SOCIETY OF ACTUARIES

Deep Learning for Liability-Driven Investment

June 2021

Deep Learning for Liability-Driven Investment

AUTHORS Kailan Shang, FSA, CFA, PRM, SCJP

SPONSORS SOA Committee on Finance Research

Caveat and Disclaimer

The opinions expressed and conclusions reached by the authors are their own and do not represent any official position or opinion of the Society of Actuaries or its members. The Society of Actuaries makes no representation or warranty to the accuracy of the information.

CONTENTS

Deep Learning for Liability-Driven Investment

Acknowledgments

The author would like to thank all members of the Project Oversight Group (POG) tasked with providing governance on this research project. This paper would not have attained its current level of relevance to practitioners without the POG's guidance, feedback and insightful input. The discussions on the challenges of LDI modeling were extremely useful for setting the direction and focus of this report.

"Deep Learning for Liability-Driven Investment" POG members are the following:

- Lindsay Allen
- David Cantor
- Naxine Chang
- Steven Craighead
- Abid Kazmi
- Zixiang Low
- Michael Niemerg
- Dennis Radliff
- Ronald Richman
- Steven Siegel
- Haofeng Yu
- Aolin Zhang

Gordon Irlam reviewed the paper and provided helpful suggestions. The author would also like to thank Barbara Scott and Steven Siegel for their effective coordination of this project, as well as the sponsorship and funding support of the Committee on Finance Research of the Society of Actuaries.

Glossary

Bellman Equation: An equation that describes the optimal value *V* given a state *s* as the sum of current payoff ($P(s,a)$) of optimal action *a* and the discounted value of the future optimal decisions in a recursive way. Here $T(s,a)$ is the new state given previous state *s* and action *a*:

$$V(s) = \max_{a}\{P(s,a) + \gamma V(T(s,a))\}$$

The Bellman equation is the function that defines the optimization problem in dynamic programming.

Deep Learning: A machine learning method that relies on artificial neural networks such as a fully connected neural network and long short-term memory (see term below) to represent the relationship between the response variable and explanatory variables in an approximated way. "Deep" refers to multiple layers of neurons usually required in deep learning to be able to have a good approximation of the relationship.

Dynamic Programming: An optimization method that breaks a multiple-period optimization problem into smaller problems that have a recursive relationship described by the Bellman equation. For example, a long-term investment strategy can be dissected into a series of quarterly rebalance decisions that depends not only on expected future conditions but also on the path that leads to the then-current state. The value of a rebalance decision is determined using backward induction starting from the end of the time horizon. To make it feasible for backward induction, simplified assumptions, especially for the liability, are usually used to limit the possible outcomes in the end and the number of possible paths of periodic decisions.

Fully Connected Neural Network: A type of artificial neural network in which neurons of adjacent layers are fully connected in a forward manner without any backward or feedback connections. It is also known as a feedforward neural network.

Grid Search: An optimization method that runs through all possible choices to determine their values and select the choice with the greatest value. It guarantees an optimal solution but is also the most computing intensive approach.

Long Short-Term Memory: A type of artificial neural network where neurons have feedback connections and can be used to learn long-term patterns. It is suitable for processing data sequences such as a time series.

Reinforcement Learning (RL): A type of machine learning to study the best action to take given a certain state in an environment to maximize the expected future reward. Usually the reward function is model free, which means it is unknown except observing the immediate reward. This is different from traditional dynamic programming where the expected future reward of each action can be calculated using backward induction. An example is AlphaGo Zero, which uses RL to study the ancient Chinese game of Go. It starts with basic rules of the game and keeps adjusting its strategy by playing the game many times and learning from the impact of each move.

Semi-supervised Learning: A learning process to study the relationships between the response variable and explanatory variables without full knowledge of the response variable. It can be the case that only part of the dataset has labeled data for the response variable. In this research, it refers to the situation where only part of the response variable is known and used to learn the relationships.

Supervised Learning: A learning process to study the relationships between the response variable and explanatory variables. The resulting model can be used to estimate the value of the response variable given the value of the explanatory variables. Linear regression is an example of a supervised learning model.

Executive Summary

Liability-driven investment (LDI) is a key investment approach adopted by insurance companies and defined benefit (DB) pension funds. However, the sophisticated liability portfolio and volatile asset market make strategic asset allocation very challenging. The optimization of a dynamic asset allocation strategy is difficult to achieve with dynamic programming, whose assumption of liability evolvement is often too simplified. On the other hand, using a grid-searching approach to find the best asset allocation or asset allocation path is too computationally intensive even with a few asset classes.

Artificial intelligence is a promising approach to address these challenges. This report uses deep learning models and reinforcement learning to construct a framework for learning the optimal dynamic strategic asset allocation plan for LDI. It designs a stochastic experimental framework of the economic system in which the program finds the optimal strategy by trying different and nonconstant asset allocation plans from time to time.

Figure E.1
LDI Experiment Design

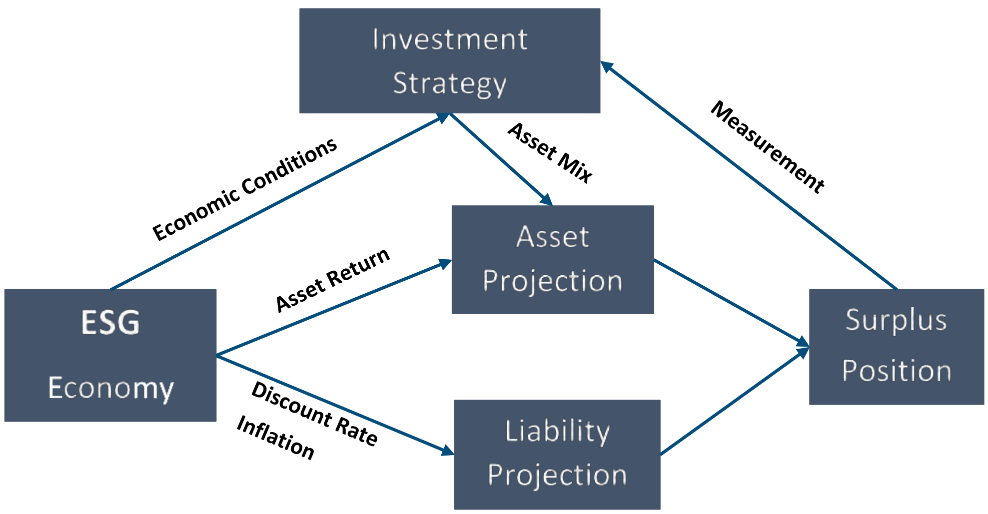

Deep learning models are trained to approximate the long-term impact of asset allocation on surplus position. Reinforcement learning (RL) is used to learn the best strategy based on a specified reward function. RL is a forward learning algorithm, which means that only the current impact is observable until the end of the experiment, and it is a semi-supervised learning. However, with enough experiments, RL is expected to learn the patterns and figure out good strategies to maximize the reward. The required calculations, however, are expected to be much less than what dynamic programming requires. The dynamic investment strategies will then choose the action that maximizes the reward at each decision point in each scenario. However, these strategies may not be the optimal ones but suboptimal ones because RL does not walk through the entire space of possible asset allocation paths.

To evaluate the effectiveness of RL compared to traditional strategic asset allocation methods, a sample DB plan is modeled with economic scenario generation, dynamic liability projection, asset allocation and projection and surplus projection. The comparison between optimal static investment strategies and RL-based dynamic strategies is performed assuming two asset classes: an AA-rated corporate bond portfolio and large-cap public equity. Efficient frontiers are built assuming fixed time horizons and static investment strategies. Both fully connected neural networks and long short-term memory models are used to approximate the reward function in the RL process. The resulting dynamic investment strategies show that RL is able to generate reasonable investment strategies and the potential for generating better risk and return tradeoff than optimal static strategies that have a target time horizon. The RL-based dynamic investment strategy also adjusts the asset allocation to be riskier with worse starting funding status.

The improvement of the risk return tradeoff is at the distribution level rather than at each individual scenario level. This means that RL is not trying to gain through market timing in this specific example, but rather to adjust asset allocation based on the then-current funding status and economic conditions.

With more asset classes coming into play, it is even more challenging for the traditional approach of LDI strategy optimization because of computational requirements. A test using four asset classes including AA-rated bonds, BBB-rated bonds, large-cap public equity and real estate investment trusts (REITs) is performed without updating optimal static investment strategies and efficient frontiers. An RL-based investment strategy, with less required computing time, can further improve the risk return tradeoff achieved with only two asset classes.

Overall, this research contributes to the existing literature in three ways. First, it introduces RL and deep learning to actuaries and tries to build some connections with actuarial concepts. Second, it provides a workable example to demonstrate the application of RL to LDI and assess its effectiveness. Without simplifying any liability and asset modeling, it shows the potential of RL to solve complex actuarial problems. The author is not aware of any previous efforts made in this area. Third, sample implementation codes are made public for educational purposes and hosted at [GitHub - Society-of-actuaries-research-institute/FP198-Deep-Learning-for-Liability-Driven-Investment].

Section 1: Introduction

Pension and insurance are important pillars that support the security of families and human society. They provide long-term guarantees of our future life. The primary investment goal of pension plans and insurance companies is to make sure that the asset value is no less than the liability value in most if not all circumstances so that long-term guarantees can be realized. This is different from many other organizations or individuals whose focus is more on asset growth.

Liability-driven investment (LDI) strategies are often used by insurance companies and defined benefit (DB) pension plans to control the mismatch between asset and liability while earning an acceptable return. However, the complexity of the liability portfolio often makes asset allocation optimization a difficult task. Both liability cash flows and liability value are sensitive to economic conditions. For example, pension benefits may be linked to the inflation rate. Lapse rates may be affected by household financial status and therefore the economic environment. Liability discount rates are affected by interest rates and credit market conditions.

The existing literature often makes simplified assumptions about the evolution of the liability portfolio. Sharpe and Tint (1990) assume that the liability portfolio evolves with a liability return that follows a Normal distribution. Asset allocation plans are determined using surplus optimization based on surplus return. The risk measure is the volatility of the surplus return. Ang et al. (2013) change the risk measure from variance to downside risk only. The optimal solution to the aggregate maximization is a weighted portfolio including both a mean-variance-efficient portfolio and a liability-hedging portfolio. Martellini (2006) uses the utility function of the funding ratio at a specified future point as the target for maximization. Like in Ang et al. (2013), the optimal asset allocation plans consist of a mean-variance-efficient portfolio, a liability-hedging portfolio and risk-free assets.

The determination of the optimal LDI strategy is often a static optimization exercise, assuming a stable development of the asset and liability portfolio in a simplified way. A target asset allocation plan, normally referred as the strategic asset allocation, is often set as the long-term plan. This approach faces two challenges:

- Economic development experiences structural changes from time to time, which makes the long-term plan not long term.
- The evolvement of the liability portfolio makes a fixed long-term asset allocation plan suboptimal.

Dynamic optimization that generates a nonfixed strategy of asset allocation is a better approach to reflect changes in economic conditions and liability portfolio in theory. Literature on applying dynamic programming to LDI exists, although the focus is more on retirement planning to maximize the utility of wealth. Brandt et al. (2005) use dynamic programming to solve the optimal dynamic asset allocations, which are path dependent. Geyer and Ziemba (2008) presented a dynamic programming approach to maximize the expected present value of terminal wealth. However, to arrive a tractable solution, simplified assumptions are usually needed, which make the strategy less useful in practice. Even though asset returns may follow some complicated processes, the liability is assumed to be fixed or follow certain nondynamic patterns.

Although dynamic programming may fail to incorporate liability details, methods used in artificial intelligence (AI) can be adjusted to handle the complexity. Reinforcement learning together with deep learning models has proved successful in strategy optimization. For example, AlphaGo Zero uses RL to study the ancient Chinese game of Go, as described in Silver et al. (2017). It starts with basic rules of the game without data from human games, and it is an improved version of AlphaGo that defeated a world champion of the original Go game. If the dynamic LDI problem can be defined in a similar way, AI models can be used to find the optimal or suboptimal investment strategy.

This study explores the possibility of using AI techniques to improve investment decision-making for pension funds and insurance companies. It develops an experiment that captures the evolution of asset portfolio and liability portfolio using actuarial models. Reinforcement learning is then used to learn the dynamic asset allocation strategy by running multiple experiments. At the same time, deep learning models are trained to link decision factors (economic conditions, liability portfolio status, risk appetite etc.) and investment choices with the outcome of whether assets are sufficient to support liabilities.

The report proceeds as follows:

- Section 2 (LDI Problem) defines the LDI problem to be solved and designs a dynamic experimental environment for RL.

- Section 3 (Reinforcement Learning Model) discusses the structure of RL and deep learning models that can be used to learn the optimal dynamic investment strategy through experiments.
- Section 4 (Dynamic LDI Example) describes an example of applying AI models to find a dynamic LDI strategy for a DB pension plan.
- Section 5 (Further Development) discusses potential extensions of this research.
- Section 6 (Conclusion) summarizes the key points of this research and concludes the main body of the report.

Section 2: LDI Problem

Both insurance and pension are important components of Social Security. The purpose of the investment activities of insurance companies and pension funds is not to achieve the highest asset return but to maximize the performance of assets relative to liability, such as maximizing a surplus or pension fund funding ratio.

2.1 LDI Analytic Framework

The complexity of the liability portfolio and its interaction with the asset portfolio make investing a very difficult task. For a full-blown LDI analysis, it is necessary to build a framework that maintains consistency between the asset side and the liability side. It is necessary to understand the key drivers of the liability risk, which are the following:

- *Discount rate:* To evaluate the liability portfolio, expected future benefit payments are discounted back using a discount curve that is highly linked with the interest rate level in the market. This is especially important for long-term payments because the impact of discounting is material. Interest rates are driven by factors including but not limited to the economic environment and monetary policy. These factors affect asset performance as well.
- *Inflation:* Inflation could affect the liability value because certain liability payments such as pension benefits and health insurance benefits may be adjusted according to inflation. Inflation measures could be the consumer price index, wage inflation, cost-of-living adjustments or medical expense inflation. These inflation measures are affected by the general economic environment and have co-movements with interest rates as well.
- *Claim rate:* Depending on the type of insurance and pension, the claim rate is the probability that an insurance claim or pension benefits will be paid. The claim rate could be the mortality rate, morbidity rate, longevity rate, accident rate or other types, depending on the type of insurance product or pension plan. Most of these claim rates are independent of the discount rate and inflation. Claim rates of certain products such as unemployment insurance are related to the economic environment.
- *Policyholder behaviors:* Many insurance products are long-term products on which policyholders' behaviors can make an enormous difference. Lapse rate, policyholder loan usage and premium payment amount are affected by householder financial conditions and are related to the economic environment to a certain degree.

Given that asset performance depends on the economy as well, the asset side is highly linked to the liability side. However, the relationship is not linear. The level of dependency usually increases in extreme events. For example, in a financial crisis, low interest rates and therefore low liability discount rates, high credit spread/default, a bear stock market and low inflation rate usually happen together. This creates adverse impacts on both the asset side (bad asset performance) and the liability side (increase in liability value). Modeling the relationship between assets and liabilities using correlation coefficients, as in many traditional dynamic optimization models, cannot capture the tail risk effectively.

Some standard LDI strategies focus on matching the sensitivity of the asset portfolio and liability portfolio to key economic risks such as the interest rate, inflation rate and foreign exchange rate. Duration matching, duration and convexity matching, cash flow matching and liability replication portfolio are all approaches to immunize the surplus from changes in interest rates. However, these approaches often focus on the current status, without fully acknowledging the evolution of the liability portfolio.

Fully matching the asset portfolio with the liability portfolio is also a costly and impractical approach of investment. Although plan sponsors or insurance companies are protected from economic risks to a high extent, the cost is high because most of the asset investments will be on fixed income securities with high credit quality. The expected return is lower than that from riskier investments such as equity and real estate. Institutions will usually trade part of the security with expected high returns from risky investments. The desired hedging ratio is determined by the institution's risk appetite.

