £163k or nearly 6 times current annual salary.
- There is a huge difference between the actual contribution rate (5%) and that needed to ensure that the fund is fully funded (22.5%) – that is to say, at a 5% contribution rate and give other parameters (projected retirement age, etc.) the individual's pension fund is massively underfunded.

Since this is a state pension scheme, either the underfunding falls on the taxpayer who makes up the difference or the scheme eventually defaults, in which case the individual does not get the guaranteed replacement ratio that he anticipated.

Note also that the shortfall is not just large, but also uncertain – the PV of the shortfall has a 90% prediction interval of [£118k £225k] - and this risk too must fall on the taxpayer if the state pension scheme is not to default.

4. Extending the analysis from individuals to cohorts

The next stage is to move from the individual to the cohort of which he is a part. The reason this matters is as follows: hitherto, we have assumed an individual and analysed outcomes conditional on their survival to the retirement age; however, from the funding point of view which is appropriate to the analysis of a state pension scheme, what is important is not the individual but the cohort, i.e., the set of all individuals of similar age and gender.

In other words, instead of considering an individual of particular age etc who is assumed to survive to retirement age, we instead consider a cohort of x similar individuals who are not assumed all to survive to retirement age. Instead, a typically small proportion of them will likely not survive and we must consider the impact of this likely pre-retirement mortality risk on the financial state of the scheme: those who die before retirement will have made contributions but will not be there to draw their retirement pensions. Thus, early deaths are good news – at least for the rest of the cohort who survive.

Working with the cohort instead of the individual-assumed-to-live-to-retirement will make a typically not massive but still important impact on the projections. Thus, Figure 3 is replaced by Figure 4 which accounts for pre-retirement mortality: the expected present value per capita shortfall at retirement falls from £163k to £136k and the bounds of the 90% prediction interval change from [£118k £225k] to [£98.1k £187k].

Figure 4:

