Institute and Faculty of Actuaries

Deloitte.

Calibrating Scenario Generators for Pensions Asset Liability Modelling

Yona Chesner
Phil Hardingham

25 May 2016

Introduction

- A key tool that has been used in the pension scheme industry, particularly by investment consultants, is the Asset Liability Model (ALM)
  - Typically this type of tool is used to inform trustees and sponsors around the level of risk and likely outcomes of following a particular investment strategy, or to compare alternative strategies in a quantitative manner
  - A typical output of such a model is shown to the right

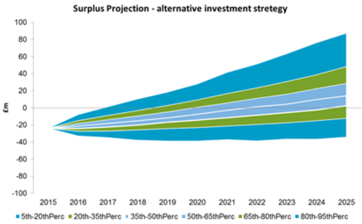

- We are aware that there are a number of very different methods currently in use for these models, however a good number of them require a set of hypothetical economic scenarios in order to base the distributions of outcomes
- As such, the output of some sort of Economic Scenario Generator (ESG) becomes a key input into such a model
  - A variety of ESGs exist, as does plenty of literature about them
  - However little of this information directly pertains to the use of such a model for this particular purpose
- The key question that we want to explore today is what approach to take to calibrate each parameter within an ESG for the purpose of an ALM described above
  - Assuming historic data is used
- Clearly parameters necessary will depend on the underlying models used, however a number of themes are likely to reoccur irrespective of the precise nature of the ESG
  - We will explore some of these by building up a basic ESG model and considering each parameter as it is introduced.

Institute and Faculty of Actuaries

25 May 2016 2

Interest rates – 1 factor Vašíček

- The first thing that we need for our model is some sort of interest rate model
- Being able to model bonds is probably the most important aspect of a model for our purpose
  - This is because changes in bond yields (of various terms) will be the driving force behind changes in the present values of pension scheme liabilities
  - Pension scheme asset portfolios tend to be bond heavy (of one type or another) as well
- Let's begin with a basic interest rate model – 1 factor Vašíček
- Some points about Vašíček
- In its most commonly quoted form, the Vašíček model is a description of the short rate over time
  - This would imply that the best method to calibrate the model to either historic or desired behaviours would be to look at the characteristics (historic or otherwise) of cash
  - Indeed in some cases this may be appropriate
- However, for our purpose it is not really the behaviour of cash that we are interested in
  - Pension schemes tend not to really be sensitive to the behaviour of the short rate
- The model can also be restated as a model of any particular bond
- So if we are most interested in behaviour of the 20 year point, for our purposes we feel that stating the model in terms of this point (and calibrating to 20 year behaviour) makes the most sense.

Institute and Faculty of Actuaries

1 factor Vašíček – continued

- The model has the following parameters that we need to consider:
  - Starting level
  - Risk premium
  - A mean reversion level
  - Speed of mean reversion
  - A volatility parameter
- All of these parameters deserve consideration as to how best to calibrate them
  - We will be concentrating on the final two of these for the next several minutes
  - Before we do, let us briefly introduce the first three
- Starting level
  - While this does deserve some thought, in practice this will likely be calibrated simply to today's market rate
  - This is the simplest starting point, is market consistent and ultimately will be the best figure to use if we want our model to be able to do some sort of cashflow discounting at time 0
- Risk premium – this is some measure of how additional volatility (i.e. longer dated bonds) are rewarded
  - A key (but not only) determinant of the slope of yield curve
  - We will consider risk premia in general in more depth once we have built the model up to include other assets.

Institute and Faculty of Actuaries

1 factor Vašíček – continued

- Reversion level
  - For Vašíček we just need a single point to revert to the long term average "ergodic" mean yield of the bond
  - For calibration the main options are historic average levels, or level implied by current long forward rate
  - Historic levels may be tempting, but do we really believe we will fully revert back to the level of yields seen in the past?
  - Forward rate is not without judgment either – what is "long end" of curve? Do we impose any risk premium into current curve?