A dynamic optimization method is needed that fully reflects the complexity of the liability evolvement, the nonlinear relationship between asset and liability and an institution's risk appetite. This research tries to solve such a complicated optimization problem using AI techniques.

2.2 Experiment

Unlike empirical analysis that relies on historical data, AI models need to be trained with many examples. Historical data present only one scenario and are not enough to train AI models. An experimental environment is needed for AI models to learn the optimal dynamic strategy. This environment can be thought of as a game with many possible paths and outcomes depending not only on the game setting but also on the strategies taken by players.

As shown in Figure 1, the experimental environment is composed of three elements: the economic scenario generator (ESG), asset and liability projection and investment strategy selection. The ESG generates the economic scenarios that govern the evolvement of the economy, asset returns and liability discount rates in a consistent way. Asset and liability projection models project future asset value, liability value and surplus value based on the economic scenario and the chosen investment strategy. The investment mix is chosen based on the economic conditions specified by the economic scenario, surplus position and a dynamic investment strategy. The investment strategy is assessed based on the resulting surplus positions and improved by running thousands of experiments.

Figure 1

LDI Experiment Design

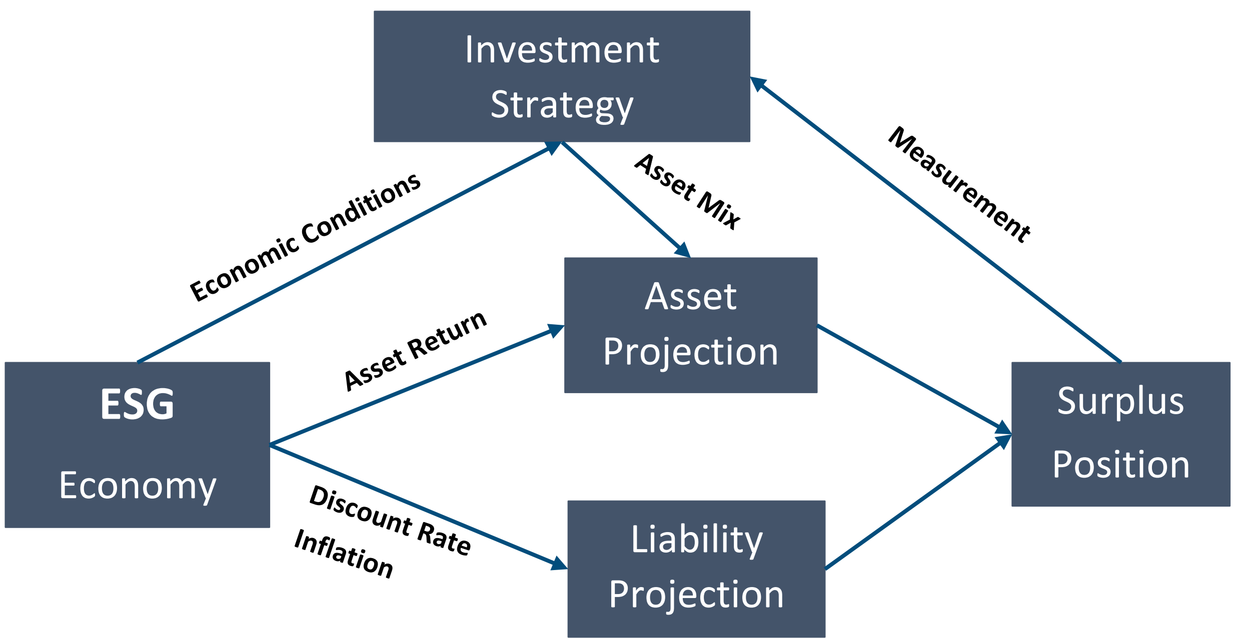

2.2.1 Economic Scenario Generation

The ESG in the experiment is not confined to any specific forms but is expected to follow a few principles:

1. The ESG needs to provide a holistic view of the economy, including major economic factors such as real gross domestic product (GDP) growth rate, inflation rate, unemployment rate, interest rate, credit spread, consumption, investment, government spending and international trading. These factors are linked in a way that represents historical patterns or forward-looking patterns that are expected to happen. The ESG generates real-world scenarios rather than risk-neutral scenarios.
2. The ESG needs to reflect a nonlinear relationship among economic factors and asset returns. For example, during economic recessions, a low interest rate, higher credit default and credit spread, high unemployment rate, lower consumption, investment and economic growth rate, and more volatile capital market usually happen together. This tail risk can be easily lost when all data are analyzed together assuming linear relationships.

3. Rather than calibrated solely based on asset market prices, asset return scenarios may be explained by economic fundamentals such as economic growth, consumption, investment, labor market shocks, monetary policies or fiscal policies. For investment strategy analysis, understanding the causes of a scenario is important because people can assess whether the scenario is reasonable for investment decision-making.
4. The cyclical pattern of the economy needs to be considered as well. Economic cycles are an important part when designing investment strategies. A market shock as in the 2008 financial crisis might be significant. But with a diversified equity portfolio and a long enough time horizon, the losses could be recovered. Unlike stochastic scenarios used in capital management that usually have a one-year time horizon, the evaluation of LDI strategies may take a much longer period covering an entire economic cycle.

Figure 2 shows the structure of a sample ESG that can be used in such an experiment.

Figure 2
Sample ESG Structure

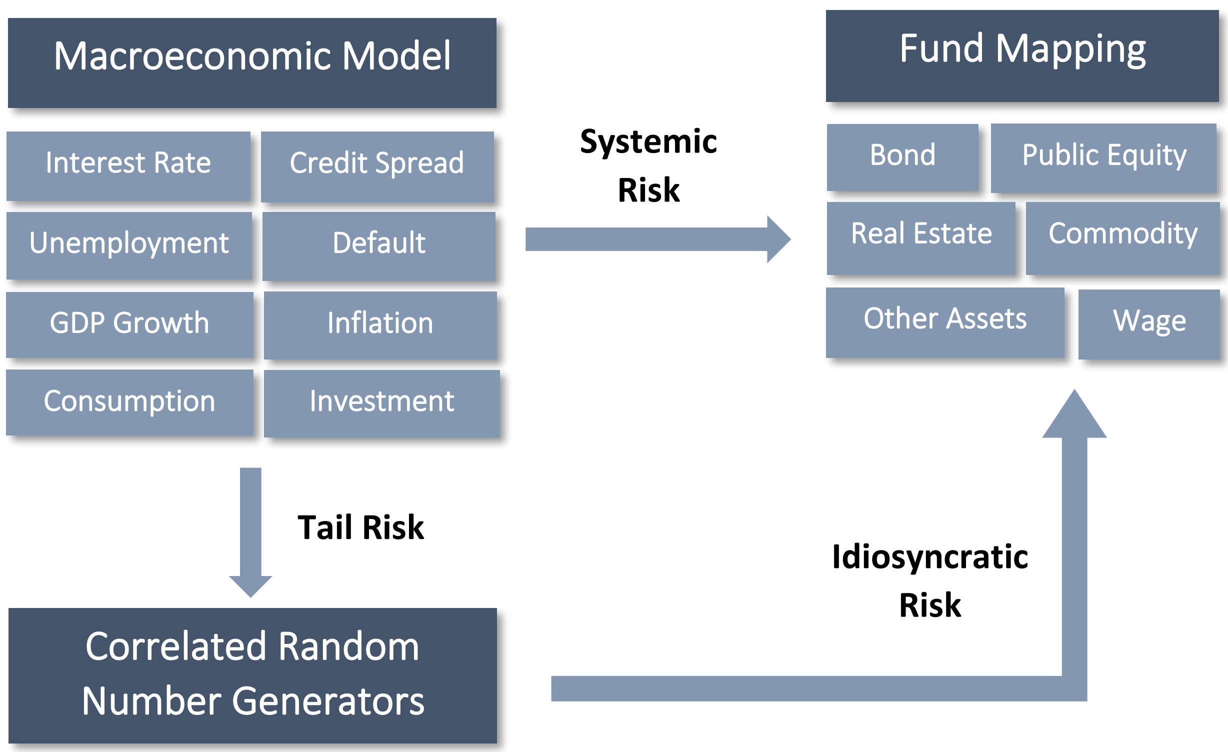

2.2.2 Asset and Liability Projection

The evolvement of an asset portfolio and liability portfolio needs to be predicted based on the generated scenarios and chosen investment strategy. The evolvement of the asset portfolio needs to consider the following elements:

1. *Asset mix of the portfolio:* This is subject to change from period to period according to the investment strategy. It involves buys and sells to rebalance the asset portfolio.
2. *Reinvestment strategy:* Cash flows could come from both the asset portfolio and the liability portfolio. Examples of cash flows include coupon payments, redemption, dividend income, premium, pension contributions, claim payments and pension benefit payments. Positive net cash flow needs to be reinvested. It may follow the asset mix or a different mix that accommodates other concerns such as liquidity.
3. *Sell orders:* When the net cash flow is negative, existing assets need to be sold to meet liability payment requirements. The selling orders of assets could be simply a proportional selling of all existing assets. It may also consider the liquidity of assets and market conditions and start from the most liquid assets.

4. *Asset value:* This is revaluated periodically based on changing economic and capital market conditions defined by economic scenarios.

The liability portfolio evolves with cash flows and revaluation:

1. Cash inflows include premium income and pension contribution. They may be linked with systemic risk factors in the ESG. For example, fewer premium payments are expected during economic recessions when customers need to reduce or terminate premium payments to meet other financial challenges. Pension contributions may be affected by wage inflation, although other factors such as changes in regulation, tax policy and company performance usually have a larger impact.
2. Cash outflows include benefit payments, policy lapses and policy loans. They are affected by systemic risk as well. For example, certain benefits such as dividend payments of participating products are adjusted according to investment performance, which is affected by systemic risk. Policy lapses may increase in two opposite situations: getting policy surrender value during financial distress to pay living expenses, or buying new investment products with higher return in the presence of interest rate spikes.
3. Liability value is revaluated periodically based on changing economic and capital market conditions defined by economic scenarios. The impact comes from changes in both discount rates and expected future cash flows.
4. Liability portfolio evolvement is affected by insurance risk such as mortality risk, morbidity risk and longevity risk. In most cases, they are independent from, if not marginally linked to, economic risk.

With the asset and liability projection, the surplus amount (Asset − Liability) or funding ratio (Asset/Liability) can be projected and used to evaluate different investment strategies.

2.2.3 Investment Strategy

The investment strategy determines the asset mix used for asset portfolio rebalancing in a dynamic way. It considers economic conditions, portfolio surplus position and risk appetite to choose the appropriate asset mix. As the economic scenario evolves, the desired asset mix changes as well. The optimal dynamic investment strategy is the target for learning. It is evaluated by the resulting surplus position, represented as either the surplus amount or funding ratio. With the same plan contributions, a higher surplus or funding ratio means a better strategy, as long as it does not cause too much downside risk that is inconsistent with the risk appetite.

The experimental environment discussed above can reflect more complicated and nonlinear relationships than traditional dynamic programming models. Although it brings the model closer to reality, solving the optimal investment strategy is a difficult task.

3. Reinforcement Learning Model

To solve the optimal dynamic investment strategy, a reinforcement learning (RL) model can be used by running thousands of experiments. It starts by trying a random asset mix to see the resulting surplus position and gradually adjusts the asset mix based on past experimental experience. As it gains more and more experience, the model can generate better asset mix suggestions. The learning problem is formalized in Section 3.1.

3.1 Reinforcement Learning

The goal is to find the optimal dynamic investment strategy $\pi^*(s)$ based on *s*, the states that decision-makers can observe at the time of decision-making. The states are economic conditions and surplus position. The optimality of an investment strategy is defined as the one that maximizes the reward function $Q^*(s,a)$ determined by *s*, the states, and *a*, the rebalancing action determined by the strategy $\pi^*(s)$:

$$\pi^*(s) = \max_a Q^*(s,a)$$

The reward function $Q^*(s,a)$ is difficult to define. The impact of an asset mix selection not only affects the current period performance but could also have long-lasting impact on the future surplus position. For example, an inappropriate asset mix may cause a deep deficit that is unlikely to be recovered solely based on future higher asset returns. Impacts on future periods are difficult to estimate given the uncertainty of future economic conditions and surplus conditions. Changes in surplus position may not always work as well. Starting from a deficit position, even when the asset return is higher than the liability return, the deficit could still increase because the starting liability value is larger than the asset value. Instead of defining the reward function directly, it can be constructed in a recursive form:

$$Q^\pi(s,a) = r + \gamma Q^\pi(s', \pi(s'))$$

where

r: current period reward that can be observed; for example, it could be defined as changes in the surplus amount or funding ratio

γ: discount factor to reflect the timing difference and

$Q^\pi(s', \pi(s'))$: reward function in the next period with new states s' and new asset rebalance action $\pi(s')$.

The reward function $Q(s,a)$ is a nonlinear function that explains nonlinear relationships. *Q* is the sum of the current period reward and discounted rewards in the future, assuming that future optimal actions will be taken based on *Q*. It is structured as the Bellman equation used in dynamic programming to describe the value given state and action, except that unlike dynamic programming, the exact function in RL is not defined but learnt by minimizing the error δ between the two sides of the recursive equation:

$$\delta = Q(s,a) - (r + \gamma \max_{a'} Q(s', a'))$$

RL is a forward process where only the current reward can be observed, and all future rewards can be estimated based on only the existing *Q*. However, traditional model training needs the true value of *Q*, which is unknown. To address this calibration challenge, RL tries to match the current reward *r* and the change in *Q*, measured as $Q(s,a) - \gamma \max_{a'} Q(s', a')$. By doing that, the observed current reward can be used to keep improving the accuracy of *Q*. Using deep learning models such as feedforward neural networks and long short-term memory models, the actual reward function can be approximated without needing to set the exact function form before model training. With enough experiments, in theory, the calibrated *Q* function will be good enough to advise a good, if not optimal, action.

Figure 3 shows the RL training process. The experimental environment is used to generate different states so that the model can try different investment strategies and find the optimal strategy. The asset rebalance action that has the highest expected reward is chosen and used to determine its impact on current period performance. After trying more and more scenarios, the deep learning model that represents the reward function is updated and is expected to move closer to the real reward function.

Figure 3
RL Process for LDI Strategy

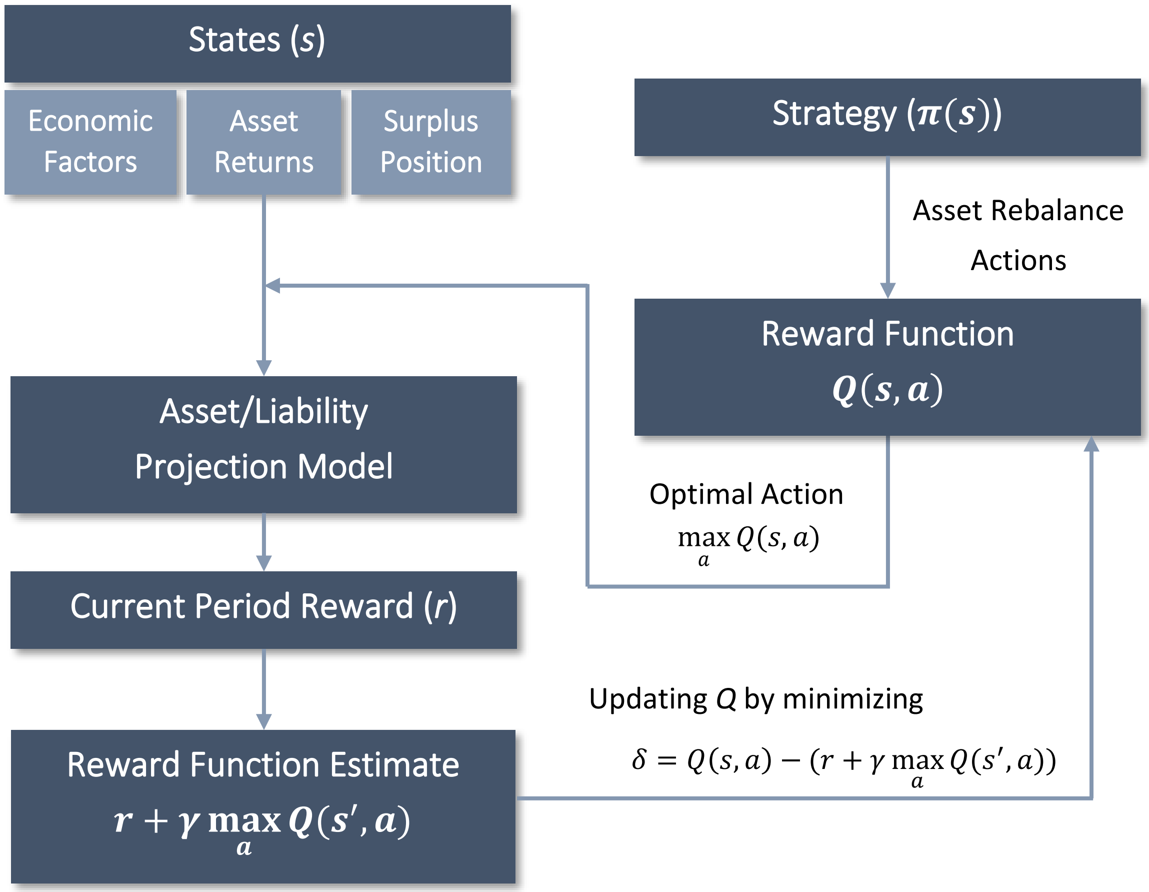

The RL process described above is also known as deep Q-learning, in which deep learning models are used to represent the reward function and the Q value is learnt and optimized with respect to action a. Many variants of RL are available and are discussed in Appendix A.

