

Pension Plan Solvency and Extreme Market Movements: A Regime Switching Approach

Niloufar Abourashchi, Iain Clacher, Mark Freeman,
David Hillier, Malcolm Kemp and Qi Zhang*

July 30, 2012

Abstract

In this paper we provide estimates of the probability that a defined benefit pension plan will become underfunded in the future. Our techniques allow both for the ‘fat-tailed’ nature of asset returns and the correlation between discount rate changes and asset returns. We show that future projections of pension plan solvency risk are highly model dependent. In particular, if discount rates are assumed to remain constant, then regardless of the choice of model for asset returns, the probability of future deficits is dramatically understated compared to forecasts where a stochastic discount rate process is used. Moreover, our results suggest that introducing leptokurtosis into asset returns greatly increases the estimated probability of a future deficit. However, it is also important to note that allowing for correlations between asset returns and the discount rate process significantly reduces the assessed probability of future underfunding.

*Mark Freeman, Author for Correspondence, School of Business and Economics, Loughborough University, United Kingdom. Email: M.C.Freeman@Lboro.ac.uk. Abouraschi, Clacher and Zhang are from the University of Leeds, Hillier is from the University of Strathclyde and Kemp is from Nematrian. **Acknowledgements:** The authors gratefully acknowledge financial support from the Rotman International Centre for Pension Management (ICPM) at the University of Toronto and from the Actuarial Profession in the United Kingdom. We would also like to thank conference participants at the October 2011 Discussion forum co-hosted by ICPM and World Bank Treasury in Washington DC, the Fifth International Conference in Mathematical and Statistical Methods for Actuarial Sciences and Finance, Venice 2012 and participants at the May 2012 Actuarial Profession Sessional Meeting in Leeds, as well as seminar participants from the University of Leeds, the University of Bradford and the University of Sheffield for their useful and insightful comments as this work has developed. The views expressed in this report and the full academic paper reflects those of the authors and not of the ICPM or the Actuarial Profession. All errors remain our own.

1 Introduction

Since the start of the credit crisis the aggregate funding position of defined benefit (DB) pension schemes in the UK has swung dramatically. A surplus of £149bn in June 2007 turned into a deficit of £209bn by March 2009. By February 2011 this had returned to a surplus of £36bn before swinging back to a deficit of £312bn in May 2012. This latest figure represents a funding ratio of only 77%.¹

In this paper we present detailed projections of how the funding position of DB schemes might evolve in the future. At the aggregate level, this work is of relevance to the Pension Protection Fund and macro-prudential regulators. At the individual fund level, our results should be of interest to those involved in pension management e.g., actuaries, trustees and sponsors.

The methods that we employ capture three key features of the data that recent market conditions have highlighted. First, we allow for the fat-tailed nature of asset returns, particularly on the downside, that has been so clearly revealed in recent years and whose relevance was highlighted by the UK Actuarial Professions Benchmarking Stochastic Models (BSM) working party that considered modelling extreme market events (Frankland et al., 2008). Second, incorporating the stochasticity of the discount rate into the model is essential because even when asset prices rise, the funding position of a DB scheme can deteriorate as a result of changes in the discount rate. This phenomena has been observed since March 2009, when the aggregate funding position of UK DB schemes worsened by approximately £100bn at a time when asset values actually increased from £837bn to £1031bn, due to even greater rises (from £972bn to £1343bn) in the present value of future liabilities as a result of a decline in gilt yields.² Finally, it is important to capture the dynamics of the correlation between the asset returns and discount rate process. Unconditionally, discount rate changes are negatively correlated with asset returns as both bond and equity prices tend to rise when interest rates are falling. This means that there is a positive correlation between the present value of liabilities and the market value of assets. This, though, is not always the case. Between June 2007 and March 2009 the aggregate value of DB assets fell by approximately 8%, from £837bn to £771bn, while simultaneously, the present value of future liabilities rose by almost 40%, from £695bn to £972bn.

Our core empirical methodology involves the use of multivariate Gaussian regime switch-

¹This is the deficit of 6,432 UK schemes covered by the Pension Protection Fund PPF 7800 index.

²The Pension Protection Fund uses a valuation process consistent with section 179 of the Pensions Protection Act. This contrasts with IAS19 and FRS17 measures that are used for company accounting purposes. Of particular relevance here is that vanilla and index-linked Gilts are used to calculate the present value of future liabilities.

ing models. This framework is well suited to modelling asset returns and extreme market events (See Ang and Timmermann, 2011 for a recent review). Not only does it capture leptokurtosis in asset returns, but it has also been shown to accurately capture the increased correlations between asset returns that often occur in bear markets (Ang and Bekaert, 2002).

Our paper is related to two predominant streams of existing literature. The first is the body of work that examines optimal contribution rates and asset allocation decisions for defined benefit pension plans. Josa-Fombellida and Roncón-Zapatero (2011), for example, consider the optimal portfolio for minimising the terminal solvency risk of a DB scheme in the presence of stochastic interest rates, asset returns and benefits in a non-regime switching environment. Our specific focus, though, is somewhat different from this literature. Rather than determining what fund managers 'should' do, we provide a framework for undertaking a risk assessment of current decisions regardless of whether or not they conform to what academics might view as being theoretically optimal. By removing the need to optimise, we are able to deal with more complex statistical representations of the data than is common in this literature. Apart from Frauendorfer et al. (2007), Markov switching models have not, to our knowledge, been widely used when determining optimal DB pension plan portfolios.³ This is despite the fact that they have been applied in more general asset allocation contexts by, for example, Ang and Bekaert (2002), Zhou and Yin (2003), Guidolin and Timmermann (2007, 2008) and Shen and Siu (2012). Even here, though, the Markov representations these authors use are simpler than our own; Guidolin and Timmermann (2008), for example, have six asset classes but only two states, in contrast to the four employed in this paper.

Our work is perhaps more closely related to Value-at-Risk (VaR) problems; we will draw out this parallel in more detail in Section 2. Here, Markov switching models have been used by a number of authors including Billio and Pelizzon (2000), Guidolin and Timmermann (2006b) and Taamouti (2009). Kawata & Kijima (2007) demonstrate that these models estimate 1% Value-at-Risk levels for portfolios better than many alternative approaches. Ferstl and Weissensteiner (2011) evaluate the optimal portfolio that will minimise the conditional VaR for a DB pension scheme, but again they do not consider regime switching within their setting.

Our main result is that the assessed risk of pension scheme underfunding is highly dependent on the precise specification of the econometric process driving asset returns and changes in the discount rate. Given that most of the recent volatility in pension scheme funding has arisen from changes in the present value of liabilities, it is perhaps no surprise

³Chen and Yang (2010) also use Markov switching models within the context of DB pension scheme funding issues, but again their specific focus, which evaluates in a theoretical manner the optimal dividend to be paid back to the fund sponsor, differs markedly from our own.

that incorporating stochasticity in the discount rate significantly increases the assessed level of future funding risk. Similarly, incorporating leptokurtosis into asset returns increases the calculated probability of a fund going into deficit in future, although this has a smaller impact than the stochasticity of the discount rate. Finally, because of the unconditional negative correlation between asset returns and interest rate changes, incorporating this relationship reduces perceived future funding risk. This is true even within our Markov switching model, as even in the crash state, this correlation remains negative. This contrasts with the situation between June 2007 and March 2009 but is representative with the more recent post-crisis period.