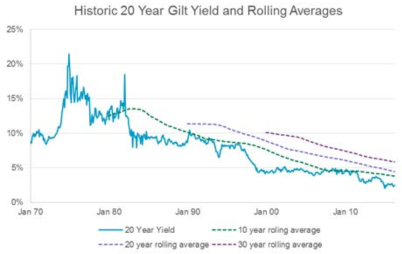

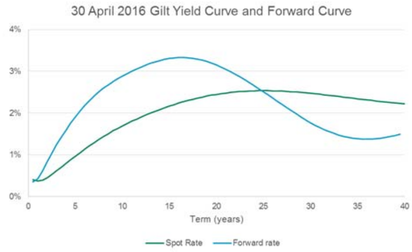

Source: BoE, Deloitte interpolation (cubic spline) used for data gaps

Institute and Faculty of Actuaries

1 factor Vašíček – continued

Speed of mean reversion

- Option 1 – we can look at the behaviour of a particular yield over time, and measure the degree to which it reverts
- However, there are a couple of statistical issues
  - Stability
  - Bias

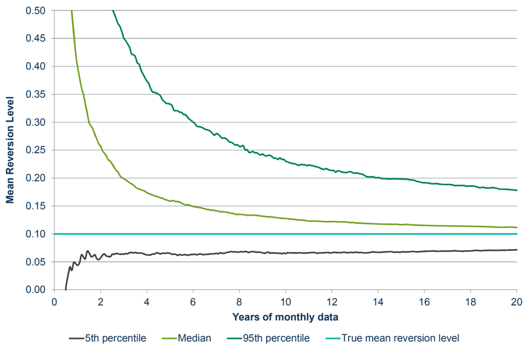

Parameter	
Simulated Term	Short rate
True mean reversion parameter	0.1
Volatility (annual)	5%
Mean reversion level	3%
Initial yield	1%
Data periodicity	Monthly
Number of sims	2000

Institute and Faculty of Actuaries

1 factor Vašíček – continued

Speed of mean reversion (continued)

- Option 2 – we can look at the relative volatilities of yields of different terms
  - If the speed of mean reversion parameter is higher, long term yields will be less volatile relative to short term yields
  - Specifically;

$$yield\ volatility \times term \propto 1 - e^{-mean\ reversion\ parameter \times term}$$

  - It is more efficient to estimate volatility than autocorrelation from limited data (as more information can be extracted by sampling more frequently)

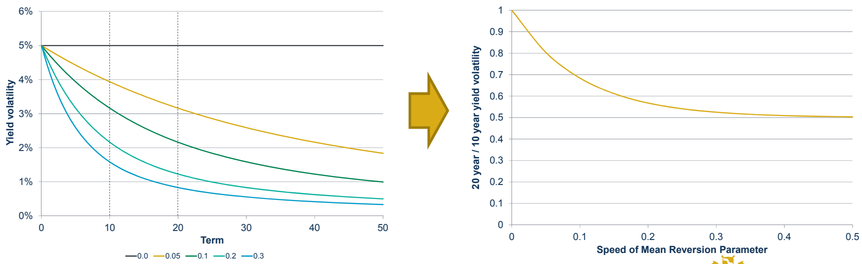

Our preference is for option 2 given the lack of bias and statistical efficiency

This is likely to calibrate weaker mean reversion, and we are more likely to generate negative rates as a result.

Institute and Faculty of Actuaries

25 May 2016 7

1 factor Vašíček – continued

Volatility

- There are two primary options for calibrating our volatility parameter

Option 1: derive from historic volatility

- We can choose a particular term of interest, calculate historic yield volatility, and calibrate our parameter so our simulations replicate this
- This is simple and easy to explain to Trustees
- Future volatility may differ significantly from past volatility

Option 2: derive from market implied volatility

- We can use market implied volatility from traded interest rate derivatives (such as swaptions)
- The implied volatility at which market participants are willing to hedge may differ from their true volatility expectations
- Implied volatility will also be impacted by many other factors (for example, transaction costs, counterparty default, capital raising costs etc)

Our preference is for option 1 for real world pension modelling.

Institute and Faculty of Actuaries

25 May 2016 8

1 factor Vašíček – continued

- Worth noting that under the preferred approach historic data on two bonds is used
  - But the model will not capture the correlation between these points
  - Only a single factor is used
  - In theory we can capture this correlation parameter by adding a second factor to the model (i.e. 2 factor Vasiček)
    - But now three bond histories are used to calibrate
    - And again not all of the correlations are captured
- Ultimately question of one vs two factors is likely to be around:
  - Materiality of curve shape changes on model
    - How does the underlying model roll forward the liabilities?
    - Client specific questions e.g. how much asset matching is involved? More matching implies residual curve risk more relevant
  - Vs. additional complexity
  - And lack of data to calibrate – e.g. limited long data on index-linked gilts
  - Insurance models often pick one factor for real vs. two for nominal
    - In our case it seems that liabilities tend to be more real in nature
    - Therefore the reverse makes more sense
    - Provided there is enough data to get a good calibration
- For our purposes today we will presume that 1 factor is being used for both nominal and real rates.

