University of Pennsylvania
ScholarlyCommons

Wharton Pension Research Council Working
Papers

Wharton Pension Research Council

9-1-2013

Implications for Long-term Investors of the Shifting Distribution of Capital Market Returns

James Moore
PIMCO, jim.moore@pimco.com

Niels Pedersen
PIMCO, niels.pedersen@pimco.com

Follow this and additional works at: https://repository.upenn.edu/prc_papers

Part of the [Economics Commons](#)

Moore, James and Pedersen, Niels, "Implications for Long-term Investors of the Shifting Distribution of Capital Market Returns" (2013). *Wharton Pension Research Council Working Papers*. 130.
https://repository.upenn.edu/prc_papers/130

The published version of this Working Paper may be found in the 2014 publication: *Recreating Sustainable Retirement: Resilience, Solvency, and Tail Risk*.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/prc_papers/130
For more information, please contact repository@pobox.upenn.edu.

Implications for Long-term Investors of the Shifting Distribution of Capital Market Returns

Abstract

The paper motivates and describes a regime-switching macro-driven simulation model for the purposes of simulating long horizon asset returns. The paths generated by this model are compared to more common approaches – multivariate normal generators and a block bootstrap simulation. Despite calibration to the same mean and variance in returns, the models display divergent behavior in the tails of long horizon return simulations. Simulations are run through representative defined contribution and defined benefit applications to examine the filtered behavior and draw inferences for future applied research and application.

Disciplines

Economics

Comments

The published version of this Working Paper may be found in the 2014 publication: *Recreating Sustainable Retirement: Resilience, Solvency, and Tail Risk*.

Recreating Sustainable Retirement

Resilience, Solvency, and Tail Risk

EDITED BY

Olivia S. Mitchell,
Raimond Maurer, and
P. Brett Hammond

OXFORD
UNIVERSITY PRESS

OXFORD

UNIVERSITY PRESS

Great Clarendon Street, Oxford, OX2 6DP,
United Kingdom

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide. Oxford is a registered trade mark of
Oxford University Press in the UK and in certain other countries

© Pension Research Council, The Wharton School, University of Pennsylvania 2014

The moral rights of the authors have been asserted

First Edition published in 2014

Impression: 1

All rights reserved. No part of this publication may be reproduced, stored in
a retrieval system, or transmitted, in any form or by any means, without the
prior permission in writing of Oxford University Press, or as expressly permitted
by law, by licence or under terms agreed with the appropriate reprographics
rights organization. Enquiries concerning reproduction outside the scope of the
above should be sent to the Rights Department, Oxford University Press, at the
address above

You must not circulate this work in any other form
and you must impose this same condition on any acquirer

Published in the United States of America by Oxford University Press
198 Madison Avenue, New York, NY 10016, United States of America

British Library Cataloguing in Publication Data

Data available

Library of Congress Control Number: 2014940448

ISBN 978-0-19-871924-3

Printed and bound by
CPI Group (UK) Ltd, Croydon, CR0 4YY

Links to third party websites are provided by Oxford in good faith and
for information only. Oxford disclaims any responsibility for the materials
contained in any third party website referenced in this work.

Contents

<i>List of Figures</i>	ix
<i>List of Tables</i>	xiii
<i>Notes on Contributors</i>	xv

1. Recreating Retirement Sustainability <i>Olivia S. Mitchell and Raimond Maurer</i>	1
---	---

Part I. Capital Market and Model Risk

2. Managing Capital Market Risk for Retirement <i>Enrico Biffis and Robert Kosowski</i>	9
3. Implications for Long-term Investors of the Shifting Distribution of Capital Market Returns <i>James Moore and Niels Pedersen</i>	30
4. Stress Testing Monte Carlo Assumptions <i>Marlena I. Lee</i>	60

Part II. Longevity Risk

5. Modeling and Management of Longevity Risk <i>Andrew Cairns</i>	71
6. Longevity Risk Management, Corporate Finance, and Sustainable Pensions <i>Guy Coughlan</i>	89
7. Model Risk, Mortality Heterogeneity, and Implications for Solvency and Tail Risk <i>Michael Sherris and Qiming Zhou</i>	113
8. The Securitization of Longevity Risk and Its Implications for Retirement Security <i>Richard MacMinn, Patrick Brockett, Jennifer Wang, Yijia Lin, and Ruilin Tian</i>	134

viii Contents

Part III. Regulatory and Political Risk

- 9. Evolving Roles for Pension Regulations: Toward Better Risk Control? 163
E. Philip Davis
- 10. Developments in European Pension Regulation: Risks and Challenges 186
Stefan Lundbergh, Ruben Laros, and Laura Rebel
- 11. Extreme Risks and the Retirement Anomaly 215
Tim Hodgson

Part IV. Implications for Plan Sponsors

- 12. Risk Budgeting and Longevity Insurance: Strategies for Sustainable Defined Benefit Pension Funds 247
Amy Kessler
- 13. The Funding Debate: Optimizing Pension Risk within a Corporate Risk Budget 273
Geoff Bauer, Gordon Fletcher, Julien Halfon, and Stacy Scapino
- The Pension Research Council* 293
- Index* 297

Chapter 3

Implications for Long-term Investors of the Shifting Distribution of Capital Market Returns

James Moore and Niels Pedersen

It's tough to make predictions, especially about the future.

(attributed to Yogi Berra)

Despite the caution from Hall of Fame baseball player *cum* philosopher Mr. Berra, few people have stopped trying to predict the future. And nowhere is prediction more common than in the financial markets. One cannot watch financial television for more than an hour without a guest being asked where he believes the level of the S&P 500, an individual stock, the price of gold, or the ten-year Treasury yield will be at year-end. Given the vagaries of the market, short-term forecasting, particularly for high volatility assets, is little more than a toss of the dice. The media soon forget the many whose predictions are off the mark. The few whose prognostications end up close to the mark are ascribed sage-like properties and develop cults of followers. While short-run forecasting has value for financial entertainment and speculation, in the United States and much of the developed world, there is perhaps no field where long-run forecasting has wider implications for personal welfare than that of forecasting asset returns.

The expected returns of stocks, bonds, and other investments play a critical role in retirement planning, budgeting, and determining future savings adequacy. Of course, practitioners in the space—plan sponsors, investment advisors, consultants, asset managers, and others versed in statistics—know that future returns are random variables. Actual returns through time are drawn from a distribution of possibilities. Given this fact, outcomes that are functions of the realized returns are themselves distributions of random variables. The specific form of the ultimate distribution in question relies on (a) the stochastic processes governing the returns themselves, (b) the functional form that these returns are ‘filtered’ through, and possibly (c) convolutions of multiple functional forms. Even when the underlying stochastic processes are known with certainty and the generating distributions are well behaved, the transforming nature of the real-world functions overlaid can lead to significant alterations of the resultant distributions. In some cases this may compress distributions; in others it may lead to tails that are exaggerated.

Yet it is important to remember that *we do not know the true statistical generating process of asset returns*. The vast majority of work done by practitioners relies on the lessons learned in undergraduate statistics courses. The insights absorbed there rely heavily on the Law of Large Numbers and asymptotic convergence to normality. These in turn rely on the stronger assumptions of stationarity and ergodicity. What if these do not hold?

To explore this question, this chapter looks at the implications of long horizon asset returns that flow from three different generating processes for stock and bond returns.¹ The first, a multivariate normal distribution, is widely used due to its familiarity and analytic tractability and has been used for Monte Carlo statistical analyses since the Second World War. The second, a block bootstrap approach, has become more common in the past few decades with increases in computing power and questions as to the appropriateness of normality given limited historical data. The third approach uses a nested structural model that links asset returns to macroeconomic fluctuations in the real economy. The core of this model relies on a non-stationary, Markov-switching evolution of real GDP as first modeled by Hamilton (1994).