The conclusion that we draw is that pension fund stakeholders need models at least as sophisticated as the ones presented here if they are to fully appreciate future solvency risk. Techniques that are easier to estimate and more pragmatic are likely to lead to materially incorrect inferences and poor decision making by those involved in pension management as well as regulators.

Our paper proceeds as follows. In Section 2, we describe the economic environment and illustrative portfolio that will be used to generate our results. In section 3, we consider the funding risk of our illustrative portfolio in a single-state model, where asset returns are lognormally distributed. In section 4, we will allow for extreme market movements by modelling asset returns as a Markov switching process. Section 5 concludes.

2 The economic environment

We consider a defined benefit pension plan that has assets under management and a future stream of liabilities that must be met. The scheme is based in the United Kingdom and payouts are sterling denominated. It invests in $N = 5$ asset classes at the index level: UK equities, US equities, European (ex UK) equities, UK gilts and Japanese equities. The assets have a combined market value at time t of V_t and we use $PV_t(L)$ to denote the present value of the future liabilities as calculated at time t . The variable that the stakeholders of this fund are interested in is z_t , which is the funding ratio of the scheme minus one:

$$z_t = \frac{V_t - PV_t(L)}{PV_t(L)}$$

The scheme is fully (under) funded if $z_t > (<)0$. Our central research question is to examine the statistical properties of z_t and, in particular, assess the probability that the scheme will be underfunded at future time t under current investment and contribution policies; $\text{Prob}(z_t) < 0$.

The problem that we face is related, but not identical, to a VaR calculation. In the latter case, for some confidence interval α the fund manager finds the value of future assets VaR_t that is the solution to the equation: $\text{Prob}(V_t < VaR_t) = \alpha$. In the situation considered in this paper, there are two adjustments to this equation. First, instead of fixing α , the value of VaR_t is determined and then α is derived as a consequence. That is, we determine probabilities of falling into deficit rather than asset values associated with pre-determined left-hand tail probabilities. Second, our research question can be rephrased as assessing $\text{Prob}(V_t < PV_t(L))$, where both the left and right hand sides of the equation are stochastic and correlated. Therefore our work can be seen as being complementary to the work of Billio and Pelizzon (2000), Guidolin and Timmermann (2006b), Kawata & Kijima (2007) and others.

In order to assess future solvency risk, we assume that the stakeholders of the scheme decide to take a frequentist approach. Historical data is used to determine the relevant underlying statistical relationships and it is then assumed that these will not change in the future. While this is a common assumption in both industry and academia, it is not uncontroversial. For example, there is an extensive literature that argues that average historic returns to equity substantially overestimate the current ex-ante equity premium (see, for example, Freeman, 2011, and the references therein). For the purposes of this paper we make no adjustments for potential limitations with the frequentist approach although it would be mathematically straightforward to do so as long as we could formulate suitable assumptions to use in its stead. For example, if one were to believe that unconditional expected equity returns are too high, then the mean estimates could be manually adjusted downwards to reflect this belief. The regime-switching model would still add value by capturing some of the complex dynamics of the volatility process.

Calibration is undertaken using monthly data from January 1970 to December 2010. For UK equities, data for the FTSE All Share total returns index are taken from Datastream. All other data for the asset side of the balance sheet are taken from Global Financial Data. These are the United Kingdom 10 year Government Bond total return index, S&P500 total returns index, Japanese Topix total returns index and MSCI Europe total returns index. As the scheme is UK-based, all returns are calculated in sterling terms.

In the table below, we present summary statistics for the total monthly returns to each series that we consider on the asset side of the balance sheet. Throughout this study we use lognormal nominal returns; $r_{nt} = \ln(I_{nt}/I_{nt-1})$, where I_{nt} is the total returns index of asset

class n :

	Mean	Standard Deviation	Skewness	Excess Kurtosis
UK Equity Returns	0.809%	5.730%	-0.710	7.436
US Equity Returns	0.895%	5.264%	-0.637	2.448
European Equity Returns	0.936%	4.716%	-1.010	3.904
UK Gilt Returns	0.784%	1.742%	0.446	2.424
Japanese Equity Returns	0.833%	6.012%	-0.054	0.517

with correlations over the period (with the assets in the same order):

$$\begin{bmatrix} 1.00 & 0.56 & 0.83 & 0.25 & 0.33 \\ & 1.00 & 0.70 & 0.02 & 0.40 \\ & & 1.00 & 0.15 & 0.47 \\ & & & 1.00 & 0.06 \\ & & & & 1.00 \end{bmatrix}$$

The pension scheme invests a proportion of its total wealth w_i in each of the asset classes with $w_1 + \dots + w_N = 1$ and assets ordered as in the table above ($i = 1$ for UK equities, $i = 4$ for UK gilts, etc.). Let w be the N -vector with elements w_i . We assume that w is fixed across time with weights rebalanced each period; $w = \{25\%, 20\%, 10\%, 40\%, 5\%\}$. This is broadly in line with current standard UK pension investment policies; see, for example, the Pension Protection Fund Purple Book. In this case, the first four moments of monthly logarithmic portfolio returns, r_p , based on our sample of data are a mean of 0.830%, standard deviation of 2.96%, skewness of -0.680, and excess kurtosis of 4.646. Again, it would be straightforward to adjust our techniques to allow for time-variation in asset allocation but this is not undertaken here for reasons of parsimony.

The liability structure of a defined benefits pension scheme is highly complex. Amongst other factors, future payouts from the scheme will depend on longevity risk, age of retirement, salary growth, additional contributions and other contribution rate changes, taxation changes, and withdrawals from and additions to the scheme. As the central focus of this paper is to understand the impact of different econometric choices on assessed future funding risk, we assume a highly stylised form for expected future liabilities of the DB scheme. At all times these stretch over the next thirty years. The first expected liability, in one year's time, is C_1 . Following that, the liabilities are expected to grow at a fixed inflation rate i .

This leads to the following schedule of expected future liabilities:

E_0 [Liabilities]	$t = 1$	$t = 2$	$t = 3$...	$t = 30$	$t = 31$	$t = 32$...
Year 0	C_1	$C_1(i + 1)$	$C_1(1 + i)^2$...	$C_1(1 + i)^{29}$	0	0	...
Year 1	0	$C_1(i + 1)$	$C_1(1 + i)^2$...	$C_1(1 + i)^{29}$	$C_1(1 + i)^{30}$	0	...
Year 2	0	0	$C_1(1 + i)^2$...	$C_1(1 + i)^{29}$	$C_1(1 + i)^{30}$	$C_1(1 + i)^{31}$...
...