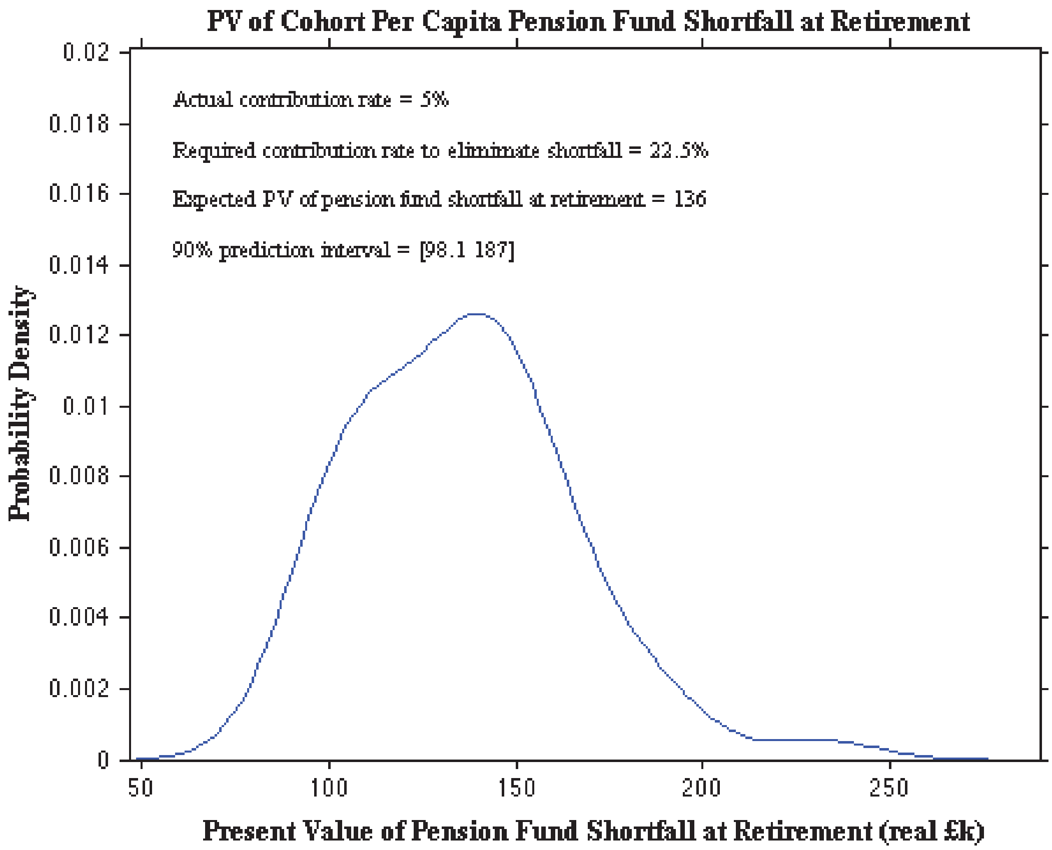

5. Extending the analysis to a larger population

The next stage is to move from the cohort to a larger population, which is in essence simply the sum of the different cohorts in the population at large.

To give a simple example, which suffices to illustrate the principles involved:

Suppose we have a population that consists of females as well as males of the age and other characteristics considered earlier.

Figure 5:

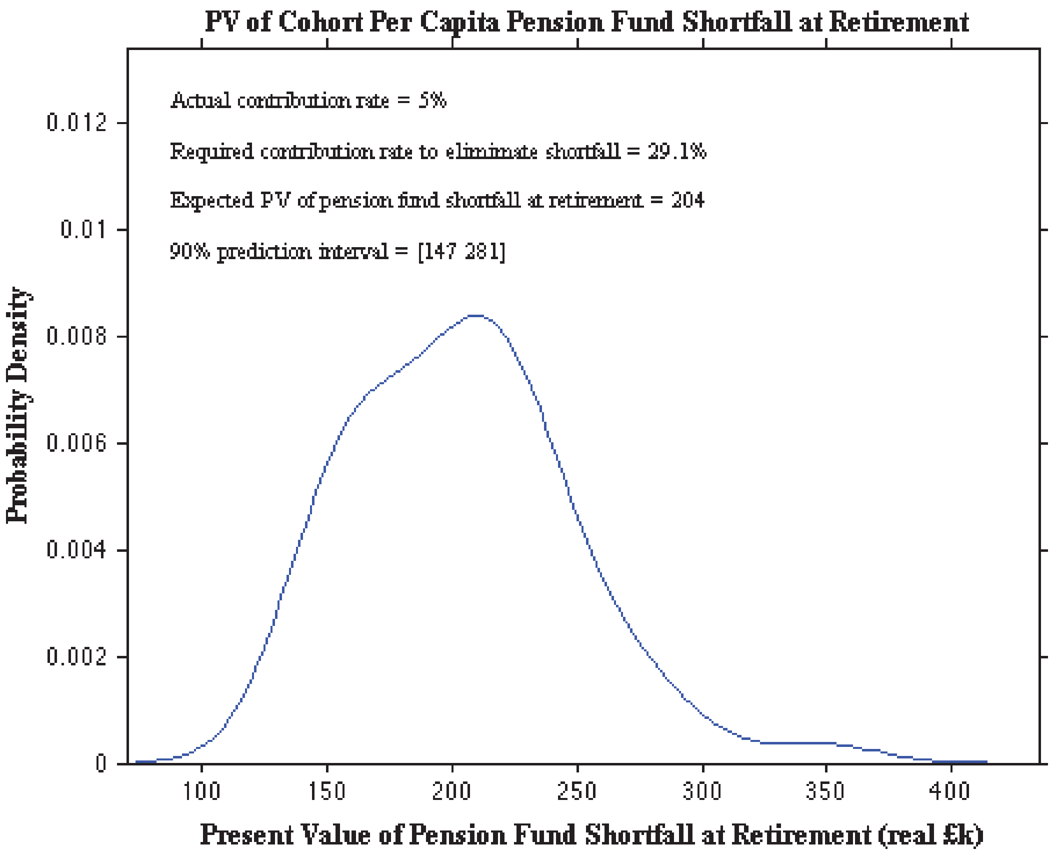

In the females case, the present value of the expected shortfall is £204k, and the bounds of the 90% prediction interval are [£147k £281k]. These are higher than for the males because of higher female life expectancy.