Each of these approaches has different pros and cons. The first approach is easy to implement. The second makes heavy use of actual historical precedence and can capture short-intermediate-horizon autocorrelation and cross-correlation structures. The third allows for the strongest linkage between economic theory and simulation results, dynamic correlation behavior, and differential, conditional sub-period dynamics. But this comes with additional complexity in model design and calibration difficulty. In what follows we explore the differences in model outcomes focusing on the behavior of distribution tails and the implications for potentially differing intra-path dynamics. This may have great importance for pensions. For defined benefit (DB) plans, this can meaningfully impact the magnitude and timing of required contributions. For defined contribution (DC) plans, it can have meaningful welfare implications—especially if there are behavioral responses to participant asset allocations around extreme performance periods.

Preliminaries

Before we elaborate on model differences, it is worth spending some time looking at the nature of uncertainty about the first moments of our return distributions. In Figure 3.1, we see three different averages of historic real equity returns—rolling ten- and 30-year geometric average returns, as well as the full sample geometric average.² Both rolling averages display quite a bit of variation. The ten-year numbers range from a low of -4.3 percent to a high of 17.9 percent. The 30-year numbers range from 3.1 percent to 9.9 percent and have deviated above or below the full sample average of 6.5 percent for periods as long as 19 years.³

32 Recreating Sustainable Retirement

Figure 3.1. Averages of S&P Composite real returns, 1871–2012.

Note: 1871–2010 data from Shiller (2013); 2010+ from Bloomberg (2013).

Source: Authors' illustration.

Standard theory would instruct us to ‘use all the data.’ Yet while this may work for a return series as long-lived and as widely followed as the S&P 500, note that even here, the standard error of the estimate is 1.55 percent. If the meandering of the rolling averages gives us reason to question stationarity, the standard error could be still wider than that. For other return series—foreign markets or new asset classes—30 years of reliable data may be difficult to obtain. Dimson et al. (2002) contains a nice discussion of additional problems due to censoring, survivorship bias, market discontinuities, and other factors that plague offshore equity markets, even for developed economies.

The economic impact of such volatility on savers can be tremendous. Taking the extremes of our 30-year rolling averages for illustration, imagine a 35-year-old putting a dollar of her 401(k) in the S&P for a planned age-65 retirement. If she assumes her dollar saved will grow at a conservative 3.1 percent *per annum*, it will be worth 2.5 times as much in real dollars. At a more robust 9.9 percent, that dollar would grow to \$17.

It should be noted that a 30-year horizon is frequently cited in corporate and public sector DB plans to justify future expected return assumptions that would imply implausibly large forward-looking equity return spreads over available risk-free debt of long maturities.⁴ Despite the market crashes of 2000–2002 and 2008–2009, the most recent 30-year real equity returns exceeded the very

long-term average by 2.0 percent *per annum*. This is due to the power of the bull run of the 1980s and 1990s. The 30-year horizon also conveniently leaves out the poor returns in the 1970s. A similar story holds for bonds.

The focus on equity alone raises to some degree the question about expected bond returns and the interplay between expected equity returns and bond returns. A serious analysis of returns should not explore the return dynamics of stocks in isolation. It should also analyze other investment choices and account for relevant conditioning variables. Bonds play a role in both cases. The body of literature on expected returns is large and growing (see Ilmanen 2011). A key implication of the literature is that expected returns vary over time. Factors such as interest rates, the spread between high yield and investment-grade debt, aggregate dividend yields and earnings yields, and book-to-market ratios, for instance, all have some predictive power in forecasting stock returns. More recently, longer-term factors such as demographic variables have been shown by Arnott and Chavez (2012), among others, to have some explanatory value.

In addition to uncertainty as to the *ex ante* mean, there is some uncertainty regarding the long-run generating process and how uncertainty compounds over time. There appears to be different behavior in the short run and the long run. Lo and MacKinlay (1988) used variance ratios to demonstrate that returns show positive serial correlation (momentum) over short horizons. Moreover, Poterba and Summers (1988), Campbell and Viciera (2005), and others have demonstrated that in long horizons there appears to be some evidence of mean reversion.

But even with conditioning variables, the specification of the *ex ante* mean expected return is an imprecise exercise. If that is the case, how much faith can we have in characterization of higher moments or distribution tails? In what follows, we first discuss the slow nature of convergence of financial variables to a true mean, even if one assumes stationarity. Next we lay out our three different simulation approaches. A macroeconomic regime-switching model is introduced as a mechanism to be able to capture both shorter-term and longer-term behavior of financial markets. We then describe the two real-world filtrations of interest: (a) a long-term glide path for a defined contribution plan, and (b) the funding and contribution impact over a more moderately long horizon for a corporate DB plan. After a summary of the results, we compare the dynamics for the three model approaches.

Slow Convergence to a True Mean

Uncertainty about the true mean or expected return is quantitatively large, even if we assume that the annual stock return can be viewed as a realization from a stationary distribution. The solid dark lines in Figure 3.2 show the 95 percent confidence interval for the expected excess stock return based on a 16 percent annual volatility and a sample average of 7 percent. The figure reminds us of the

34 Recreating Sustainable Retirement

Figure 3.2. Slow convergence to true mean.

Note: 1871–2010 data from Shiller (2013); 2010+ from Bloomberg (2013).

Source: Authors' illustration.

unnerving statistical uncertainties associated with some of the key parameters that most simulation techniques take as given.

Of course, if the true process is not stationary, convergence can be even slower still. Figure 3.2 also shows the convergence in estimation for another process with true mean of 7 percent and volatility of 16 percent. In this case, the generating process is a Markov-switching mixture of normal distributions where the two sub-sample means are set at a level ± 3.5 percent around our true mean (i.e. 3.5 percent and 10.5 percent). The parameters are chosen to correspond approximately to the widest deviations of the 30-year rolling averages from the full sample mean in the Shiller dataset (2013) shown previously.

A total of 10,000 150-year paths were run with starting draws equally split between the high and low conditional mean return states. State transitions occurred according to a two-state Markov chain set with $p = q$ so that each state is equally probable over the course of the simulation and the distribution of returns is symmetric.⁵ Values of p and q are set at 0.90 (dashed line) and 0.95 (dotted line) corresponding to average durations of the conditional states of ten and 20 years. Bootstrap standard errors are calculated from the 10,000 simulation trails to construct confidence intervals.

We see that this generates consistently wider confidence intervals. At the 30-year sample point, the ten-year average regime duration process would yield a confidence interval with a 1.93 percent wider spread (3.32 percent for the 20-year

average duration process). Convergence is also slower: comparing the ratio of the confidence intervals as the sample increases, at 30 years, the 20-year, slow switch process is 29 percent wider than our baseline case; at 150 years it is 36 percent wider.

Three Different Return-generating Models and Their Results

We compare three different return-generating models—a multivariate normal distribution model, a block bootstrap model which resamples from past historical data, and a third structural model that is designed and calibrated to capture a number of economic features. The more involved macroeconomic long horizon simulation model is described in some detail below, with a more detailed technical appendix at the end of the chapter. The other two models, multivariate normal and block bootstrap, are commonly used by both academics and practitioners and are presented later as counterparts for comparison and evaluation of the more expansive macroeconomic approach. Limitations of both of these models provided the impetus for the development of this modeling approach.