3.2 Fully Connected Neural Networks

The reward function $Q(s,a)$ is learnt from experiments without knowing its exact form. Using a generic model that can approximate any functions can avoid the burden of specifying the form of the function. This is especially useful when the relationships are complicated and difficult to represent with a concise mathematic function. Deep learning models, such as fully connected neural networks (FCNNs), can be used to represent the reward function in RL (see Figure 4).

Figure 4
Fully Connected Neural Network

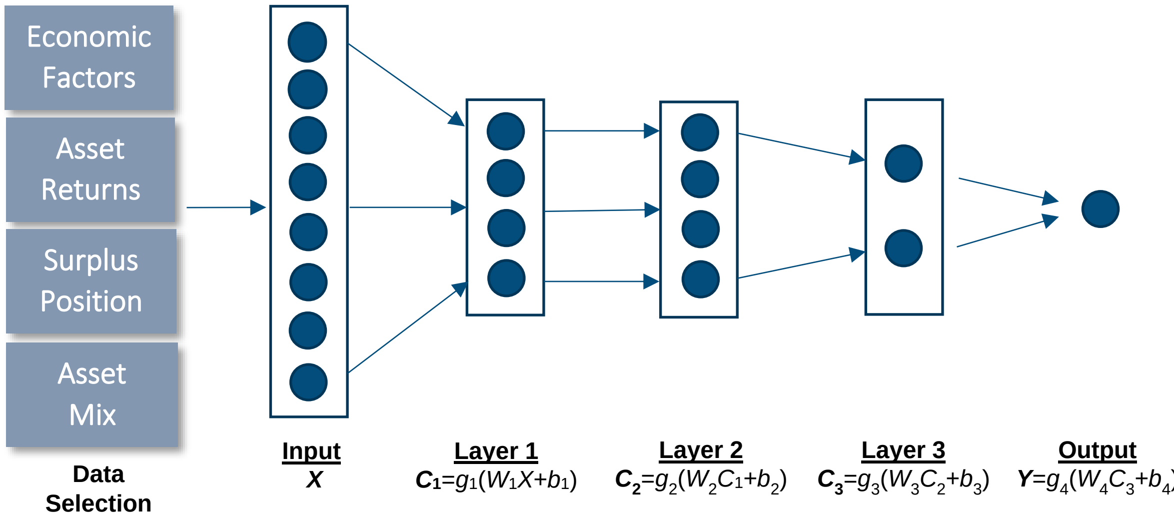

C_i : the value of neurons in hidden layer i

g_i: the activation function for hidden layer i. g_4 is the function for the output layer. A common activation function is the sigmoid function $\frac{1}{1+e^{-x}}$ (a.k.a. logistic function). Many other activation functions are available as well.

The neural network can still be represented as a function of the input X. The function is not a simple linear function, a polynomial function, a generalized linear function or other nonlinear function but is a few linear layers (*WX+b*) and nonlinear layers (activation function) stacked together:

$$Y = f(X; W_1, b_1, W_2, b_2, W_3, b_3, W_4, b_4)$$

3.3 Long Short-Term Memory

Another choice of reward function is long short-term memory (LSTM) models, which have been widely used in text mining. A fully connected neural network treats all the inputs as independent variables, which means no specific relationship among inputs is assumed. For LDI strategy decision-making, it is unlikely that only the current status matters. In other words, it is not a standard Markov process with only the current status that matters. The path leading to the current status is useful information as well. For example, a GDP growth rate of 2.5% may be misleading without knowing previous GDP growth rates. It could have increased from 1% or dropped from 4%, which tells totally opposite stories about the economy. If previous statuses are used together with the current status as inputs for the reward function, it makes sense to tell the model that they are time series.

An LSTM model is composed of a sequence of standard modules. Each module corresponds to a step in the sequence. A step can be considered as a time point when the asset portfolio needs to be rebalanced in LDI analysis. Figure 5 shows the structure of an LSTM model.

Figure 5
LSTM Structure

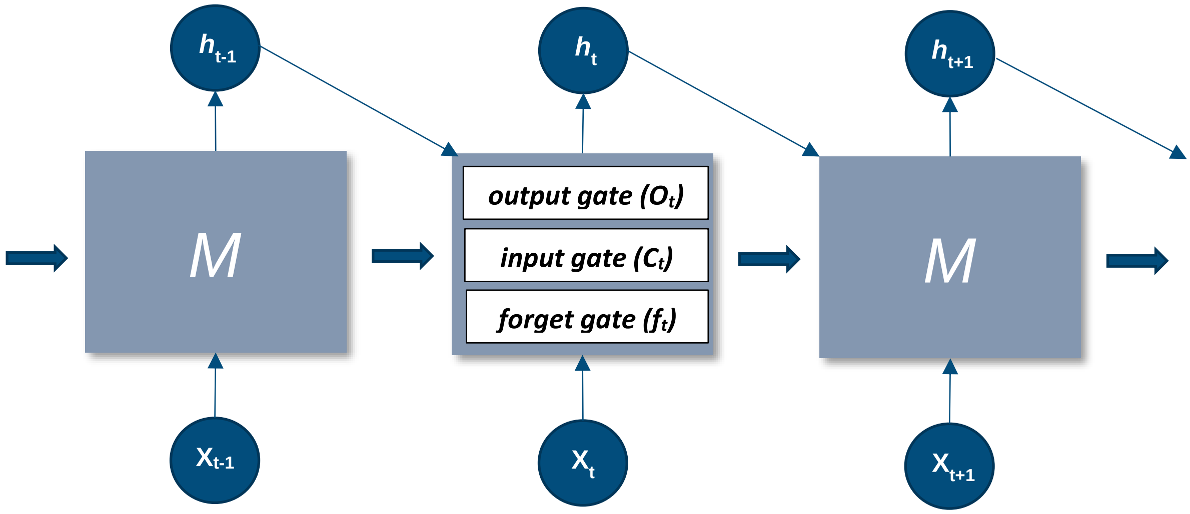

Each module *M* has the same structure. At time *t*, input vector X_t contains states including current economic conditions, asset returns and surplus position. The module consists of three gates, which are similar to layers in neural networks:

1. Forget gate $f_t = \sigma\left(W_f \cdot [X_t \quad h_{t-1}] + b_f\right)$. It defines the information to be forgotten.
2. Input gate $C_t = f_t \cdot C_{t-1} + i_t \cdot \tilde{C}_t$
   where
   $i_t = \sigma(W_i \cdot [X_t \quad h_{t-1}] + b_i)$
   $\tilde{C}_t = \tanh(W_C \cdot [X_t \quad h_{t-1}] + b_C)$
   It is the information to be kept. It calculated as a weighted average of the old kept information C_{t-1} and the new information $\tilde{C}_t$.
3. Output gate $o_t = \sigma(W_o \cdot [X_t \quad h_{t-1}] + b_o)$. The hidden state $h_t = o_t \cdot \tanh\,(C_t)$ is then calculated and used for the next module at time *t*+1.

LSTM models can maintain short-term memory for a long time, which may be helpful for recognizing and maintaining time-related patterns such as economic cycles. Fully connected neural networks allow only forward connections between adjacent layers. LTSM models include a sequence of hidden states that can contain information from nonadjacent earlier layers. LSTM models are also less insensitive to the length of time lags, which is good for modeling cycles with an uncertain length.

Fully connected neural networks and LSTM models are just two options to represent the reward function. As long as the models are flexible enough to approximate unknown reward functions, they can be used in RL.

4. Dynamic LDI Example

To evaluate AI models in learning dynamic optimal investment strategies, a simple DB pension plan is used to test the models' effectiveness. The example is built purely for testing and illustrating the AI framework described in Section 3. Results are intended to show the effectiveness of the AI framework instead of providing investment suggestions. Different liabilities, current asset mix, investment goals and views of the future economy and capital markets could change assumptions in the analytical framework and generate different conclusions. However, they are not included in this paper but reserved for future research.

To avoid reinventing the process of economic scenario generation and dynamic asset liability projection, the benchmark model in Shang and Hossen (2019) is used. The model includes a comprehensive and complicated system of generating fundamental economic factors, determining asset returns based on fundamental economic factors and idiosyncratic risk factors, projecting assets and liabilities dynamically, and measuring the success of investment strategies.

By reflecting real-world economic patterns, the benchmark model provides a good experimental environment for deep learning models. Deep learning models can try different investment strategies and learn the optimal rules based on the resulting financial outcomes predicted by the benchmark model. A simplified version of the benchmark model is used for this research and described below. In addition, the EXC EL tool build by Shang and Hossen (2019) is transformed into Python programs to train deep learning models easily.

4.1 Simplified LDI Benchmark Model

The LDI benchmark model is composed of three elements: economic scenario generation, plan projection and investment strategy evaluation. It can project the funding status of a pension plan given an asset allocation plan and an economic scenario. The majority of Section 4.1 is an excerpt of Section 4 in Shang and Hossen (2019), except for the static optimization part, which will be replaced by dynamic strategies optimized using RL. It is included here so that this report is self-contained. Figure 6 shows the structure of the simplified model.

Figure 6
Simplified LDI Benchmark Model Structure

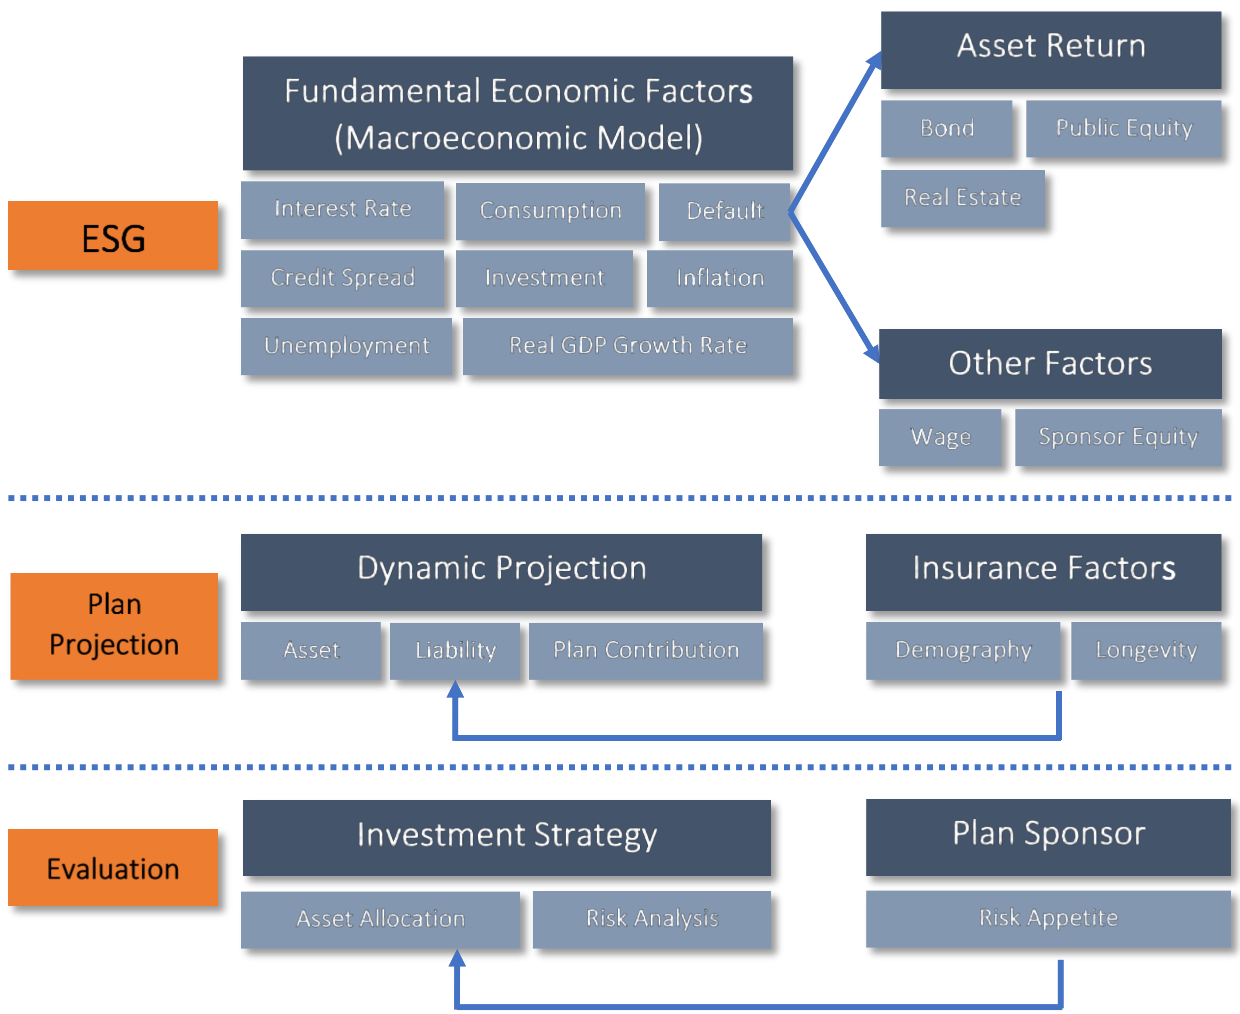

4.1.1 Economic Scenario Generation

The LDI benchmark model starts from an economic scenario generator, which generates fundamental economic factors, asset returns and wage inflation. The economic scenario generator relies on macroeconomic models to represent the economic system.

Fundamental economic factors such as GDP growth rates, interest rates, credit defaults, credit spread, unemployment rate and inflation rate can be chosen to govern the systemic risk in the economic scenario generator. These fundamental economic factors can be considered the roots of our economic and financial system. They define the current economic status and predict the future status, whether the economy is in expansion or recession. Macroeconomic models study both the contemporary and intertemporal relationships among fundamental economic factors to predict the future. They ensure the consistency and reasonableness of generated economic scenarios.

A vector autoregression (VAR) model is used to describe the intertemporal relationships among eight fundamental economic factors for the U.S. economy: the real GDP growth rate, inflation rate, unemployment rate, short-term interest rate, long-term interest rate, credit spread, consumption and investment. These fundamental economic factors govern the period-to-period changes in the status of the economy and indirectly affect the asset returns and other model outcomes in a consistent way:

$$\mathbf{F}_{\boldsymbol{t}} = \mathbf{c} + \mathbb{A}_{\mathbf{1}}\mathbf{F}_{\boldsymbol{t}-\mathbf{1}} + \mathbf{e}_{\boldsymbol{t}}$$

where

$\mathbf{F}_{\boldsymbol{t}}$ is a column vector with fundamental economic factors at time t or during period t
$\mathbf{c}$ is a column vector representing the constant terms of the fundamental economic factors
$\mathbb{A}_{\mathbf{1}}$ is a square matrix containing the model parameters describing the linear dependence of fundamental economic factors and
$\mathbf{e}_{\boldsymbol{t}}$ is a column vector to store the error terms of fundamental economic factors that cannot explained by linear models.