Given this schedule, the present value of liabilities that will be calculated at any time t is given by the growth annuity formula $PV_t(L) = C_1(1 + i)^{t-1} G(r_{ft}, i, 30)$ where $G(r_{ft}, i, 30)$ is the thirty year growth annuity value based on the rate, r_{ft} , that is used at time t to determine the present value of future liabilities:⁴

$$z_t = \frac{V_t - C_1(1 + i)^{(t-1)}G(r_{ft}, i, 30)}{C_1(1 + i)^{(t-1)}G(r_{ft}, i, 30)}$$

$$G(r_{ft}, i, T) = \frac{1}{r_{ft} - i} \left(1 - \left(\frac{1 + i}{1 + r_{ft}} \right)^T \right)$$

The initial expected cash flow, C_1 , is set so that $C_1G(r_{f0}, i, 30) = 1/(1 + z_0)$, where z_0 represents the initial solvency level of the scheme expressed so that, for every £1 of current liabilities the fund has market assets valued at £1 + z_0 . For our illustrative example we set $z_0 = 15\%$. We also assume that there are no net contributions to the asset side of the balance sheet after time zero as new contributions exactly offset fund payouts. It should be stressed that these assumptions are for simplicity of exposition only. The processes that we describe could easily be extended to more complex cash flow dynamics and net inflows / outflows from the fund.

In order to calculate the present value of the liabilities, we use the 10-year UK gilt yield, rather than the AA corporate bond rate, as given by Datastream. While we do not necessarily advocate risk-free discounting of pension liabilities, the choice of this rate is broadly consistent with section 179 of the Pension Protection Act that was used by the Pension Protection Fund to construct the summary statistics presented in the introduction to this paper. Again, the methods that we use are broadly insensitive to the choice of yield that is used for discounting purposes.

Our central interest is in calculating the probability that $z_t < 0$ for all t up to a horizon

⁴We ignore term structure issue here. Many financial economists would argue that each separate cash flow should be discounted at an individual rate that reflects the shape of the term structure that prevails at that horizon. While we agree with this point, for simplicity, and in keeping with many practitioners, at time t we discount cash flows of all horizons at the same discount rate r_{ft} .

of thirty years based on six different econometric specifications of the data. The first three take a simple one-state lognormal process for asset returns while the last three incorporate Markov regime switching. In Models 1 and 4 it is assumed that the discount rate is fixed at its current level: $r_{ft} = r_{f0}$ for all t . All other models incorporate an AR(1) or, equivalently, discrete-time Ornstein-Uhlenbeck (O-U) model for the discount rate. In Models 2 and 5 the discount rate process is assumed to be independent of asset returns. Models 3 and 6, however, incorporate the observed correlation between discount rate changes and portfolio returns.

3 Modelling the risk of underfunding: 1-state

Let μ_t denote the N -vector with elements $E_t[r_{nt+1}]$ and Σ_t denote the $N \times N$ matrix with elements $Cov_t(r_{nt}, r_{mt})$. Throughout this section it is assumed that $\mu_t = \mu$ and $\Sigma_t = \Sigma$ for all t . This assumption will be relaxed in Section 4. In addition, it is assumed that log returns are multivariately normally distributed.

In this case, the T -period portfolio logarithmic return has expected value and variance $m_T = Tw'\mu$ and $\sigma_T^2 = Tw'\Sigma w$. This characterisation of the data accurately captures the observed mean and standard deviation of historical asset returns, and the contemporaneous correlation between any two asset classes. It does, though, have a number of weaknesses. First, it does not capture any higher moments of asset returns — in particular their ‘fat-tailed’ properties that are clearly present in the summary statistics of the data presented above. This means that this model underestimates extreme market events and therefore underestimates the true spread of possible future portfolio values. Second, it is well-known that the variance-covariance matrix of asset returns is time-varying, and this is not captured within this setting.

The variable of interest is the probability that the DB scheme becomes underfunded: $\text{Prob}(z_t < 0)$ or equivalently after taking logarithms of both sides of the numerator of z_t , $\text{Prob}(\ln(V_t) < \ln(C_1(1+i)^{(t-1)}G(r_{ft}, i, 30)))$. As $\ln(V_t) \sim N(m_t, \sigma_t^2)$, this becomes a z -score problem and therefore solutions will be expressed in terms of the cumulative distribution function for a standard normal distribution, $\Phi(\cdot)$.

Model 1 is our simplest. We add the assumption that the risk-free rate will remain constant at its current value, $r_{ft} = r_{f0}$ for all t , where $r_{f0} = 3.53\%$ is the gilt yield in December 2010. In this case $\ln(C_1(1+i)^{(t-1)}G(r_{ft}, i, 30))$ is non-stochastic and:

$$\text{Prob}(z_t < 0) = \Phi\left(\frac{\ln(C_1(1+i)^{(t-1)}G(r_{f0}, i, 30)) - m_t}{\sigma_t}\right)$$

Model 2 relaxes the constant discount rate assumption by now letting it follow an independent AR(1), or discrete-time Ornstein-Uhlenbeck, process. The continuous-time form of this process, as suggested initially by Vařicek (1977), has been used in a related context by Josa-Fombellida and Ronc3n-Zapatero (2011):

$$r_{ft} - r_{ft-1} = \begin{cases} a + br_{ft-1} + e_t \\ \theta(r_{ft-1} - \bar{r}) + e_t \end{cases}$$

where $\theta = b$, $\bar{r} = -a/\theta$ and $e_t \sim N(0, \sigma_e^2)$. This gives the AR(1) process an economic interpretation; \bar{r} represents the long-run interest rate value to which the discount rate mean-reverts and θ is the parameter value that determines the speed and strength of the mean reversion. While more sophisticated interest rate models are available, this choice introduces analytical tractability for a number of reasons. First, the probability density function (pdf) of r_{ft} , $f(r_{ft})$, conditional on the current discount rate, r_{f0} , is normally distributed with mean $E[r_{ft}]$ and variance $\sigma^2(r_{ft})$:

$$E[r_{ft}] = (1 + b)^t r_{f0} - \frac{1 - (1 + b)^t}{b} a, \quad \sigma^2(r_{ft}) = \frac{1 - (1 + b)^{2t}}{1 - (1 + b)^2} \sigma_e^2$$

This is a feature we will exploit in Model 3 below, where we allow for the correlation between asset returns and the discount rate. Second, the transformed variable x_t :

$$x_t = r_{ft} - (1 + b)r_{ft-1} = a + e_t$$

is independent and identically normally distributed. This will be helpful when we come to our most sophisticated model below, Model 6, which co-estimates the discount rate and asset returns processes simultaneously within a Markov switching environment. Basing our parameterisations on the deannualised monthly 10-year UK gilt yield for the interval January 1970 to December 2010, we derive estimates of $a = 0.0000165$, $b = -0.00359$ and $\sigma_e = 0.000285$.