Macroeconomic Long Horizon Simulation (LHS) Model

The premise of the macro-driven long horizon simulation model is that the dynamic processes for macroeconomic growth, inflation, and risk aversion jointly determine both the short-term dynamics as well as the variations in discount factors applied to these cash flows in financial markets to form asset prices and valuations for both risk-free and risky investments. When we impose restrictions and assumptions on the dynamics of these fundamental variables that are based on theory and academic research, we indirectly imply a set of ‘structural’ restrictions for the long-term valuations of asset prices, such as bond prices and equity prices, which restrict the plausible range of average investment returns and volatility over a given investment horizon. Cochrane (2011) contains a discussion of some of these issues.

The secular dynamics of productivity growth or ‘potential’ GDP growth, inflation rates, and equilibrium risk premiums shape the distributions of real interest rates and nominal interest rates and equity yields (dividends and earnings yields) over a long-term investment horizon. Similarly, the business cycle dynamics of unemployment, output gap, and central bank policy, as well as the accompanying temporary bouts of extreme uncertainty and risk aversion in financial markets, shape the distribution of asset returns and yield curves in the short term.

A structural macroeconomic regime-switching (LHS) model combines these guiding principles and ideas within a unified structural framework designed to

36 Recreating Sustainable Retirement

Figure 3.3. Schematic of structural LHS process.

Source: Authors' illustration.

remain highly empirically and quantitatively relevant. Figure 3.3 summarizes the structure of the model at a high level.

This LHS model explicitly incorporates the main dynamic linkages between economic growth and inflation, monetary policy, risk aversion, and realized asset returns. In this way it provides a suitable framework for assessment of both long-term and short-term distributions of asset returns and yields curves. The framework allows us to quantify both the short-term and long-term tail risks that strategic long-term investors face based on inputs for the main drivers of returns (long-term growth, long-term inflation, and equity risk premium) and allows us to explore the impact of parameter uncertainty.

A key feature of the model is the regime-switching component, which generates realistic business cycle dynamics in the model (see Hamilton 1989). The regime-switching dynamics directly translate into realistic cyclical fluctuations in inflation, GDP growth, and risk premiums. They also enable the model to more closely match the higher frequency moments of asset returns distributions, as well as the general correlations with macro variables over the business cycle. More specifically, such regime-switching models can produce the type of rapid changes in risk premiums required to generate the ‘boom–bust’ characteristics of asset returns experienced historically. The regime-switching behavior ultimately also generates a strong tendency for mean reversion in average returns, and it reduces volatility of those returns over long-term investment horizons. A more detailed description of each component of the model is provided next.

Macro Regimes

The business cycle dynamics are determined by a regime-switching model that places the macro economy in either a ‘contraction’ or an ‘expansion,’ similar to the specification first introduced by Hamilton (1989). One extension is that the transitions between these two regimes are governed by time-varying probabilities modeled with duration-dependent hazard functions. The duration dependence is such that the hazard rate increases exponentially with the time the economy has been in a specific regime. The probability that a specific recession or expansion comes to an end consequently increases with the duration of the current stage of the business cycle. The parameters of the hazard function for contractions and expansion are calibrated to match the characteristics of National Bureau of Economic Research (NBER) business cycle durations, both in terms of their average length and their standard deviation. Since contractions tend to be shorter-lived than expansions, the transition probabilities are inherently asymmetric.

Real Activity, Growth, and Output Gaps

The business cycle dynamics of real GDP growth are derived from the specification of a dynamic process for the output gap. In a recession, the output gap converges to a negative level, whereas it converges to a positive level in an expansion. The specification that is used within regimes ensures that real GDP growth rate is generally most negative at the beginning of a recession and most positive at the beginning of an expansion. Overall the parameters are calibrated to match the distribution of real GDP growth and the distribution of output gap realizations within both expansions and contractions. The long-term potential GDP growth is set to an exogenous rate in the model, which can reflect forward-looking views on productivity growth or simply the historical growth rate of about 2 percent per year. In principle the long-run growth rate can have its own dynamics.

Inflation

Inflation in the model has both a cyclical business cycle-driven component and a persistent, long-term component. The business cycle component of inflation is assumed to be driven by cyclical fluctuations in aggregate demand and, as such, it is based on the simulated level of the output gap. A positive output gap is associated with high cyclical levels of inflation, whereas a large negative output gap will reduce the inflation rate to a level below the long-term ‘structural’ inflation rate that is prevailing at a given point in time. This feature of the model resembles the well-known ‘Phillips’ curve from theoretical economic models. The long-run level of realized inflation, and hence also the long-run level of inflation expectations, is determined by shocks to a stationary but highly persistent process for non-business cycle-related inflation fluctuations. Shocks to this process have a very significant

38 Recreating Sustainable Retirement

impact on the expected inflation rates many years into the future, which are important in shaping the long-term distribution of nominal yields.

Monetary Policy

The central bank is assumed to respond to the cyclical components of inflation and real activity. The response function is calibrated with a Taylor rule that implies positive real rate responses to both the cyclical component of inflation and the output gap. Both effects will tend to push the Federal Reserve to reduce short-term real rates in recessions (absent any short-term supply shocks to inflation) and increase the short-term real rate in expansions. The current unusual monetary policy stance, however, warrants an explicit adjustment to this general description of monetary policy. The impact of quantitative easing (QE) is therefore explicitly modeled. Specifically, for the two years until the end of 2014, our model assigns only a small probability that the Fed exits its current stance (with a fixed fed funds rate at 25 basis points) and raises the fed funds rate toward the level implied by the Taylor rule. After that, it is assumed that the probability of exiting the regime increases over time.

Yield Curve and Bond Returns

The nominal yield curve is derived as the expected future short rate plus an inflation risk term premium and a real rate risk term premium. The ‘expectations component’ of the yield curve is implicitly derived from the expected dynamics of the output gap and inflation because they ‘pass through’ the Taylor rule to the expected future fed funds rate. As a consequence, the expectations component of the yield curve responds to the current state of the business cycle and becomes logically consistent with the specified dynamics for monetary policy, inflation, and real activity. The business cycle dynamics of short maturity yields are dominated by business cycle-driven fluctuations in expectations about monetary policy, which creates interesting and very realistic endogenous dynamics in the yield curve simulations. For instance, at the beginning of an expansion the market will price additional rate increases resulting in a steepening yield curve, whereas at the beginning of a recession the market will anticipate further easing of monetary policy over the course of the recession and yield curves will potentially be inverted for some time as the economy enters a contraction. Overall, the yield curve will be ‘steepest’ in the middle of or at the end of a recession, whereas the curve will be ‘flattest’ in the middle of an expansion. The current shape of the yield curve is consistent with our modeling of QE.

The cyclical dynamics of long maturity yields in the model are, on the other hand, mostly driven by fluctuations in inflation and real rate risk premiums, since long-term expectations for real rates and inflation (and hence expectations for

distant future policy rates) are fairly stable across the business cycle. This is consistent with empirical evidence suggesting that survey-based measures of long-term expectations and real activity are quite constant over time and simply do not fluctuate enough to generate the observed volatility in long-term real and nominal yields at annual frequencies.

While risk premiums fluctuate in the short run and drive dynamics in the short run for long maturity yields, they remain fairly constant and are 'bounded' in the long run. The main driver of any significant dispersion in the simulated long-run dispersion of nominal yields is therefore gradual accumulation of small but persistent shocks to inflation. Generally speaking, significant changes in long-term interest rate changes must be accompanied by persistent changes in the level of inflation.