Institute and Faculty of Actuaries

Inflation

Introduction

- We have implicitly calibrated breakeven inflation by calibrating real and nominal rates, but not RPI itself
- Does RPI differ from the short breakeven inflation rate?
- We have recorded the cumulative return on realised inflation minus breakeven inflation over time

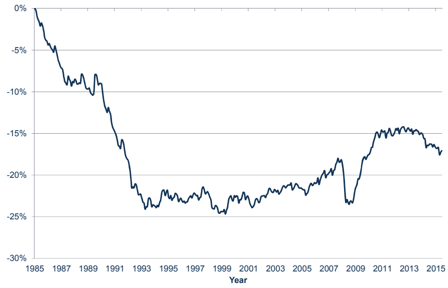

- What does this mean?
  - The return clearly isn't flat – RPI and short breakeven inflation are not identical
  - This implies another risk factor exists for RPI.

Institute and Faculty of Actuaries

Inflation - continued

Significance

- How important is this in a Pension Scheme context?
  - Breakeven inflation rates are potentially more significant than realised RPI
  - However, we may require RPI for some Asset Liability Modelling (for example, if we are using a cashflow framework)

Calibration

- We can model the difference between RPI and the short breakeven rate as a random walk, with a volatility and risk premium parameter
- However, in practice the process might not have independent increments (it may trend)
- Therefore, the sampling frequency for calibrating volatility becomes significant.

Institute and Faculty of Actuaries

25 May 2016 11

Inflation - continued

Term Structure of Volatility

- There is a tradeoff between calibrating using the shortest data available whilst avoiding the distorting impact of autocorrelation
- The term structure of volatility for actual minus breakeven inflation is shown below

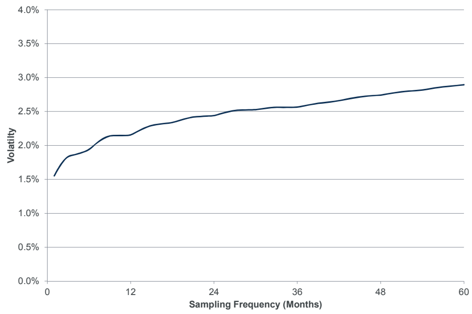

- This is a wider consideration when calibrating volatility, particularly where pricing is stale
- As in the case of gilts, we will defer discussion of the associated risk premium.

Institute and Faculty of Actuaries

25 May 2016 12

Risk asset(s) – Introduction

- To complete our basic model we need to add risk assets
  - In practice we would want several of these given the current nature of pension scheme investing
  - However for our purposes today lets just pick one – say UK equities
- We will model the risk asset class as geometric Brownian motion
  - Independent increments imply an efficient market
  - We will assume there is a static Equity Risk Premium (ERP) that is relative to the short rate
    - i.e. avoiding negative ERP in scenarios with high interest rates
- We have the following parameters to calibrate:
  - Volatility
  - Risk premium

Institute and Faculty of Actuaries

Risk asset(s) – Volatility

- As with bonds and RPI we want to know what period and sampling frequency to use
- For our dataset (FTSE All Share) short term volatility has significant fluctuations, but longer term fairly stable
  - And term structure less of an issue, so using as frequent sampling as possible (within reason) makes sense

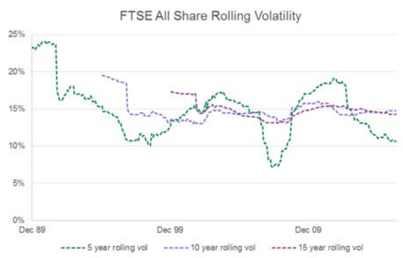

- Any other asset classes used would require similar analysis

Institute and Faculty of Actuaries

Risk Premia

Introduction

- Let's discuss risk premia in a little more detail
- We have four risk premia to calibrate
  - Nominal, real, inflation and equity risk factors
- Why are risk premia important, and what properties do we want to capture?
  - Modelling likelihood of being fully funded requires asset class growth rates
  - Portfolio optimisation exercises require arbitrage free scenarios.

Institute and Faculty of Actuaries

Risk Premia (continued)

Option 1: Use historic data for calibration

- If we simulate returns from a known distribution, how quickly do the annualised sample returns converge?

