ISSN 1745-8587

Birkbeck Working Papers in Economics & Finance

Department of Economics, Mathematics and Statistics

BWPEF 1901

"Sustainable and Affordable"? Actuarially Fair Contribution Rates for the USS Pension Scheme

Kenjiro Hori
Birkbeck, University of London

Stephen Wright
Birkbeck, University of London

January 2019

▪ Birkbeck, University of London ▪ Malet Street ▪ London ▪ WC1E 7HX ▪

“Sustainable and Affordable”? Actuarially Fair Contribution Rates for the USS Pension Scheme

Kenjiro Hori and Stephen Wright*

Birkbeck, University of London

January 3, 2019

Abstract

We compute actuarially fair contribution rates (aggregating both employers' and employees' contributions) for the USS pension scheme, using UK life tables and market yield curves. The fair rate is sensitive to life expectancy and the level of real yields, neither of which appears stationary. So any scheme predicated on a constant contribution rate is inherently unstable. We therefore argue that, to survive, defined benefit schemes such as USS must explicitly incorporate time variation in contribution rates, ideally along with some dependence on individual characteristics. Our formulae in principle provide an objective, verifiable and implementable methodology to calculate such fair contribution rates.

*All at Department of Economics, Mathematics and Statistics, Birkbeck, University of London, Malet Street, London, WC1E 7HX, UK. Email: k.hori@bbk.ac.uk and s.wright@bbk.ac.uk, respectively

JEL Classification: J32

Keywords: defined benefit, pension contribution rate

1 Introduction

In May 2018, UK universities and trade unions agreed to set up an expert panel to consider what, if any, changes were required to the existing defined benefit scheme, the Universities Superannuation Scheme (USS), expressing the hope that "... the expert panel's deliberations will lead to... a defined benefit scheme which is sustainable and affordable..." (USS, May 2018). In this paper we consider whether this objective is achievable.

The USS scheme, like any defined benefit scheme, makes a set of promises to its members. These can be summarised in terms of four components:[1]

1. For a scheme member who has R years until retirement, and stays in the scheme, USS will progressively accumulate a liability to pay a pension equal to $\frac{1}{75}$ of their salary for each year they[2] contribute to the scheme,[3] for the remainder of their lifetime. This pension is (more or less) indexed to the cost of living.[4] So someone who contributes for example $R = 20$ years will get a pension equal to $\frac{20}{75} = 26.7\%$ of their average salary over the 20 years.

2. USS also pays a lump sum at retirement, equal to 3 additional years' worth of pension.

3. After the member's death the scheme will continue to pay 50% of the pension to

[1] Note that USS also has a defined contribution element for some scheme members; this paper however focuses solely on the defined benefit components.

[2] Throughout the paper we use the term "they" in the singular, gender-neutral sense.

[3] This was previously $\frac{1}{80}$, until 2016.

[4] There are exclusion clauses which introduce partial indexation if inflation exceeds some critical level. We assume that the impact of this partial indexation is minimal, and therefore ignore it in our analysis.

any surviving spouse or civil partner.[5]

4. Finally, the scheme offers a death-in-service grant of 3 times annual salary at the time of death.

In order for the scheme to be just viable, the typical scheme member must pay just enough into the scheme, as a deduction from salary, to fund the payouts the scheme will expect to make. These contributions are made both by the employer and the employee, but since both are taken out of the gross salary cost of employing the scheme member, we simply consider the total amount, rather than who pays it. The "actuarially fair" contribution rate, which we refer to as c^*, is then the rate of total contributions, as a percentage of salary, that ensures this is the case.

Technically speaking, this requires that the present value of contributions into the scheme must match the present value of the scheme's liabilities. In order to calculate these present values we need to take account of the nature of the risk (or lack of it) associated with both inflows and outflows to the scheme. However, this is *all* we need to know: specifically we do *not* need to know the investment policies of the scheme, since these have no impact on these present value calculations.

The exact calculation of the fair contribution rate is quite complicated. In the main body of the paper we carry out this task in a range of environments of varying degrees of complexity. However, we can get quite a long way by first doing the calculation for what might appear to be a highly restricted special case, but which turns out to have a surprising degree of empirical relevance. This special case has the advantage that it is easy to present in non-technical terms, so we start by presenting this as an example, before summarising the key features of our results. We then proceed to the technical details.

[5] There is also some limited provision for pensions to be paid to dependent children, but these are time-limited and discretionary so we exclude these from our analysis.

1.1 The suitcase calculation

Imagine a world in which both inflation and growth of salaries in nominal terms were precisely zero, as was the nominal rate of interest. By implication growth of salaries in real terms would also be zero, as would be the real interest rate. Not only does this simplify the calculation considerably, but - despite its highly restrictive assumptions, it turns out to provide a very relevant benchmark case.

In this highly simplified hypothetical case, we can do some very simple and intuitive illustrative calculations. Consider the example of a 27 year old who joins the scheme and continues to contribute until retirement at the age of 67. At this point, the scheme is committed to pay the first three benefits listed above.[6] On current historic UK data for life expectancy, at this point the typical scheme member (averaging across genders) can expect to live until the age of 85, so a further 18 years. So, for the average scheme member the scheme will expect to pay $\frac{40}{75} = 53.3\%$ of their career average salary for 18 years. (Of course not all scheme members will stay in the scheme for so long, but we show below that in this simple case this does not affect the fair contribution rate).

On top of this the scheme will need to pay the lump sum equal to three years of pension, which (in this simplified case) has an identical effect to the scheme member living for 3 extra years. Finally, the promise to continue to pay half the pension to a surviving spouse also has an impact that is equivalent to an increase in the scheme member's life expectancy. Assuming that there is roughly a 50% chance that the spouse will outlive the scheme member,[7] and that the spouse will live for a further six years (again, roughly in line with the data), this is equivalent to an increase in the scheme member's life expectancy of $\frac{1}{2} \times \frac{1}{2} \times 6 = 1.5$ years. So we can simplify by ignoring items

[6] In this example we are thus neglecting the contributions required to fund the death-in-service grant. We show below that when we calculate the contribution of this element to the fair contribution rate it plays a very small role. This is unsurprising since if the death-in-service grant were itself actuarially fair - ie, if it was simply equal to the present value of benefits accrued up to that date - then it would have no impact at all on the contribution rate. In practice it increases the fair contribution rate somewhat, implying that the amount paid out is more than fair value for the typical member.

[7] The probability is in fact somewhat less than 50%, since we are conditioning on the scheme member surviving to 67 (since the spouse may not survive to that point) but we ignore this for expositional simplicity. Using the ONS figures for $2015 - 17$, the actual probability is 45.8%.

2 and 3, and simply considering the cost to the scheme of paying the pension to someone with an “effective life expectancy” of $18 + 3 + 1.5 = 22.5$ years.

A calculation of the fair contribution rate in this simple case reduces to the requirement that the cumulative total payments made by the scheme member must be enough to pay for the benefits the scheme expects to pay. Under the restrictive assumptions given above, this is equivalent to the scheme member simply putting some fixed share, c, of their salary into a suitcase until retirement, and then handing the suitcase over to the scheme to pay their pension. The cash in the suitcase may be too much or too little for any individual pensioner - the scheme pools the risk that some members will die early or late - but an actuarially fair contribution rate ensures that on average there will be just enough money in the suitcase to pay the benefits the scheme promises to the average scheme member.

We shall refer to the fair contribution rate for the suitcase calculation as c_S^*. In the case of our example of someone who stays in the scheme for 40 years, with an “effective life expectancy” at retirement (as defined above) of 22.5 years we must have

$$\begin{aligned} \text{total payments into the suitcase} &= \text{total expected payments out of the suitcase} \\ \overset{\text{fair contribution rate}}{c_S^*} \times \text{ salary} \times \overset{\text{time in scheme}}{40} &= \overset{\text{pension}}{\frac{40}{75}} \times \text{salary} \times \overset{\text{effective life expectancy at retirement}}{22.5} \end{aligned}$$

Here the salary refers to career average salary. Simplifying this expression gives,

$$c_S^* = \frac{22.5}{75} = 30\%$$

i.e., in this special case we can think of the fair contribution rate just in terms of a ratio of years: the top of the ratio is the effective number of years the pension will on average need to be paid, and 75 is the notional number of years a scheme member would need to contribute to receive a pension exactly equal to their average salary. Furthermore, note that because both salary and time to retirement, R, cancel out, it means that in this special case they have no impact on the fair contribution rate: c_S^* is the same at all salaries and for any length of time in the scheme.

While the suitcase calculation requires some drastic simplifying assumptions, we shall show that it provides a very helpful benchmark under much more general conditions.

1.2 The suitcase becomes a bank account

Indeed our first adjustment towards more realistic assumptions turns out to make absolutely no difference at all to the answer. Consider a somewhat more realistic case where inflation is typically positive, so that cash left in a suitcase would progressively lose value in real terms. But money regularly paid into a deposit account (which we need to assume is completely risk-free, hence has a government guarantee) accrues interest. If we assumed that the nominal interest rate on such an account was precisely equal to the inflation rate, then the *real* interest rate would be precisely zero.

In this special case any contribution paid into the bank account would maintain its value in real terms, so total payments into such a deposit account would simply equal the contribution rate times total salary over the period that the scheme member pays into the account, expressed in real terms. Such payments may increase over the working life, if they get promoted, or at least move up their salary scale, but in this special case, the pattern of salary over time is irrelevant: all that matters is the total amount paid in, in real terms. For a scheme member contributing for 40 years, the pension payments out of the scheme are calculated as $\frac{40}{75}$ of average salary, and are also guaranteed in real terms. But this pension obligation is in turn simply equal to $\frac{1}{75}$ $\times$ total real earnings while they are in the scheme. We can then repeat the calculation of the fair contribution rate for the bank account case, c_B^* as

$$\begin{aligned} \text{total real payments in} \quad &= \quad \text{total expected real payments out} \\ \overset{\text{fair contribution rate}}{c_B^*} \times \text{ total earnings} \quad &= \quad \overset{\text{pension}}{\frac{1}{75}} \times \text{total earnings} \times \overset{\text{effective life expectancy at retirement}}{22.5} \end{aligned}$$

which gives exactly the same answer, i.e.: $c_B^* = c_S^* = \frac{22.5}{75} = 30\%$. In this case it is transparent that, with a zero real interest rate, the amount of time in the scheme, and any real growth of earnings, are both irrelevant to this calculation.[8] The suitcase calculation above still applies exactly, and we can again think of the fair contribution rate solely in terms of a ratio of years. Furthermore the amended version of the formula makes it clear that we do not need to assume that the scheme member stays in the scheme throughout their working life, since a shorter period of contributions affects both sides of the formula symmetrically.[9]

The extreme simplicity of the suitcase calculation immediately brings out a key feature of the fair contribution rate: for a given accrual rate for the pension (i.e., holding the notional number of years constant at 75), and for a given retirement date, the fair contribution rate must increase with effective life expectancy. This feature is always present even when we consider much more complicated cases.[10]

Figure 1 shows that, roughly from the start of the twentieth century onwards, there have been steady rises in UK life expectancy at age 65, with nontrivial differences by gender.

[8]In the more general case we show below that c^* is barely affected by real growth of wages during the lifetime. In contrast, calculating the pension on the basis of the final salary (as was the case for pre-2016 version of the scheme) raises the fair contribution rate nontrivially (see discussion in Section 3.3 below).

[9]This feature also holds, to a fairly good approximation, even in more complex cases: the member can be assumed to quit the scheme randomly, with minimal, if any, impact on the contribution rate. See discussion in Section 3.4 below.

[10]Only in the resrictive case of the suitcase calculation is it precisely *proportional* to life expectancy at retirement..

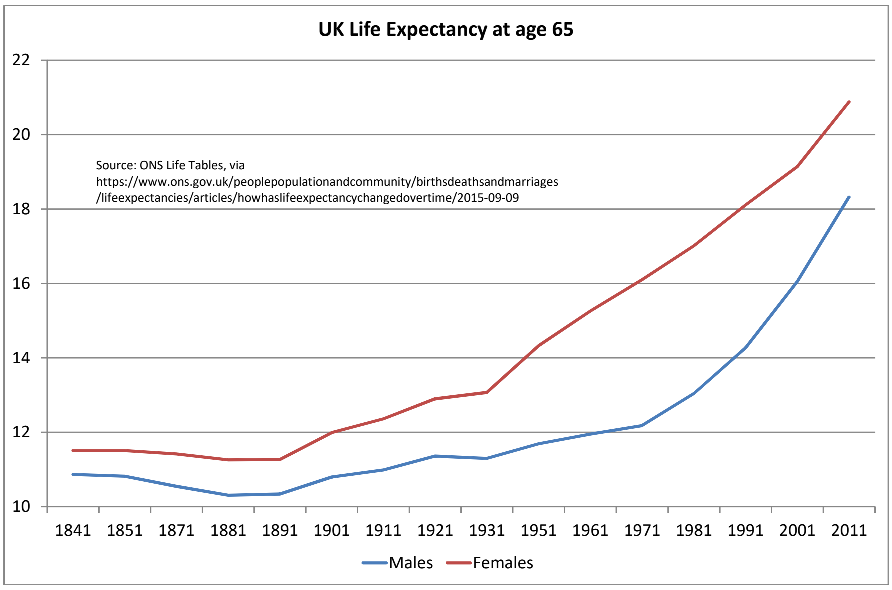

Figure 1: UK Life Expectancy at Age 65

The only way that such rises would *not* impact on the fair contribution rate would be if the age of retirement within the scheme adjusted automatically by precisely enough to offset the impact of rising life expectancy. While retirement ages have risen they have not risen enough to do so. Thus the clear implication is that on this basis alone, fair contribution rates should be on a secular upward trend.

Despite its simplicity, and the apparently restrictive assumptions made in deriving the suitcase calculation, it turns out to provide a very useful, and empirically relevant benchmark case, when we turn to consider the impact of real interest rates *not* being equal to zero.