Thus, if we have a 'population' of one 25 year old male and one 25 year old female, then the expected shortfall based on the above results would be £136k+£204k=£340k, and lower and upper bounds of the 90% prediction interval would be £98.1k+£147k=£245.1k and £187k+£281k=£468k respectively.

In practice, however, it would probably be more convenient (and the numbers easier to digest) if one worked with their per capita equivalents. Hence, the per capita shortfall would have an expected value of £170k and a 90% prediction interval of [£122.55k £234k] respectively.

6. Using the model for the analysis of long-term fiscal planning

The model has many applications in the analysis of long-term fiscal planning, and I briefly mention three. For convenience, assume we are only interested in the cohort of 25 year old males considered earlier.

Long-term solvency

The first of these is the analysis of the long-term solvency of the state pension system. In the case of our cohort of 25 year old males, we had a shortfall with an expected PV of £136k. As I have noted already, this is a very large amount – in fact, it is 136/24 is nearly 6 times the assumed annual salary. This figure alone should be ringing alarm bells …

Furthermore, it will generally be the case that the shortfall associated with any cohort will increase as the age of the cohort rises. For example, if we have an otherwise similar cohort of 45 year old males, the expected PV of the shortfall rises to £173k, which is over 7 times the assumed starting salary. Thus, the underfunding problem generally becomes more aggravated, the higher the age of the cohort.

There is, in short, a major (to put it mildly) shortfall problem, and there are only three possible outcomes: either something is done to reduce the shortfall, or the taxpayer makes it up, or the scheme eventually defaults.

In an ideal world, fiscal policy makers would confront this shortfall problem and propose solutions. Indeed, some (e.g., in the UK) are attempting to do so, even though they have little real sense of the magnitudes involved. Others, such as in France, still have their heads in the sand – as the recent decision there to lower the state pension age back to 60 clearly illustrates. Unfortunately, the problem will not go away just because many policy makers would prefer it went away (who wouldn't?) or, even worse, continue to be in denial about it.

It would therefore make sense to consider alternatives such as raising the required contribution rate, postponing retirement, reducing the guaranteed replacement ratios or moving the state system to a fully-funded DC model.

To assess these different options, we need a properly calibrated model that gives us quantifiable answers (however limited those answers might be).

The results from our model are not especially encouraging:

- If we wish to eliminate the shortfall for our 25 year old male cohort by raising the contribution rate, then our model tells us that we would have to raise this from 5% to 22.5% (!).
- We could eliminate the shortfall by raising the retirement age, but my projections indicate that this would have to be raised to unrealistic levels before the shortfall was eliminated.[12]
- If we wish to eliminate the shortfall by eliminating the guaranteed replacement rate, then the given inputs (5% contribution rate, given AAS, etc.) would produce a replacement ratio of a little over 21%.

[12] The 'correct' interpretation of this latter finding is *not* that one realistically expects the average worker to work till to very high ages. Instead, it should be interpreted in terms of common sense and indicates that there is no realistic way in which the shortfall can be eliminated by raising the retirement age *alone.*

A partial alternative is to increase the equity weighting of the investment portfolio, taking on more risk to produce a higher expected return. This, however, is rarely the panacea it is sometimes portrayed to be. In the context of our cohort, increasing the equity weight to 50% leads the expected PV shortfall to fall from £136k to £117k and leads the required contribution rate to fall from 22.5% to 16.5%. The scheme's financial position is improved, but even so, the scheme is still very much underfunded.[13]

Transition

A related issue is transition: if a decision is made to reform the system, then how do we implement the move from here to there? There are two ways in which this can be done.

The first involves a managed transition. The keys to this are to ensure that the promised outputs are consistent with the inputs (that is to say, to ensure that the new scheme is inherently solvent) and that existing accumulated 'rights' are at least acknowledged, albeit not honoured outright. So, for example, one might say that a new worker starting at 25 has no built-up rights, but a 55 year old might have an accumulated pension fund of, say, £120k (or whatever), and this would be carried forward into the new system. One way to do this is to scrap the entitlement to a guaranteed replacement ratio, but to give the 55 year old worker the credit of having accumulated a fund of £120k (or whatever) which would then be credited into his pension fund account. Some decision would also have to be made about the payments to be made to current pensions (e.g., should their pensions be cut, and if so, by how much, etc.). This process is akin to a managed insolvency, in which the creditors get some but not all of their entitlements under the old system. The process itself is inevitably somewhat arbitrary, but at least it enables thought-out decisions to be made and there is the chance, at least, of a not-too-inequitable allocation of losses amongst interested parties (e.g., taking account of contributions already made etc.).