With the fundamental economic factors, asset returns can be generated through the process of fund mapping. The return of each asset class is generated as the sum of two factors: a systemic factor and an idiosyncratic factor. The systemic factor is determined based on the relationship between the asset return and the general economy governed by fundamental economic factors. The idiosyncratic factor is determined by the unique features of each asset class. The systemic factor can be determined based on historical data or adjusted with forward-looking views of the future economy. The idiosyncratic factor is independent of the systemic factor on average but usually has higher volatility and nonzero correlation with the systemic factor during economic recessions.

In addition to asset returns, some other factors, such as wage inflation and the plan sponsor’s equity return, may be important for LDI analysis. Wage inflation is important for projecting plan liability and benefit payments, which are usually linked to the future wage level. The plan sponsor’s equity return is a useful indictor of whether sufficient plan contributions will be available in the future. These factors can be generated using the same fund-mapping process as is used for asset returns.

Both contemporary and intertemporal relationships are represented by the linear functions

$$\mathbf{y}_{\boldsymbol{t}} = \boldsymbol{\alpha} + \boldsymbol{\phi}_{\mathbf{1}}\boldsymbol{y}_{\boldsymbol{t}-\mathbf{1}} + \boldsymbol{\phi}_{\mathbf{2}}\mathbf{y}_{\boldsymbol{t}-\mathbf{2}} + \mathbb{B}_{\mathbf{0}}\mathbf{F}_{\boldsymbol{t}} + \mathbb{B}_{\mathbf{1}}\mathbf{F}_{\boldsymbol{t}-\mathbf{1}} + \mathbb{B}_{\mathbf{2}}\mathbf{F}_{\boldsymbol{t}-\mathbf{2}} + \mathbf{e}_{\boldsymbol{t}}$$

where

$\mathbf{y}_{\boldsymbol{t}}$ is a column vector containing the returns of all asset classes and other factors during period t
$\boldsymbol{\alpha}$ is a column vector containing the constant terms of all asset classes
$\boldsymbol{\phi}_{\boldsymbol{i}}$ is a column vector containing parameters to govern the relationship between the current return and the return for all asset classes during period $t-i$
$\mathbb{B}_{\boldsymbol{i}}$ is a matrix that contains all asset classes’ model parameters for the fundamental economic factors during period $t-i$ and
$\mathbf{F}_{\boldsymbol{t}}$ is a column vector including all the fundamental economic factors at time t or during period t.

Model outcomes used in this report include the U.S. Treasury bond yield curve (for one-, two-, three-, five-, seven-, 10-, 20- and 30-year yields); the AAA-, AA-, A- and BBB-rated corporate bond credit spread and default rate; large-, mid- and small-cap equity index dividend yields and capital returns; high-dividend equity index dividend yields and capital returns; equity, mortgage and REIT cap rates and capital returns; wage inflation; and plan sponsor equity return.

The last step of the economic scenario generation process is to project bond fund returns based on bond fund investment strategy, term mix and reinvestment strategy. Simulated yield curve movements, defaults and credit-rating migration need to be reflected when projecting the bond fund returns.

Both the VAR(1) and linear models are calibrated using historical data from 1991 to 2016, where available. The ESG and calibration have been fully vetted in Shang et al. (2019).

4.1.2 Liability Model

With the economic scenario generator ready, the next part is plan projection. Future pension benefit payments can be projected based on demographic details of plan participants (age, gender, occupation, length of service, expected retirement age, current salary etc.), mortality assumptions, wage inflation, benefit rate, cost-of-living adjustment (COLA), COLA limit, lump sum payment option and other factors. Except for wage inflation and COLA, which are linked to the model inflation rate, all other assumptions are deterministic. Nonfinancial risk such as mortality risk and demographic changes can have a large impact on liability as well. However, they are likely to be addressed by insurance risk pooling and transferring rather than investment strategies. These risks are not considered in this research. With the projected benefit payments, the liability value (projected benefit obligation) can be calculated as

$$L_t = \sum_{i=1}^{\infty} \frac{B_{t+i}}{(1 + \mathrm{DR}_{t,i})^i}$$

where $\mathrm{DR}_{t,i}$ is the discount rate applicable to benefit payments due i time periods after the liability value is determined, at time t. In this report, it is the AA-rated i-year corporate bond spot rate at time t.

The liability value is projected dynamically. At the end of each period, the accrued benefit payments will increase because of extra service accrued during the period for active plan participants. Accordingly, contributions will be made to the plan to fund the extra benefit payments as follows:

$$C_t = \sum_{i=1}^{\infty} \frac{\Delta B_{t,i}}{\left(1 + \mathrm{CR}_{t,i}\right)^i} [1 + (\mathrm{SER}_t - \overline{\mathrm{SER}}) \times \mathrm{adj}]$$

where

$\Delta B_{t,i}$ = extra benefit payments at time $t + i$ caused by the extra length of service during period t
$\mathrm{CR}_{t,i}$ = discount rate of term i at time t used to determine the amount of plan contribution (this could be the expected investment return based on actual asset allocation or the liability discount rate)
SER_t = plan sponsor's equity return during period t
$\overline{\mathrm{SER}}$ = expected plan sponsor's equity return and
adj = adjustment factor to reflect the impact of the sponsor's equity return on plan contributions (adjusted to mitigate the impact of financial stress experienced by the plan sponsor).

In practice, the discount rates are often set to be the expected rate of investment return for the purpose of determining the plan contribution. The scenarios generated by the economic scenario generator include volatilities of all periods until the end of the projection. During the dynamic projection, the future volatilities are unknown. The expected return needs to reflect the current economic state but exclude the unknown future volatilities. As shown in Figure 7, the outer scenario is the entire scenario, including the volatility of all periods. The inner scenarios are the expected scenarios at each time point. At time *t*, the then-current plan assets and plan liabilities are determined based on the outer scenario till time *t*. The inner scenario at time *t* is constructed with the expected future economic fundamental factors and asset returns till the end of projection at time T.

Figure 7
Dynamic Projection

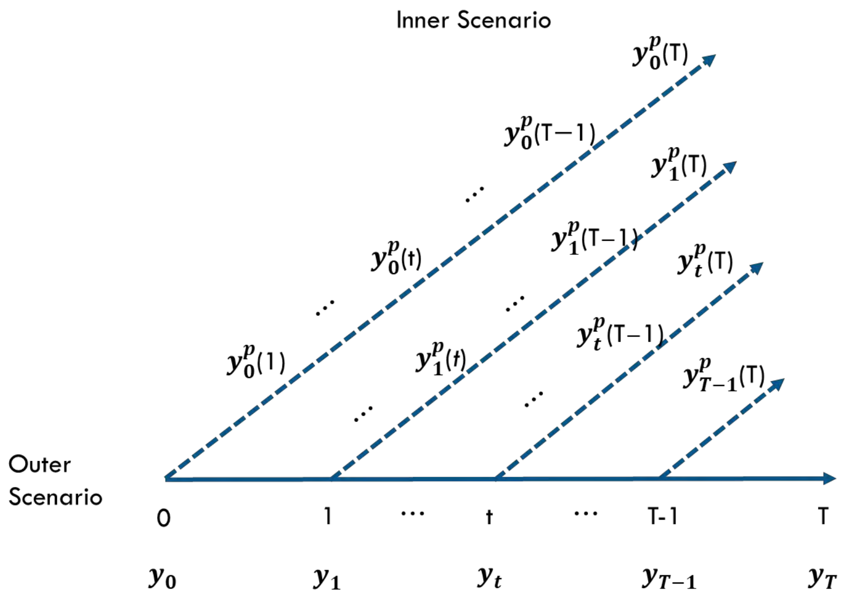

The discount rate $\mathrm{CR}_{t',i}$ can be derived using the following steps:

Step 1. Generate the fundamental economic factors with the random term $\mathbf{e_t}$ till time t' (simulated economic factors) and without the random term thereafter (expected economic factors):

$$\mathbf{F}_{t'}^{p}(t) = \begin{array}{ll} c + \mathbb{A}_1\mathbf{F}_{t-1} + \mathbf{e}_t & \quad t \le t' \\ c + \mathbb{A}_1\mathbf{F}_{t-1} & \quad t > t' \end{array}$$

where $\mathbf{F}_{t'}^{p}(t)$ is a column vector that contains expected fundamental economic factors, given the history till time t'.

Step 2. Generate the asset returns with the random term $\mathbf{e_t}$ till time t' (simulated asset returns) and without the random term thereafter (expected asset returns):

$$\mathbf{y}_{t'}^{p}(t) = \begin{array}{ll} \alpha + \phi_1\mathbf{y}_{t-1} + \phi_2\mathbf{y}_{t-2} + \mathbb{B}_0\mathbf{F}_t + \mathbb{B}_1\mathbf{F}_{t-1} + \mathbb{B}_2\mathbf{F}_{t-2} + \mathbf{e}_t & \quad t \le t' \\ \alpha + \phi_1\mathbf{y}_{t-1} + \phi_2\mathbf{y}_{t-2} + \mathbb{B}_0\mathbf{F}_t + \mathbb{B}_1\mathbf{F}_{t-1} + \mathbb{B}_2\mathbf{F}_{t-2} & \quad t > t' \end{array}$$

where $\mathbf{y}_{t'}^{p}(t)$ is a column vector that contains expected asset returns and other factors, given the history till time t'.

Step 3. Calculate the expected aggregated asset return based on the asset mix:

$$\mathrm{CR}_{t',i} = \mathbf{w}_{t'+i}\mathbf{y}_{t'+i}^{p}$$

where $\mathbf{w}_{t'+i}$ is a row vector that contains the asset mix at time $t' + i$. If the plan has a target asset mix and the portfolio is rebalanced at the end of each period, the vector will be constant through time.

With this method, the discount rate used to determine contributions is not fixed but varies according to circumstances. The expected return on a bond portfolio would be sensitive to changes in the average yield to maturity on the bonds, while the expected return on equities and other investments could be sensitive to anticipated real GDP growth rates or other fundamental economic factors.

With the projected benefit payments and plan contributions, asset balances can be dynamically projected under each scenario, considering the net cash flow from the liability side.

4.1.3 Asset Model

With the economic scenario generator and liability cash flow projection in place, asset projection is straightforward. A simplified version is used here, assuming that the benefit payment B_t and plan contribution C_t occur at the end of each period:

$$A_t = A_{t-1}(1 + c_t + r_t) - B_t + C_t$$

where

A_t = total asset value at time t
c_t = aggregate income rate (coupon, cap rate or dividend yield) of the asset portfolio during period t, equal to the weighted average income rate of all asset classes with weight equal to the investment amount at the beginning of period t and
r_t = aggregate capital return (price return) of the asset portfolio during period t, or the weighted average capital return of all asset classes with weight equal to the investment amount at the beginning of period t.

Based on the asset allocation plan, the asset portfolio may be rebalanced, which will affect the calculation of future c_t and r_t. This is where the LDI strategy plays its important role in matching liabilities with assets.

With projected assets and liabilities, the economic funding ratio FR_t and economic funding surplus/shortfall FS_t can be calculated accordingly:

$$\mathrm{FR}_t = \frac{A_t}{L_t}$$

$$\mathrm{FS}_t = A_t - L_t$$

Unlike regulatory valuation that may use a smoothed discount curve, economic valuation uses a mark-to-market approach that depends on the market condition at the valuation date.

4.1.4 Dependency

The dependency in the LDI benchmark model is governed by the economic scenario generator and modeled in five places:

1. Intertemporal dependency among fundamental economic factors is built in the economic scenario generator with the VAR(1) model. The fundamental economic factors jointly determine the status of the economy: recession or expansion.
2. The error terms (nonsystemic component) of the fundamental economic factors in the VAR(1) model are not independent from one another. Their contemporary correlations are reflected when generating the random part of the fundamental economic factors.
3. Both contemporary and intertemporal dependency between asset returns and fundamental economic factors is governed by the linear models used to generate asset return scenarios.
4. The error terms (idiosyncratic factors) of asset return variables or other factors are not independent from one another. Their contemporary correlations are reflected when generating the random part of the asset returns.
5. To address the issue of nonlinear relationship, the volatility of idiosyncratic factors, their interdependency and their dependency on fundamental economic factors vary by the status of the simulated economy. Higher volatility and dependency are used when the economy is projected to be in recession.

Each scenario generated by the economic scenario generator will bear the appropriate relationships. Assets and liabilities projected based on these scenarios will provide a consistent, realistic and holistic view of the pension plan's future.

4.1.5 Evaluation

Assessment of the asset allocation plans requires choosing return and risk measures. Different measures may be used and tested in this report:

- *Return measure:* This equals the average funding ratio minus the target funding ratio, or the average funding surplus to assess the dollar impact.
- *Risk measure:* This equals the average funding ratio minus the left-side *x*th percentile of the funding ratio (that is, the lowest [100 – *x*]% of the funding ratio). Alternatively, volatility of the funding ratio may also be used.

The plan contributions under different asset allocation plans may be different. If expected return is used when determining the plan contribution, high-risk plans will have fewer plan contributions. The difference could be material and needs to be reflected in the return and risk measures. For ease of comparison, it is assumed that plan contributions are the same for all asset allocation plans, with the amount determined as the discounted value of additional benefit payments attributed to current period employment using a 100% Treasury bond investment. Adjustments can be made to test the impact of plan contributions on investment strategies, which by itself can be a research topic.

To assess overall performance under multiple scenarios, a utility function is used with degree of risk aversion as an input:

$$\text{Utility} = \text{Return} - \text{Degree of Risk Aversion} \times \text{Risk}^2$$

Risk can be measured by the volatility of the return measure, value at risk (VaR) or conditional tail expectation (CTE) at a chosen confidence level. The most appropriate asset allocation plan may be determined as the one with the highest utility value. Qualitative considerations are likely to be involved in the final decision as well.

Table 1 lists the basic information of the sample pension plan for which deep learning models are used to find the best dynamic investment strategy.