1.3 The fair contribution rate when real interest rates are not zero

1.3.1 A general definition of the fair contribution rate

The calculation of the fair contribution rate is in general more complicated than in the suitcase calculation because we need to compare amounts in present value terms, i.e. for the general case we calculate, for any given scheme member the fair contribution rate, c^*, that equates the present value of payments into the scheme to the present value of scheme liabilities to the scheme member, i.e. that satisfies

$$c^* \times (\text{present value of salary payments}) = \text{present value of liabilities}$$

which in turn requires

$$c^* = \frac{\text{present value of liabilities}}{\text{present value of salary payments}}$$

1.3.2 Which discount rate?

As is the case whenever we calculate a present value, we need to pick an appropriate rate at which to discount both future payments into the scheme and the scheme's liabilities. We assume that the appropriate interest rate to use is a real interest rate on bonds of an appropriate maturity that are free of both default risk and inflation risk: thus in our calculations we use yields on index-linked UK government bonds.[11]

We argue that if anything this may result in an under- rather than over-estimate of the fair contribution rate, since the use of a risk-free discount rate is more obviously applicable to scheme liabilities than it is to contributions.

There are two distinct rationales for the use of such risk-free rates in calculating the present value of liabilities.

The first rationale arises from a standard assumption in the analysis of insurance markets. Liabilities to any individual scheme member are risky, from the scheme's point of view, since they depend on the member's date of death, which is unknown. But a standard assumption in insurance is that the law of large numbers applies, such that for payments to a large number of scheme members with the same mortality risk, individual risks cancel each other out, so that aggregate payments are close to risk-free.

[11] We discuss below one adjustment that is required to these yields for the fact that scheme liabilities are indexed to CPI rather than RPI inflation.

An alternative rationale, which is arguably more widely applicable, focuses on the *nature* of mortality risk at the individual level. In standard finance theory, it is not risk, *per se* that matters; what does matter is whether the risk associated with any payment or receipt is correlated with *market* risk factors that have associated risk premia. We are not aware of any evidence of such a correlation. In the absence of any correlation, the present value of any payment with purely idiosyncratic risk is calculated in terms of the *expected* payment, discounted using a risk-free rate. The advantage of this rationale in terms of USS is that it applies at the level of any individual scheme member, rather than simply at the aggregate scheme level.

In contrast, the use of risk-free rates to calculate the present value of salary payments is arguably less defensible on the basis of either rationale. While liabilities to any individual scheme member are of fixed amounts, conditional upon mortality status, salary payments are not fixed, either at the individual or aggregate level, even conditional upon the scheme member being alive. To the extent that salaries are correlated with standard measures of market risk that have associated risk premia (most obviously, stock market risk) salary payments may have "positive beta", which would imply a higher discount rate than we use, and hence a lower present value than in our calculations (which simply assume that salary payments follow a fixed, known, path until retirement). Thus in the estimates of c^* we provide below, this would point to these being under-estimates.

A crucial point to note is that the choice of discount rate in calculating the fair contribution rate is entirely unaffected by the nature of the scheme's investments. What the scheme does with the contributions made by the scheme member will have a significant impact on risks to the scheme's sponsors, but it does not affect the present values that feed into the contribution rate, so it does not affect the contribution rate.[12] It *would* make a difference if the scheme's commitments to members were contingent on the outcome of investments, but they are not - indeed this is the key distinction between

[12] If this is not obvious, consider the answer to the question: what is the present value today of £100 invested in safe assets for ten years, compared to £100 invested in equities? The answer is the same in both cases: £100.

defined benefit schemes like USS and defined contribution schemes.

1.3.3 The impact of different real interest rates

For a given choice of interest rates to calculate present values, we now turn to the impact on the fair contribution rate.

In the conventional textbook case, when real interest rates are positive, the further a payment is into the future, the less it is worth today, in real present value terms. Since payments into the scheme come first, and the promised benefits stretch out into the distant future, with positive real interest rates, for any given contribution rate, this raises the present value of contributions relative to benefits. As a result, we show below that a positive real interest rate means that the actual fair contribution rate c^* can be significantly lower than in the suitcase calculation.

Figure 2 shows that a commonly used measure of real interest rates, the yield on an RPI-linked UK government bond with maturity of 10 years, was indeed typically positive, albeit on a steadily declining trend, from 1985 until 2011, but has been consistently negative since then.

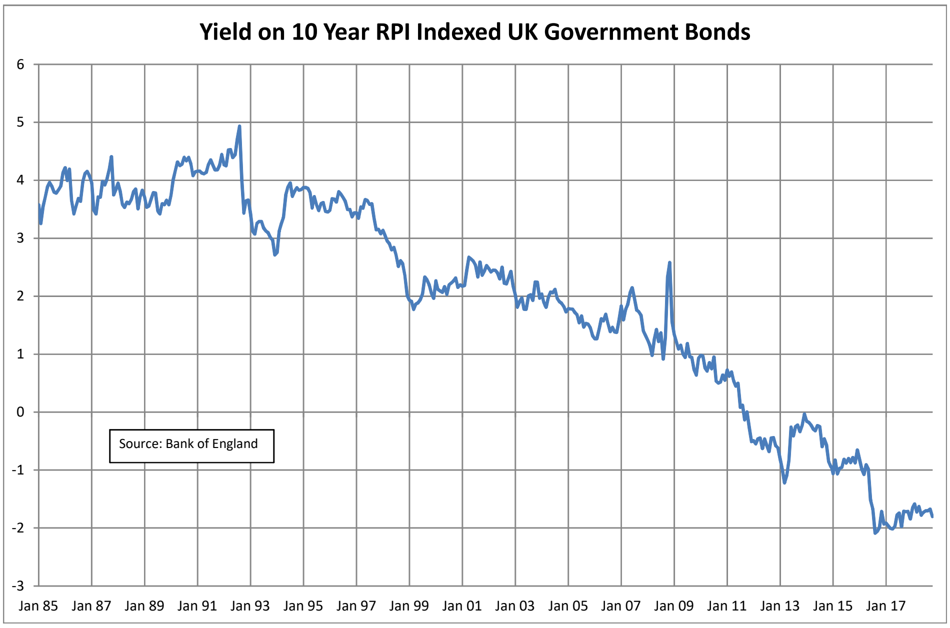

Figure 2: Yield on 10 Year RPI Indexed UK Government Bonds

Since the obligations of the USS are now indexed to CPI inflation, in our calculations below we adjust RPI-indexed yields for the gap between CPI and RPI inflation. This makes the appropriate measure of real rates in recent years somewhat less than as shown in the chart, but it is still currently markedly negative at all maturities.

When real rates are negative, as at present, the further a payment is into the future, the *more* it is worth in present value terms, so c^* needs to be *higher* than in the suitcase calculation.

The impact on c^* of quite small differences in the average level of interest rates at different maturities can be very significant. We show below that to a fairly good approximation, when real interest rates are close to zero a one percentage point fall in real rates raises the fair contribution rate c^* by around 6 percentage points for someone who contributes to the scheme for 20 years, and by 9 percentage points for someone who contributes for 40 years.[13] (This greater sensitivity for younger members reflects the longer maturities of the scheme's liabilities).

Thus even when we allow for the impact of real interest rates being different from zero, our suitcase calculation provides a crucial benchmark. For as long as real interest rates remain negative, the fair contribution rate in the suitcase calculation, $c_S^* = 30\%$ represents a lower bound for the true contribution rate, c^*. But the *extent* to which the true contribution rate differs from c_S^* is determined by the length of the period to retirement. Thus the fair contribution rate is higher for younger scheme members.

A second look at Figure 2 also conveys a crucial message. The history of real interest rates since 1985 gives no reassurance that real interest rates have any tendency to revert to a stable average level.[14] We have already seen, in Figure 1, that the other key determinant of the fair rate, life expectancy at retirement, has also clearly drifted over

[13] We show below (see Section 3) that to a fairly good approximation the sensitivity to the real rate is equal to c_S^* times one half the *current* effective life expectancy of the worker. We also show that this can be related straightforwardly to the difference between the "durations" of the scheme's liabilities vs assets. Thus in our numerical example of someone who pays into the scheme for 40 years this coefficient is given by $-0.30 \times \left(\frac{40+22}{2}\right) = -9.3$ for a unit change in yield, or -9.3 percentage points for a 1% change in yield.

[14] This feature is borne out by long-run international data.

time. Thus the fair contribution rate must clearly also vary over time; and in the recent past at least, this drifting has clearly been in an upward direction.

1.4 The embedded life insurance policy

Our discussion thus far has focussed only on benefits paid after retirement. This ignores the impact on the fair contribution rate of two elements of the scheme. First, as noted under item 3 of the list of benefits at the start of the paper, *whenever* the scheme member dies, the USS will pay to the surviving spouse, for the remainder of their lifetime, half of the pension that would have been paid to the member, had they survived to retirement. For those scheme members who die before retirement, the expected duration of the spouse's pension is clearly considerably higher, and hence the implied scheme liability is distinctly higher. Second, as noted under item 4, the USS pays a death-in-service grant of 3 times salary at the time of death. The total cash value of both of these benefits can be viewed as is, in effect, an embedded life insurance policy.[15]

In principle, it should be noted that some element of life insurance need not necessarily affect fair contribution rates. This can be illustrated by considering two specific cases.

First consider a case in which the three post-retirement benefits we illustrated with our suitcase calculation were the only benefits offered by the scheme.[16] In this case, to the extent that some scheme members died before retirement, this would reduce the scheme's expected payouts, and we would therefore over-estimate the fair contribution rate.

Now consider a second case in which, if the scheme member dies before retirement, the USS pays the spouse's pension from that date onwards, together with some additional amount as a death-in-service grant. Both will push the actuarially fair contribution

[15] The cash value of the total payout would be equal to 3 years' salary plus the cost of buying an annuity equal to the spouse's pension: the latter is likely to be the larger component.

[16] Specifically, if we assume that there was no death-in-service grant, and that the spouse's pension was only paid for as long as would be expected for a scheme member whose life expectancy after retirement was as assumed above.

rate back up again; and there will be some critical value of the death-in-service grant at which the total costs of the insurance policy exactly counterbalance the savings on post-retirement benefits, leaving the fair contribution rate unaffected.

Any value of the death-in-service grant larger than this critical value will of necessity imply a higher fair contribution rate. We calculate below[17] that, for virtually all scheme members, the value of the death-in-service grant is higher than this critical value. For a member with 20 years to retirement its net effect is to boost the fair contribution rate by around 1.6 percentage points.

1.5 A Summary of the Key Results of the Paper

In the remainder of the paper we set out a formal mathematical framework that allows us to calculate estimates of the fair contribution rate, c^*, at prevailing market interest rates, and given the typical pattern of life expectancy over life, as given by UK life tables. While these calculations are complicated, the results are actually quite easy to interpret in light of the arguments set out above, so we summarise them here.

Figure 3 shows the results of these calculations for scheme members joining the scheme at the ages of 27 and 47, respectively, and hence having time to retirement of 40 and 20 years respectively. The calculation shows what would have been the hypothetical required fair contribution rate over the period since 1985, for each of these individuals at then-prevailing market interest rates and life expectancies, given the terms of the USS scheme at the time. Thus the calculations also reflect the change in the scheme from providing a pension which until 2011 was defined in relation to final salary to one based on career average revalued earnings (CARE), as described above. A final salary scheme implies that growth rate of earnings through the working life increases fair contribution rates nontrivially, whereas in the current scheme the impact is minimal: thus a fair contribution rate based on current scheme benefits would have been distinctly lower before 2011.[18] The estimates also take into account changes in UK-wide growth rates of

[17] See Table 7.
[18] See Section 3.3 below.

real earnings,[19] as well as the 2016 change in the annual accumulation rate from $\frac{1}{80}$ to $\frac{1}{75}$. The table of results are given in Appendix A.

We should stress that, while the calculations involved are complicated, they are conceptually fairly straightforward, require minimal assumptions and exploit readily available data that would, to a very close approximation have been available in real time.[20] In particular the calculations do not require any assumptions about the future path of interest rates, so there is no "look-ahead" bias. We do, of course, exploit the market forecasts that are embedded in yields, and hence prices, of government bonds; but when these prices change, our calculated fair contribution rate changes.

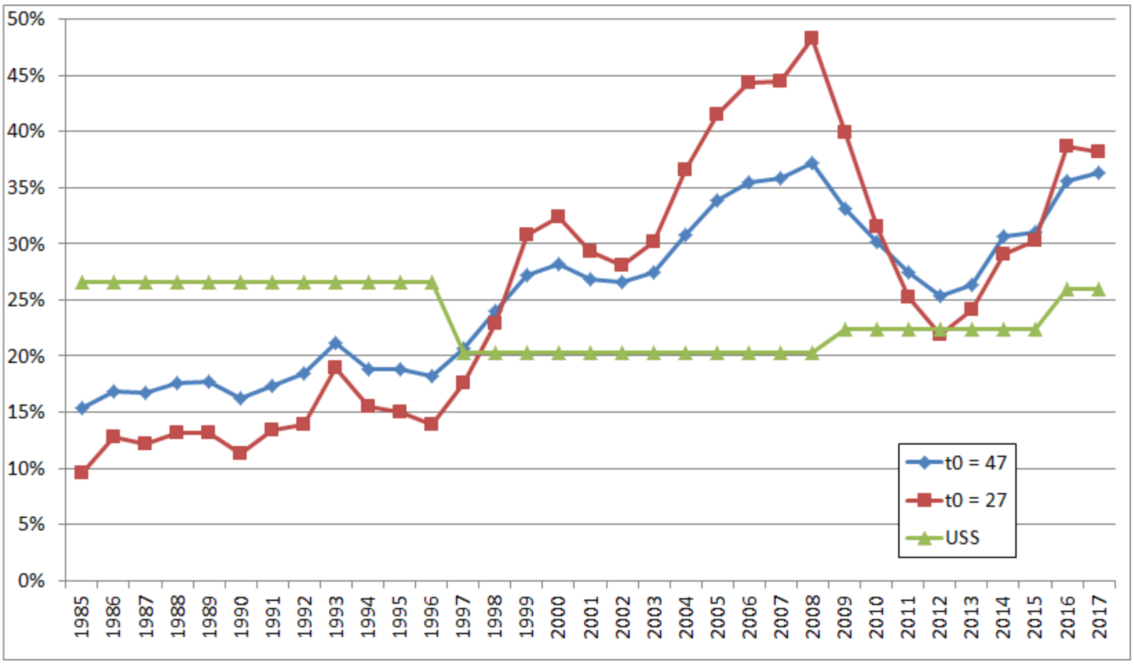

Figure 3: Actuarially Fair Contribution Rates (1985-2010)

Figure 3 shows that in the most recent data (as of August 2018) our estimated fair contribution rates were just under 37% for a scheme member with 20 years to retirement, and around 39% for a member with 40 years to retirement.