By the law of total probability, $\text{Prob}(z_t < 0) = E_{r_{ft}}[\text{Prob}(z_t < 0 | r_{ft})]$. When the portfolio return is independent of the discount rate process, this is equivalent to:

$$\text{Prob}(z_t < 0) = \int_{-\infty}^{\infty} \Phi \left(\frac{\ln(C_1(1+i)^{(t-1)}G(r_{ft}, i, 30)) - m_t}{\sigma_t} \right) f(r_{ft}) dr_{ft}$$

This can be evaluated by numerical integration. In Model 3, we continue to assume that asset returns are one-state lognormal and also discount rates are characterised by the AR(1)

process, but now allow for the observed correlation between e_t and the portfolio return at time t , r_{pt} . Notice that, by backward iteration:

$$\begin{aligned} r_{ft} &= (1+b)r_{ft-1} + x_t \\ &= (1+b)^t r_{f0} + \sum_{i=1}^t (1+b)^{t-i} x_i \end{aligned}$$

Therefore:

$$\begin{aligned} Cov\left(\sum_{i=1}^t r_{pt}, r_{ft}\right) &= Cov\left(\sum_{i=1}^t r_{pt}, \sum_{i=1}^t (1+b)^{t-i} e_i\right) \\ &= \sum_{i=1}^t (1+b)^{t-i} Cov(r_{pi}, e_i) \end{aligned}$$

When there is only one state, $Cov(r_{pi}, e_i) = \rho\sigma_p\sigma_e$, where ρ is the correlation between r_{pt} and e_t which is a constant. Therefore:

$$Cov\left(\sum_{i=1}^t r_{pt}, r_{ft}\right) = \rho\sigma_p\sigma_e b^{-1} [(1+b)^t - 1]$$

The variance of $\sum_{i=1}^t r_{pt} = t\sigma_p^2$ where σ_p^2 is the constant period variance of the portfolio return and therefore the correlation between $\sum_{i=1}^t r_{pt}$ and r_{ft} , ρ_t , is:

$$\rho_t = \rho \frac{(1+b)^t - 1}{b\sqrt{t} \frac{1 - (1+b)^{2t}}{1 - (1+b)^2}}$$

Applying a well-known result for the conditional distribution of bivariate normal variables:

$$\begin{aligned} f\left(\sum_{i=1}^t r_{pt} | r_{ft}\right) &\sim N(m_{pt}, s_{pt}^2) \\ m_{pt} &= m_t + \rho_t \frac{\sigma_t}{\sigma(r_{ft})} (r_{ft} - E[r_{ft}]) \\ \sigma_{pt}^2 &= (1 - \rho_t^2) \sigma_t^2 \end{aligned}$$

and so:

$$\text{Prob}(z_t < 0) = \int_{-\infty}^{\infty} \Phi\left(\frac{\ln(C_1(1+i)^{(t-1)}G(r_{ft}, i, 30)) - m_{pt}}{\sigma_{pt}}\right) f(r_{ft}) dr_{ft}$$




Figure 1: This graph presents the probability that the scheme will be underfunded at time t ; $\text{Prob}(z_t < 0)$. Results are presented for three models, all of which assume that portfolio returns are independently and identically normally distributed. In Model 1, the discount rate is fixed at its current level. In Model 2, the discount rate is stochastic but independent of the portfolio return. Model 3 allows for the correlation between the discount rate and the portfolio return process.

Again, this can be evaluated by numerical integration.

In figure 1, we present the results from the three models when $i = 4\%$. It can be seen that the probability of default first rises and then dissipates quickly in each case. This is because, under a frequentist approach the expected return to the portfolio is substantially above the assumed fixed inflation rate of 4%. Therefore the expected returns effect quickly dominates the stochasticity of the asset returns and discount rate processes.

Under Model 1, where the discount rate is fixed, the maximal probability of becoming underfunded is relatively low; around 4% after two years and four months. This then drops away quite rapidly.

Under Model 2 the results are markedly different with a much higher maximal probability of underfunding, which now lies at approximately 17% after one year and nine months. There are three key causes for this effect, two of which amplify the difference and the third of which reduces it. First, there is now a probability that discount rates will fall substantially. To illustrate this, suppose momentarily that the portfolio will deliver a non-stochastic return of $E[r_p]$ over the next twelve months. For the pension scheme now to be underfunded in one

year:

$$G(r_{f1}, i, 30) > G(r_{f0}, i, 30) (1 + z_0) e^{E[r_p]} / (1 + i)$$

Given $E[r_p] = 12 \times 0.830\%$, $z_0 = 15\%$, $i = 4\%$ and $r_{f0} = 3.53\%$, this corresponds to $r_{f1} < 2.27\%$. Based on the calibrated 1-state Ornstein-Uhlenbeck model this movement is not unlikely with a p -value of over 12%. Therefore even if there is no uncertainty over asset returns, the interest rate stochasticity effect is highly significant.

Second, there is a Jensen's inequality effect that also increases the perceived solvency risk. When interest rates are stochastic the expected net present value is greater than the present value calculated at the expected discount rate. To illustrate this, set $i = 4\%$ and $T = 30$ and suppose that the discount rate in five years, r_{f5} might equal 8% or 2% with equal probability. In this case $E[G(r_{f5}, i, T)] = 0.5 [G(8\%, 4\%, 30) + G(2\%, 4\%, 30)] = 28.236$. By contrast, $G(E[r_{f5}], i, T) = G(5\%, 4\%, 30) = 24.955$. This effect has been very heavily documented in environmental economics, where it leads to a declining schedule of discount rates with increasing time horizon (Weitzman, 2001) and, in turn, a higher social cost of carbon. In the context given here, after 12 months the Ornstein-Uhlenbeck process leads to estimates $E[r_{r1}] = 3.62\%$ and $\sigma(r_{f1}) = 1.16\%$. The growth annuity value at the expected discount rate, $G(E[r_{f1}], i, T) = 30.558$. However, when we run a simulation of 1,000,000 drawings of r_{f1} from the relevant normal distribution, $E[G(r_{f1}, i, T)] = 31.215$. In figure 2, we present a comparison of these two types of annuity value for $T = 30$ and for t up to 360 months. This demonstrates how economically significant this effect can be. Therefore, ignoring uncertainty about the discount rates at future time t will lead to an underestimate of the expected growth annuity value at time t , and thus an understatement of the expected present value of future liabilities that will prevail at time t .

The third, offsetting, effect, is that the long-run equilibrium discount rate, $\bar{r} = -a/\theta = 5.53\%$ on an annualised basis, is considerably higher than the current risk-free rate of 3.53%. This implies that the term structure of expected discount rates is upward sloping, reducing the growth annuity value estimates in the future. This feature is clearly evident from the "Known rf" line in Figure 2. We would expect the differences between Models 1 and 2 to be greater than presented here where the term structure is flat.