The second approach is a mismanaged transition or mismanaged insolvency aka chaos, in effect the collapse of the pension system and its replacement with a new system in which no credit is given for earlier contributions. This is exactly what happened in some of the ex-Soviet systems after the collapse of the Soviet Union.[14] In these cases, the creditors lose everything and the results are catastrophic. Those approaching retirement would find themselves looking at negligible pensions having been promised very generous ones – guaranteed replacement ratios of 100% were not uncommon in the Soviet Union, not that the guarantees were actually honoured in the end, but that is another matter - and would have no time to build up significant entitlements under the new system. Even worse affected would be those who had already retired, who would be left pauperised. The youngsters, at least, would be able to build up their

[13] In this particular case, the lower bound of the 90% prediction interval also falls (from £98.1k to £59.3k), indicating a reduction in lower tail risk, but this result cannot generally be taken for granted.

[14] Leontjeva (2012) has a very good account of post-Soviet pension systems.

entitlements and have some prospects under the new system – assuming, that is, that the new system was insolvent, which this cannot be taken for granted in any system with defined benefits, and leaving aside the fact that they would also have the interim burden of looking after their parents and grandparents whose pension were close to worthless

Analysis of alternative scenarios

Analysis of long-term state pensions also requires a sensible selection of alternative scenarios. To state the obvious: any state pension scheme can only be judged as solvent if analysis shows it is able to honour its obligations under any reasonable set of alternative scenarios.

Consider two examples.

In our earlier projections, we assumed that the underlying economic growth rate was on average 2% a year. But what happens if, say, the underlying mean growth rate was 4%?

In this case, the PV of the expected shortfall would rise from £136 to £334 and the required contribution rate – the contribution required to eliminate the shortfall - would rise from 22.5% to 32.1%: the underfunding is now revealed as much worse than it previously looked. This – at first sight counterintuitive result – arises because the higher economic growth rate means higher projected salary and hence greater entitlements (i.e., a higher pension). Thus, the good news of higher economic growth is bad news for the pension scheme.

Conversely, a lower economic growth rate of 0% would lead the PV of the expected shortfall to fall from £136 to £50.5 and the required contribution rate to fall from 22.5% to 15.1%. In this case, the bad news of lower economic growth rate translates into good news for the pension scheme, because accumulated entitlements would be lower.

Another important scenario is longevity. In our previous projections, we assumed that the projected longevity was constant – that is, that life expectancy would not change over the horizon period considered. However, there is strong evidence of sharply rising life expectancy in western countries[15] and we might reasonably ask what would happen to our pension scheme scenario if this rise was to continue instead of stop.

Going back to our earlier assumption of a 2% mean real economic growth rate, but adding in longevity risk, we then find that the PV expected shortfall rises from £136 to £218 and the required contribution rate rises from 22.5% to 31.4%. Higher life expectancy is bad news – at least for the pension scheme – because it has to pay out for longer post-retirement periods.

[15] See, e.g., Dowd et alia (2010).

Conversely, lower life expectancy is good news for the pension scheme, and the sad example of Russia comes readily to mind: the Russian state pension system would be in much worse shape than it already is but for the fact that life expectancy in Russia has deteriorated sharply over recent decades.

7. Conclusions

This paper has outlined a methodology to project outcomes from state pension schemes. Of these the most important is the projected shortfall. The paper also provides some illustrative results based on a model calibrated to a UK data and based on a hypothetical and not particularly generous state pension system. One would imagine, however, that the results calibrated for other western European countries would not be much different.