[19] We set the growth rate of earnings equal to 10 year rolling average annual growth rates (source: ONS)

[20] The market yield curve data which largely drive short term movements are available on a daily basis for the previous day; mortality statistics are available only with a lag of a year or two but evolve only slowly so the calculations would be only minimally altered if we used data that would have been available at the time.

Despite the complexity of the underlying calculations, these figures are relatively easy to interpret in terms of our analysis above. We showed above that if real interest rates were zero at all maturities, the "suitcase" calculation would point to a fair contribution rate of around 30%. But real rates across relevant maturities are currently all negative, which implies that the suitcase calculation must understate the current fair contribution rate. We also noted that if real interest rates across all maturities shift away from zero by one percentage point this must, to a reasonable approximation, raise the fair contribution rate by around 6 percentage points for a member with 20 years to retirement, and by around 9 percentage points for a member with 40 years to retirement. We show below that in the most recent data real rates at the (more relevant) longer maturities (and adjusted for the expected gap between CPI and RPI inflation) are around -0.9%. If, finally, we allow for the additional costs of the embedded insurance policy (the death-in-service grant), which we estimate adds around $1\frac{1}{2}$ percentage points to the fair contribution rate, we get very similar answers to those shown in the chart.

For example for a scheme member with 20 years to retirement:

$$\begin{aligned}
\overset{\text{Fair contribution rate}}{36.6\%} \approx \; & \overset{\text{"suitcase" calculation}}{30\%} \\
& \overset{\text{sensitivity to real interest rate}}{-6} \times \overset{\text{real interest rate}}{-0.9\%} \\
& \overset{\text{impact of death-in-service grant}}{+1.6\%}
\end{aligned}$$

For comparative purposes the chart also shows the actual historical contribution rate for USS members. This shows that in the most recent data, there is a very substantial gap between our estimated fair contribution rates and the current contribution rate.

What do we learn from these figures?

Probably the most immediately striking feature of our calculations is the very substantial gap between our estimated fair contribution rates and those in the current scheme. It is not surprising, therefore, that the scheme's sponsors have recently proposed substantial increases in contribution rates.

However, we would argue that this is *not* the most crucial feature of our results. We would lay more stress on the more general implications of our analysis. Most crucially, we have shown that the fair contribution rate is, to a very good approximation, driven by just two factors: life expectancy at retirement and the real interest rate. Both of these factors share two key features.

First, as we have noted above, they have both tended to drift over time: in statistical terms they are not mean-reverting. This in turn implies that the fair contribution rate must itself drift over time;[21] and Figure 3 showed that the implied shifts have been very large. We show below that the dominant element in these shifts has been the path of real interest rates: the secular downward trend in real rates shown in Figure 2 is largely reflected in a secular upward trend in fair contribution rates.[22]

Second, and even more crucially, they are both exogenous processes, beyond the control of anyone involved in the scheme. By implication, the same must apply to the fair contribution rate itself.

Thus it is clear that any scheme based on a constant contribution rate must be inherently unstable. This property is implicitly acknowledged in the periodic adjustments of contribution rates seen historically, but such adjustments usually involve both painful processes of negotiation and are also typically delayed responses to changes in the underlying determinants.

At the time of writing (December 2018) negotiations are continuing between the scheme's sponsors and the trade unions who represent a substantial proportion of the scheme members as to whether, and to what extent, the scheme should change. But the logic of our analysis is that, if a pension scheme is to be viable, the associated contribution rates that ensure this viability, for a given set of scheme benefits, are not a fit subject for negotiation - any more than we expect to negotiate over the impact of, say, the price of oil, or the impact of rising longevity on the price of life insurance. We cannot negotiate away these features; all we can choose to negotiate over is the necessary

[21] Strictly speaking, in the absence of some fortuitous offsetting movements in the two factors.

[22] We show below that the pattern of fair rates is also affected nontrivially by shifts in the slope of the yield curve.

adjustments to the terms of agreements that are affected by them. Thus we could in principle neutralise, or at least offset, the impact of further increases in longevity or falls in market yields by some combination of higher contribution rates, reduced scheme benefits and/or adjustments to the scheme's retirement age. The precise combination of adjustments *could* in principle be a subject of negotiation, but the requirement for an adjustment is not.

While the primary focus of our analysis is on the determinants of fair contribution rates for the "typical" scheme member, a further implication of our analysis is that no such individual exists. While some individual features (most notably the initial salary of the scheme member) have absolutely no impact on fair contribution rates, others, most notably age (and hence time to retirement) only drop out of our formulae in restrictive cases (when real interest rates are zero at all maturities at the time of joining the scheme); while yet others (marital status, gender, number of dependent children) must always impact on fair contribution rates. As such, any scheme that sets a common contribution rate for all scheme members implies some element of cross-subsidy across scheme members. Since the attributes that affect fair contribution rates are typically eminently measurable, it is not immediately obvious why a common contribution rate for scheme members with different individual characteristics should necessarily be desirable.[23]

Our analysis suggests that scheme contribution rates should vary both over time (as exogenous determinants like interest rates and life expectancy vary) *and* across individuals. While these conclusions may appear unpalatable, it is worth noting that, in many other contexts, such variations are taken for granted. Arguably the most relevant analogy is the market for annuities. Anyone purchasing an annuity will face a price that varies from day to day, as interest rates vary, and also in secular terms as insurance companies assess evidence on life expectancy. The price paid will also depend on a range of personal characteristics - most notably, of course, age.

[23] The usual economic arguments for pooling apply only to unobservable characteristics such as health and longevity risk.

The clear difference between purchasing an annuity and being a member of a defined benefit scheme is that the former is delivered by a reasonably transparent market, whereas the latter is run by a set of administrative rules and procedures, changes in which have in the past typically been both slow, and painful. But a further implication of our analysis is that regular adjustments of contribution rates to (best estimates of) actuarially fair rates could in principle be entirely formulaic, with formulae based on clear, implementable, and defensible principles. We would argue that our analysis represents a useful first step in that direction.

1.6 A summary of the remainder of the paper

The rest of the paper is organised as follows. In Section 2 we derive an expression for the fair contribution rate based on a general market yield curve, an assumed deterministic path for wages, and a hazard rate function that determines life expectancy at any given age. While the yield curve and hazard rate functions are known at any point in time, they will change over time, hence so must the fair contribution rate. In Section 3 we consider a simplified example, with constant real yields and a known date of death, and use this to derive a first-order Taylor Series approximation for the fair contribution rate (which we exploit in the discussion above). In Section 4, we specify a particular functional form for the market yield curve, consistent with the Vasicek model, which allows us to consider subtler variants on the current market yield curve. In Section 5 we use UK data to estimate fair contribution rates for a range of values of key determinants. Section 6 concludes the paper.

2 The Fair Contribution Rate: a formal derivation

A scheme member joins a defined benefit pension scheme at time 0. If they survive, they will retire at $R > 0$.

2.1 The market environment

The scheme member's real wage income w, from which contributions to the scheme are drawn, is assumed to grow deterministically at a constant continuously compounding annual rate $g \geq 0$. The wage w_0 is normalised to 1 at $t = 0$, and therefore the wage at time t, $w_t = e^{gt}$.

At time t there is a real market yield curve described by the function $y(t, m)$ the yield (continuously compounded risk-free real return to maturity) on a zero-coupon bond issued at time t, and maturing with a risk-free real payoff at time m. This function evolves stochastically through time, but at any point t is a known deterministic function in m.

2.2 Mortality Risk

There is a single source of uncertainty, which is the time of death of the scheme member and their spouse. The probability of death in a small interval $(t, t + \Delta t)$ is given by $\lambda(t)\Delta t$, where $\lambda(t)$ is the *hazard rate* (or the *force of mortality*). Denote by $\overline{\lambda}(t_1, t_2)$ the average hazard rate for period $[t_1, t_2]$. The probability of surviving from $t = 0$ until $t = T$ is then $e^{-\int_0^T \lambda(t)dt} = e^{-\overline{\lambda}(0,T)T}$.[24] We assume identical but independent probabilities of death for the scheme member and spouse.[25]

2.3 The Scheme

The following assumption simplifies our analysis:

[24] Strictly speaking the hazard rate function is itself a function of time, ie $\lambda = \lambda_t(t)$. Hence, analogously to the market yield curve, at any point in time t we assume there is a known probability $\exp\left(-\overline{\lambda}_t(t, T)\right)$ of surviving until time $T > t$, but new information on mortality will cause this probability to adjust over time. We exploit such shifts in the hazard rate function in the numerical calculations in Figure 3, but suppress the dependence on t in our notation for expositional purposes.

[25] In the data, as Figure 1 showed, they are clearly not identical, so the calculations below can also be viewed as being calculated under a veil of ignorance on the gender of the scheme member. If we calculated gender-specific fair contribution rates, these would clearly be higher for women than for men. Anti-discrimination law rule this out in practice, but as a result a gender-neutral fair contribution rate, as calculated below, clearly implies a cross-subsidy within the scheme from men to women.

Assumption A1: *The Law of Large numbers applies to all assets and liabilities of the scheme.*

Assumption A1 implies that idiosyncratic probabilities of death at any time t during the life of the scheme translate to exact proportions of scheme participants, of any given age, who die at time t. While this is a standard assumption in calculating actuarially fair magnitudes, we show below that our calculated fair contribution rate can also be calculated under the alternative, and arguably less restrictive assumption that both assets and contingent liabilities of the scheme have zero risk prices.

The scheme member's instantaneous contribution to the pension scheme at time $t \in [0, R)$ is assumed to be a fixed time-invariant proportion c of their wage w at time t. The present value at $t = 0$ of the scheme member's total payments into the scheme A_0 (the assets of the fund) is simply proportional to the present value of the scheme member's lifetime earnings within the scheme, W_0

$$A_0 = c.W_0 \tag{1}$$

where

$$W_0 = \int_0^R e^{\left(g-y(0,t)-\overline{\lambda}(0,t)\right)t} dt. \tag{2}$$

Note that here, as in the remainder of the analysis we apply Assumption A1, such that expected earnings of any individual scheme member of a given age, at a given time, translate to the actual earnings of scheme members in aggregate of the same age and time.

Note also that, because the wage is assumed to grow deterministically, this is the only correction for risk required in the present value calculation. Thus, to the extent that actual observed wage growth may covary with market risk factors, we ignore this.[26]

[26]To the extent that there is covariation between wage growth and market risk factors, we would normally expect this to be positive. If so the value of A_0 given here represents an upper bound to the present value of the scheme member's contributions. We also ignore the (empirically nontrivial) probability that the scheme member may quit the scheme before retirement; however since this has almost precisely offsetting impacts on scheme assets and liabilities, allowing for this probability would have a negligible impact on our results.

In return for these contributions, the scheme member receives two kinds of payouts if they survive beyond R:

- A constant pension p, payable from R until the scheme member's death, given by

$$p = \alpha \int_0^R e^{gt} dt = \alpha \left(\frac{e^{gR} - 1}{g} \right) \tag{3}$$

  where α is the annual accrual rate to pension for the years that they contribute to the pension scheme (e.g. in the USS scheme $\alpha = \frac{1}{75}$).

- A lump-sum payout at $t = R$ of Z years' worth of the pension, i.e. of $Z.p$.

In addition the scheme also makes payments to spouses or other beneficiaries. There are two kinds of payouts:

- If the scheme member dies before reaching retirement, the scheme makes a lump-sum death-in-service payment to a named beneficiary of B times the annual wage at the time of death, $\tau \in [0, R)$, i.e. $Be^{g\tau}$. We assume that the scheme member always names such a beneficiary.[27]

- Additionally, in the case that the scheme member dies before their spouse, the spouse receives a proportion s of what the pension would have been, had the scheme member survived until R (irrespective of whether the scheme member does in fact die before R) i.e. the spouse receives $s.p = s\alpha \left(\frac{e^{gR}-1}{g} \right)$. We assume that (in contrast with the death-in-service grant) any such payments can only be made to the spouse, and thus any such payment is contingent on the spouse surviving the scheme member.[28].

[27] Note that in the current USS scheme $Z = B = 3$.

[28] In USS there are in practice some additional pension payments to dependent children but these are of limited duration, and are discretionary, so we ignore them in our calculations.

2.4 The Fair Contribution Rate

Let L_0 denote the present value of all the scheme's contingent liabilities to the scheme member. Define the *fair contribution rate* c^* as the contribution rate that equates the present values of the scheme's assets and liabilities. Thus, given (1) and (2), we have

$$c^* = \frac{L_0}{W_0} \tag{4}$$

The fair contribution rate is thus simply equal to the present value of the scheme's liabilities as a proportion of the present value of the scheme member's lifetime earnings. Note that we follow standard academic practice in ignoring the administrative costs of the scheme in calculating this fair rate.

We now find the present values of each of the components of the scheme's contingent liabilities. Under Assumption A1, all such contingent liabilities, as outlined above, are risk-free.