Table 1
Sample DB Plan Information

Item	Assumption
Initial plan assets	\$10,000,000
Initial plan liabilities	\$10,000,000
Valuation date	Dec. 31, 2016
Pension benefit	Five-Year Average Salary before Retirement × 1% × Length of Service × COLA
COLA	$\prod_{i=1}^{\text{Current Age} - \text{Retirement Age}} \left(1 + \min\left(5\%, \quad 80\% \times \text{Inflation Rate}_i\right)\right)$
Payment option	1. Retirees can choose the lump sum payment option at retirement. The lump sum amount is equal to the expected future benefit payments discounted by 4%, and 10% of retirees are assumed to use this option. 2. The remaining 90% of retirees are assumed to choose life payment option.
Administration, investment and tax expenses	These expenses are not modeled explicitly in this example. It is assumed that administration expenses are included in benefit payments, and investment/tax expenses are deducted from gross investment return.
Standard mortality rate	Example ultimate mortality rate till age 110 by gender (life expectancy: 77 years)
Mortality improvement rate	1% per year
Vesting period	Three years

<table>
<tr><td>Vesting period turnover rate</td><td>
<table>
<tr><th>Occupation Type</th><th>Turnover Rate</th></tr>
<tr><td>I</td><td>5%</td></tr>
<tr><td>II</td><td>4%</td></tr>
<tr><td>III</td><td>3%</td></tr>
<tr><td>IV</td><td>2%</td></tr>
<tr><td>V</td><td>1%</td></tr>
</table>
</td></tr>
<tr><td>Salary growth rate</td><td>
Basic salary growth rate is linked to the scenarios of wage inflation. Different occupations have different multiples of the basic growth rate.
<table>
<tr><th>Occupation Type</th><th>Turnover Rate</th></tr>
<tr><td>I</td><td>0.9</td></tr>
<tr><td>II</td><td>1</td></tr>
<tr><td>III</td><td>1.1</td></tr>
<tr><td>IV</td><td>1.2</td></tr>
<tr><td>V</td><td>1.5</td></tr>
</table>
</td></tr>
<tr><td>Plan participant mix</td><td>
<table>
<tr><th>ID</th><th>Retired</th><th>Sex</th><th>Date of Birth</th><th>Date of Hire</th><th>Annual Salary</th><th>Date of Retirement</th><th>Occupation</th><th>Weight</th></tr>
<tr><td>1</td><td>Y</td><td>F</td><td>June 30, 1950</td><td>April 30, 1979</td><td>50,000</td><td>July 31, 2014</td><td>II</td><td>15%</td></tr>
<tr><td>2</td><td>Y</td><td>M</td><td>June 30, 1955</td><td>May 31, 1980</td><td>45,000</td><td>June 30, 2015</td><td>I</td><td>15%</td></tr>
<tr><td>3</td><td>N</td><td>F</td><td>June 30, 1981</td><td>January 31, 2005</td><td>80,000</td><td>June 30, 2045</td><td>IV</td><td>8%</td></tr>
<tr><td>4</td><td>N</td><td>M</td><td>June 30, 1991</td><td>January 31, 2015</td><td>40,000</td><td>June 30, 2055</td><td>II</td><td>8%</td></tr>
<tr><td>5</td><td>N</td><td>M</td><td>June 30, 1971</td><td>January 31, 2005</td><td>120,000</td><td>June 30, 2035</td><td>V</td><td>8%</td></tr>
<tr><td>6</td><td>N</td><td>F</td><td>June 30, 1961</td><td>December 31, 2000</td><td>80,000</td><td>June 30, 2020</td><td>II</td><td>8%</td></tr>
<tr><td>7</td><td>N</td><td>F</td><td>June 30, 1981</td><td>January 31, 2005</td><td>100,000</td><td>June 30, 2045</td><td>IV</td><td>8%</td></tr>
<tr><td>8</td><td>N</td><td>M</td><td>June 30, 1986</td><td>January 31, 2005</td><td>55,000</td><td>June 30, 2050</td><td>II</td><td>10%</td></tr>
<tr><td>9</td><td>N</td><td>M</td><td>June 30, 1966</td><td>January 31, 1991</td><td>95,000</td><td>June 30, 2030</td><td>IV</td><td>10%</td></tr>
<tr><td>10</td><td>N</td><td>F</td><td>June 30, 1971</td><td>January 31, 1995</td><td>50,000</td><td>June 30, 2035</td><td>II</td><td>10%</td></tr>
</table>
We use compressed model points to represent the entire portfolio. The weight stands for the portion of the liability portfolio that a model point represents. They are scaled so that the total liability value equals the plan’s initial liability value.
</td></tr>
</table>

4.1.6 Description of Simulated Results

The simulated scenarios and projected funding status are used for RL model training and validation. One thousand scenarios with 10 years of quarterly projections are generated, with 800 scenarios used for training purposes and 200 scenarios used for validation purposes. Additional scenarios may be needed if risk measures such as VaR and CTE are calculated at an extremely high confidence level, such as 99.9%.

Before applying RL models, it is helpful to have a glimpse of the distribution of simulated results, as shown in Figures 8 to 10. Deterministic scenarios are generated by the ESG assuming a zero random component. The average and volatility of simulated annual returns/yields are also presented. Two asset classes, an S&P 500 index and AA-rated corporate bond portfolio, are used as a starting point. The bond portfolio is assumed to be rebalanced to maintain the current credit rating with a target duration of 10 years. The discount curve of the liability valuation is assumed to be AA-rated corporate bond yield curves.

Figure 8
Equity Return Projection

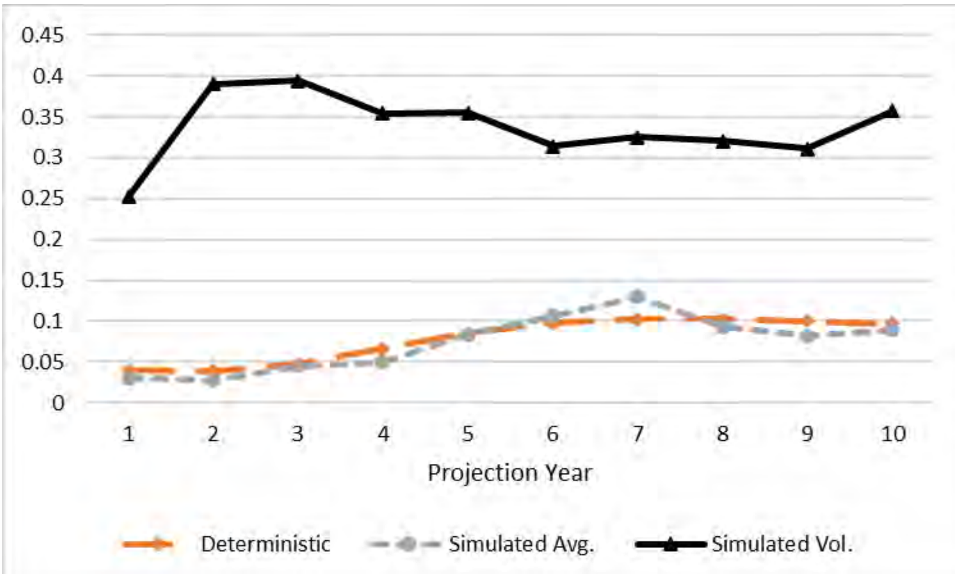

Figure 9
AA-Rated Corporate Bond Portfolio Return Projection

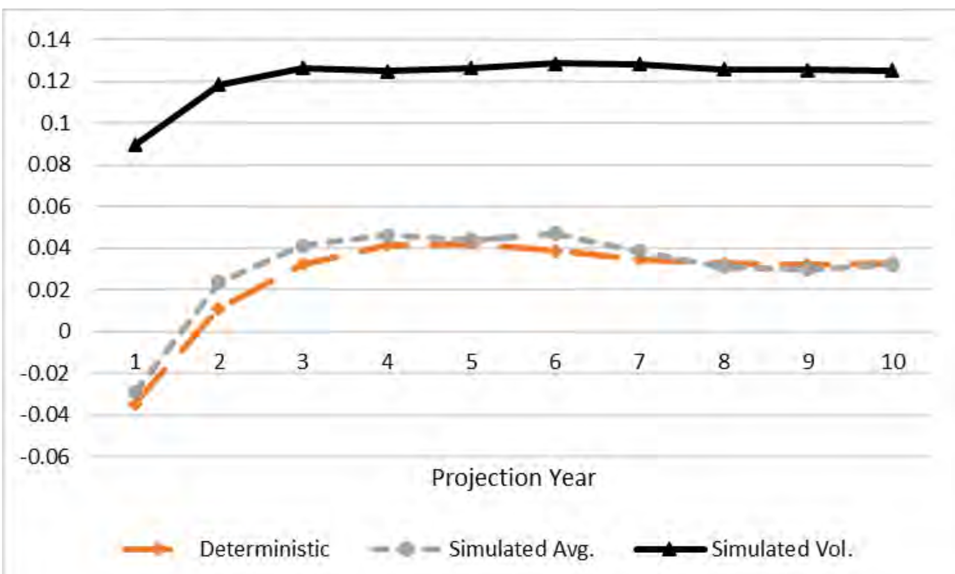

Figure 10
Liability Discount Rate: 10-Year AA-Rated Corporate Bond Yield Projection

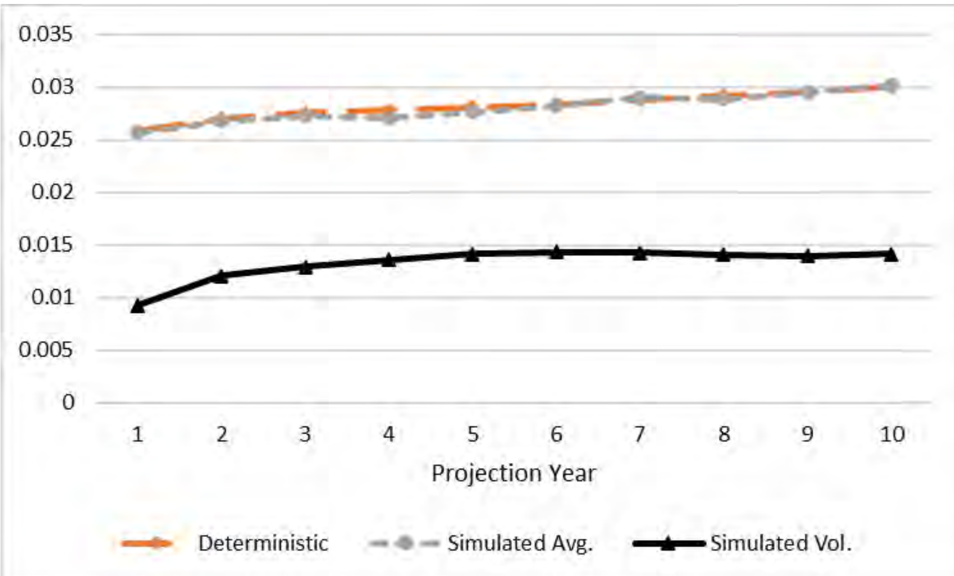

The interest rates are expected to increase gradually in the long term based on the calibrated ESG, which leads to low projected bond portfolio returns in general. Although different views of future returns exist, this report focuses on using simulated scenarios to test the effectiveness of RL, rather than choose forward-looking views. Table 2 lists the statistics of the quarterly returns.

Table 2
Statistics of Simulated Returns

			Correlation		
Asset Class	Mean	Volume	Equity Return	Bond Return	10-Year Bond Yield
Quarterly Equity Return	1.8%	16.2%	1.000	0.038	0.147
Quarterly Bond Return	0.7%	6.0%	0.038	1.000	0.124
10-Year Bond Yield	2.8%	1.3%	0.147	0.124	1.000

The plan assumes a fully funded status at the beginning, with the projected liability values shown in Figure 11.

Figure 11

Dynamic Liability Projection

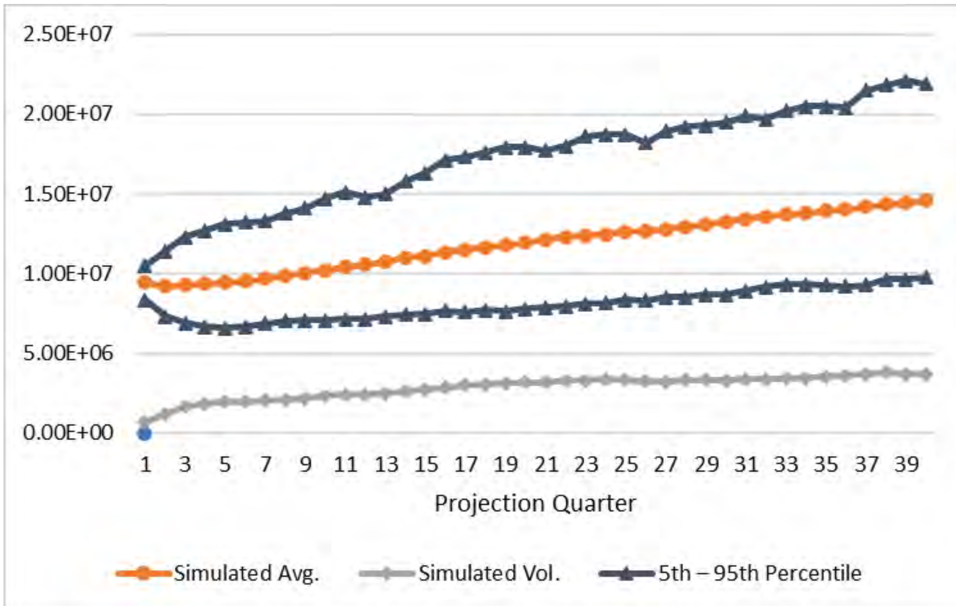

It is an open pension plan with liability growing steadily over 10 years. Figure 12 shows the distribution of benefit payments with the spikes due to the exercise of lump sum payment options. The payments over the next 10 years are almost deterministic except the inflation adjustment part.

Figure 12

Benefit Payment Projection

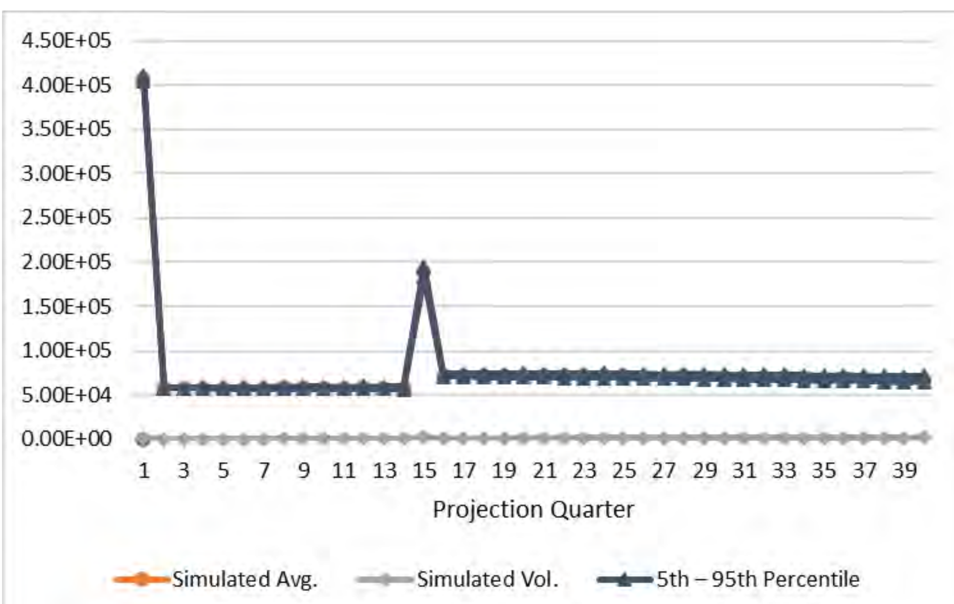

Figure 13 shows the plan cost projection, which is the discounted value of additional future benefits accrued during each period. It can be considered as a conservative estimation of plan contribution and is not contingent on the funding status. In actual projection, it can be adjusted according to real practices.

Figure 13
Plan Cost Projection

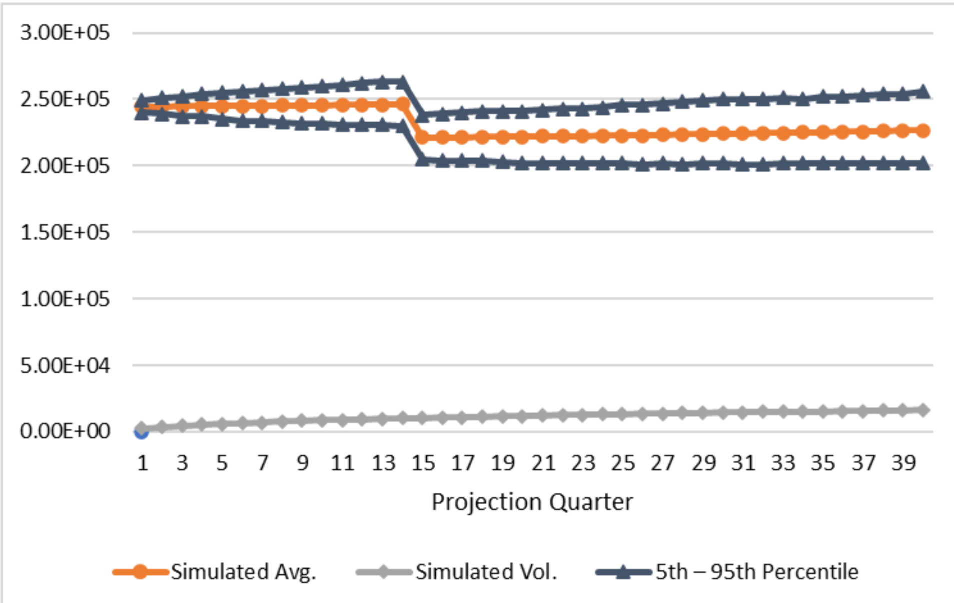

Figure 14 shows the projected average funding ratios with five constant asset allocation plans over 10 years. Equity investment is preferred in terms of its higher expected return. The average funding ratio is expected to stay above 1 because of a fully funded status at initial and plan contributions calculated using government bond yield curves for discounting. The interest rates are also projected to increase gradually, which helps increase funding ratios as well.