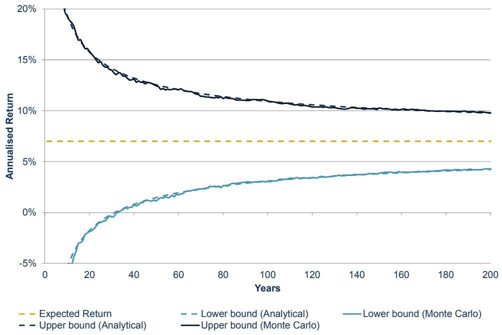

Parameter	
Expected return	7%
Volatility	20%
Significance	5%

- Convergence takes a long time
- Past performance is not an indicator of future results
- This may generate arbitrage or near arbitrage opportunities.

Institute and Faculty of Actuaries

Risk Premia (continued)

Option 2: Derive from fundamentals

- We can use economic theory or judgement to set a particular risk premium
- However, if we do this in isolation for many asset classes we are increasingly likely to introduce near arbitrage opportunities
- We can set some risk premia relative to others to avoid these opportunities
  - Black-Litterman
  - The tradeoff is losing direct control over calibrating particular risk premia

Our preference is option 2, deriving from risk premia from fundamentals, whilst limiting the maximum efficiency an optimised portfolio can achieve

Modelling in a Pension Scheme context

- We need to focus on important risk premia and let secondary risk premia vary to avoid unrealistic portfolio efficiency
- What risk premia do we care about most?
  - Equity risk premium
  - Term premium
  - We can let the inflation risk premium float

Institute and Faculty of Actuaries

Risk Premia (continued)

Equity Risk Premium

- What do we use the equity risk premium (ERP) for?
  - Model the projected asset growth rate in excess of gilts
  - Determine the discount rate spread
  - The discount rate spread will depend on covenant strength; from a modelling perspective we only want a single ERP
- We can estimate this from fundamentals
  - Estimate the expected return on equity from a fundamental model
  - Subtract the risk free rate

Term Premium – Introduction

- Within the model, the term premium determines the expected returns on short versus long bonds
  - Are investors compensated for taking interest rate exposure?
  - Or do they give up return to get it?
- Before considering the term premium, we need to be more precise about what we mean by:
  - Arithmetic risk premium
  - Geometric risk premium

Institute and Faculty of Actuaries

Risk Premia (continued)

Geometric versus arithmetic risk premium

Property	Arithmetic Risk Premium	Geometric Risk Premium
Definition 1	The arithmetic return of an asset above the risk free rate	The geometric return of an asset above the risk free rate
Definition 2	The instantaneous expected return above the risk free rate	The difference between the long run growth rate of an asset and the risk free rate
Can be blended…	Cross-sectionally	Longitudinally

- If an asset has variable returns, the geometric return will be lower than the arithmetic return
- Specifically;

$$G \approx A - 0.5\sigma^2$$

- In the case of risk assets, there is a one to one pairing between risk factors and assets
  - We can target either the arithmetic average or the geometric risk premium (if we recognise the difference between the two)
- However, in the case of rates, we have only one (or two) factors which determine the return of an infinite number of assets (bonds of different terms)
  - Therefore, the risk premium for each bond will be related to all other bonds to avoid arbitrage
  - What does this term structure look like for both arithmetic and geometric risk premia?

Institute and Faculty of Actuaries

Risk Premia (continued)

- In the Vašíček model, the risk premium parameter is arithmetic
- Therefore, the arithmetic risk premium of a bond will be proportional to its volatility

$$A \propto Bond\ Volatility$$

- This implies:

$$G = Constant \times Bond\ Volatility - 0.5 \times (Bond\ Volatility)^2$$

- The expected long run growth rates of different term bonds (rolled continuously) will differ
  - This is the case even if the arithmetic risk premium is set equal to zero
- If geometric risk premia were constant across term, this would imply arbitrage
  - Portfolios of many bonds of different terms would have a higher geometric risk premia than each standalone bond
  - If geometric risk premia are constant, arithmetic risk premia would be non linear across term
  - A long position of long term bonds with a short position in short term bonds could eliminate risk whilst generating excess returns.