Finally, we note that, on average, long maturity bonds have higher expected returns due to the maturity-dependent term premium that is specified in the model. As a result of this risk premium, the average yield curve tends to be upward sloping at maturities out to about 20–25 years. After that convexity, effects in the yield curve flattens the curve and cause even longer-term yields to decline gradually with maturity.

Equity Returns

Equity prices are based on expected dividends discounted with the term structure of risk-free yields as well as the equity risk premium. The equity return is composed of dividend yields plus capital appreciation (re-pricing of equities). Earnings growth is based on future GDP growth and mean reversion in corporate profit margins. We assume a constant payout ratio for dividends. Consequently, short-term earnings growth expectations will rationally fall when the probability of entering a recession increases in the model, but long-term earnings growth and dividend growth expectations do not fluctuate a lot over time. The dominant component of equity returns and return volatility is the 're-pricing' component of returns and the associated changes in the price to 'stabilized' earnings ratio. The dynamics of the P/E ratio are inherently linked to the behavior of both long-term real rates and, especially, the dynamics of the equity risk premium. The re-pricing return in a given period is approximately equal to the change in the equity discount factor times the equity price duration with respect to yield and equity risk premium changes. The equity price duration can be inferred from the discounted cash flow model but will be about 20–30 years depending on current valuations. In the recessionary regime, equity risk premiums will tend to widen (capital loss) with higher volatility, but will narrow (capital gain) in an expansion. Similarly, the volatility of risk aversion levels and hence equity premiums are assumed to be higher in recessions than expansions.

40 Recreating Sustainable Retirement

To summarize, most of the annual volatility in equity returns is due to the fluctuations in the discount rate and hence fluctuations of the equity risk premium and long-term interest rates. The volatility is not due to shocks to expected future cash flows. These features are consistent with the extensive body of academic research associated with Robert Shiller.

Comparison to Alternative Simulation Approaches

To highlight some of the major advantages of the macro-based approach to long-run simulations we compare our simulation results with two simpler and more commonly used approaches to simulation of both the yield curve and the equity market returns. These approaches do not incorporate the same inter-temporal structural relationships between equity returns and bond returns and macro variables that the macro model ‘enforces.’

One model which we can compare results with is a normal approximation to the LHS model dynamics, where we match the average annual return, volatility, and correlations of two-, five-, ten-, and 30-year maturity points on the yield curves, as well as equity returns. A second model is a ‘bootstrap’ model, similar in that we match the average return and volatility to the structural model, but this instead samples historical quarterly data to also match the higher moments of the empirical distribution that a normal approximation may be missing. The bootstrap generally produces fatter tails in equity returns, since they are non-normal on an annual basis. But for long-term investment horizons, the difference between the bootstrap and normal approximation is small.

The structure of the LHS model serves to limit the plausible range of outcomes over longer periods of time. It induces mean reversion in excess bond and equity returns over time, which means that the long-run volatility of real returns decreases with the investment horizon. For instance, following an increase in the equity risk premium, investors are ‘hit’ by an immediate capital loss, but they face higher returning investment opportunities afterwards. Similarly, a shock to interest rates causes negative bond returns in the short run, but investors then face an environment with higher yields subsequently. Figure 3.4 shows how the term structure of volatility is downward-sloping in our model, whereas it is flat in both the normal model and the bootstrap model.

The downward-sloping term structure of volatility that the macro-based model generates means that there are very large differences between the long-term predictions of the macro-based structural model and the two ‘naïve’ memoryless simulations of returns. Over time the cumulative volatility of returns explodes in both of the non-structural, reduced-form approaches. This gives rise to exaggerated tail behaviors in these two simulation approaches. This can be seen in Figure 3.5,

Figure 3.4. Volatility term structure: equity returns.

Source: Authors' illustration.

which shows the distribution of equity and bond returns from the three different models.

For example, both the normal and bootstrap processes give roughly 5 percent likelihoods of negative annualized equity returns over a 20-year horizon. It is difficult to imagine such persistence in an economy without a fundamental structural shift in markets, given the implied long-term real capital decimation. Similarly, the 95th percentile of 20-year annualized compounded equity returns is nearly 14 percent for both simple models. Here either growth or inflation would have to be materially higher for a sustained economic period than we have witnessed, or price-earnings ratios would eclipse historic levels. Similarly implausible relationships would have to hold to generate bond return behavior as seen at the outer percentiles.

Importantly, the specific paths that generate the tail events in pure ‘engineering’ models, such as the independent and identically distributed (i.i.d.) normal distribution model, of asset price returns cannot be linked to a specific economic environment or assumption. It is not possible to assess whether the ‘tail’ outcomes are reasonable from an economic standpoint as the fundamental economic parameters that shape the distribution of returns are hidden from the visible eye. In the macro model, it is much easier to pinpoint the economic assumptions that have to be made to generate a given tail scenario, and hence to assess whether it is ‘plausible’ or not.

Panel A. Annualized horizon equity returns.

Panel B. Annualized horizon bond returns.

Figure 3.5. Models' distributions of equity and bond returns.

Note: Panel A: Annualized horizon equity returns; Panel B: Annualized horizon bond returns; Panel C: Horizon annualized 60/40 returns.

Source: Authors' illustration.

Panel C. Horizon annualized 60/40 returns.

Figure 3.5. (Continued)

Implications for Projections of DC and DB Plans

To get a sense of the impact of the three different return-generating processes in real-world settings, we construct two examples: one for a DC plan and another for a DB plan. The returns and yield curves generated by our models can then be used to assess the impact on key decision variables used by plan sponsors. While the annualized distributions of return dynamics are informative, without the overlay of the pension-specific models, key inferences may be missed. In addition, the long-term nature of the pension saving/funding dynamic introduces a series of other effects due to the changing dynamics of contributions. In the case of the DC plan, the representative saver is increasing contributions though time as she ages and presumably has a higher income from which to save. For the DB plan, there is an endogenous, path-dependent relationship between contributions, asset returns, and changes in the discount rate that can be highly non-linear through time.

For our DC case, we look at the cumulative savings dynamics of our return process over a 40-year glide path.⁶ With the passage of the Pension Protection Act (PPA) in 2006 and its codification of target date funds as Qualified Default Investment Alternatives (QDIAs), there has been explosive growth in these target date funds. According to Morningstar (Furman 2013),⁷ total target date fund assets reached \$475 billion by November 2012—a four-fold increase since the passage of the PPA. In practice each fund family has its own glide path, and most have individual nuances. Yet as a rule, the glide paths are invariably designed so that financial market risk in portfolios is decreasing in time-to-retirement. This is

44 Recreating Sustainable Retirement

akin to the financial planning heuristic that one saving for retirement should hold $(100 - \text{age})$ in equities or other similarly risky assets. The essential motivation here is that as the DC plan participant is moving through his career toward retirement, he is replacing the present value of future human capital with financial assets. In the early years, with many years to retirement, the reservoir of this store of human capital is large relative to financial assets and provides a buffer against market shocks. As the individual approaches retirement, the relative size of the combined portfolio = $PV(\text{Human Capital}) + \text{Financial Savings}$, tilts toward an increasing fraction on the financial assets side. To maintain a relatively balanced risk position, the mix in the financial portfolio (here the target date fund) decreases in risk as one moves closer to retirement.

In practice, target date portfolios may have many different asset classes and gradations within each asset class. MarketGlide calculates indices of weighted-average glide paths from 37 fund families (see Abidi and Quayle 2010). For our purposes we use a simplified version of the glide path that maps assets into cash, fixed income, and equities. This provides a relatively accurate representation of key risks as the indices are dominated by U.S. aggregate fixed income and U.S. large cap equities. As a general rule, asset classes with distinct dynamics, principally commodities and real estate, comprise less than 2 percent each at any point along the industry average glide path. The glide path used is shown in Figure 3.6.