Figure 14
Average Funding Ratio Projections

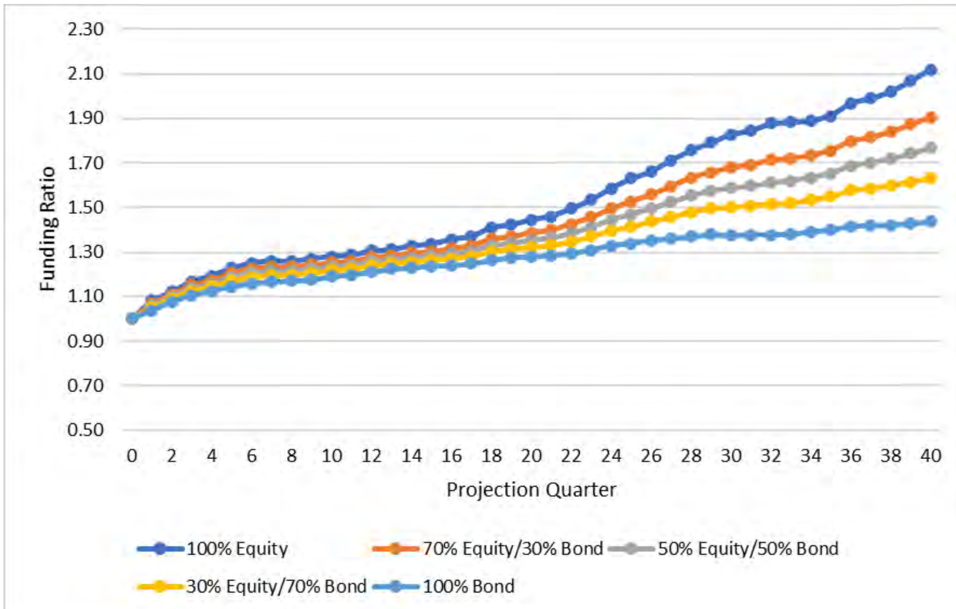

Figure 15 shows the range of possible funding ratios with a two-asset mix: 100% equity investment and 100% bond portfolio investment. Equity investment comes with higher expected value and higher risk.

Figure 15
Range of Funding Ratio Projections

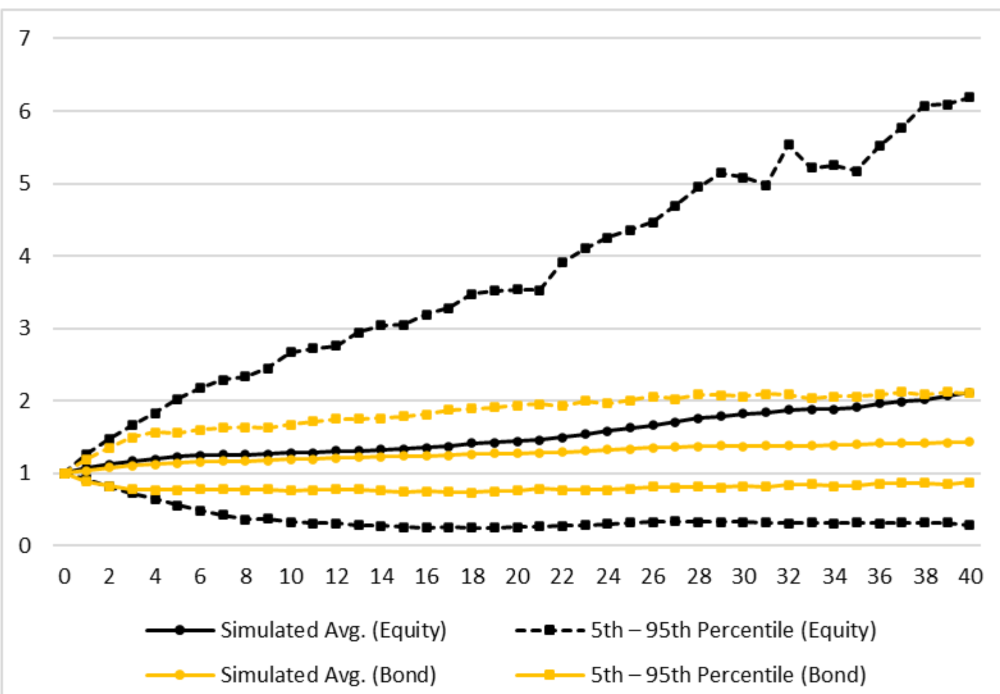

Before RL is used to find a dynamic strategy, it is helpful to evaluate different asset allocation plans based on degree of risk aversion, time horizon and initial funding status. Figure 16 shows the efficient frontier of asset allocation plans for the funding ratio in 10 years using the average funding ratio minus 1 as the return measure and the volatility of the funding ratios as the risk measure. Fifty-one asset allocation plans are tested, with bond investment increasing from 0% to 100% with a step of 2%.

Figure 16
Efficient Frontier Sample

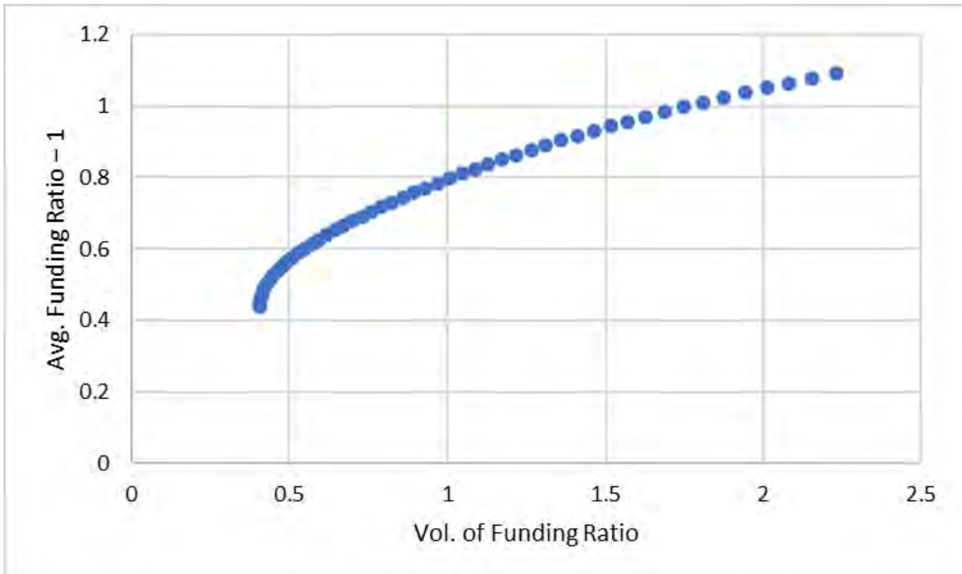

Figure 17 shows the optimal asset allocation plans with funding ratio volatility as the risk measure, using a different time horizon and degree of risk aversion. With an increasing degree of risk aversion, more bond investment is desired. For a short time horizon less than about two years, a bond portfolio is less preferred, which is likely to be caused by the projection of increasing interest rates in the early years till they reach a plateau in terms of average values. For some degrees of risk aversion, the desired bond investment gradually decreases by time horizon after four years. This may be caused by lower long-term volatility.

Figure 17

Static Optimal Strategy Sample: Funding Ratio Volatility

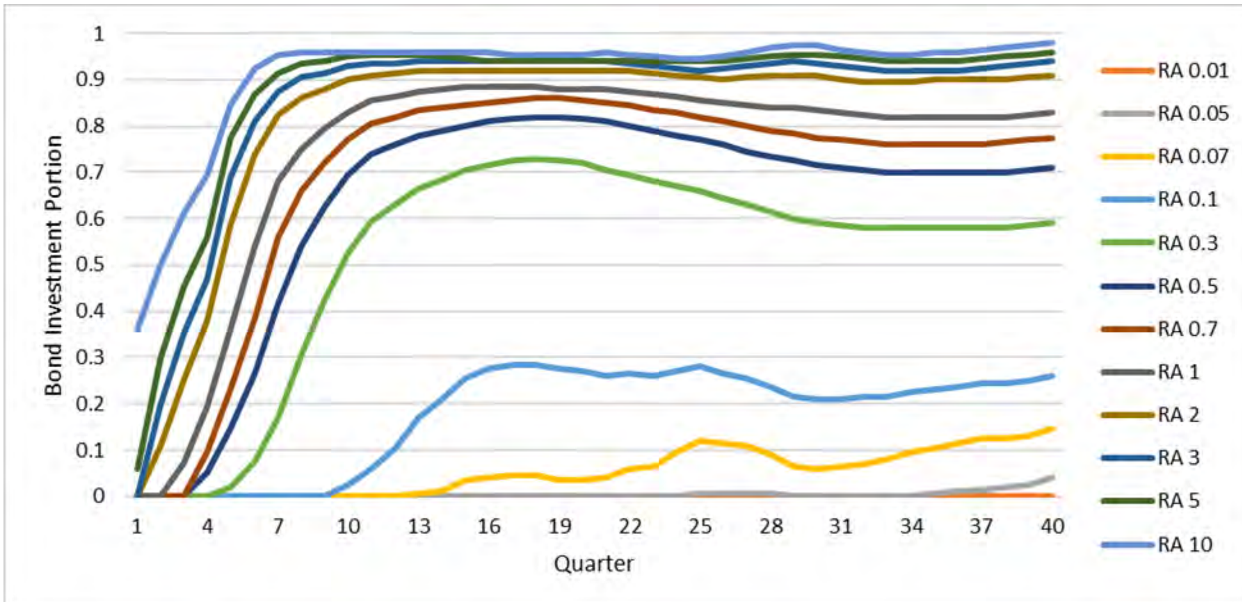

Note: RA: degree of risk aversion

Using the difference between the left-side VaR and average value as the risk measure, the optimal asset allocation plans show a similar pattern as those using volatility to measure risk. Figures 18 to 20 show the bond investment portion using VaR90, VaR95 and Var99, respectively.

Figure 18

Static Optimal Strategy Sample: Funding Ratio VaR90

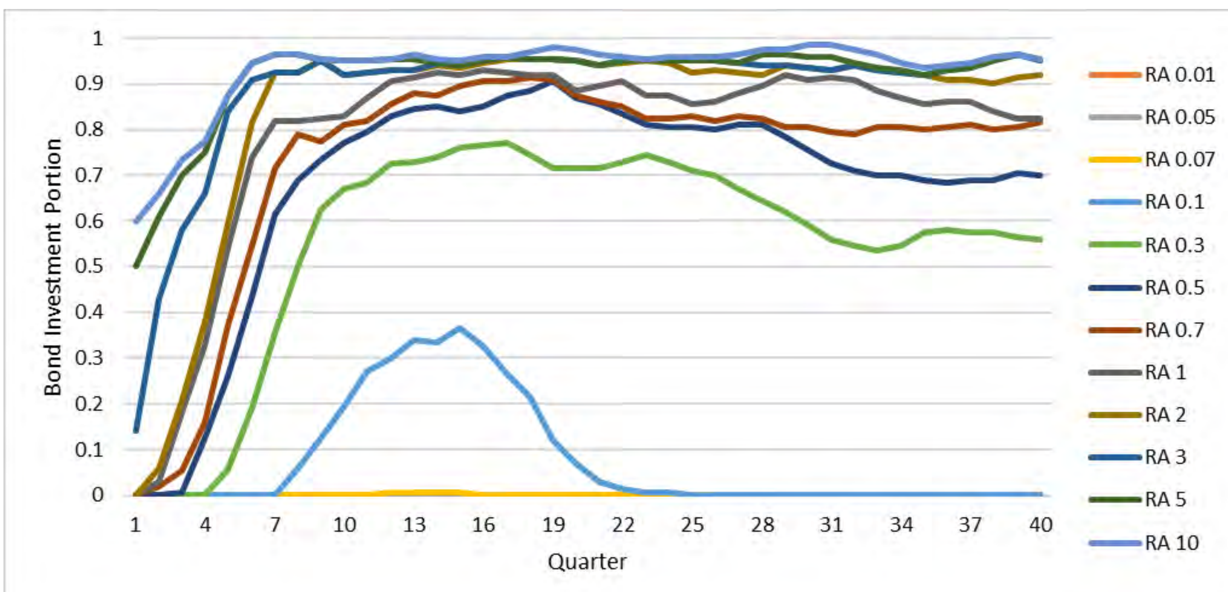

Figure 19

Static Optimal Strategy Sample: Funding Ratio VaR95

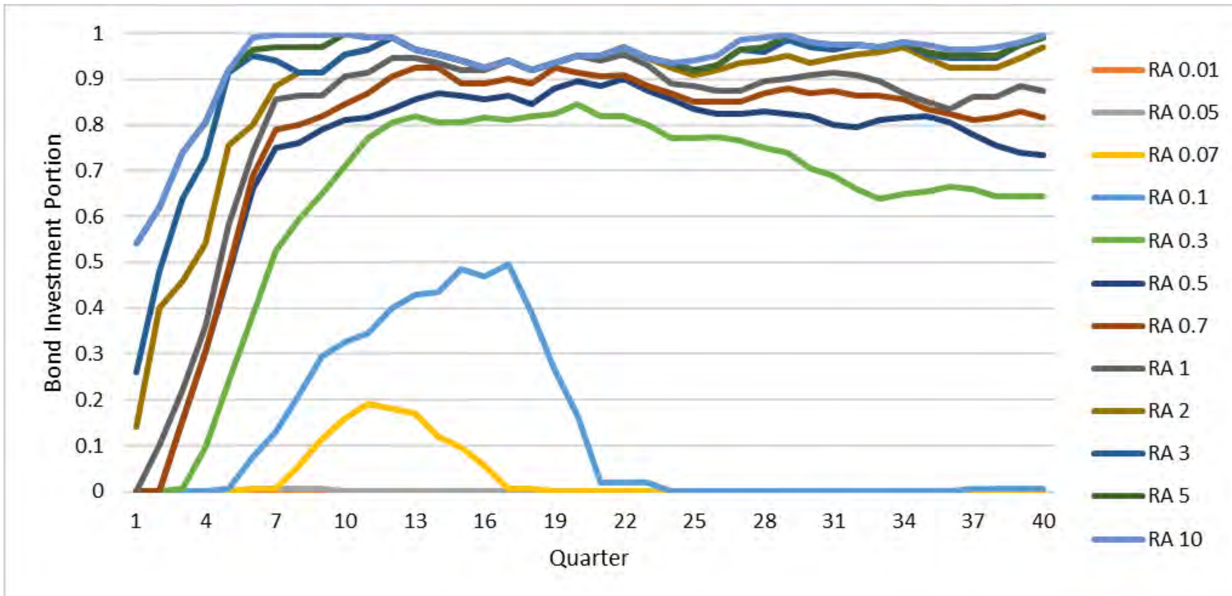

Figure 20

Static Optimal Strategy Sample: Funding Ratio VaR99

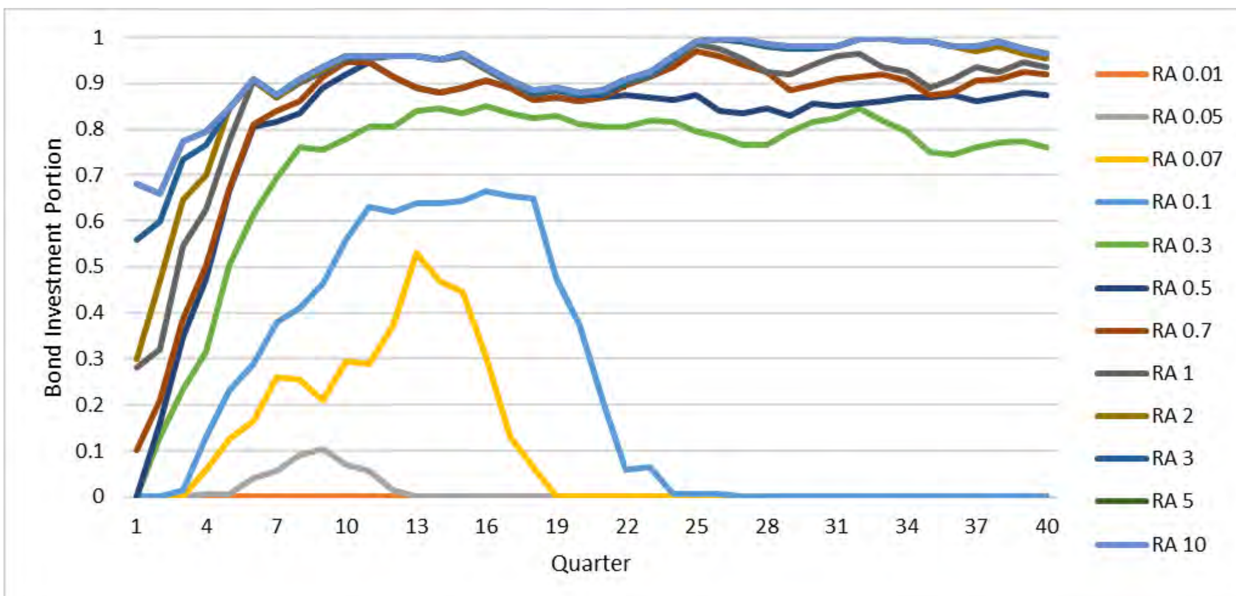

Figure 21 shows the optimal strategies using the ending funding surplus as the target. For comparison purposes, volatility of the funding surplus is used. The pattern is similar to Figure 17 with a slightly lower desired bond investment. However, the degree of risk aversion has different implications between the funding ratio and funding surplus.