Figure 1 also presents results for Model 3, which allows for the covariance between discount rates and asset returns. This somewhat reduces the perceived pension fund solvency risk, with a maximal probability of underfunding now being at about 13% at one year and nine months. The reason for this is that, in general, asset returns are negatively correlated with the discount rate, with unconditional correlation between r_{pt} and e_t of -40.7% . This is most obviously driven by UK gilts, which make up 40% of the overall portfolio. As might be expected, the total return on this asset class is highly negatively correlated with the discount

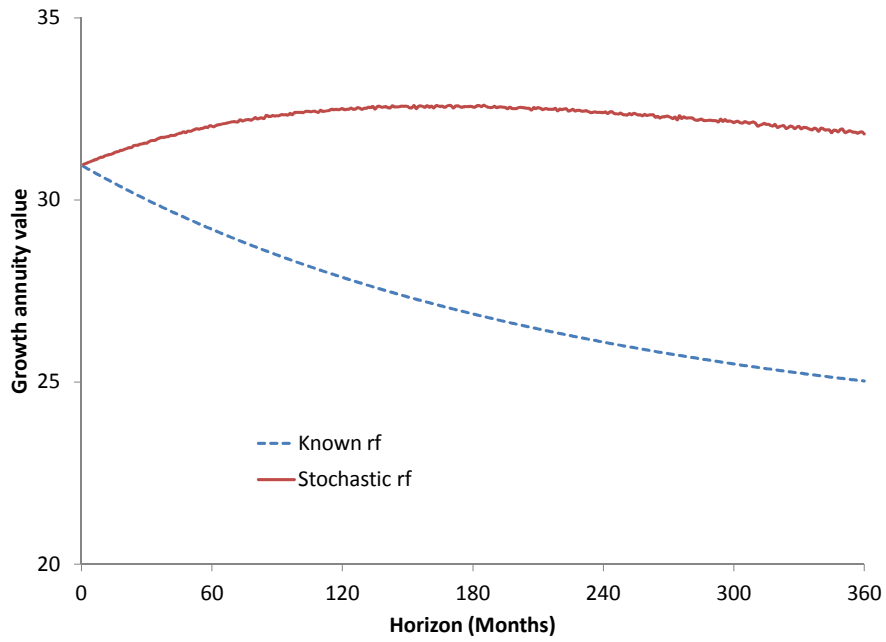


Figure 2: This graph presents the growth annuity value for $T = 30$ years for t up to 360 months. The “Known rf” line assumes that $r_{ft} = E[r_{ft}]$ with certainty, where $E[r_{ft}]$ is determined by the discrete-time Ornstein-Uhlenbeck process described in the body of the paper. The “Stochastic rf” is calculated as the average growth annuity value determined over 1,000,000 simulations when r_{ft} is drawn at random from this distribution.

rate (-86%). In addition, UK equity returns are also negatively correlated with bond yields (-32%). Therefore the market value of assets and the present value of liabilities tend to move up and down in tandem, providing a natural hedging effect for pension fund stakeholders.

4 Modelling the risk of underfunding: 4-states

To overcome the weaknesses of the traditional one-state lognormal model of asset returns we extend our analysis to incorporate extreme events into the analysis. There are a number of techniques available to do this — see, for example, Kemp (2011) in the context of pension fund management — of which perhaps the two most obvious are extreme value theory (EVT) including copulas and Markov regime switching models. We have chosen to use the latter for a number of reasons, although ultimately there is some subjectivity in this modelling choice. The most practical is that, under the UK Pension Regulator’s Code of Practice 7, there is a requirement for trustees to understand the advice that they are given. While it is difficult for a non-specialist to appreciate all the details of Markov models, we believe that they are easier to explain intuitively to pension scheme stakeholders than the EVT alternatives. Specifically, as we will show below, the different states have clear economic interpretation and can be relatively easily mapped to known historical periods to which we believe practitioners will naturally respond. Second, as argued above, our topic of interest shares many features with Value-at-Risk problems and Kawata & Kijima (2007) have shown how well Markov models perform in such context. Finally, we are dealing with a multi-asset problem while EVT tends to be applied within more restrictive environments with lower number of asset classes.

Markov models have been heavily used in economics and finance. For example, early economics studies used Markov switching models to capture structural breaks in the economy (see Hamilton, 1989; Lam, 1990; Raymond and Rich, 1997; Storer and van Audenrode, 1995). In addition, these models have also been used in a wide range of other settings including combining Markov Switching with other models such as GARCH (Hamilton and Susmel, 1994); error correction (Psaradakis, Sola and Spagnolo, 2004); and causality, (Ravn, Psaradakis and Sola, 2005). In addition to VaR and asset allocation, finance Markov Switching has been used in option pricing (Boyle and Draviam, 2007), bond pricing (Elliott and Siu, 2009), foreign exchange modelling (Dueker and Neeley, 2007) and elsewhere. Ang and Timmermann (2011) provide a recent review.

In contrast to the traditional model it is no longer assumed that μ_t and Σ_t are constant across time. Instead, the Markovian assumption is that, at any time t , the world lies in one of S states. We use the dummy variable $\delta_{st} \in \{0, 1\}$ for $s \in [1, S]$ to denote the state

that occurs at time t . The probability we assign at time t to the world being in state s at time $t + 1$, $\text{Prob}_t(\delta_{st+1} = 1)$, depends only on the state at time t . A simple, fixed, transition probability matrix, M , can then be used to fully describe the stochastic way in which the prevailing state changes over time. The eigenvectors of M give the ergodic probabilities, π , associated with each state, which can be interpreted as the proportion of time that the economy spends in each of the states over very long time-periods.

Within each state, asset returns are modelled as multivariately independently and identically normally distributed with a vector of expected returns μ_s and variance covariance matrix Σ_s . Again, modelling alternatives are available. Elliott and Miao (2009), for example, allow the error terms to be Student-t distributed within a Markov Switching evaluation of Value-at-Risk problems, thus introducing leptokurtosis within each state as well as by switching between states. Okimoto (2008) and Chollette et al. (2009) have recently presented combined regime switching with copula models. These more sophisticated models, though, increase the estimation difficulty and therefore allow for less rich state processes and reduce the number of asset classes that might reasonably be included within the estimation process. In unreported results we have also run estimations with a first-order Vector Autoregressive specification of the Gaussian Markov process for some of our models. This has no substantive impact on our results and comes at considerable computational cost.

Even with the Gaussian specification of a Markov switching process, there are a large number of parameters to be estimated. For each of the S states it is necessary to estimate N values of μ_s , $N(N - 1)/2$ values of Σ_s , plus the $S(S - 1)$ values of M , giving a total of $S(0.5N(N + 1) + S - 1)$. We use a four-state ($S = 4$) process throughout, which is consistent with the choice of Guidolin and Timmermann (2006a) for jointly capturing US stock and bond dynamics. In addition, both the Akaike and Bayesian Information Criteria show a preference for a four-state specification of the data employed here over either two or three states. This is sufficiently sophisticated to capture many of the broad statistical properties of the historical data, but is sufficiently limited so that, at 72, the number of parameters for estimation remains parsimonious.

In order to estimate this multivariate regime switching environment, we invoke the MSVARlib package in GAUSS (<http://bellone.ensae.net/download.html>) written by Benoit Bellone. This code uses maximum likelihood methods to estimate regimes in a vector autoregressive framework. We estimate two Markov environments, both under the assumption that asset returns are AR(0) processes. The first incorporates the five asset classes alone and this characterisation of the data is used for Models 4 and 5 below. The second incorporates the variable x_t , which is also AR(0), into the estimation process, giving six variables in all. This parameterisation is used for Model 6, our most sophisticated as it captures the correlation

between the discount rate process and asset returns within a Markov environment.

4.1 Markov estimates: 5 asset classes

Based on the five asset classes, the empirical estimates of the transition probability matrix, M , is:

$$M = \begin{bmatrix} 0.9791 & 0.0010 & 0.0353 & 0.0444 \\ 0.0010 & 0.9817 & 0.0010 & 0.0967 \\ 0.0120 & 0.0010 & 0.9641 & 0.0010 \\ 0.0080 & 0.0163 & 0.0000 & 0.8578 \end{bmatrix}$$

where element $M_{ij} = \text{Prob}(\delta_{it+1} = 1 | \delta_{jt} = 1)$. The transition matrix allow us to estimate the expected period of time in any one state before transition into an alternate state. This is given by $\sum_{t=1}^{\infty} t M_{ii}^{t-1} (1 - M_{ii}) = (1 - M_{ii})^{-1}$. This has values of approximately 7 months for state 4 and over 2 years for all the other states.