Figure 3.6. Simplified market average DC glidepath.

Source: Authors' illustration.

Our hypothetical participant starts his participation in our DC plan with a salary of \$30,000 per year, which grows at a rate of 1 percent per year in excess of inflation.⁸ He contributes 6 percent of pay and has a match of 4 percent from his employer. Moreover, he is assumed to invest solely in the target date fund corresponding to the 40-year glide path.

The variables of interest are the plan account balances at various points along his progression to retirement. Note that for each individual simulation, there will be some path-dependency as the values grow by the contributions as well as asset returns and the asset allocation is changing through time, so timing and order of specific return draws can be very important. Table 3.1 shows the characteristics of two hypothetical DB plans that we examine.

We assume both plans are currently 80 percent funded, approximately equal to the funding level (81.7 percent) estimated for the Milliman Pension Funding Index as of January 31, 2013 (Milliman 2013). The funding rules for contributions are those as set out in the PPA.⁹ Unfunded amounts are subject to a seven-year amortization basis. Each year the amount of underfunding is compared to the rolled-forward amortization bases. If the new underfunding exceeds the value of the unamortized prior-year bases, a new basis is created with a seven-year amortization. Once a plan is fully funded, all existing bases are erased. The required contributions are the sum of the amortization charges from these bases and the plan's normal (service) cost.¹⁰ For simplicity, we handled the plan's current underfunding by assuming that equal amortization charges were established in the current and each of the three prior years. Funding levels are a function of both the plan's assets and liabilities. In reality, plan sponsors have some latitude in the yield curve used to determine liabilities—this is currently even more the case given the rules promulgated by MAP-21 in 2012. For our purposes we use an approach closer to the mark-to-market liability valuation as set forth by the FASB for accounting purposes and use a point-in-time yield curve rather than one that is a moving average through time.

Assets are rolled forward assuming asset returns per a 60/40 mix of stocks and bonds rebalanced quarterly, less current year benefit payments. Liabilities are revalued each year given changes in the discounting yield curve. We examine two

TABLE 3.1 Characteristics of DB plans

	Frozen plan	Accruing plan
PV (liabilities)	\$1.25 billion	\$1.25 billion
0–10 years	42%	32%
10–20 years	35%	36%
20–30 years	16%	20%
30+ years	7%	12%
Duration of liabilities	13.8 years	16.5 years
Normal cost (%) of liabilities	—	5%

Source: Authors' tabulations.

46 Recreating Sustainable Retirement

cases: a frozen plan and an open accruing plan. For the frozen plan, benefits are paid and the $t+1$ set of cash flows becomes the current year's benefit payments (we assume that the actuaries have perfect foresight into plan demographics). For the accruing plan, we make the simplifying assumption that normal cost is applied *pro rata* in future years. This yields a stationary distribution of liabilities which, as a first order approximation, only shifts duration in response to changes in the yield curve.

The principal variable of interest is the value of required plan contributions. For the DB case, we expect scenario results to be highly path-dependent. Unlike the DC case where we had a time-varying asset allocation, here the path-dependency is a result of the interplay between assets and liabilities (stocks and the yield curve) and the contribution rules. Strong equity markets and/or large increases in the discount rate can abruptly halt required contributions, and they may stay at zero for some time. Poor equity markets/flat-to-declining rates, combined with benefit outflows, can cause required contributions to increase and stay persistently high.

Results for DC Simulations

Results for our 40-year DC glide path are presented in Figure 3.7 and Table 3.2. Over time, we see a pattern emerging. The spread of the block bootstrap

Figure 3.7. Total savings in DC plan.

Source: Authors' illustrations.

TABLE 3.2 Total savings in DC plan

Structural (\$)	2018	2023	2028	2033	2038	2043	2048	2053
1st percentile	15,122	30,537	51,002	83,016	128,019	188,575	260,682	335,884
5th percentile	17,341	36,168	60,584	100,360	151,868	216,685	295,950	397,964
10th percentile	18,513	39,179	66,852	108,428	164,470	237,206	325,632	432,843
25th percentile	20,363	44,014	77,175	124,257	188,668	273,212	382,817	514,899
Median	22,490	50,207	88,966	144,981	220,990	324,748	458,374	629,008
75 percentile	24,564	56,930	103,199	168,051	259,750	385,915	549,521	778,129
90th percentile	27,003	64,669	117,039	192,388	298,291	458,437	664,364	952,725
95th percentile	28,518	68,642	128,872	206,481	326,813	504,899	744,054	1,076,167
99th percentile	31,366	77,712	147,174	245,714	382,644	590,109	904,104	1,431,140
Mean	22,364	50,103	88,809	144,339	221,569	326,856	462,825	638,618
Bootstrap (\$)	2018	2023	2028	2033	2038	2043	2048	2053
1st percentile	14,079	25,777	41,732	65,719	96,981	141,697	184,930	231,906
5th percentile	16,619	32,288	52,603	81,677	121,999	176,682	240,333	310,006
10th percentile	17,911	35,756	58,710	93,381	140,008	201,478	278,002	361,897
25th percentile	20,103	42,704	72,530	115,284	173,771	255,539	350,723	467,323
Median	22,639	50,872	90,233	145,922	225,470	332,590	467,660	641,948
75 percentile	25,180	60,403	112,530	186,935	293,122	440,524	637,955	918,742
90th percentile	27,688	69,510	134,961	235,436	366,077	562,444	830,329	1,238,562
95th percentile	29,366	75,592	151,949	267,700	434,956	649,153	1,016,471	1,588,341
99th percentile	32,485	89,988	187,266	335,609	552,169	854,683	1,349,915	2,263,013
Mean	22,399	50,388	89,801	146,925	226,834	336,252	477,879	661,902
Normal (\$)	2018	2023	2028	2033	2038	2043	2048	2053
1st percentile	14,820	27,836	45,457	70,568	106,233	153,494	218,312	276,610
5th percentile	16,603	32,775	54,421	86,520	133,322	194,091	267,094	348,871
10th percentile	17,725	36,096	60,757	96,957	146,713	213,807	297,960	395,908
25th percentile	19,907	42,079	71,883	116,552	177,523	263,153	367,440	487,140
Median	22,464	50,538	89,817	145,566	222,230	327,427	464,200	634,796
75 percentile	25,252	60,067	110,930	182,323	282,972	418,014	599,582	856,484
90th percentile	27,955	70,093	134,154	227,201	354,872	531,657	761,865	1,105,678
95th percentile	29,672	77,677	149,767	252,908	397,478	604,754	878,447	1,320,910
99th percentile	33,480	94,584	188,617	333,923	552,279	814,311	1,248,614	1,945,710
Mean	22,402	50,408	89,822	146,573	225,607	333,430	472,565	652,603

Source: Authors' tabulations.

48 Recreating Sustainable Retirement

distributions of account balances is wider than the spread of the multivariate normal distributions, which in turn is wider than those seen for the structural LHS model. A similar pattern is seen when we examine the means and medians of the distribution as we go out to the far years of the distribution. We might expect differences in the width of the distributions given the different generating processes, but one must recall that the models were calibrated to have the same mean and volatility of returns for any single year. Hence, *ex ante*, most readers would probably expect the means and the medians of the distributions to stay quite close for our entire horizon. But note that we are showing the means and medians of the distributions at each point in time, not the values corresponding to compounding the mean or median single-year returns.