As stated, the scheme member receives an instantaneous pension $p = \alpha\left(\frac{e^{gR}-1}{g}\right)$ with probability $e^{-\overline{\lambda}(0,\tau)\tau}$ at time $\tau \geq R$. The present value of this pension payout stream is then,

$$L_0^p = \alpha\left(\frac{e^{gR}-1}{g}\right)\int_R^\infty e^{-\left(y(0,\tau)+\overline{\lambda}(0,\tau)\right)\tau}d\tau. \tag{5}$$

The scheme member also receives a lump-sum payout of $Z.p$ at $t = R$, which happens with probability $e^{-\overline{\lambda}(0,R)R}$. The present value of this at time 0 is

$$L_0^Z = Z\alpha\left(\frac{e^{gR}-1}{g}\right)e^{-\left(y(0,R)+\overline{\lambda}(0,R)\right)R}. \tag{6}$$

The scheme pays a death-in-service grant equal to $Bw_\tau = Be^{g\tau}$ if the scheme member dies before retirement at $\tau < R$. This happens with probability $e^{-\overline{\lambda}(0,\tau)\tau}\lambda(\tau)\Delta\tau$ in the period $(\tau, \tau+\Delta\tau)$. The present value of this lump-sum payout is

$$L_0^B = B\int_0^R \lambda(\tau)\, e^{\left(g-y(0,\tau)-\overline{\lambda}(0,\tau)\right)\tau}d\tau. \tag{7}$$

Finally, when the scheme member dies at τ, with probability $e^{-\overline{\lambda}(0,\tau)\tau}$ if the scheme member's spouse is still alive, they will receive an instantaneous pension payout of $s\alpha\left(\frac{e^{gR}-1}{g}\right)$ until their death. Under our assumption this is true irrespective of whether the scheme member dies before or after R. The present value of this payout is,

$$L_0^s = \int_0^{\infty} \lambda(\tau)\, e^{-\overline{\lambda}(0,\tau)\tau} \left\{ s\alpha \left(\frac{e^{gR}-1}{g} \right) \int_{\tau}^{\infty} e^{-\left(y(0,t)+\overline{\lambda}(0,t)\right)t} dt \right\} d\tau. \tag{8}$$

Here the integrals over τ and t capture the probabilities of the scheme member's and spouse's deaths, respectively at a given time horizon.

Thus we have

Proposition 1 *The fair contribution rate c^* is given by,*

$$c^* = \frac{1}{\underbrace{\int_0^R e^{\left(g-y(0,t)-\overline{\lambda}(0,t)\right)t} dt}_{\textit{present value of earnings, } W_0}} \left[\underbrace{B \int_0^R \lambda(\tau)\, e^{\left(g-y(0,\tau)-\overline{\lambda}(0,\tau)\right)\tau} d\tau}_{L_0^B \textit{ (death-in-service grant)}} + \underbrace{Zpe^{-\left(y(0,R)+\overline{\lambda}(0,R)\right)R}}_{L_0^Z \textit{ (lump-sum)}} \right. \tag{9}$$

$$\left. + \underbrace{p \int_R^{\infty} e^{-\left(y(0,\tau)+\overline{\lambda}(0,\tau)\right)\tau} d\tau}_{L_0^p \textit{ (pension to scheme member)}} + \underbrace{sp \int_0^{\infty} \lambda(\tau)\, e^{-\overline{\lambda}(0,\tau)\tau} \left\{ \int_{\tau}^{\infty} e^{-\left(y(0,t)+\overline{\lambda}(0,t)\right)t} dt \right\} d\tau}_{L_0^s \textit{ (pension to surviving spouse)}} \right].$$

$$\textit{where } p = \alpha \left(\frac{e^{gR}-1}{g} \right)$$

We immediately note that the expression for c^* in Proposition 1 can be generalised. Consider the following alternative to Assumption A1.

Assumption A2: *Assume that all contingent payoffs of the scheme have zero risk prices.*

We then have

Corollary 1 *Proposition 1 also holds, at the level of an individual scheme member, if Assumption A1 is replaced with Assumption A2.*

Corollary 1 follows from standard asset pricing theory. If we abandon Assumption 1 and consider the nature of the payoffs underlying both the scheme's assets, and liabilities relating to an individual scheme member, these are clearly stochastic, due to mortality risk. However if the payoffs are uncorrelated with market risk factors they will have zero risk prices so the no-arbitrage prices of both assets and liabilities for an individual scheme member can be calculated using expected payoffs, discounted by the risk-free rate, resulting in an identical formula for c^*.

2.5 Interest rate sensitivity of the fair contribution rate: a duration-based approach

In the introduction we gave an intuitive argument for the impact of changes in real interest rates on the fair contribution rate, c^*. The formula in Proposition 1 does not make this relationship transparent. However, analysis of the durations (somewhat loosely defined) of the scheme's assets and liabilities brings this out more clearly and precisely.

Both assets and liabilities of the schemes can be interpreted as portfolios of zero coupon bonds. Practitioners typically calculate the duration of a bond portfolio by analysing the sensitivity to a parallel shift in the yields of all the bonds in the portfolio. We take a similar approach here by considering the impact of a parallel shift in the zero coupon yield curve. We decompose the zero coupon yield at any maturity m as

$$y(t,m) = r(t) + \gamma(t,m) \tag{10}$$

where r is the instantaneous rate (defined such that $r(t) \equiv y(t,0)$), and $\gamma(t,m)$ is the slope of the yield curve. We then have:

Proposition 2 *Defining (spot rate) durations at $t = 0$ of the scheme's assets ($A_0 = c^* W_0$) and scheme liabilities, L_0 respectively by*

$$\delta_W \equiv -\frac{\partial \ln W_0}{\partial r_0};\ \delta_L \equiv -\frac{\partial \ln L_0}{\partial r_0},$$

the notional duration of c^* is given by

$$\delta_{c^*} = -\frac{\partial \ln c^*}{\partial r_0} = \delta_L - \delta_W. \tag{11}$$

Then for $\frac{d\lambda(t)}{dt} > 0$ for $t \in [0, R]$ and $g < g^- >> 0$,

$$\frac{\partial c^*}{\partial r_0} = -c^* \left(\delta_L - \delta_W\right) < 0.$$

Proof. See Appendix B. ■

The intuition for this result is fairly straightforward. Since c^* is simply the ratio of liabilities to assets, its notional duration (the semi-elasticity with respect to the interest rate) is simply equal to the difference between the two durations. Under minimal restrictions (which we show in Appendix B are satisfied for all plausible parameter values) this difference must be positive, hence an upward parallel shift in the yield curve must lower c^*.

3 A simplified example and an approximation

3.1 A simplified example

To understand eqn (9) consider a special case under the following assumptions:

1. Deterministic time of death D for the scheme member, assumed after retirement: $D > R$. The assumption means that we lose the term for the lump-sum payout if the scheme member dies before R.

2. The spouse outlives the scheme member with probability one half (given independence), and, conditional upon surviving the scheme member, the spouse's date of death is given by $D^S > D$, with certainty.

3. The yield curve is flat, so that $y(t, m) = r \; \forall m$

Under these assumptions we have

$$c^* = \frac{\alpha\left(\frac{e^{gR}-1}{g}\right)\left(\overbrace{\left\{\frac{e^{-rR}-e^{-rD}}{r}\right\}}^{\text{capitalisation factor for pension}} + \overbrace{\frac{s}{2}\times\left\{\frac{e^{-rD}-e^{-rD^S}}{r}\right\}}^{\text{capitalisation factor for spouse's pension}} + Z\times\overbrace{\left\{e^{-rR}\right\}}^{\text{discount factor for lump sum}}\right)}{\underbrace{\frac{e^{(g-r)R}-1}{g-r}}_{\text{Wealth}}}. \tag{12}$$

The denominator of this expression is simply equal to W_0, the present value of lifetime earnings. In the numerator the common term outside the brackets is simply the value of the pension, $p = \alpha\left(\frac{e^{gR}-1}{g}\right)$. The first two terms in curly brackets represent capitalisation factors for the member's and spouse's pensions; the third term is simply a discount factor for the lump sum payment.

Using this relatively simple expression we can obtain some simple and intuitive analytical results.

We first consider a useful special case, which, as discussed in the introduction is not too far from current data: namely we evaluate the limit of the fair contribution rate as r approaches zero (noting that c^* is undefined at this point). This gives, letting $\alpha = \frac{1}{Y}$,

$$\lim_{r\to 0} c^* = \frac{E}{Y}$$

where E is "effective life expectancy at retirement", defined by

$$E = D - R + \frac{s}{2}\left(D - D^S\right) + Z$$

since in this limiting case the capitalisation factor for the pension is simply equal to $D - R$ (the number of years the pension will be paid); the capitalisation factor for the spouse's pension is simply equal to $\frac{s}{2}\left(D - D^S\right)$ (the effective number of additional years the pension will be paid to the spouse, in expectation) and both enter symmetrically with Z, the number of years' extra pension provided as a lump sum. (This feature is visible by inspection of the exact formula above).

We thus find that this limiting case yields identical results to the "suitcase" calculation as set out in the introduction. Thus setting $Z = 3$, $s = \frac{1}{2}$ (which are the current parameters in the actual USS scheme) and assuming $D - R = 18$, $D - D^S = 6$, we then have effective life expectancy, $E = 18 + 3 + \frac{6}{4} = 22.5$ implying

$$\lim_{r \to 0} c^* = \frac{22.5}{75} = 30\%$$

as calculated above.

We also have

$$\begin{aligned} \lim_{r \to 0} \frac{\partial c^*}{\partial R} &= 0; \\ \lim_{r \to 0} \frac{\partial c^*}{\partial g} &= 0 \end{aligned}$$

The first result follows because in this limiting case a longer period of contributions to the fund increases both fund assets and fund liabilities proportionally. The second follows because in the absence of discounting the timing of salary payments, and hence the growth rate of salaries, is irrelevant.

By taking the double limit as both r and g approach zero, and setting $Z = s = 0$ for expositional simplicity, we can also derive an intuitive expression for the limiting value of the derivative with respect to r.[29] From Proposition 2, given the flat yield curve, we have

$$\frac{\partial c^*}{\partial r} = -c^* \left(\delta_L - \delta_W\right)$$

and in this limiting case we have

$$\begin{aligned} \lim_{r \to 0, g \to 0} \delta_W|_{Z=s=0} &= \frac{R}{2} \\ \lim_{r \to 0, g \to 0} \delta_L|_{Z=s=0} &= R + \frac{E}{2} \end{aligned}$$

[29] We take the double limit for expositional simplicity. If we only take the limit as $r \to 0$, the terms below are all functions of g, since higher values of g impact on duration of both assets and liabilities. Non-zero values of s and Z introduce some cross-product terms but qualitative results are similar.

both of which features are fairly intuitive. With w constant and equal to 1, $W_0 = R$, and duration of wealth at time zero is simply the average maturity of all wage payments, given by $\frac{R}{2}$. The duration of liabilities at $t = R$ will be $\frac{E}{2}$, but a zero coupon bond that would provide this amount will have duration R at time zero, so its total duration is $R + \frac{E}{2}$. Hence

$$\lim_{r \to 0, g \to 0} \frac{dc^*}{dr}|_{Z=s=0} = -\frac{E}{Y}\left(\frac{R+E}{2}\right).$$

or equivalently

$$\lim_{r \to 0, g \to 0} \delta_{c^*}|_{Z=s=0} \equiv \frac{\partial \ln c^*}{\partial r} = \frac{R+E}{2}.$$

i.e., in this doubly limiting case the notional duration of c^* itself (the semi-elasticity with respect to the interest rate) is equal to the duration of a zero coupon bond with maturity $\frac{R+E}{2}$, where $R + E$ is life expectancy at time zero.

3.2 A linear approximation for the fair contribution rate

The doubly limiting case allows us to derive the first-order Taylor approximation to (12)

$$c^* \approx c_A^* = \frac{E}{Y}\left(1 - \frac{(R+E)}{2} r\right) \tag{13}$$

As is to be expected this provides a good approximation close to $g = r = 0$, but is modestly biased downwards with respect to larger deviations of both g and r from zero, as Figure 4 shows for a reasonably wide range of values of r.

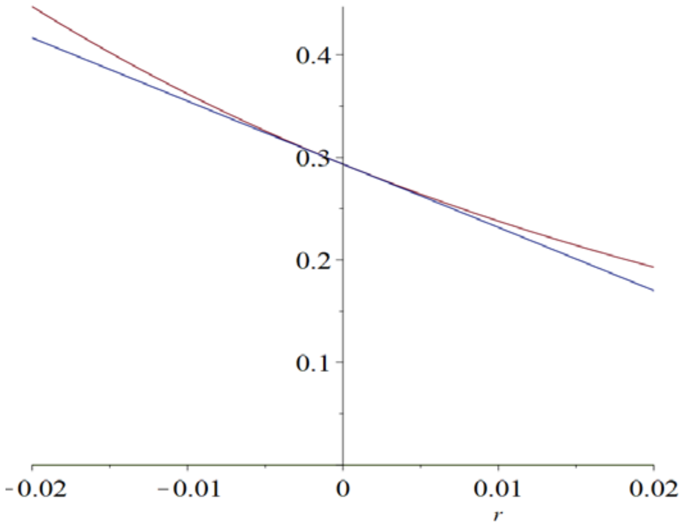

Figure 4: (c^* brown, c_A^* blue)

Note that the approximation has the feature

$$\frac{\partial^2 c_A^*}{\partial R \partial r} < 0$$

For $E = 22.5$ as above, and e.g., $R = 40$, we have

$$c^* \approx \frac{22.5}{75}\left(1 - \frac{1}{2}(40 + 22.5)\,r\right) = 0.3 - 9.38r$$

so a one percentage point fall in r increases c^* by roughly 9 percentage points in the neighbourhood of $r = g = 0$.

3.3 Comparison with a final salary scheme

Until 2011 for new members, and 2016 for all, the USS scheme provided a pension based on final, rather than average salary. The simplified version of the model also allows a straightforward comparison of the two approaches.

Inspection of the simplified expression (12) for c^* above shows that it can be factored into three components as

$$c^* = \frac{p}{W_0} F(r) = \frac{\alpha \left(\frac{e^{gR}-1}{g}\right)}{\frac{e^{(g-r)R}-1}{g-r}} F(r)$$

where $F(r)$ is a weighted sum of the capitalisation and discount factors in curly brackets in (12), that depend only on r. By inspection, for $r = 0$, the terms in g in the ratio p/W_0 cancel precisely, for any g, implying the result that, as shown above, c^* is near-invariant to g for r close to zero.