The table below presents summary statistics for the returns to our illustrative portfolio over each of the four states, with the ‘‘Avg Correl’’ column presenting the average correlation between the five asset returns in each state. Of course in a Markov model where the investment opportunity set is time-varying, the theoretically optimal portfolio is dependent on the underlying state. However, we continue to keep the weights in our illustrative portfolio fixed through time and identical to those described above. This makes the impact of different modelling choices on our results more transparent. Again, though, it would be straightforward to extend this analysis to state dependent weights.⁵

	Mean	Standard Deviation	π	Avg Correl
State 1	0.66%	2.83%	40.8%	37.3%
State 2	1.31%	2.43%	37.8%	29.8%
State 3	0.90%	1.40%	14.8%	34.3%
State 4	-1.16%	6.43%	6.6%	54.0%

The four states have clear economic representation. State 1 has average returns slightly below the unconditional expectation (0.66% vs. 0.83%) but is otherwise unexceptional. State 2 is the bull state, with annualised average returns of almost 16%. The average correlation between asset classes is relatively low in this state, giving portfolio volatility that

⁵This is the main justification for estimating the Markov switching process at the individual asset class level rather than the portfolio level. Had we estimated switching for portfolio returns alone, then it would not be possible to dynamically change the weightings within future simulations should a fund manager wish to do so.

is somewhat below the unconditional estimate. In total, the economy spends almost 80% of its time in one of these two states. State 3 has close to average expected returns and correlations, but very low volatility. State 4, the most infrequent, is the most interesting for our purposes. It is a clear “crash” state with very high volatility, very high correlation between different asset classes and an annualised expected return of approximately -14%.

To check the accuracy with which this four-state representation of the data captures the unconditional dynamics of portfolio returns, we simulate returns over 50,000 months. At time $t = 0$, the initial state is chosen according to the ergodic probabilities. For each future time period, the state is calculated at random using the transition probability matrix, M . Within each state, the mean and standard deviation of portfolio returns are calculated by $\mu_{ps} = w'\mu_s$ and $\sigma_{ps}^2 = w'\Sigma_s w$ respectively, which are reported in the table above. A single-period return is then constructed using a random number generator based on these parameter values and the first four moments of the returns across the 50,000 months are then calculated. This process is repeated 500 times, and the mean values of the first four moments across those 500 iterations are presented in the table below, together with lower 2.5% and upper 97.5% estimates from these simulations. These are compared against the estimated first four moments from the data with associated 95% confidence intervals⁶:

	Model	Data
Mean	0.824% [0.798%, 0.851%]	0.830% [0.757%, 0.903%]
Standard deviation	2.98% [2.95%, 3.00%]	2.96% [2.79%, 3.16%]
Skewness	-0.562 [-0.648, -0.475]	-0.680 [-0.896, -0.464]
Excess Kurtosis	3.683 [3.292%, 4.080%]	4.646 [4.213, 5.079]

As is well-known, conditional heteroskedasticity leads to unconditional fat-tailed distributions when the process has a constant mean.⁷ This is clearly observed here, where the excess kurtosis is statistically significantly above zero. This characterisation of the data still

⁶Approximate sample standard deviations for the mean, skewness and kurtosis statistics were calculated as σ/\sqrt{n} , $\sqrt{6/n}$ and $\sqrt{24/n}$ respectively, where $n = 491$ is the number of observations and σ the estimated volatility of the data. Approximate 95% confidence intervals for each statistic are then calculated as the sample mean ± 1.96 standard deviations. The standard deviation confidence interval is calculated from the Chi-squared statistic with 490 degrees of freedom.

⁷See, for example, <http://www.nematrion.com/MixturesOfNormalDistributions.aspx>.

falls some way short of capturing all the leptokurtosis in the data, primarily because within each state returns continue to be generated by Gaussian distributions.

A priori, we would expect that the presence of fat-tails from this Markov Switching environment would lead to wider 95% confidence intervals for the future asset value of the portfolio, V_T , than the one-state model. This would also be consistent with the findings of Kawata & Kijima (2007). To test this, we simulate forward the value of the portfolio, V_T . We do this 10,000 times for 360 months. We can then compare these value for $T = 5, 10$ and 30 years with those from the 1-state model.

	4-state			1-state		
	$T = 5$	$T = 10$	$T = 30$	$T = 5$	$T = 10$	$T = 30$
V_T						
Mean	1.71	2.93	26.98	1.69	2.84	23.16
Lower 2.5%	0.87	1.13	4.21	1.04	1.44	6.53
Upper 97.5%	2.80	6.05	98.28	2.58	5.13	60.17

As can be seen, there are significant differences. At a maturity of five years, the traditional model predicts with 97.5% confidence that the value of the portfolio will be at least 4% more than the initial value. By contrast, the AR(0) Markov Switching model gives a 2.5% chance of a fall of 13%. The differences between the models become ever greater with increased time horizons.

4.2 Markov estimates: 5 asset classes and x_t

In Model 6 presented below we include x_t as a sixth variable to be estimated in the Markov switching process. x_t is constructed by first estimating the parameter b over the whole sample and then treating this as the “true” and time-invariant value. Within an Ornstein-Uhlenbeck interpretation of the interest rate process, this is equivalent to having the speed of mean reversion, θ , constant across states but allowing both the long-run interest rate value \bar{r} and the volatility of the discount rate process, σ_e^2 , to be state dependent. Since asset returns and the variable x_t are all assumed to be AR(0) processes, there are now 96 variables that need estimating.

In this case, the transition probability matrix changes to

$$M = \begin{bmatrix} 0.5049 & 0.0086 & 0.0752 & 0.0010 \\ 0.0593 & 0.9906 & 0.0010 & 0.0010 \\ 0.0419 & 0.0010 & 0.9118 & 0.1754 \\ 0.3939 & 0.0000 & 0.0121 & 0.8226 \end{bmatrix}$$

The summary statistics for the four states are:

	Mean	Standard deviation	π	$\text{Corr}(r_p, x_t)$	Avg Correl
State 1	-1.90%	5.67%	6.0%	-23.1%	45.4%
State 2	0.53%	2.48%	43.6%	-11.8%	34.2%
State 3	1.11%	2.39%	34.6%	-52.3%	41.1%
State 4	1.77%	3.10%	15.7%	-53.9%	30.7%

There are clear similarities between States 2, 3 and 1 in this case with States 1,2 and 4 respectively in the previous parameterisation. What is particularly noticeable is that, even in the “crash” state (now State 1), asset returns and the discount rate process remain negatively correlated, although at a lower absolute level than the unconditional average. This gives fund managers a natural hedge during market falls. While this is consistent with the post-crash period since March 2009, it does not reflect the simultaneous falls in discount rates and asset values that occurred between June 2007 and March 2009 in the UK. State 4 is noticeably different to state 3 in the previous parameterisation, although their ergodic probabilities are similar. Now this state gives very high expected returns that are also highly negatively correlated with changes in the discount rate. Although this state has lower than average correlations between different asset class returns, the overall standard deviation of portfolio returns in this state is higher than the unconditional average.