The same forces that give rise to longer-run mean reversion in returns in the structural model affect the means and medians of the distributions as well. For shorter horizons, these effects are small—even out to 20 years, values of medians stay within 1 percent or less (means within 2 percent), but this drift compounds out to 40 years where there is as much as a 4 percent discrepancy in mean account balances in the difference between the bootstrap simulations and the structural model results.

Both tails of the account balances are also quite different. After 20 years, the saver has combined employee and employer contributions of just over \$92,000. In both the bootstrap and normal models, he would expect that, almost 10 percent of the time, his account balance would be less than cumulative contributions. Under the LHS model, this would happen less than 3 percent of the time. The LHS model produces downside accumulation results that are 26 percent better in the lowest percentile of the distribution, and 23 percent better at the 5th percentile. On the flip side, the LHS approach also produces smaller account balances in the good scenarios. Balances are approximately 18–23 percent (26–27 percent) lower at the 95th (99th) percentiles than in the bootstrap and normal distributions.

It is also instructive to look at the implied geometric average returns in the upper tails. For the normal and bootstrap cases, these are approximately 11 percent and 13.3 percent. These do not seem too extreme when one considers that the average glide path weight in equity in the first 20 years is roughly 85 percent. Yet they may appear wider when viewed against the backdrop of current bond yields of less than 2 percent on ten-year Treasuries, and also given that higher equity returns have tended to coincide with periods of persistently moderate or declining inflation. A back-of-the-envelope calculation assuming he gets 2.5 percent from bonds over the period would imply either 10.0 percent or 12.7 percent excess returns for equities. This seems inordinately high for such a long period.

At the 40-year horizon, the bootstrap-driven model has the widest distribution, with both more negative and more positive outcomes. It should also be noted here that the asset allocation changes the most dramatically in the last 20 years of the glide path. Equities drop from 78 percent with 20 years to go, to 44 percent at retirement. The bond allocation rises from 19 percent to 45 percent.

The lack of mean reversion in the bootstrap and normal models may be most distortionary here, since there is natural reason to expect negative serial correlation of returns over the medium to longer term in fixed income markets. If a riskless bond, which is held to maturity, returns less than its yield to maturity it *must* be followed by returns greater than the initial yield to maturity. This mathematical fact is lessened somewhat if we think about rolling portfolios and portfolios containing spread product, but for the investment-grade fixed income assumed here, these factors would be insufficient to overcome the rationale for longer-term mean reversion in returns.

After 40 years, the account balances in the bootstrap simulations are on average 1.5 percent higher than those for the normal simulations, and 3.5 percent higher than for the structural model. At the upper extremes of the distributions (99th percentile) the bootstrap gives account balances nearly 40 percent greater than those in the normal model, and nearly 60 percent greater than those seen in the structural model. The implied annualized return is approximately 10.3 percent. Again this number seems inordinately large over a 40-year horizon, given the mix of assets and our current starting point. In the bottom percentile, the bootstrap model lags the normal model by 16 percent and the structural model by 31 percent. The persistent negative returns required to get approximately 1.8 percent *per annum* seem implausibly low for a 40-year implied compounded average.

The spread for the normal model seems more reasonable by comparison, yet it also spans a range where the 99th percentile outcome is almost seven times that seen in the 1st percentile results. Its median and mean results are 1–2 percent higher than those seen in the structural simulation model. For the extreme percentiles, the normal model is about 20–25 percent wider than the structural model. In the worst percentile case, we again see negative nominal returns over a 40-year horizon. Upside effective annual asset appreciations are 8.0 percent annualized at the 95th percentile (9.7 percent annualized for the 99th percentile). This does not seem terribly unreasonable given our horizon length and the fact that these are for the extremes of the distribution.

The LHS model produces an annualized return of 0.8 percent in the worst percentile. This would be depressing, but much less so than for the other two models. Results to the upside are more muted as well producing effective annualized returns of 7.1 percent at the 95th percentile (and 8.4 percent for the 99th). It should be noted, however, that annualized returns can be a bit deceiving in this context when one thinks about the asset-weighted averages. Time-weighted, the averages have substantially more exposure to equities, but when viewed against the construct of the glide path and increasing reliance on bonds in the future, the average effective weight is likely more closely balanced.

Results for DB Plans

Turning our attention to the DB plans, we first analyze the behavior of contributions across the three models for the frozen plan over a ten-year horizon. Recall that

50 Recreating Sustainable Retirement

asset allocation is kept at 60 percent equity and 40 percent fixed income with quarterly rebalancing. Under the contribution rules, median minimum required contributions drop to zero by 2017 (see Figure 3.8, Panel A). This is largely driven by the upward drift in the discount curve, assumed to increase by 100 basis points on average at the ten-year point of the yield curve and somewhat more for short–intermediate maturities with the two-year rate rising by some 240 basis points. Average required contributions stay positive in all years given the funding rules. For the first three years, there is little difference across the models given the nature of the amortization bases. Average contribution results differ a bit as we move forward in time, but even at their widest they amount to approximately \$6 million in ten years' time, or less than 1 percent of current liabilities. At the 99th percentile, the potential contributions display greater and widening variance. For 2014, the spread between the bootstrap and LHS models is a little over \$9 million. In ten years' time, this grows to \$61 million.

Panel B in Figure 3.8 shows a distribution of the present value of contributions for the three models. Interestingly, all three models have average required contributions less than current shortfalls, given the central tendency for rates to rise. The present value of average contributions ranges from a low of \$170 million in the LHS model to a high of \$210 million in the bootstrap model. The bootstrap and normal models display significantly higher distribution in the possible future contributions. This difference is again largely attributable to the lack of mean reversion of returns in the bootstrap and normal distribution models and to the propensity for a greater left skew in the historic return distribution for equities.

Contribution requirements for the plan still accruing benefits are more interesting (see Figure 3.9, Panel A). There is a similar pattern of mean and median required contributions for the first few years, but then there is some divergence. Means and medians stay positive, as additional benefits are earned while the combined increase in the yield curve and current asset levels are not sufficient to award 'free' additional benefit accruals. Interestingly, the medians and means diverge most for the bootstrap simulations, with the bootstrap ultimately having the lowest median contribution requirements, but the highest mean requirements. After the first few years, the bootstrap consistently produces the largest mean and 99th percentile contributions.

Over the next ten years, the average present value of contributions for the open plan ranges from \$789 million in the structural model to \$866 million in the bootstrap model. The normal model is in between, at \$841 million. The bootstrap model also displays the widest variation in contributions and a substantial tail (Panel B, Figure 3.9).

Conclusion

This chapter examines three alternative methods to simulate long horizon yield curves and asset returns: a 'block bootstrap' simulation, a normal 'Monte Carlo'

Panel A. Minimum required contribution (\$ millions): frozen plan.

Panel B. Distribution of ten year cumulative contributions.

Panel C. Distribution of five year cumulative contributions.

Figure 3.8. Contribution patterns under alternative scenarios: frozen plan.

Notes: Panel A: Minimum required contribution (\$mm): frozen plan. Panel B: Distribution of ten-year cumulative contributions. Panel C: Distribution of five-year cumulative contributions.

Source: Authors' illustration.

Panel A. Minimum required contributions (\$ millions): accruing plan.

Panel B. Distribution of ten-year cumulative contributions.

Panel C. Distribution of five-year cumulative contributions.

Figure 3.9. Contribution patterns under alternative scenarios: accruing plan.

Notes: Panel A. Minimum required contribution (\$mm): accruing plan. Panel B. Distribution of ten-year cumulative contributions. Panel C. Distribution of five-year cumulative contributions.