In contrast, given our assumption on the wage process, a final salary scheme provides a pension given by

$$p_{FS} = \alpha e^{gR}. \tag{14}$$

which implies that the associated fair contribution rate can be factored as

$$c^*_{FS} = \frac{p_{FS}}{W_0} F(r) = \frac{\alpha e^{gR}}{\frac{e^{(g-r)R}-1}{g-r}} F(r).$$

If $r = 0$, we then have

$$\frac{p_{FS}}{W_0} = \frac{\alpha g e^{gR}}{e^{gR}-1}$$

which is a convex, increasing function of g. The corresponding linear approximation in the neighbourhood of $r = g = 0$ is given by

$$c^*_{FS} \approx \frac{E}{Y}\left(1 + \left(\frac{R}{2}\right) g - \left(\frac{R+E}{2}\right) r\right)$$

so higher growth of earnings and a longer period to retirement interact to boost the fair contribution rate in a final salary scheme.

3.4 The impact of early leavers

The calculation underpinning the formula for c^* in Proposition 1 is predicated on the apparently restrictive assumption that the scheme member will stay in the scheme, with certainty, until retirement at $t = R$. However, it is straightforward to show that the impact of this assumption is minimal, at least in the case of the simplified example.

Consider an extension to the example, in which the scheme member has a constant probability q of quitting the scheme over the interval $t + \Delta t$ over a small interval Δt. If the scheme member quits the scheme's pension liability will be fixed at the level accrued at the time they quit, but so will accrued assets. Thus the pension that will be paid is now a random variable, and we now need to calculate liabilities in terms of the expected pension, given by

$$E_0(p) = \int_0^R e^{(g-q)t} dt = \frac{e^{(g-q)R} - 1}{g - q}$$

At the same time the relevant measure of wealth, out of which contributions are accrued also needs to be adjusted for the quit probability, so this becomes

$$W_0 = \frac{e^{(g-q-r)R} - 1}{g - q - r}$$

but we can immediately note that the formulae for $E_0(p)$ and W_0 are identical to those for p and W_0 calculated above if we replace g with $\widetilde{g} = g - q$. Since c^* is near-invariant to g for r close to zero, it must also be near-invariant to $\widetilde{g}$ and hence to q.

4 The Actuarially Fair Contribution Rate with a Vasicek Model of the Yield Curve

We now extend the model by implementing the Vasicek model (1977) of the yield curve. This allows us two extensions: (1) we can both extrapolate and interpolate market yields where market data is not available,[30] and (2) simulations of the actuarially fair contribution rates under different scenarios for the shape of the yield curve.

[30] We use yields form the Bank of England database, for which the maximum maturity is 40 years; at various points yields are also missing at some intermediate maturities.

The Vasicek model of the instantaneous rate $r(t)$ is

$$dr = a(b - r)\,dt + \sigma dz \tag{15}$$

where a is the rate of mean reversion, b is the long-run mean interest rate, σ is the rate's volatility and $dz \sim N(0, dt)$. Integrating both sides, we have

$$r(t) = r_0 e^{-at} + b\left(1 - e^{-at}\right) + \sigma e^{-at} \int_0^t e^{aS} dz_S. \tag{16}$$

This process in turn implies a particular function form for the market yield curve at any time t, which we can exploit in the definition of c^*, as follows:

Proposition 3 *If the spot rate $r(t)$ follows a Vasicek (1977) process, the actuarially fair contribution rate c^* is as given in Proposition 1, setting*

$$y(0, m) = \frac{F}{m} r_0 + (1 - \frac{F}{m})\widetilde{b} + \frac{\sigma^2}{4a}\left(\frac{F}{m}\right)^2$$

throughout, where $\widetilde{b} = \lim_{m \to \infty} y(0, m)$ *is given by*

$$\widetilde{b} = b - \frac{\sigma^2}{2a^2}$$

Proof. See Appendix C. ■

In Section 5 we calibrate the initial parameters (r_0, a, b, σ) to match current market zero coupon bond prices. In doing so, we show in Section 5 below that the resulting estimate of the actuarially fair contribution rate c^* is extremely close to that using the full set of market yields. Having done this, we can then simulate the impact on c^* of shifts in the key Vasicek parameters.

5 Numerical Estimates and Simulations

We apply UK data to estimate the actuarially fair contribution rates in Propositions 1 and 3. These are shown in Figure 3 in the introduction to the paper. For this,

- The hazard rate data for UK are published by the Office of National Statistics, for which the latest data available are for 2015-17 (Appendix D).

- For interest rates, the Bank of England publishes the RPI instantaneous real zero coupon yields and forward rates (https://www.bankofengland.co.uk/). The historical year-end rates for $1985 - 2017$, as well as the rates for 10 August 2018 are used in the following numerical simulations. In 2011 (Appendix E) USS changed the inflation index used to adjust pension payouts from RPI to CPI. In line with this, the historical real interest rates are also adjusted from 2011 using a RPI-CPI wedge. For example a wedge of $83bp$ is applied for the 30 August 2018 (see Appendix E for a discussion). The historical spot (r_0), 10-year (y_{10}) and 20-year (y_{20}) CPI adjusted yields are depicted in Figure 5.

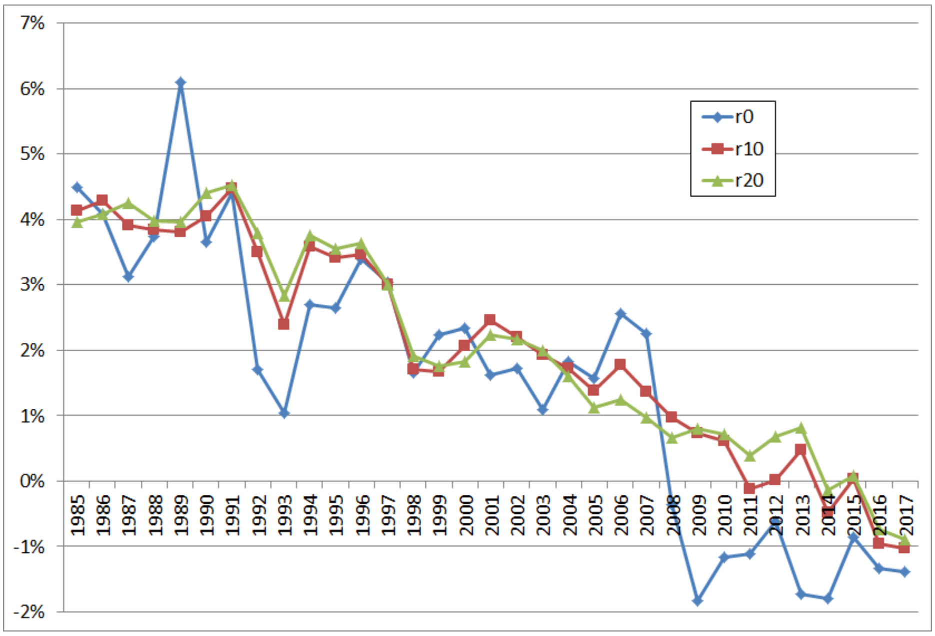

Figure 5: Real Yields (spot, 10-year and 20-year), $1985 - 2010$, RPI; $2011 - 2017$, CPI

- The term structure of real yields for 20 August 2018 is shown in Figure 6. As shown, real are currently negative at all maturities.

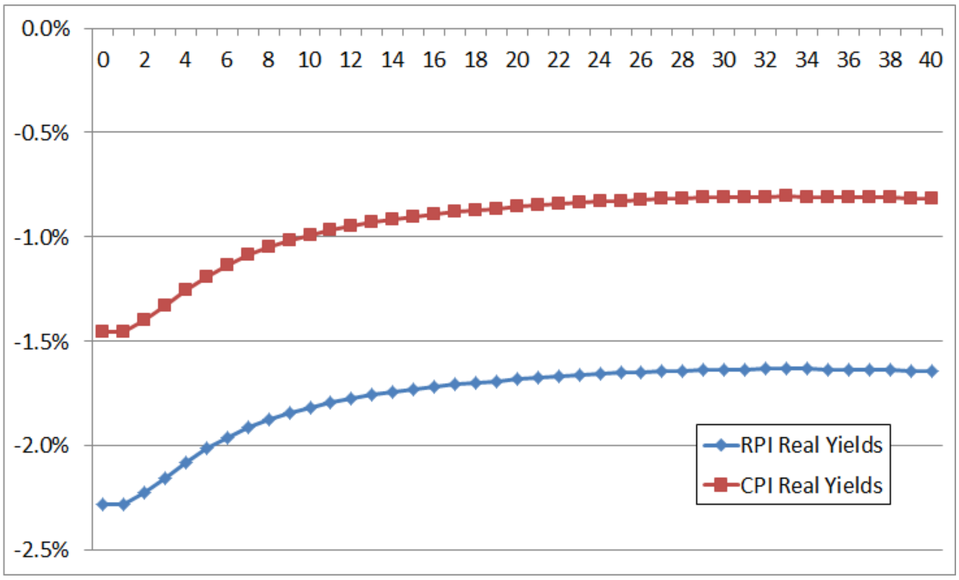

Figure 6: RPI and CPI Instantaneous Real Yields, 20 August 2018 (source: Bank of England; RPI-CPI wedge estimated by authors)

- For the Vasicek model, parameters (a, b, σ) are calibrated to fit the yields of zero coupon bonds $y(0, m)$ given for 10 August 2018. This yields the values of $(0.050, 1.291\%, 1.086\%)$. (Appendix F)

Note that b estimated here is the long-run mean of the instantaneous spot rate, $\lim_{t\to\infty} E[r(t)]$, where $r(t)$ was given in (16). The long-run means for zero coupon bond yields $y(0, m)$, and the corresponding instantaneous forward returns $f(0, m) = \frac{\partial}{\partial m}[y(0, m)]$, include a convexity term, such that $\lim_{m\to\infty} y(0, m) = \lim_{m\to\infty} f(0, m) = \widetilde{b}$, where $\widetilde{b} = b - \frac{\sigma^2}{2a^2}$. The calibrated value of $\widetilde{b}$ on the market yield data used here is -1.088%. (See Appendix F for further discussion.)

5.1 Calculating historical Fair Contribution Rates

The actuarially fair contribution rates in Proposition 1 are computed and plotted in Figure 3, in the introduction to this paper, using the year-end real instantaneous forward rates (and the CPI-adjusted rates from 2011), and hazard rate data as described above. Two cases are shown, where the worker joins the pension scheme at the ages of 47 and 27, where the retirement age is assumed to be 67, hence $R = 20$ or 40. Before 2011, when the scheme defined the pension in relation to final salary, the expression for p in (3) is replaced with $p_{FS} = \alpha e^{gR}$ as in (14). α is set as $\frac{1}{80}$ and $\frac{1}{75}$ for up to 2015 and 2016 and after, respectively.

Also plotted are the historical USS contribution rates (employees and employers combined). The breakdown between the employees and employers are given in Appendix A.

5.2 Fair Contribution Rates using Market Yields and Counterfactual Vasicek Parameters

Using the interest rates for 10 August 2018, the fair contribution rates in Propositions 1 and 3 are estimated and compared for $g = 0\%$:

Years in Scheme	Market Forward Rates	Vasicek Model
20	36.8%	36.8%
40	39.3%	39.2%

(Table 1)

The comparison makes clear that the Vasicek model allows sufficient flexibility in the calibrated parameters to give virtually identical results to those calculated using the full set of market forward rates/yields.

We can also compare the effects of a change in the wage growth rate g on the fair contribution rate for 20 and 40 years of contributions, using market yields:

g	0%	1%	2%	3%	4%
20 years in scheme	36.8%	36.6%	36.4%	36.3%	36.1%
40 years in scheme	39.2%	38.9%	38.6%	38.3%	38.0%

(Table 2)

This demonstrates that $\frac{\partial c^*}{\partial g}$ is relatively small and negative. This is consistent with our simplified analytical example, with $Z = s = 0$, where we have $\text{sign}\left(\frac{\partial c^*}{\partial g}\right) = \text{sign}(r)$.

The Vasicek model allows us to analyse different scenarios for the crucial model parameters.

We first consider the impact of a shift in b, the implied mean of the instantaneous spot rate in the Vasicek model. We represent the impact in terms of a shift in the limiting value of the yield and forward curve, $\widetilde{b} = b - \frac{\sigma^2}{2a^2}$, since this magnitude is directly comparable with market yields. For $g = 0$ we have

$\widetilde{b}$	-3.0%	-2.0%	-1.0%	0.0%	1.0%	2.0%	3.0%
20 years in scheme	46.6%	41.1%	36.4%	32.3%	28.9%	25.9%	23.3%
40 years in scheme	60.6%	48.2%	38.4%	30.7%	24.6%	19.7%	15.9%

(Table 3)

This demonstrates that, compared to the current Vasicek estimate of $\widetilde{b} = -1.088\%$, the long-run mean rate would need to rise to between $0 - 1\%$ if workers contribute for 40 years, and to nearly 2% if they only contribute for 20 years, in order for the current aggregate contribution rate of 26% to be actuarially fair.

Note that in this initial exercise the current interest rate r_0 is kept fixed while the long-run mean rate is shifted, in effect rotating the yield curve pivoted at r_0. In Table 4 we instead simulate parallel-shifts in the yield curve:

r_0	-3.4%	-2.4%	-1.4%	-0.4%	0.6%	1.6%	2.6%
$\widetilde{b}$	-3.0%	-2.0%	-1.0%	0.0%	1.0%	2.0%	3.0%
20 years in scheme	53.7%	43.9%	36.1%	29.9%	24.8%	20.7%	17.3%
Implied "duration"	–	20.0	19.5	19.0	18.5	18.0	17.5
40 years in scheme	67.7%	51.0%	38.2%	28.5%	21.2%	15.7%	11.6%
Implied "duration"	–	28.9	29.3	29.8	30.2	30.5	30.7

(Table 4)

The implied duration figures shown here have a direct correspondence to the definition of (spot rate) durations applied in Proposition 2, expressed in terms of discrete changes. Thus the implied durations for c^* are approximately given here by the change in $\ln c^*$ moving right across the table (since yields at all maturities are shifting by one percentage point). These correspond quite well to the implied figures for δ_{c^*} given by our linear approximation for c^* in Section 3, which implies $\delta_{c^*} \approx \frac{R+E}{2}$ for r close to zero. For higher values of c^* the notional duration falls off, given the convexity of the true function for c^*.