To understand how well this new representation of the data captures the unconditional summary statistics of the data, we run simulations of the portfolio returns process in the same way as was reported in the previous subsection:

	Model, r_p	Data, r_p	Model, x_t	Data, x_t
Mean	0.780% [0.754%, 0.803%]	0.830%	0.00120% [0.00098%, 0.00140%]	0.00165%
Standard deviation	2.961% [2.937%, 2.986%]	2.963%	0.02613% [0.02589%, 0.02635%]	0.02858%
Skewness	-0.467 [-0.531, -0.408]	-0.680	—	—
Excess Kurtosis	2.366 [2.091, 2.673]	4.646	—	—
$\text{Corr}(r_p, x_t)$	-37.7% [-38.7%, -36.7%]	-40.7%	—	—

While there is a clear improvement in capturing the excess kurtosis compared to the one-state model, this parameterisation performs less well than the five-asset only model against this criterion. There remains a danger, therefore, that this calibration still somewhat underestimates the true funding risk.

4.3 Results

In figure 3, we present results for $\text{Prob}(z_t < 0)$ for all models. This graph is identical to Figure 1 for Models 1–3, which are shown here for comparison purposes. For Models 4–6, we are no longer able to construct closed-form solutions to the problem and therefore resort to Monte Carlo methods instead. Over 100,000 simulations at time $t = 0$ the initial state is chosen at random according to the ergodic probabilities.⁸ The state at each future time, until $T = 360$ months, is then calculated at random according to the transition probability matrix, M . At each time, portfolio returns, r_{pt} , are drawn at random from a normal distribution with mean μ_{ps} and variance σ_{ps}^2 . This allows for the calculation of V_t within each simulation at all intervals over a horizon of 30 years.

The difference between Models 4, 5 and 6 is in the treatment of the discount rate process. For comparison with Model 1, in Model 4 it is assumed that $r_{ft} = r_{f0}$ for all t . For comparison with Model 2, in Model 5 values of x_t are drawn at random and independently from the portfolio return from a normal distribution with mean a and variance σ_e^2 . Therefore, although Model 5 includes Markov switching for the asset returns process, discount rates are driven as if there were only one state in the economy. From x_t , values of the discount rate are constructed iteratively by $r_{ft} = (1 + b)r_{ft-1} + x_t$. In Model 6, the variable x_t is estimated as part of the Markov switching process and therefore its mean and variance are both state dependent, as is its correlation with the portfolio returns process. We use the standard Cholesky decomposition approach to construct correlated random variables r_{pt}, x_{pt} at each time t within each simulation to capture the properties of the underlying state at that time.

The comparison of results of Models 4–6 are similar to those of Models 1–3. Assuming that the discount rate is constant leads to a substantial underestimation of the solvency risk, while assuming that the interest rate process is independent of the asset returns process leads

⁸In unreported results, we also run simulations for Models 1–3 to check the sensitivity of our results to the method of their derivation. The simulation results are, in all cases, highly similar to those reported in Figure 1. An issue also arises as to whether to start the simulations according to the ergodic probabilities or the estimated smoothed probabilities in December 2010 that are an output from the Markov estimation process. The former is more suitable for our purposes as we are comparing modelling choices, and are therefore reporting unconditional estimates. For a fund manager assessing his or her own fund risk in December 2010, it would be more appropriate to use the smoothed probability estimates for that date.




Figure 3: This graph presents the probability that the scheme will be underfunded at time t ; $\text{Prob}(z_t < 0)$. For Models 1–3, the results are as reported in Figure 1. For Models 4–6, it is assumed that asset returns are generated by a 4-state Markov regime switching model. In Model 4, the discount rate is fixed at its current level. In Model 5, the discount rate follows a one-state Ornstein-Uhlenbeck process that is independent of the portfolio return. Model 6 allows for the correlation between the discount rate and the portfolio return process within a Markov switching environment.

to an overestimate of the risk because asset returns and the discount rate are negatively correlated.

By incorporating Markov switching, three new features emerge. First, the impact of introducing leptokurtosis into the asset returns process noticeably increases the assessed funding risk in all cases, but the impact is less dramatic than allowing for stochasticity in the discount rate. This is consistent with the summary statistics presented in the introduction to this paper as most of the volatility comes from changes in the calculated present value of liabilities rather than asset values. Second, in Models 4–6, the maximal probability of fund risk, at months 43, 29 and 33 respectively, occurs approximately one year later than in Models 1–3 (months 28, 21 and 21 respectively). Finally, the funding risk decays much more slowly under Models 4–6 than under Models 1–3. The $\text{Prob}(z_t < 0)$ at 30 years is under 1% of the maximal risk for Models 1 and 3 and approximately 4% for Model 2. By contrast, it is almost 9% for Models 4 and 6 and over 11% for Model 5. At longer horizons, the estimated risks from Model 6 are greater than those from Model 2.

In addition to revealing the probability that $z_t < 0$, the simulations also present broader statistical information about this variable. The following table presents summary statistics for z_t at horizons of 1 and 5 years for the four models with stochastic interest rates:

z_t	Mean	Median	Lower 2.5%	Upper 2.5%
	One year horizon			
Model 2	0.259	0.237	-0.188	0.834
Model 3	0.250	0.238	-0.127	0.699
Model 5	0.252	0.228	-0.197	0.843
Model 6	0.231	0.217	-0.141	0.680
	Five year horizon			
Model 2	0.762	0.636	-0.343	2.638
Model 3	0.723	0.645	-0.227	2.094
Model 5	0.741	0.595	-0.386	2.771
Model 6	0.610	0.500	-0.285	2.119

Again, the effects of more accurately modelling the underlying dynamics are dramatic. Despite starting with an initial funding surplus of $z_0 = 15\%$, the most sophisticated model predicts that the lower 2.5% of funds will be more than 14% in deficit after 12 months and double that after five years. This again emphasises not only the probability of falling into a funding shortfall but also, in a Value-at-Risk sense, how severe such a fall might be.

5 Conclusion

This paper has examined the impact of extreme market movements on pension fund solvency. This issue was forcefully brought home by the credit crunch in 2008 which resulted in significant falls in asset values, and the subsequent policy response of QE, which dramatically reduced gilt yields. For pension funds this presented the perfect storm. Asset values were depressed, while the discount rate for estimating the present value of the liabilities were pushed lower, resulting in inflated liabilities. Jointly, the effect was to leave huge deficits on the balance sheet of defined benefit pension funds. As a result, to better estimate future solvency scenarios it is critical that both asset returns and discount rates are modelled jointly.

To undertake this analysis we use Markov Regime Switching Models. The modelling choice seems appropriate given its widespread use in both academic and practitioner work. Moreover, this type of modelling addresses some of the key issues raised by the Actuarial Profession BSM Working Party in capturing extreme market movements. Markov Regime Switching Models can capture, rare but extreme market events, time-varying asset return volatility and the fat-tailed nature of stock returns.