Institute and Faculty of Actuaries

Risk Premia (continued)

Positive Arithmetic Risk Premium

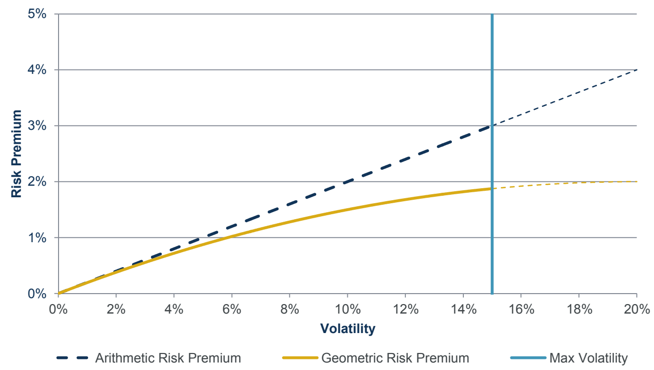

Institute and Faculty of Actuaries

Risk Premia (continued)

Negative Arithmetic Risk Premium

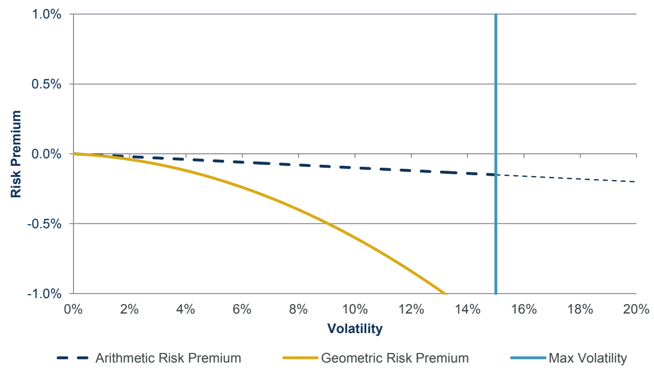

Institute and Faculty of Actuaries

Risk Premia (continued)

Term premium – Calibration

- Setting the (arithmetic) risk premium will be judgemental
  - We can set to zero, implying the expected wealth at any given point in time from investing in gilts will be equal regardless of the specific mixture of gilts held
  - We can choose an expected growth rate of bonds of two terms over the long run, and set the arithmetic risk premium so this is reflected in our model (which will in also determine the growth rate of all other standalone bonds)
  - We can choose a non – zero arithmetic risk premium to broadly target the slope of the initial yield curve
    - Taking into account the impact of the mean reversion level on the initial slope of the curve

Our preference is to set the arithmetic risk premium to zero (or a low number) as a default

In general, at the 20 year point in the curve market participants (predominantly Pension Schemes) are buying and selling to hedge rather than to generate excess growth. With no strong rationale to set this premium significantly above or below zero, we think it is safer to take no credit for a term premium and to discuss the potential return implications of leveraged LDI (for example) outside the model.

Institute and Faculty of Actuaries

Correlations

- For the sake of completeness it is worth briefly considering correlations
- We will require an nxn matrix where n is the number of factors in the model
  - For our current model we have n = 4
- Once a methodology has been chosen for all other parameters we can calibrate this historically
- As with volatility some consideration around sampling frequency is required
- For Vašíček, as noted earlier we have more correlation parameters than are actually required
  - In practice this means discarding one bond term when calibrating correlations
  - Term to discard would be less "useful" term, e.g. keep 20 year correlation but ignore longer term bond
- Finally due to sampling error comparison of model output correlations with inputs helps identify stability issues

Institute and Faculty of Actuaries

Wrap up

- We believe that simply utilising an "off the shelf" ESG as an input to a pension ALM is likely to come with certain drawbacks
  - Where the ESG has not been calibrated specifically for this purpose there are likely to be many judgement "gaps"
  - Even where calibration has taken these issues into account it is imperative that the model end user understands the process that has been used
- Ultimately there are always trade-offs between various possible calibration methods, and this should be acknowledged in the way that any client advice is presented
- Most imperative is that neither consultants nor clients slip into the belief that the model output is a fully fair representation of the "real world", but rather is a tool to aid comprehension of the possible ramifications of a given approach to investment strategy
  - If we do not understand the simplifications and approximations within the model calibration process, we may not understand which risks are vulnerable to being understated or assumed away
- Not considered in today's presentation are many other issues including:
  - The pros and cons of a wide range of alternative models
  - Dealing with illiquidity risk premiums
  - Output validation.

Institute and Faculty of Actuaries

The views expressed in this presentation are those of invited contributors and not necessarily those of the IFoA. The IFoA do not endorse any of the views stated, nor any claims or representations made in this presentation and accept no responsibility or liability to any person for loss or damage suffered as a consequence of their placing reliance upon any view, claim or representation made in this presentation.

The information and expressions of opinion contained in this publication are not intended to be a comprehensive study, nor to provide actuarial advice or advice of any nature and should not be treated as a substitute for specific advice concerning individual situations. On no account may any part of this presentation be reproduced without the written permission of authors.

Institute and Faculty of Actuaries