With such high implied durations for c^* it is unsurprising that the range of different values of c^* is wide, and highly sensitive to real interest rates. To illustrate the impact of different input values, in Table 5 we simulate the fair contribution rates for different hypothetical yield curves depicted in Figure 7 using hazard rates for 10 August 2018. To show the impact of different hazard rate functions, in Table 6 we recompute the numbers in Table 5 using the hazard rates for 1985 (see Figure 8).

R	g	Up (< 0)	Flat (< 0)	Down (< 0)	Up (> 0)	Flat (> 0)	Down (> 0)
20	0%	34.0%	38.4%	43.7%	23.6%	26.3%	29.5%
20	2%	33.7%	38.0%	43.1%	23.7%	26.4%	29.6%
40	0%	32.2%	41.0%	52.4%	18.0%	23.0%	29.3%
40	2%	32.0%	40.1%	50.4%	18.9%	23.7%	29.7%

(Table 5: Impact of different yield curve shapes given latest mortality data.)

R	g	Up (< 0)	Flat (< 0)	Down (< 0)	Up (> 0)	Flat (> 0)	Down (> 0)
20	0%	28.5%	31.4%	34.8%	20.9%	22.8%	25.0%
20	2%	28.5%	31.4%	34.6%	21.2%	23.1%	25.2%
40	0%	25.8%	32.1%	39.9%	15.2%	18.9%	23.5%
40	2%	26.1%	31.8%	38.9%	16.2%	19.8%	24.2%

(Table 6: Impact of different yield curve shapes given mortality data for 1985.)

Tables 5 and 6 show that:

1. The higher the long-run yield, the lower the fair contribution rate.

2. For downward (> 0) where the long-run yield is 0%, the fair contribution rate is insensitive to changes in R or g.

3. The fair contribution rates are lower for longer R if the yields are negative, while they are higher for positive yields.

4. The increase in hazard rates since 1985 (i.e. longer life expectancy) have increased the fair contribution rate, with the increase being larger for longer R, lower g and negative yields.

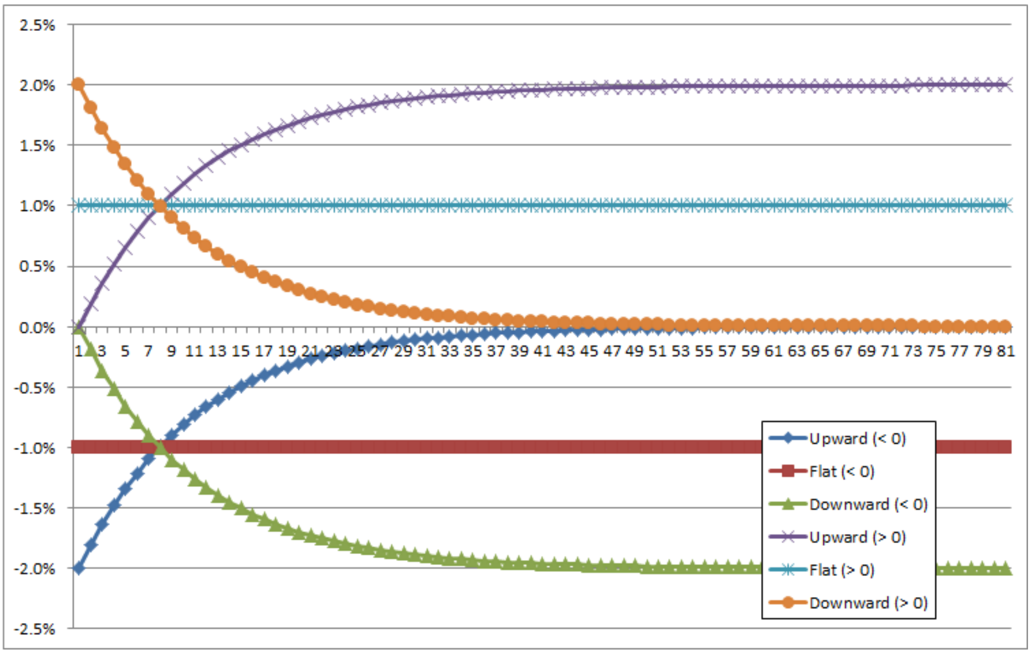

Figure 7: Different Yield Curve used for Simulation

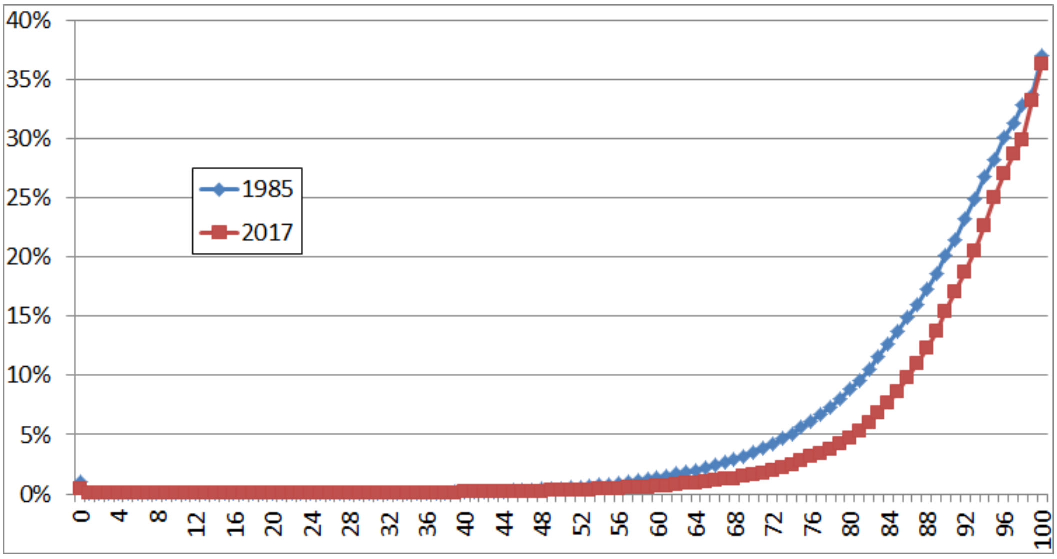

Figure 8: Hazard Rates for 1985 and 2017

Finally, we investigate the relative importance of the four payout terms in Proposition 1, i.e. the contribution rate required to cover (i) the lump-sum payout to the beneficiary if the worker dies before R, (ii) the lump-sum payout to the worker at R, (iii) the pension payout to the worker from R until their death, and (iv) the pension payments to the spouse if they outlive the worker. For a 20 year contribution and $g = 0$, we have

$\widetilde{b}$	-3.0%	-2.0%	-1.0%	0.0%	1.0%	2.0%	3.0%
Lump-sum if die before R	1.7%	1.7%	1.6%	1.6%	1.6%	1.6%	1.6%
Lump-sum to worker at R	4.4%	4.2%	4.0%	3.8%	3.7%	3.5%	3.3%
Pension payout to worker	33.6%	29.3%	25.6%	22.5%	19.8%	17.5%	15.5%
Pension payout to spouse	7.0%	6.0%	5.1%	4.4%	3.8%	3.3%	2.9%
Total	46.6%	41.1%	36.4%	32.3%	28.9%	25.9%	23.3%

(Table 7)

Unsurprisingly the pension payout to the worker is the largest element (around 70% of the total payout). The relative weight of the three remaining terms corresponds reasonably well to their respective contributions to "effective life expectancy", as discussed in Section 3 in relation to our simplified example. The only exception is that the contribution of the spouse's pension is distinctly higher, since the calculations in that discussion only relate to benefits payable after retirement.

6 Concluding Remarks

Since we summarise the key conclusions of the paper in the Introduction, we do not repeat them here; we focus instead on some qualifications, and possible objections, to our modelling approach.

Clearly, even in our more complicated calculations described below, we have made important simplifications. If anything, however, we would argue that allowing for these would increase, rather than decrease our estimates of typical fair contribution rates. Specifically:

- We treat wage growth as if it had no risk associated with it. This has two important effects on our calculations. First, c^* is assumed constant throughout the scheme member's working life. If, for example, the member's salary rises unexpectedly above or below what was factored into the original c^* calculation, the scheme cannot change its contribution rate on these higher contributions. But by the time they occur market interest rates may have changed, implying a different value of c^*. This means that the scheme carries interest rate risk - it is in effect providing a free option to members. Second, to the extent that wage growth is correlated with market returns, it has "positive beta", which means that the present value of contributions is less than that implied by the calculation using yields on risk-free bonds.

- We make no adjustment for the transactions costs of running the scheme.

- We do not attempt to extrapolate the steady rise in life expectancy observed in the data.

- Nor do we take account of any statistical evidence of links between life expectancy and socioeconomic factors, which would almost certainly point to USS members having higher life expectancies than the average population.

As against these arguments, a counter-argument that is often proposed is that the contributions scheme members invest in the scheme can be (and indeed typically have been) invested in risky assets with higher expected returns than those on risk-free bonds. We would argue strongly that this argument is mistaken. Our calculations are based purely on present value calculations that only require us to know the properties of contributions into, and contingent liabilities of, the scheme. These are invariant to the scheme's investment strategy, and thus the investment strategy has no impact on fair contribution rates.

We could of course imagine a scheme that would not have this invariance property; but it would not be a defined benefit scheme. Thus, at the opposite extreme, in a pure defined contribution scheme the liabilities of the scheme are *entirely* determined by the investment strategy.

It is possible, in principle, to imagine a hybrid scheme in which liabilities were not risk-free (contingent on the scheme member's mortality status), but instead had some form of explicit dependence on investment returns.[31] To the extent that liabilities in such a scheme then had "positive beta" with respect to market risk factors, which in turn had positive market risk premia, this would straightforwardly imply a lower fair contribution rate, albeit at the cost of a riskier income in retirement. However even such a hybrid scheme would almost certainly share the key characteristics we have stressed throughout the paper, that fair contribution rates would be increasing in life expectancy and decreasing in *expected* investment returns. Thus even such a hybrid scheme would face the same inherent instability as a pure defined benefit scheme like USS if contribution rates were held fixed. Furthermore there would be a sharp contrast with the relatively straightforward calculation of fair contribution rates presented in this paper, which require only information which is readily available on a virtually real-time basis.

[31] Indeed some attempts have been made to develop such hybirid schemes, see, for example, Munnell & Sass (2013).

Calculation of fair contribution rates for such hybrid schemes would be both considerably more complex, and distinctly more contentious, since it would require assumptions on expected market returns and risk premia, all of which are the subject of extensive debate and continued disagreement, both amongst academics and practitioners.

References

[1] Bank of England (2006), "New information from inflation swaps and index-linked bonds", *Bank of England Quarterly Bulletin 2006*

[2] Bank of England (2015), "The informational content of market-based measures of inflation expectations derived from government bonds and inflation swaps in the United Kingdom", *Staff Working Paper* No.**551**

[3] Hull, J. (2012), *Options, Futures, and Other Derivatives*, 8th edition, Pearson

[4] Mamon, R.S. (2004), "Three Ways to Solve for Bond Prices in the Vasicek Model", *Journal of Applied Mathematics and Decision Sciences*, **8**(1), 1-14

[5] Munnell, A.H. and Sass, S.A. (2013), "New Brunswick's new shared risk pension plan", Centre for Retirement Research, Boston College, *Working Paper*

[6] Office for Budget Responsibility (2015), "Revised assumption for the long-run wedge between RPI and CPI inflation", *https://obr.uk/box/revised-assumption-for-the-long-run-wedge-between-rpi-and-cpi-inflation/*

[7] Vasicek, O.A. (1977), "An Equilibrium Characterization of the Term Structure", *Journal of Financial Economics*, **5**, 177-188

7 Appendix

A Historical USS and Simulated Fair Contribution Rates

The historical USS contribution rates for employees and employers are sourced from Wikipedia. The period between 1974 and 1997 includes a 2% surcharge aimed at covering benefits for service prior to the scheme's inception in 1974. Between 2011 and 2016 the employee contribution rate was 7.5% for existing workers with final salary scheme and 6.35% for new workers with career average revalued earnings (CARE).