Figure 21
Static Optimal Strategy Sample: Funding Surplus Volatility

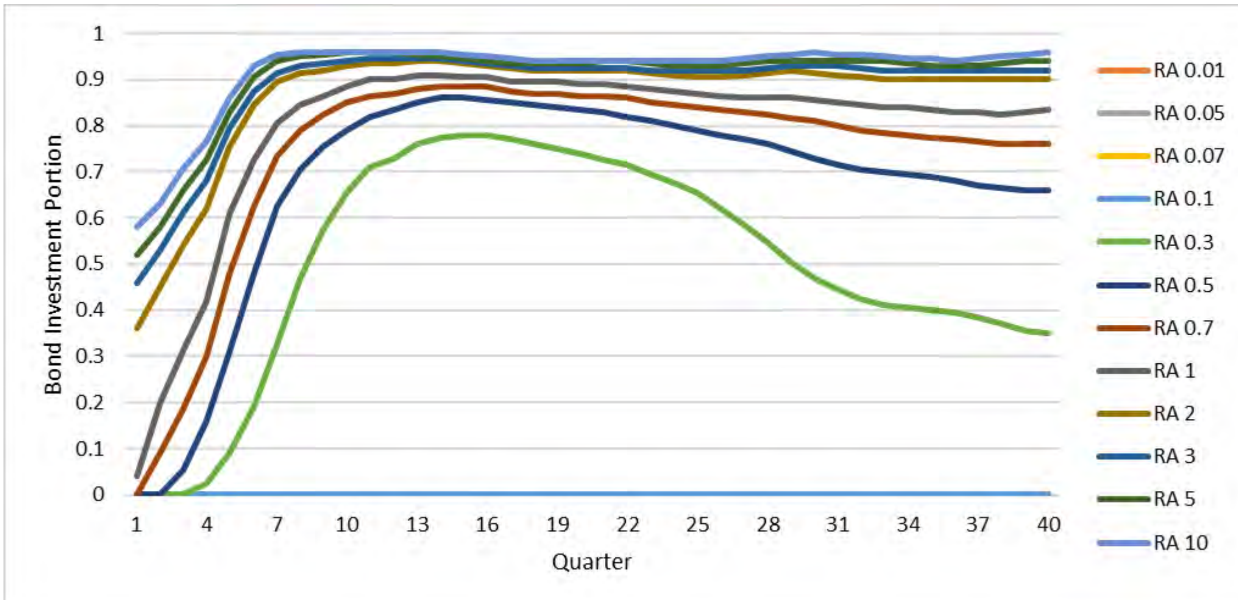

As expected, the desired bond investment increases by percentile. The bond investment may also decrease with a longer time horizon. Although this analysis is helpful for benchmarking, it is still based on static strategies:

- The optimal investment strategies shown above assume a constant mix from the start. They do not include periodic readjustments. By observing real performance in each period, the ability to adjust the asset allocation plan may provide extra gains that are not captured by static strategies. This gain may be insignificant for an open DB plan as in this research with a steady growth of liability shown in Figure 11, but the gain can be meaningful for closed DB plans.

- Although the static strategy analysis can be reperformed periodically, given the then-current funding status and economic conditions, it provides only part of the picture. It is very challenging to run through all the possible paths with varying asset mixes during different time periods given the burden on dynamic liability projections, the continuous space of funding status and the available asset mixes, especially when more than two asset classes are used. For example, Figure 22 shows the optimal static strategy with an initial funding ratio of 80% instead of 100% as assumed in Figure 17. The optimal strategies show less bond investment with the underfunded status.

Figure 22
Static Optimal Strategy Sample with Underfunding: Funding Ratio Volatility

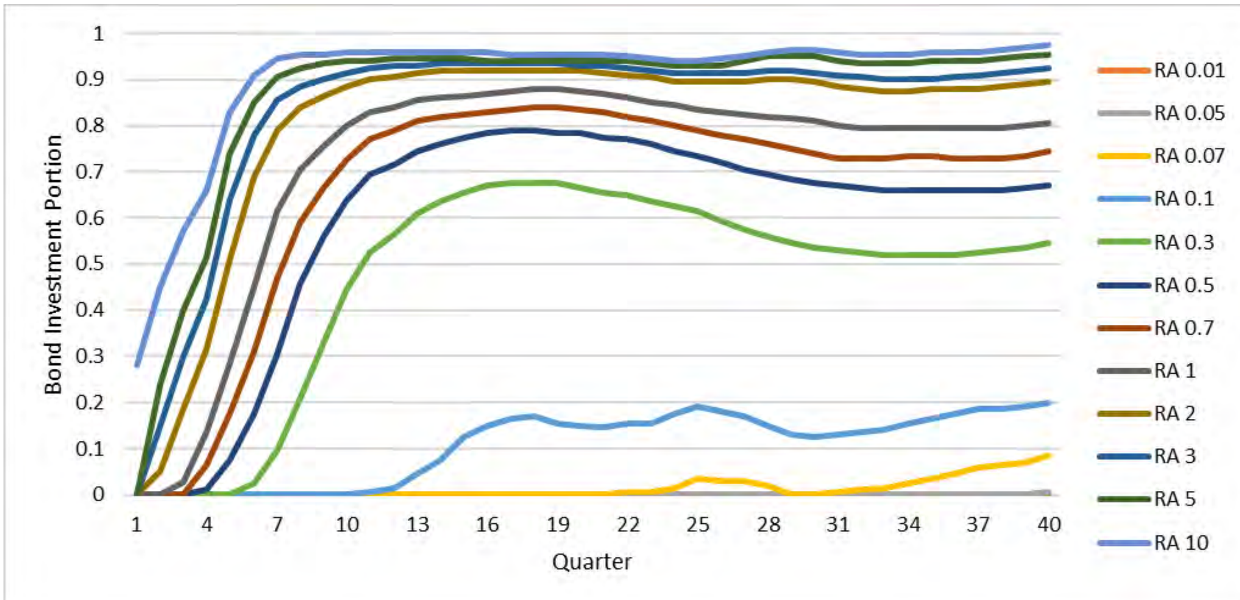

4.2 Model Training

With the experimental environment set up, RL models can be trained to find appropriate but not necessarily optimal dynamic investment strategies. As explained in Section 3, the key is to find the reward function $Q^{\pi}(s,a)$, which may be approximated by feedforward fully connected neural networks or LSTM models.

4.2.1 State

Input s is the information that decision makers can observe in the real world. A complete set is all the historical fundamental economic factors, asset returns and pension plan surplus position. A subset is usually selected depending on data availability, model type and relevance of historical patterns in the future. In this example, for fully connected neural networks, s includes not only the current status but also the status of the previous two quarters. It is assumed that given three quarters' information, the model is able to detect the status of the economy and predict future movements of economic factors and asset returns. For LSTM models, each module in the time series corresponds to a quarter's information. Information about previous quarters is included in the hidden state. Figure 23 illustrates the inputs to LSTM and FCNN models. LTSM models include the full sequence as one training example. NN models consider each time point in the sequence as a separate training example.

Figure 23
Reward Function Inputs

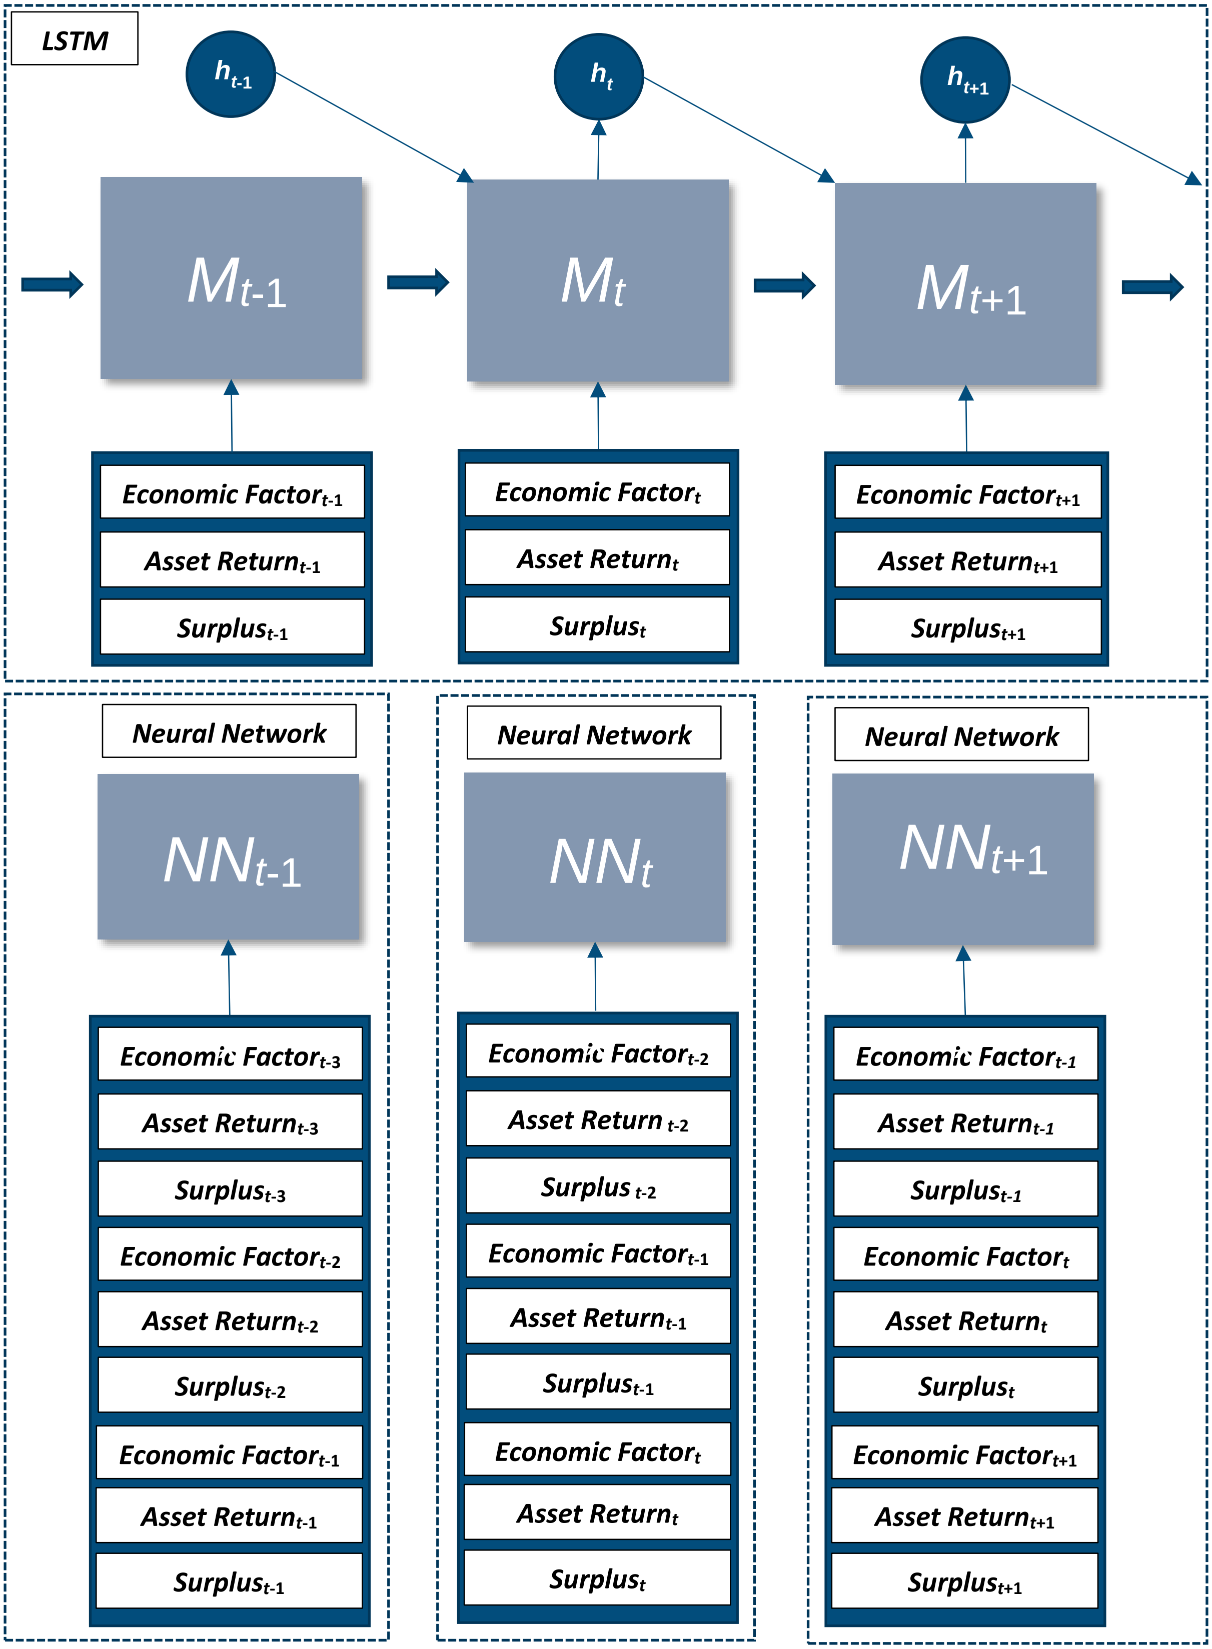

4.2.2 Action

Input *a* is the asset rebalance action to take. Many ways can be used to define possible actions. For example, if the asset portfolio is invested evenly in two asset classes, bonds and public equity, possible actions could be (a) keep the current mix (b) increase bond investment and (c) increase equity investment. To facilitate model training, constraints can be put on possible actions:

- Each asset class may have a range of allocation percentage. For example, equity allocation may not be allowed to exceed 60% according to the risk appetite.
- The magnitude of the asset mix change can be limited. For example, for quarterly rebalancing, the maximum allowed change of one asset class's allocation is 2% to avoid high transaction cost, market impact of massive buy and sell, and the adverse impact of wrong decisions.