There is of course scope for much improvement/refinement in the modelling itself. What is possible in this regard depends in part on the data available, however, and this would vary from one country to another; nonetheless, any such improvements would lead to greater accuracy/realism in projected outcomes, which is to be welcomed. Possible improvements include:

- Allowing for scheme charges and extra entitlements (such as entitlements to extra healthcare, when extending the analysis to broader social security), etc.
- Allowing for differences in income amongst the cohorts considered, and allowing for the impact of career salary progression as workers move through their working life.
- Allowing for changes in entitlement: for example, instead of the member getting a guaranteed replacement ratio, he might get a pension determined by the final salary and the number of years worked or the mean salary and the number of years worked. Alternatively, one might have a state pension scheme on the DC model, in which there is no fixed entitlement but only an accumulated pension pot from the retirement income will be derived.
- Allowing for unemployment and other periods of break from the labour force (e.g., to raise a family): these impact on pension fund contributions both directly, since there would be no contributions when the individual is not working, and also indirectly via their impact on human capital and subsequent earnings and hence subsequent contributions when the individual returns to the labour force. [16]
- Allowing for alternative decumulation strategies, e.g., other forms of annuitisation, an income drawdown strategy, or a more realistic representation of the defined benefits e.g. in the existing UK state pension system.
- Allowing for demographic dynamics: we have hitherto assumed static demographics, but a more realistic model would take account of the effects of immigration and emigration, which would affect either or both

[16] Simulations not reported here suggest that this latter effect be very substantial, too.

of the contributions going into the pension system and or the retirement pension claims on it.

Limited as the modelling might be, the illustrative results presented here are striking, and what jumps out is the massive degree of underfunding in our hypothetical state pension scheme[17] – and, by implication, of any western state pension scheme with DB characteristics. [18]

This leaves policy makers with a simple choice: they can either address the underfunding issue up-front and try to reduce the shortfall; or they can kick the can down the road like their predecessors. In the short term, such a response – or rather lack of it – would make the problem even greater for later policy makers. But if continued in the long term, the mathematics indicate that the pension scheme will default: no insolvent scheme can last forever.

References

Blake, D., A. J. G. Cairns and K. Dowd. "Pensionmetrics: Stochastic Pension Plan Design and Value-at-Risk during the Accumulation Phase." *Insurance: Mathematics and Economics*, Vol. 29, No. 2, October 2001, pp. 187-215.

Cairns, A. J. G., D. Blake, K. Dowd, G. D. Coughlan, D. Epstein, A. Ong, and I. Balevich, "A Quantitative Comparison of Stochastic Mortality Models Using Data from England & Wales and the United States." *North American Actuarial Journal*, 2009, Volume 13, No. 1, pp. 1-35.

Currie, I. D. "Smoothing and Forecasting Mortality Rates with P-Splines." Paper given at the Institute of Actuaries, June 2006.

Dowd, K., D. Blake and A. J. G. Cairns. "Facing Up to Uncertain Life Expectancy: The Longevity Fan Charts." *Demography*, Volume 47, Number 1, February 2010, pp. 67-78.

Leontjeva, K. (2012) "Old–Age State Social Insurance: May Its Failure be Averted?."Lithuanian Free Market Institute.

[17] To give a crude ballpark estimate of the total shortfall, let's suppose that, say, 90% of all adults work in the labour force, and that each of these has the PV of the expected shortfall of a 25-year-old male, which we earlier estimated to be £136k. Then the PV of the expected shortfall per capita would be 0.9 times £136k = £122.4k.. Assuming a 2% real salary growth rate, initial salary would have grown to £55.6k and so the burden of the shortfall in PV terms would be over twice national income. Moreover, I would assert that these assumptions bias the scale of the problem downwards, i.e., so twice national income is easily an under-estimate.

[18] An interesting question is whether these sorts of results apply to poorer countries as well. Other things equal, a country that is half as wealthy as another will have half the expected shortfalls and so on, but in terms of burden (expected shortfall etc.) vs. capacity to bear that burden (annual income, etc.), there would be no substantive difference. Perhaps the biggest substantive differences in fact life expectancy and economic growth rate: other things equal, a country with a lower life expectancy will have a smaller pension 'problem' – for obvious reasons, and (as we saw on p. 13) a country with a higher growth rate will have a bigger such problem.

Pensions Commission (2004) *Pensions: Challenges and Choices. First Report of the Pensions Commission*. London: HMSO.

Pensions Commission (2005) *A New Pension Settlement for the Twenty-First Century. Second Report of the Pensions Commission*. London: HMSO.