Year	Scheme	Alpha	RPI/CPI	Employee	Employer	USS	g	R = 20	R = 40
1985	Final	1/80	RPI	8.00%	18.55%	26.55%	2.0%	18.1%	14.1%
1986	Final	1/80	RPI	8.00%	18.55%	26.55%	2.0%	18.4%	15.5%
1987	Final	1/80	RPI	8.00%	18.55%	26.55%	2.0%	16.6%	12.0%
1988	Final	1/80	RPI	8.00%	18.55%	26.55%	2.0%	17.3%	12.6%
1989	Final	1/80	RPI	8.00%	18.55%	26.55%	2.0%	17.5%	12.7%
1990	Final	1/80	RPI	8.00%	18.55%	26.55%	2.0%	15.9%	10.8%
1991	Final	1/80	RPI	8.00%	18.55%	26.55%	2.0%	16.3%	11.6%
1992	Final	1/80	RPI	8.00%	18.55%	26.55%	2.0%	17.1%	11.7%
1993	Final	1/80	RPI	8.00%	18.55%	26.55%	2.0%	19.9%	16.6%
1994	Final	1/80	RPI	8.00%	18.55%	26.55%	2.0%	17.9%	13.8%
1995	Final	1/80	RPI	8.00%	18.55%	26.55%	2.0%	18.3%	14.0%
1996	Final	1/80	RPI	8.00%	18.55%	26.55%	2.0%	18.1%	13.6%
1997	Final	1/80	RPI	6.35%	14.00%	20.35%	2.0%	20.5%	17.3%
1998	Final	1/80	RPI	6.35%	14.00%	20.35%	2.0%	24.3%	23.6%
1999	Final	1/80	RPI	6.35%	14.00%	20.35%	2.0%	27.2%	30.9%
2000	Final	1/80	RPI	6.35%	14.00%	20.35%	2.0%	27.5%	30.9%
2001	Final	1/80	RPI	6.35%	14.00%	20.35%	2.0%	25.0%	25.4%
2002	Final	1/80	RPI	6.35%	14.00%	20.35%	2.0%	24.4%	23.5%
2003	Final	1/80	RPI	6.35%	14.00%	20.35%	2.0%	25.2%	25.2%
2004	Final	1/80	RPI	6.35%	14.00%	20.35%	2.0%	28.1%	30.4%
2005	Final	1/80	RPI	6.35%	14.00%	20.35%	2.0%	31.0%	34.9%
2006	Final	1/80	RPI	6.35%	14.00%	20.35%	2.0%	32.2%	36.9%
2007	Final	1/80	RPI	6.35%	14.00%	20.35%	2.0%	33.4%	38.7%
2008	Final	1/80	RPI	6.35%	14.00%	20.35%	2.0%	36.8%	47.2%
2009	Final	1/80	RPI	6.35%	16.00%	22.35%	2.0%	34.0%	42.1%
2010	Final	1/80	RPI	6.35%	16.00%	22.35%	2.0%	33.1%	38.0%
2011	CARE	1/80	CPI	6.35%	16.00%	22.35%	2.0%	27.5%	25.6%
2012	CARE	1/80	CPI	6.35%	16.00%	22.35%	2.0%	25.4%	22.4%
2013	CARE	1/80	CPI	6.35%	16.00%	22.35%	2.0%	26.5%	24.7%
2014	CARE	1/80	CPI	6.35%	16.00%	22.35%	2.0%	30.5%	29.2%
2015	CARE	1/80	CPI	6.35%	16.00%	22.35%	2.0%	31.0%	30.4%
2016	CARE	1/75	CPI	8.00%	18.00%	26.00%	2.0%	35.4%	38.3%
2017	CARE	1/75	CPI	8.00%	18.00%	26.00%	2.0%	36.0%	37.7%

Table 8: Fair and USS Actual Contribution Rates

B Proof of Proposition 2

Proof. (11) follows immediately from the definition:

$$c^* = \frac{L_0}{W_0} \Rightarrow -\frac{\partial \ln c^*}{\partial r_0} = -\frac{1}{c^*}\frac{\partial c^*}{\partial r_0} = -\left(\frac{\partial \ln L_0}{\partial r_0} - \frac{\partial \ln W_0}{\partial r_0}\right) = \delta_L - \delta_W.$$

For the sign of δ_{c^*}, recall L_0 is,

$$\begin{aligned} L_0 &= B\int_0^R \lambda(\tau)\, e^{\left(g-y(0,\tau)-\overline{\lambda}(0,\tau)\right)\tau} d\tau + sp\int_0^\infty \lambda(\tau)\, e^{-\overline{\lambda}(0,\tau)\tau}\left\{\int_\tau^\infty e^{-\left(y(0,t)+\overline{\lambda}(0,t)\right)t} dt\right\} d\tau \\ &\quad + Zpe^{-\left(y(0,R)+\overline{\lambda}(0,R)\right)R} + p\int_R^\infty e^{-\left(y(0,\tau)+\overline{\lambda}(0,\tau)\right)\tau} d\tau \\ &= L_0^B + L_0^s + L_0^Z + L_0^p. \end{aligned}$$

Then δ_L is a weighted average of the duration of its components:

$$\delta_L = \frac{L_0^B}{L_0}\delta_{L^B} + \frac{L_0^s}{L_0}\delta_{L^s} + \frac{L_0^Z}{L_0}\delta_{L^Z} + \frac{L_0^p}{L_0}\delta_{L^p}. \tag{17}$$

Compare this with δ_W. Note first that,

$$\delta_W = \frac{1}{W_0}\int_0^R te^{\left(g-y(0,t)-\overline{\lambda}(0,t)\right)t} dt < \frac{1}{W_0}\int_0^R R.e^{\left(g-y(0,t)-\overline{\lambda}(0,t)\right)t} dt = \frac{1}{W_0} RW_0 = R$$

as $t \in [0, R]$ in the integrand. Analogously for the latter two terms of (17),

$$\begin{aligned} \delta_{L^Z} &= \frac{1}{L^Z} RZpe^{-\left(y(0,R)+\overline{\lambda}(0,R)\right)R} = \frac{1}{L^Z} RL^Z = R \\ \delta_{L^p} &= \frac{1}{L^p} p\int_R^\infty \tau e^{-\left(y(0,\tau)+\overline{\lambda}(0,\tau)\right)\tau} d\tau > \frac{1}{L^p} p\int_R^\infty R.e^{-\left(y(0,\tau)+\overline{\lambda}(0,\tau)\right)\tau} d\tau = \frac{1}{L^p} RL^p = R \end{aligned}$$

i.e. both terms are greater than or equal to δ_W. As for the first term in (17),

$$\delta_{L^B} = \frac{1}{L^B} B\int_0^R \tau\lambda(\tau)\, e^{\left(g-y(0,\tau)-\overline{\lambda}(0,\tau)\right)\tau} d\tau = \frac{\int_0^R \tau\lambda(\tau)\, e^{\left(g-y(0,\tau)-\overline{\lambda}(0,\tau)\right)\tau} d\tau}{\int_0^R \lambda(\tau)\, e^{\left(g-y(0,\tau)-\overline{\lambda}(0,\tau)\right)\tau} d\tau} = \int_0^R \tau w^{L_B}(\tau)\, d\tau \tag{18}$$

where

$$w^{L_B}(\tau) = \frac{\lambda(\tau)\, e^{\left(g-y(0,\tau)-\overline{\lambda}(0,\tau)\right)\tau} d\tau}{\int_0^R \lambda(\tau)\, e^{\left(g-y(0,\tau)-\overline{\lambda}(0,\tau)\right)\tau} d\tau}.$$

In other words δ_{L^B} is a weighted average of the times of death-in-service payouts τ, $\tau \in [0,R]$, with the weights given by $w^{L_B}(\tau)$. Compare this with

$$\delta_W = \frac{\int_0^R \tau e^{\left(g-y(0,\tau)-\overline{\lambda}(0,\tau)\right)\tau} d\tau}{\int_0^R e^{\left(g-y(0,\tau)-\overline{\lambda}(0,\tau)\right)\tau} d\tau} = \int_0^R \tau w^W(\tau)\, d\tau \tag{19}$$

where

$$w^W(\tau) = \frac{e^{\left(g-y(0,\tau)-\overline{\lambda}(0,\tau)\right)\tau} d\tau}{\int_0^R e^{\left(g-y(0,\tau)-\overline{\lambda}(0,\tau)\right)\tau} d\tau}.$$

Note for constant $\lambda(\tau) = \lambda\ \forall \tau$, $w^{L_B}(\tau) = w^W(\tau)$. Then for $\frac{d\lambda(\tau)}{d\tau} > 0$, $w^{L_B}(\tau)$ gives higher weights to larger τ than $w^W(\tau)$, and hence $\delta_{L^B} > \delta_W$. This leaves the second term in (17),

$$\delta_{L^s} = \frac{sp}{L_0^s}\int_0^\infty \lambda(\tau)\, e^{-\overline{\lambda}(0,\tau)\tau}\left\{\int_\tau^\infty te^{-\left(y(0,t)+\overline{\lambda}(0,t)\right)t} dt\right\} d\tau = \frac{\int_0^\infty \lambda(\tau)\, e^{-\overline{\lambda}(0,\tau)\tau}\left\{\int_\tau^\infty te^{-\left(y(0,t)+\overline{\lambda}(0,t)\right)t} dt\right\} d\tau}{\int_0^\infty \lambda(\tau)\, e^{-\overline{\lambda}(0,\tau)\tau}\left\{\int_\tau^\infty e^{-\left(y(0,t)+\overline{\lambda}(0,t)\right)t} dt\right\} d\tau}.$$

This can be rewritten as,

$$\delta_{L^s} = \frac{\int_0^\infty te^{-\left(y(0,t)+\overline{\lambda}(0,t)\right)t}\left\{\int_0^t \lambda(\tau)\, e^{-\overline{\lambda}(0,\tau)\tau} d\tau\right\} dt}{\int_0^\infty e^{-\left(y(0,t)+\overline{\lambda}(0,t)\right)t}\left\{\int_0^t \lambda(\tau)\, e^{-\overline{\lambda}(0,\tau)\tau} d\tau\right\} dt}.$$

In splitting D_L^s into two,

$$\delta_{L^s} = \frac{L_0^{s<R}}{L_0^s}\delta_{L^s_{<R}} + \frac{L_0^{s\geq R}}{L_0^s}\delta_{L^s_{\geq R}}$$

where

$$\begin{aligned} L_0^s &= sp\int_0^\infty e^{-\left(y(0,t)+\overline{\lambda}(0,t)\right)t}\left\{\int_0^t \lambda(\tau)\, e^{-\overline{\lambda}(0,\tau)\tau} d\tau\right\} dt \\ &= sp\int_0^R e^{-\left(y(0,t)+\overline{\lambda}(0,t)\right)t}\left\{\int_0^t \lambda(\tau)\, e^{-\overline{\lambda}(0,\tau)\tau} d\tau\right\} dt + sp\int_R^\infty e^{-\left(y(0,t)+\overline{\lambda}(0,t)\right)t}\left\{\int_0^t \lambda(\tau)\, e^{-\overline{\lambda}(0,\tau)\tau} d\tau\right\} dt \\ &= L_0^{s<R} + L_0^{s\geq R}, \end{aligned}$$

once again it can be seen that,

$$
\begin{aligned}
\delta_{L^s_{\geq R}} &= \frac{1}{L_0^{s\geq R}} \int_R^\infty te^{-\left(y(0,t)+\overline{\lambda}(0,t)\right)t} \left\{ \int_0^t \lambda(\tau) e^{-\overline{\lambda}(0,\tau)\tau} d\tau \right\} dt \\
&> \frac{1}{L_0^{s\geq R}} \int_R^\infty R.e^{-\left(y(0,t)+\overline{\lambda}(0,t)\right)t} \left\{ \int_0^t \lambda(\tau) e^{-\overline{\lambda}(0,\tau)\tau} d\tau \right\} dt = \frac{1}{L_0^{s\geq R}} RL_0^{s\geq R} = R.
\end{aligned}
$$

Finally for

$$
\delta_{L^s_{<R}} = \frac{\int_0^R te^{-\left(y(0,t)+\overline{\lambda}(0,t)\right)t} \left\{ \int_0^t \lambda(\tau) e^{-\overline{\lambda}(0,\tau)\tau} d\tau \right\} dt}{\int_0^R e^{-\left(y(0,t)+\overline{\lambda}(0,t)\right)t} \left\{ \int_0^t \lambda(\tau) e^{-\overline{\lambda}(0,\tau)\tau} d\tau \right\} dt}
$$

to be greater than δ_W in (19), the sufficient (but not necessary) condition is that $e^{g\tau}$ grows more slowly than $\int_0^t \lambda(\tau) e^{-\overline{\lambda}(0,\tau)\tau} d\tau$ (again such that larger weights are assigned to larger t in $\delta_{L^s_{<R}}$ than in δ_W), or that g is less than some critical value g^- such that $\delta_{L^s_{<R}} = \delta_W(g^-)$. ■

We note that $\frac{d\lambda(t)}{dt} > 0$ for working age for all historical data attained from ONS $(1982 - 2017)$, and similarly that $\delta_{L^s_{<R}} > \delta_W$ comfortably for an extremely wide range of values for g.

C Proof of Proposition 3

Analogous to (2), the expectation of the present value of the scheme member's lifetime earnings within the scheme is now

$$
W_0 = \widehat{E}_0 \left[\int_0^R e^{\left(g-\overline{\lambda}(0,t)\right)t} e^{-\int_0^t r_S dS} dt \right], \tag{20}
$$

where $\widehat{E}_0[\,]$ is the risk-neutral expectation over stochastic r_t taken at $t = 0$. Similarly the expectation of the total present value of the scheme member's pension payouts is,

$$
\begin{aligned}
L_0 = \; & \widehat{E}_0 \left[B \int_0^R \lambda(\tau) e^{\left(g-\overline{\lambda}(0,\tau)\right)\tau} e^{-\int_0^\tau r_S dS} d\tau + Zpe^{-\overline{\lambda}(0,R)R} e^{-\int_0^R r_S dS} \right. \qquad (21) \\
& \left. + p \int_R^\infty e^{-\overline{\lambda}(0,\tau)\tau} e^{-\int_0^\tau r_S dS} d\tau + Sp \int_0^\infty \lambda(\tau) e^{-\overline{\lambda}(0,\tau)\tau} \left\{ \int_\tau^\infty e^{-\overline{\lambda}(0,t)t} e^{-\int_0^t r_S dS} dt \right\} d\tau \right],
\end{aligned}
$$

where p is given by (9).