Source: Authors' illustration.

simulation, and a structural economic regime-switching model. We explore how the choice of modeling approach affects the simulated distribution of future returns and outcomes (such as plan contributions and retirement wealth) for both DC and DB plans.

The bootstrap model has desirable properties for shorter simulations of a few years or less, but it produces very questionable results for long horizons. Examination on a path-by-path basis yields individual scenarios that strain credibility, both in individual variables and between variables. The normal model with independently, identically distributed returns also fails along several dimensions. Over a longer horizon, fundamental tenants of economic rationality dictate bounds on valuations that would almost certainly rule out a true memoryless process, which the normal and bootstrap models assume. The weak linkages between variables that correlation provides are not sufficient to ensure reasonable long-term relationships. A simple correlation matrix cannot capture the complex dynamic relationship between stock returns, bond returns, and macro variables. These relationships are better characterized by a model that generates correlations which vary with the macroeconomic regime, such as our proposed macro-driven regime-switching model (LHS). This model captures many desirable properties at both the secular and cyclical frequencies and overcomes many of the flaws of simpler models. The model also makes it possible to directly link investment returns with simulated macroeconomic time series for inflation and GDP growth. For this reason we would argue that the structural model should be preferred for longer horizon simulations. The model can be used to evaluate different asset allocations, market or macro-contingent dynamic asset allocations, and hedging strategies.

A key implication of the LHS model is mean reversion of returns at longer investment horizons, which implies a downward-sloping term structure of volatility. This feature or idea is supported by the work of Campbell and Viceira (2005) and by Siegel's well-known work (1994) on long-run stock returns. We initially described the statistical uncertainty surrounding estimates of risk premiums and expected returns that may even vary over time. So, what is the equity risk premium we are supposed to be converging to in the simulation? The task ahead of us is to incorporate some essence of the uncertainty about the mean processes (risk premiums) into our long-term simulation models (see Pastor and Stambaugh 2012).

Our results and discussion should provide some cautionary lessons. Characterizing distributions of outcomes ten or 40 years hence is an exercise that should be taken with more than a single grain of salt. One should look at such models (no matter how sophisticated) only as one among a number of guides to answers, rather than the sole guide—or, worse yet, the answer in and of itself.

Appendix: The Regime-Switching Long Horizon Simulation Model

Regime Dynamics

The transition probabilities between regimes are a function of the time spent in the regime. Formally, it is assumed that the transitions are governed by a Weibull distribution. The survival function (the probability of staying in a given regime past time t is given by $S(t) = e^{-(t/a)^b}$, with associated probability density function $f(t) = ba^{-b}t^{b-1}e^{-(t/a)^b}$, and hazard rate function $h(t) = (f(t))/(S(t)) = ba^{-b}t^{b-1}$. As discussed above, we calibrate the parameters to match the empirical mean and standard deviation of the durations of economic expansions and contractions.

Output Gap Dynamics

The output gap is assumed to follow an Ornstein-Uhlenbeck process within regimes. In a regime-switching setting the key feature of this specification is that changes in GDP growth are going to be more pronounced at the business cycle turning points.

$$dy = -\kappa_{y,s} [y - \theta_{y,s}] dt + \sigma_{y,s} \sqrt{dt} \cdot \varepsilon_t$$

Mechanically, the reason why the most dramatic changes occur around turning points is that it is typically at the business cycle turning points in the simulation (where the regime changes from one to the other) that the difference between the current output gap and the new ‘target level’ for the regime is the greatest. As a result, extreme levels of GDP growth are going to be realized at the beginning of a recession (from positive to highly negative) and the beginning of an expansion (from negative/zero to highly positive).

Inflation Dynamics

Inflation is assumed to have two components.

$$\pi = \pi_{cyclical} + \pi_{persistent}$$

$$d\pi_{cyclical} = -\kappa_{cyclical,s} \pi_{cyclical} dt + \sigma_{cyclical,s} \sqrt{\delta\tau} \cdot \varepsilon_\pi$$

The first component is meant to capture excess ‘demand’-driven inflation. Its behavior is strongly linked to the output gap. It is in other words a short-term Phillips-type inflation vs. aggregate demand relationship.

$$d\pi_{\text{persistent}} = -\kappa_{\pi,s}(\pi_{\text{persistent}} - \theta_{\pi})dt + \sigma_{\pi,s}\sqrt{dt} \cdot \varepsilon_{\pi}$$

The second component is a very low frequency/almost permanent component of inflation. This component captures secular changes in inflation levels in the economy.

The Taylor Rule for Monetary Policy

Monetary policy is assumed to follow a modified Taylor rule, such that the targeted short-term real rate responds to deviations of inflation from its long-run level and the output gap.

$$r = (r^* + \pi) + \beta_{\pi}(\pi - \pi_{\text{long}}) - \beta_{y,s}[y - y_{\text{long}}]$$

Monetary policy is neutral when the current inflation equals the current value for the time-varying long-term component of inflation. This captures the notion that a substantial increase in long-run inflation expectations and average future realized inflation has to be a monetary phenomenon. On a path toward higher inflation, which would correspond to a shock in the permanent component of inflation, we would expect to see ‘unsustainable’ or overly accommodating levels of (low) short-term real rates of interest. It is implausible for real policy rates to remain very high during inflationary periods. The general cyclical dynamics of monetary policy in the model are however determined by the business cycle. In recessions the central bank cuts real rates, and in expansions the central bank tightens, increasing real short-term rates.

The Yield Curve Model

The nominal yield at maturity T is given by the average expected future short rate plus a nominal risk premium minus a nominal rate volatility convexity adjustment. The real yield for maturity T is given by the average expected future short real rates

56 Recreating Sustainable Retirement

plus a real risk premium (which may be negative) plus a liquidity risk premium—real rate volatility convexity adjustment. Mathematically, the nominal yield at time T can be written as

$$y_T = E \left[\frac{1}{T} \sum_{j=1}^T r_j \right] + \delta_T + \theta_T - \gamma_T$$

where r is the short rate (three-month T-Bill), δ is the real rate risk premium for maturity T , θ_T is the duration/inflation risk premium, and γ_T is the convexity adjustment at maturity T . The real rate risk premium is driven by the business cycle in that we relate it to the output gap and is given by:

$$\delta_T = \alpha_{\delta,T} + \beta_{\delta,T} y_{\text{gap}} + \epsilon$$

$\beta_{\delta,T}$ is assumed to be positive, which captures the positive correlation between the real risk premium and the output gap. This captures the ‘flight-to-safety’ characteristics of U.S. securities in general and U.S. Treasuries in particular. The nominal duration (inflation) risk premium is given by:

$$\theta_T = \alpha_{\theta,T} + \beta_{\theta,T} \pi_{\text{inf}} + \epsilon$$

Since $\beta_{\theta,T}$ is assumed to be positive, we generally expect a positive correlation between the inflation risk premium and the level of inflation. The convexity adjustments to the nominal and real yield curves are simply given by the conventional expression of:

$$\gamma_T = \frac{1}{2} T^2 \sigma^2$$

The convexity adjustment plays a significant role in shaping the longer-term maturities of the yield curve, but does not in itself affect the expected term premium.

Overall the dynamics of the shorter end of the curve (maturities inside five years) are heavily influenced by the business cycle and the expected federal funds rate.

The Equity Model

Equities are priced according to discounted dividend model with a terminal condition:

$$P = E \left[\sum_{t=1}^{30} \frac{D_t}{(1+r_t + \lambda)^t} \right] + E \left[\frac{1}{(1+r_{30} + \lambda)^{30}} P_T \right]$$

The equity risk premium is assumed to follow a regime-switching process (Ornstein–Uhlenbeck process). Each regime is characterized by a mean, volatility, and mean reversion parameter.