Our results show that future projections of fund solvency are extremely sensitive to modelling choices. In particular, our analysis shows the importance of estimating a stochastic interest process that is allowed to vary with asset returns. If interest rates are allowed to remain constant through time, then both standard one-state model and the multi-state Markov Regime Switching model significantly underestimate the likelihood of future pension plan underfunding. Where interest rates are estimated as a stochastic process that varies with asset returns, then our results suggest that both the traditional one-state model and the multi-state model predict a much higher proportion of underfunded pension plans in the future. Although the difference between the most sophisticated multi-state model and the traditional one-state model graphically does not appear to be large, the multi-state model predicts between approximately 2.5% and 4.5% more schemes underfunded and the increased number of underfunded schemes is persistent through time.

These results have a number of pension management and policy considerations. From a micro pension management perspective, the advice that is given to trustees and sponsors must be carefully explained. The presentation of one result, or one type of result (i.e. a fixed discount rate), does not present an accurate picture for decision making. In particular a much broader and clearer discussion about potential outcomes would allow for more effective decision making around issues such as short-term and long term funding plans, potential risk management strategies and asset allocation decisions.

Understanding the potential impact of these different outcomes in both the short and long-run has huge implications for macro-prudential pension regulation. For example, if our models were re-calibrated to incorporate a sustained period of low asset returns, then the percentage of funds that are likely to have deficits would increase at the further out projection times. Consequently, a much richer data set of future pension outcomes could be estimated and better decisions could be made in terms of pension funding at a macro level. Moreover, from the perspective of The Pensions Regulator scheme specific sensitivities to potential future outcomes could be considered, and so the identification of ‘at risk’ schemes may become better, while for the PPF a richer set of future scenarios could be projected that may help with the identification of potential funding pressures.

6 References

Ang, A. and G. Bekaert, International asset allocation with regime shifts, *Review of Financial Studies*, 2002, 15, 1137–1187.

Ang, A. and A. Timmermann, Regime changes and financial markets, Netspar discussion paper #DP 06/2011-068, 2011.

Billio, M. and L. Pelizzon, Value-at-Risk: a multivariate switching regime approach. *Journal of Empirical Finance*, 2000, 7, 531–554.

Boyle, P.P., and T. Draviam, Pricing exotic options under regime switching. *Insurance: Mathematics and Economics* 2007, 40, 267–282.

Chen, P., and H. Yang, Pension funding problem with regime-switching geometric Brownian motion assets and liabilities, *Applied Stochastic Models in Business and Industry*, 2010, 26, 125–141.

Chollette, L., A. Heinen, and A. Valdesogo, Modeling international financial returns with a multivariate regime-switching copula, *Journal of Financial Econometrics*, 2009, 7, 437–480.

Dueker, M. and C.J. Neeley, Can Markov switching models predict excess foreign exchange returns?, *Journal of Banking and Finance*, 2007, 31, 279–296.

Elliott, R.J., and H. Miao, VaR and expected shortfall: a non-normal regime switching framework. *Quantitative Finance*, 2009, 9, 747–755.

Elliott, R.J., and T.K. Siu, On Markov-modulated exponential-affine bond price formulae. *Applied Mathematical Finance*, 2009, 16, 1–15.

Ferstl, R., and A. Weissensteiner, Asset-liability management under time-varying investment opportunities, *Journal of Banking and Finance*, 2011, 35, 182–192.

Frankland, R., A.D. Smith, T. Wilkins, E. Varnell, A. Holtham, E. Biffis, S. Eshun, and D. Dullaway, Modelling extreme market events: A report of the benchmarking stochastic

models working party, Institute of Actuaries and Faculty of Actuaries, 2008.

Fraudendorfer, K., U. Jackoby, and A. Schwendener, Regime switching based portfolio selection for pension funds, *Journal of Banking and Finance*, 2007, 31, 2265–2280.

Freeman, M.C., The time-varying equity premium and the S&P500 in the twentieth century, *Journal of Financial Research*, 2011, 34, 179–215.

Guidolin, M., and A. Timmerman, An econometric model of nonlinear dynamics in the joint distribution of stock and bond returns, *Journal of Applied Econometrics*, 2006a, 21, 1–22.

Guidolin, M., and A. Timmerman, Term structure of risk under alternative econometric specifications. *Journal of Econometrics*, 2006b, 131, 285–308.

Guidolin, M., and A. Timmerman, Asset allocation under multivariate regime switching, *Journal of Economic Dynamics and Control*, 2007, 31, 3503–3544.

Guidolin, M., and A. Timmerman, International asset allocation under regime switching, skew, and kurtosis preferences, *Review of Financial Studies*, 2008, 21, 889–935.

Hamilton, J.D., A new approach to the economic analysis of non-stationary time-series and the business cycle. *Econometrica*, 1989, 57, 357–3840.

Hamilton, J.D., and R. Susmel, Autoregressive conditional heteroskedasticity and changes in regime. *Journal of Econometrics*, 1994, 64, 307–333.

Josa-Fombellida, R. and J.P. Roncón-Zapatero, Optimal asset allocation for aggregated defined benefit pension funds with stochastic interest rates, *European Journal of Operational Research*, 2011, 201, 211–221.

Kawata, R., and M. Kijima, Value-at-risk in a market subject to regime switching, *Quantitative Finance*, 2007, 7, 609–619.

Kemp, M., *Extreme events: Robust portfolio construction in the presence of fat tails*, John Wiley, 2011.

Lam, P. S., The Hamilton model with a general autoregressive component. *Journal of Monetary Economics*, 1990, 26, 409–432.

Okimoto, T., New evidence of asymmetric dependence structure in international equity markets. *Journal of Financial and Quantitative Analysis*, 2008, 43, 787–815.

Psaradakis, Z., M. Sola, and F. Spagnolo, On Markov error-correction models, with an application to stock prices and dividends, *Journal of Applied Econometrics*, 2004, 19, 69–88.

Psaradakis, Z., M. O. Ravn, and M. Sola, Markov switching causality and the money-output relationship. *Journal of Applied Econometrics*, 2005, 20, 665–683.

Raymond, J.E., and R. W. Rich, Oil and the macroeconomy: A Markov state-switching approach. *Journal of Money, Credit and Banking*, 1997, 29, 193–213.

Shen, Y., and T.K. Siu, Asset allocation under stochastic interest rate with regime switch-

ing. *Economic Modelling*, 2012, 12, 1126–1136.

Storer, P.A., and M. Van Audenrode, Unemployment insurance take-up rates in Canada: Facts, determinants and implications. *Canadian Journal of Economics*, 1995, 28, 822–835.

Taamouti, A., Analytical Value-at-Risk and Expected Shortfall under regime-switching, *Finance Research Letters*, 2009, 6, 138–151.

Vařicek, O., An equilibrium characterisation of the term structure, *Journal of Financial Economics*, 1977, 5, 177–188.

Weitzman, M., Gamma discounting, *American Economic Review*, 2001, 91, 260–271.

Zhou, X.Y., and G. Yin, Markowitz’s mean-variance portfolio selection with regime switching: A continuous-time model. *SIAM Journal on Control and Optimization*, 2003, 42, 1466–1482.