As before c^* is given by $\frac{L_0}{W_0}$, where, from (20) and (21),

$$
\begin{aligned}
W_0 &= \int_0^R e^{\left(g-\overline{\lambda}(0,t)\right)t}\widehat{E}_0\left[e^{-\int_0^t r_S dS}\right]dt \\
L_0 &= B\int_0^R \lambda(\tau)\, e^{\left(g-\overline{\lambda}(0,\tau)\right)\tau}\widehat{E}_0\left[e^{-\int_0^\tau r_S dS}\right]d\tau + Zpe^{\left(g-\overline{\lambda}(0,R)\right)R}\widehat{E}_0\left[e^{-\int_0^R r_S dS}\right] \\
&\quad +p\int_R^\infty e^{-\overline{\lambda}(0,\tau)\tau}\widehat{E}_0\left[e^{-\int_0^\tau r_S dS}\right]d\tau + Sp\int_0^\infty \lambda(\tau)\, e^{-\overline{\lambda}(0,\tau)\tau}\left\{\int_\tau^\infty e^{-\overline{\lambda}(0,t)t}\widehat{E}_0\left[e^{-\int_0^t r_S dS}\right]dt\right\}d\tau,
\end{aligned}
$$

where the price of a zero coupon bond with maturity t at $t=0$ is

$$P(0,t) = \widehat{E}_0\left[e^{-\int_0^t r_S dS}\right].$$

To find the expression for this substitute from the Vasicek model (16),

$$
\begin{aligned}
P(0,t) &= \exp\left\{-\int_0^t \left(r_0 e^{-aS} + b\left(1-e^{-aS}\right)\right)dS\right\}\widehat{E}_0\left[\exp\left\{-\int_0^t\left(\sigma e^{-aS}\int_0^S e^{au}dz_u\right)dS\right\}\right] \\
&= \exp\left\{-bt-(r_0-b)\frac{1-e^{-at}}{a}\right\}\widehat{E}_0\left[\exp\left\{-\int_0^t\left(\sigma e^{-aS}\int_0^S e^{au}dz_u\right)dS\right\}\right].
\end{aligned}
$$

Consider the term $\widehat{E}_0[.]$. Noting that for a stochastic variable X with $\widehat{E}_0[X]=0$, $\widehat{E}_0\left[e^X\right] = e^{\frac{1}{2}Var[X]}$, we have,

$$\widehat{E}_0\left[\exp\left\{-\int_0^t\left(\sigma e^{-aS}\int_0^S e^{au}dz_u\right)dS\right\}\right] = \exp\left\{\frac{\sigma^2}{2}Var\left[\int_0^t\left(e^{-aS}\int_0^S e^{au}dz_u\right)dS\right]\right\}.$$

Now,

$$
\begin{aligned}
Var\left[\int_0^t\left(e^{-aS}\int_0^S e^{au}dz_u\right)dS\right] &= \int_0^t\int_0^t Cov\left[e^{-aS}\int_0^S e^{au}dz_u, e^{-av}\int_0^v e^{au}dz_u\right]dSdv \\
&= \int_0^t\int_0^t e^{-a(S+v)}\widehat{E}_0\left[\int_0^S e^{au}dz_u\int_0^v e^{au}dz_u\right]dSdv \\
&= \int_0^t\int_0^t e^{-a(S+v)}\left(\int_0^{S\wedge v} e^{2au}du\right)dSdv \\
&= \frac{1}{2a}\int_0^t\int_0^t e^{-a(S+v)}\left(e^{2a(S\wedge v)}-1\right)dSdv,
\end{aligned}
$$

where $S \wedge v = \min[S, v]$. Continuing,

$$
\begin{aligned}
&= \frac{1}{2a}\int_0^t \left[\int_0^v e^{-a(S+v)}\left(e^{2aS}-1\right)dS\right]dv + \frac{1}{2a}\int_0^t \left[\int_v^t e^{-a(S+v)}\left(e^{2av}-1\right)dS\right]dv \\
&= \frac{1}{2a^2}\int_0^t \left(1+e^{-2av}-2e^{-av}\right)dv + \frac{1}{2a^2}\int_0^t \left(-e^{-a(t-v)}+e^{-a(t+v)}+1-e^{-2av}\right)dv \\
&= \frac{1}{2a^3}\left(2at-3+4e^{-at}-e^{-2at}\right).
\end{aligned}
$$

Thus,

$$
\begin{aligned}
P(0,t) &= \exp\left\{-bt-(r_0-b)\frac{1-e^{-at}}{a}+\frac{\sigma^2}{4a^3}\left(2at-3+4e^{-at}-e^{-2at}\right)\right\} \\
&= \exp\left\{-r_0\left(\frac{1-e^{-at}}{a}\right)+\left(b-\frac{\sigma^2}{2a^2}\right)\left(\frac{1-e^{-at}}{a}-t\right)-\frac{\sigma^2}{4a}\left(\frac{1-e^{-at}}{a}\right)^2\right\} \\
&= E(0,t)\exp\{-F(0,t)\,r_0\}. \qquad (22)
\end{aligned}
$$

where

$$
\begin{aligned}
E(0,t) &= \exp\left\{\left(b-\frac{\sigma^2}{2a^2}\right)(F(0,t)-t)-\frac{\sigma^2}{4a}F(0,t)^2\right\} \qquad (23) \\
F(0,t) &= \frac{1-e^{-at}}{a}. \qquad (24)
\end{aligned}
$$

If we then define

$$P(0,m) = e^{-y(0,m).m}$$

we then have

$$y(0,m) = \frac{F}{m}r_0 + \left(1-\frac{F}{m}\right)\widetilde{b} + \frac{\sigma^2}{4a}\left(\frac{F}{m}\right)^2$$

as given in the proposition. ■

D Hazard Rates

The Office of National Statistics releases the national life table (source: https://www.ons.gov.uk/). The mortality rate using the data for 2015-17 is shown in Figure 9.

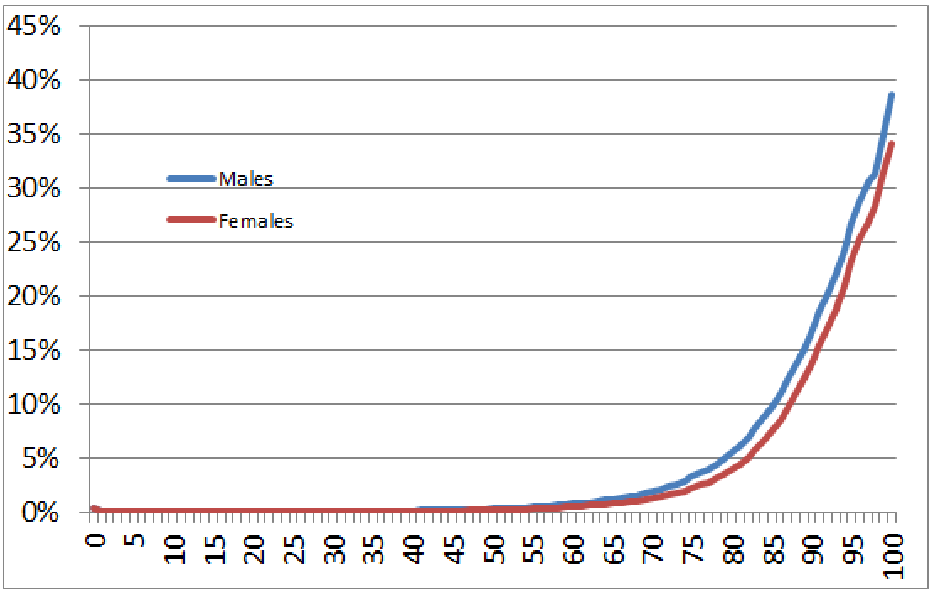

Figure 9: UK mortality rate, 2015-17 (source: ONS)

Using these data the average age of death can be calculated as 79.2 years for male, 82.9 years for female and 81 years for the whole population.

E UK Real Yields and Forward Rates

The Bank of England publishes real yields and instantaneous forward rates for RPI inflation

(https://www.bankofengland.co.uk/statistics/yield-curves/) for years 2.5 to 40. As of 10 August 2018 these rates are negative as shown in Figure 2. As USS changed the inflation index applied to adjust pension payouts from RPI and CPI in 2011, it is necessary to adjust this to real rates reflecting CPI. To do this the RPI-CPI wedge is estimated using a 20 year rolling average. The wedge applied is between 52 and 92bp, with 82bp for the most recent data (August 2018).

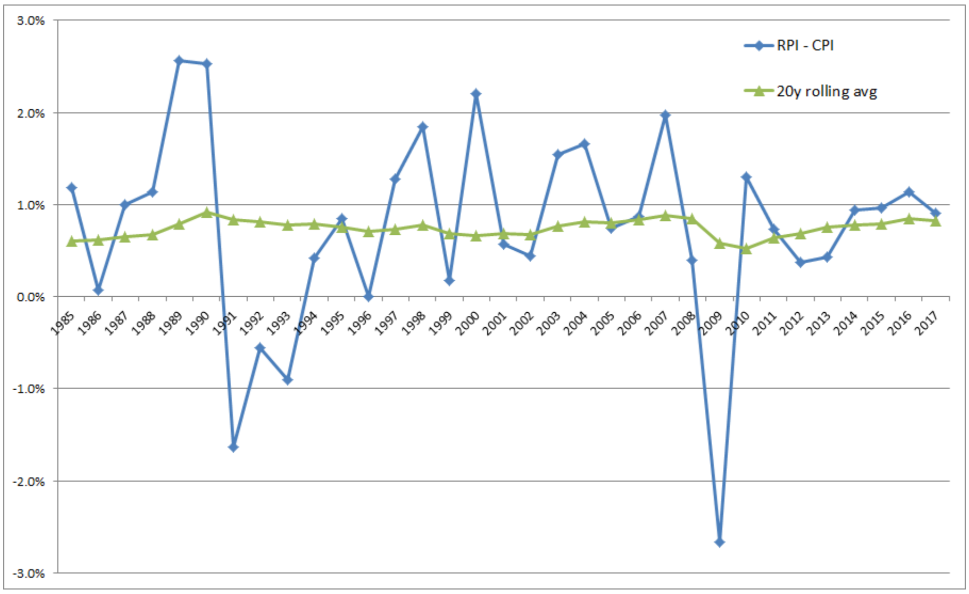

Figure 10: RPI - CPI Wedge and 20 year Rolling Average

This contrasts with the “long-run average difference between RPI and CPI inflation” of $80bp$ applied in Bank of England Quarterly Bulletin (2006), or the statement in their Staff Working Paper No.551 (2015) that, “[o]ver longer horizons, the expected RPI/CPI wedge appears fairly stable at around 66 basis points”. On the other hand, OBR (2015) quote their “new estimate of the long-run wedge between RPI and CPI inflation of 1.0 percentage points”. Our estimate of $82bp$ for the most recent data lies within the appropriate range of these estimates.

F Vasicek Parameter Estimation

Given the market price of zero coupon bonds with maturity t, $D(0,t)$, the Vasicek parameters (a,b,σ) were calibrated to minimised the following error function,

$$\min_{(a,b,\sigma)} \sum_t \left[P(0,t) - D(0,t)\right]^2,$$

where $P(0,t)$ is given by Eqn (22). For 10 August 2018, this resulted in the estimates $(0.050, 1.291\%, 1.086\%)$. In the Vasicek model, the future interest rates have the expectation of,

$$E\left[r\left(t\right)\right]=r_0e^{-at}+b\left(1-e^{-at}\right),$$

which is derived by taking the expectation of Eqn (16). On the other hand, the forward rates seen at time 0 includes a convexity term:

$$f\left(0,t\right)=r_0e^{-at}+b\left(1-e^{-at}\right)-\frac{\sigma^2}{2a^2}\left(1-e^{-at}\right)^2.$$

In other words the market forward rate curve is lower than the path of future spot rates estimated by the model by the convexity term,

$$f\left(0,t\right)=E\left[r\left(t\right)\right]-\frac{\sigma^2}{2a^2}\left(1-e^{-at}\right)^2. \tag{25}$$

The reason for this is because bond prices are convex in rates, which means that the average price is higher than the price at the average rate (Jensen's inequality applied to convex functions), or equivalently, the average bond returns are lower than the average rate by the convexity term. Specifically, the average bond return is,

$$R\left(0,t\right)=-\frac{1}{t}\ln P\left(0,t\right)=-\frac{1}{t}\ln E\left(0,t\right)+\frac{1}{t}F\left(0,t\right)r_0.$$

In the limit then,

$$\lim_{t\to 0}R\left(0,t\right)=r_0 \text{ and } \lim_{t\to\infty}R\left(0,t\right)=\widetilde{b}\text{, where } \widetilde{b}=b-\frac{\sigma^2}{2a^2}.$$

Figure 11 shows the calibrated forward curve fitting closely with the market forward curve. The difference between the $E\left[r\right]$ curve and the fitted forward curve is the convexity term $\frac{\sigma^2}{2a^2}\left(1-e^{-at}\right)^2$ in (25).

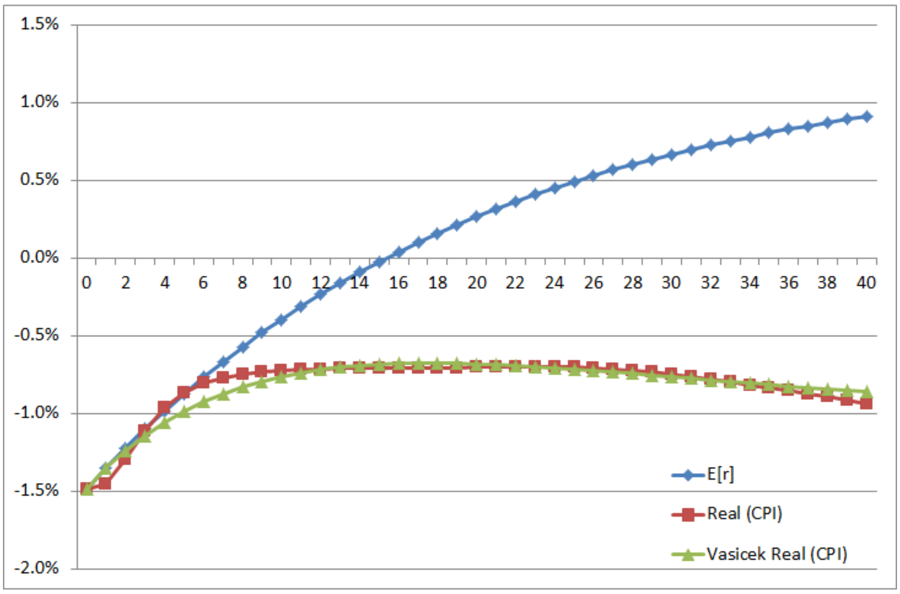

Figure 11: Market and Vasicek-fitted Real (CPI) Instantaneous Forward Rates, 10 August 2018