Equity Risk Premium

The dynamics of the equity risk premium are specified as an Ornstein–Uhlenbeck process that depends on the regime (contraction or expansion) that the economy is currently in. Formally,

$$d\lambda_{erp} = -\kappa_{s,erp} [\lambda_{erp} - \theta_{j,erp}] dt + \sigma_{s,erp} \sqrt{dt} \cdot \varepsilon_{erp}$$

The parameters of the model are calibrated to generate higher volatility in recessions and more violent changes in equity risk premiums around business cycle turning points.

This means that the parameters satisfy the following conditions: $\theta_{contraction,erp} > \theta_{expansion,erp}$, $\kappa_{contraction,erp} > \kappa_{expansion,erp}$, $\sigma_{contraction,erp} > \sigma_{expansion,erp}$. The first condition ensures that the P/E ratio will tend to decline when the economy goes into a recession. The second condition means that changes in P/E ratios are asymmetric in that they decline faster in recessions than they improve in expansions. The final condition gives more idiosyncratic volatility in the recession phase of the business cycle. This is consistent with the skewness of empirical equity returns where the left tail of equity returns is ‘fatter’ than the right tail and consistent with higher market volatility and uncertainty in recessions.

The Terminal Condition on Equity Valuations

In steady state the expected return on equity should equal the required rate of return on equity. We impose this condition to pin down the expected terminal valuation of equities, beyond the 30-year forecasting horizon at each ‘node’ in the simulation.

$$P_T = E_T / (r + \lambda)$$

The steady state condition can be derived as follows. The terminal expected price level is $P_T = d_T / (r + \lambda - g_D)$, where d is dividends, r is long-term interest rate, and roe is the return on equity. Now use the fact that $d_T = p_{0T} e_T$, $g_D = (1 - p_{0T}) roe$, and impose $roe = r + \lambda$ in equilibrium.

A Note on Expectations in the Model

The complicated regime-switching dynamics of the model do not admit a closed-form solution for expected future fed funds rate, real GDP growth and inflation at any given ‘simulation node’ in the simulation. To deal with this issue, the expected future values of a specific variable are set to model consistent linear

58 Recreating Sustainable Retirement

functions (projections) onto the most relevant current state variables in the model. The expected future fed funds rate at a given point in time in a given regime is for instance based on the linear projection of future simulated realizations of the fed funds rate, the current fed funds rate, the current level of cyclical and persistent inflation, and the current output gap for a given set of model parameters. To get a term structure of expectations at a given point in time, the parameters for a given variable vary at different forecast horizons, which are one, two, three, four, five, seven, ten, 20, and 30 years. The expectations are in this way by construction ‘unbiased,’ conditional on the current state of the economy.

Notes

1. In practice, most modern implementations would extend to a broader array of asset classes (e.g. real estate, commodities, private equity, hedge funds, etc.). The reduced set of asset classes is chosen to reduce complexity of computation and to focus development of intuition.
2. Data are taken from Robert Shiller’s (2013) website for 1871–2010 and from Bloomberg (2013) thereafter.
3. In nominal terms, the corresponding average ranges are ten-year rolling: (1.7 percent)–20.4 percent; 30-year rolling: 5.2 percent–13.6 percent; Full Sample Average: 8.7 percent. On a cumulative basis using Shiller’s CPI series, nearly 60 percent of the cumulative inflation has occurred in the past 30 years.
4. Discussed in Moore (2011).
5. p is the probability if in the high state to stay in that state, $1-p$ is the probability of transition to the low state. Symmetric results hold for q with respect to the low and high states.
6. Schaus (2010) provides a thorough treatment of the design and motivations behind DC glide paths.
7. Furman (2013) cites sources at Morningstar in her January 13, 2013, *New York Daily News* article.
8. For the Multivariate Normal and Bootstrap models inflation is assumed to be 3 percent over the period. For the Nested Structural model, inflation is a stochastic process which drives yield curve dynamics and feeds into equity valuations. Here the process is calibrated to a long-run mean of 3 percent.
9. This analysis assumes zero funding balances (Prefunding Balance and Funding Standard Carryover Balance), which impacts the funding ratio and minimum required contribution.
10. Provided the plan is underfunded or 100 percent funded. If the plan is overfunded, the required contribution is $\text{Max}(\text{Normal Cost} - \text{Overfunding}, 0)$.

References

- Abidi, N., and D. Quayle (2010). ‘Fixing the Flaws with Target Date Funds,’ *Journal of Indexes*, 13(2): 46–51.
- Arnott, R. D., and D. B. Chavez (2012). ‘Demographic Changes, Financial Markets, and the Economy,’ *Financial Analysts Journal*, 68: 1–24.

- Bloomberg (2013). *S&P 500 Index* [website]. <<http://www.bloomberg.com/quote/SPX:IND>>.
- Campbell, J. Y., and L. M. Viceira (2005). 'The Term Structure of the Risk-return Tradeoff,' *Financial Analysts Journal*, 61(1): 34–44.
- Cochrane, J. H. (2011). 'Discount Rates,' *Journal of Finance*, 66(4): 1047–1108.
- Dimson, E., P. Marsh, and M. Staunton (2002). *Triumph of the Optimists*. Princeton, NJ: Princeton University Press.
- Furman, P. (2013). 'Target-date Funds as a Retirement Investment Option are Growing in Popularity, But Do Your Homework,' *New York Daily News*, January 13.
- Hamilton, J. D. (1989). 'A New Approach to the Economic Analysis of Non-stationary Time Series and the Business Cycle,' *Econometrica*, 57(2): 357–384.
- Hamilton, J. D. (1994). *Time Series Analysis*. Princeton, NJ: Princeton University Press.
- Ilmanen, A. (2011). *Expected Returns: An Investor's Guide to Harvesting Market Rewards*. Hoboken, NJ: Wiley Finance.
- Lo, A., and A. C. MacKinlay (1988). 'Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test,' *Review of Financial Studies*, 1(1): 41–66.
- Milliman (2013). 'Milliman Analysis: Pension Plans Off to a Roaring Start in 2013 as Funded Status Improves by \$106 Billion,' Milliman 100 Pension Funding Index [website]. <<http://www.milliman.com/expertise/employee-benefits/products-tools/pension-funding-index/pdfs/pfi-february-2013.pdf>>.
- Moore, J. (2011). *PIMCO Viewpoints: Prediction? Pain* [website]. <<http://www.pimco.com/en/insights/pages/prediction-pain.aspx>>.
- Pastor, L. and R. Stambaugh (2012). 'Are Stocks Really Less Volatile in the Long Run?' *Journal of Finance*, 67(2): 431–478.
- Poterba, J. M., and L. H. Summers (1988). 'Mean Reversion in Stock Returns: Evidence and Implications,' *Journal of Financial Economics*, 22(1): 27–60.
- Schaus, S. L. (2010). *Designing Successful Target-Date Strategies for Defined Contribution Plans*. Hoboken, NJ: John Wiley & Sons.
- Shiller, R. J. (2013). 'Shiller, R., U.S. Stock Price Data, Annual, with Consumption, both Short and Long Rates, and Present Value Calculations.' *Online Data Robert Shiller* [website]. New Haven, CT: Yale University Department of Economics. <<http://www.econ.yale.edu/~shiller/data.htm>>.
- Siegel, J. J. (1994). *Stocks for the Long Run: The Definitive Guide to Financial Market Returns & Long Term Investment Strategies*. New York, NY: McGraw-Hill.