Possible actions to test at each time point would be the following:

- Keep current mix
- Increase equity allocation by 2% of total asset and decrease bond investment by the same amount
- Decrease equity allocation by 2% of total asset and increase bond investment by the same amount.

The action setup described above may be reasonable in practice for an existing pension plan with a starting asset mix and concerns about transaction costs and market impact. At the same time, it explores only a part of possible actions. An alternative approach is to expand the action space to all possible asset allocation plans. For example, 51 possible actions can be defined as 0% bonds/100% equity to 100% bonds/0% equity with a step of 2% in the allocation percentage. At each decision point, asset allocation can be changed to the best action that satisfies any constraints.

With the states and action defined, the current period reward r, as in the recursive representation of the reward function, needs to be projected:

$$Q^{\pi}(s,a) = r + \gamma Q^{\pi}(s', \pi(s'))$$

With the economic conditions, asset returns and surplus position, the current period reward of each possible action can be estimated using an asset and liability projection model. In this report, the following options are used:

- Asset return

  $$R_t^{ar} = \frac{A_t}{A_{t-1}} - 1$$

  The ultimate goal is to find a dynamic strategy that minimizes the chance of being underfunded or maximizes the ultimate overfunding. However, instead of trying to reach the ultimate goal in one step using RL, a gradual approach is used to better understand the advantages and disadvantages of RL.

  Ideally, the reward measure is a function of both asset and liability. But, as the first step, a simpler reward measure is needed so that it can be validated using backward induction. Discounted multiple-period asset returns are used as the total reward with asset return as the current reward. No matter what the past path, given the current state, the optimal strategy of the future can be determined precisely. Here the optimal strategy is scenario/state dependent.

  A more appropriate asset-only measure is the discounted asset returns in a dollar amount. Although the optimal strategy is likely to be close to the one using discounted asset returns in percentages, it runs into the issue that semi-supervised learning is needed because the optimal value is not explicit and cannot be calculated using backward induction, or at least not in a computationally efficient way. The reason is that asset value depends on the entire path of the asset allocation rather than only the current asset allocation. With the past asset allocations incorporated in the reward function, the calculation will be about 20-fold of using asset returns in percentage, given a 10-year time horizon with quarterly adjustments. The main reason for using an asset-based measure is to determine the effectiveness of reward function calibration using deep learning models, and so asset gains/losses in a dollar amount are not tested.

- Change in funding ratio

$$R_t^{fr} = \frac{A_t}{L_t} - \frac{A_{t-1}}{L_{t-1}}$$

The second reward function is the discounted changes in future funding ratios. Like the discounted asset gains/losses in a dollar amount, this reward function depends on the entire asset allocation path. It is unknown and scenario dependent. Therefore, the reward function needs to be estimated using semisupervised learning. Deep learning models are used to estimate the reward function in an unsupervised way because only the current reward is observable.

- Change in funding surplus

$$R_t^{fs} = A_t - A_{t-1} - (L_t - L_{t-1})$$

The third reward function is the discounted changes in the future funding surplus. Like the discounted changes in future funding ratios, this reward function depends on the entire path of both asset allocation and scenario. It may be a more meaningful measure for a runoff plan.

Theoretically, the dynamic investment strategy is chosen based on the value of the reward function:

$$\pi^*(s) = \max_a Q^*(s, a)$$

However, to add uncertainty in the training process, a rebalance action also may be determined randomly so that it will not stay at a local optimal point. At the initial phase of training, the reward function has not been updated with many training examples. It makes sense to choose the rebalance action randomly with a high probability. As the training process goes on, more and more rebalancing action should be chosen based on maximum reward:

$$\mathrm{prob}(n) = l + (h - l)e^{-n-1/rate}$$

where

$\mathrm{prob}(n)$: The probability that the action is chosen randomly for the *n*th example
l: The lowest probability of random selection; 5% is used in the example
h: The highest probability of random selection; 75% is used in the example
$rate$: The decay rate that controls the speed of reducing probability; the maximum of 200 and batch size is used in the example.

4.2.3 Reward Function Specification

Two models are used to approximate the reward function, with the detailed structure shown in Figure 24.

Figure 24
Reward Function Model Specification: Neural Network and LSTM

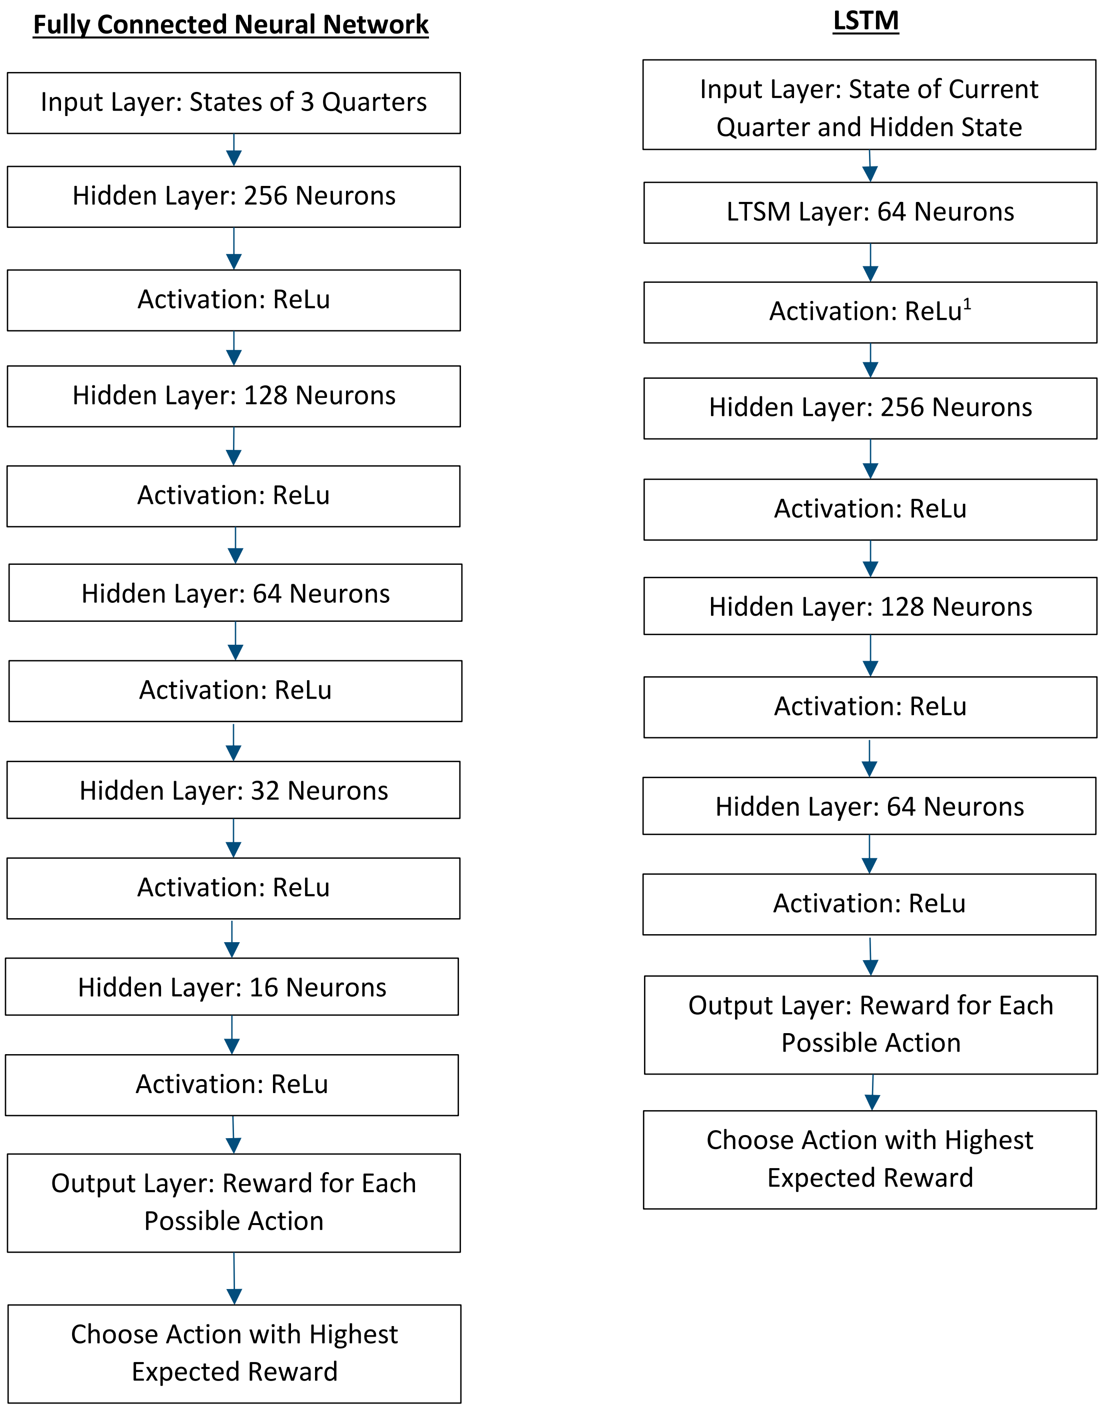

Notes:

1. ReLu: An activation function to determine which neurons are active. $f(x) = \begin{matrix} 0 & for\ x < 0 \\ x & for\ x \geq 0 \end{matrix}$

The activation function is applied to a neuron's value before fed into the next layer. A neuron's value is a linear function of the neurons in the previous layer ($WX + b$) with W the weights applied to the neurons and b the bias term. Depending on its value, the next step is to determine whether the neuron should be activated, like the way our brains work. However, the value of the neuron could range from negative infinity to positive infinity and make the choice hard. Activation functions are used to add nonlinearity into neural networks and bring the values down to a desired range.

4.2.4 Control Parameters

Table 3 shows the control parameters used in model training. Other options were tested without observing material improvement on model accuracy.

Table 3
Model Training Parameters

Parameter	Value
Batch Size	1,000
Epoch	200
Update Frequency	10
Initial Learning Rate	0.01
Optimization Method	RMSprop proposed by Hinton et al. (2012)
Loss Function	Mean squared error

Batch Size

Optimization with the entire dataset of experiments is demanding on both computing capacity and speed of convergence. Rather, the optimization is done many times with smaller datasets referred to as batches. In this example, 800 training experiments are conducted, with each covering 10 years of quarterly projections. For neural networks, this means 32,000 training examples. If each batch contains 1,000 training examples, running through the entire training dataset means running 32 batches. For each batch, the model tries to update model parameters and reduce the loss.

Epoch

One epoch means the entire training dataset will be run through one time. Usually the training dataset needs to be used multiple times for the optimization process to converge. With 200 epochs, the model will perform optimization 6.4 million times.

Update Frequency

During training, instead of updating model parameters each time after optimizing based on one batch, additional stability can be incorporated to defer a model parameter update till a certain number of batches have been run. Here an update frequency of 10 means that model parameters will be updated every 10 simulations, although optimization will be performed every time. It is important to know that each time that model parameters are updated, the reward function will be changed as well. Allowing this gives the model a chance to reduce errors before changing to another optimization target.

Learning Rate

The learning rate controls the parameter updating speed during training. A high learning rate means that the gradient information will be reflected at a higher speed.

Optimization Method

In this example, the RMSprop method is used for minimizing the loss function. It is an adaptive learning rate method that uses the moving average of squared gradients to adjust parameter updating:

$$\theta_{t+1} = \theta_t - \frac{\eta}{\sqrt{MeanSquare_t}} g_t$$

$$MeanSquare_t = 0.9MeanSquare_{t-1} + 0.1g_t^2$$

where

θ: Model parameter to update
η: Learning rate
g_t: Gradient of the loss function w.r.t. parameter θ at step t
$MeanSquare_t$: Moving average of squared gradients.

Other adaptive learning rate schedules, such as Adam, were also tested, but no material difference was observed regarding the resulting dynamic investment strategy.

Loss Function

The mean squared error (MSE) is the average of squared errors. Errors are defined as

$$\delta = Q(s, a) - (r + \gamma \max_a Q(s', a))$$

The control parameters described above were used to produce the results of this research. The selection of these parameters along with other modeling choices is mostly a trial-and-error process considering the impact on model accuracy and the required computing resources and model training time.

4.3 Model Validation

Unlike traditional supervised learning problems that have actual values and predicted values, it is not straightforward to tell if an investment decision is optimal. To assess when the RL is effective, a few checks can be made:

1. Using a reward function where supervised learning can be applied to check if deep learning models is effective for optimal action selection. However, the choices of the reward function are limited, an example being a discounted asset return.

2. Even though the reward function is learned with a deep learning model, which means the true reward value is unknown, it is still possible to check the convergence of the trained model by comparing MSE to the estimated reward function. For a good model convergence, it is expected that MSE as a percentage of the estimated reward decreases to the expected level. A good benchmark can be derived by conducting analysis based on the discounted asset return, as described above.

3. Another comprehensive validation measure is the utility function described in Section 4.1.5. By running the trained model with multiple scenarios, a utility function that reflects the tradeoff between risk and return can be calculated and compared with other static asset allocation plans at different time horizons. A higher value of the utility function means a better investment strategy.

Using discounted asset returns as the reward function, supervised learning using deep learning models is performed to check whether it is possible to gain from active investment based on the projection of asset returns for each asset class. Two asset classes are used: an AA-rated corporate bond portfolio and large-scale public equity. The starting asset mix is an even allocation to both asset classes. Three actions are available: keeping the then-current asset mix, increasing the bond allocation by 2% and increasing the equity allocation by 2%. Given the limited space of possible actions and therefore asset mix, the value of the reward function can be precisely calculated for each scenario and each action at each time step. Two deep learning models explained in Section 4.2.3 are calibrated to represent the reward function. The resulting strategy from both models is the 2% shift from bonds to equity investment each quarter until the investment is fully allocated to equity. This is somehow understandable because the deep learning models will choose the asset class with the higher expected returns if the time horizon is long enough. For example, with a 10-year time horizon, a quarterly model will likely smooth out any high volatility, which is deemed a disadvantage for pension funding. However, it also proves that it is not capable of cherry picking asset classes for each period under each scenario. The random component is large enough to dissuade any market-timing efforts in this example. Table 4 lists the calibration results.

Table 4
Supervised Learning Results

	FCNN		LSTM	
	Validation Data	Training Data	Validation Data	Training Data
MSE	24%	22%	24%	22%
Avg. Optimal Reward	0.66	0.75	0.66	0.75
MSE/Avg. Optimal Reward	23.9%	21.7%	23.6%	21.7%
R^2	22%	21%	22%	21%

Note: Reward is measured as the sum of discounted future asset returns.

Given that it does not favor market timing, economic factors except the recession status are removed from model inputs. With simplified inputs, the two deep learning models generate the same optimal strategy as before.

Using the RL method, both FCNN and LSTM generate the same optimal strategy. This confirms that the possible optimization is mainly based on funding status, return and risk of asset classes. It also shows that RL can find the same optimal strategy without using full-scale supervised learning. The ratio of MSE and average optimal reward can be used as a benchmark to evaluate the convergence of FCNN and LSTM using other reward functions.

An assessment of model convergence and utility functions against other strategies is given together with results in Section 4.4.

4.4 Results

The following model calibrations are done to test the efficiency of RL in dynamic asset allocation:

- Simple calibration with rebalance constraints that allows only investment in an AA-rated bond fund and large-scale public equity. In the simple cases, a maximum 2% shift in allocation is allowed on a quarterly basis. This means only three options are available: increasing equity investment by 2%, keeping the same target asset mix from the previous period and decreasing equity investment by 2%. Short selling is not allowed.
- Simple calibration without rebalance constraints that allows only investment in an AA-rated bond fund and large-scale public equity. In the second set, no constraints are applied to reallocation except short selling. Fifty-one different allocation plans ranging from 0% to 100% bond investment with a step of 2% are available and can change from quarter to quarter.
- Comprehensive calibration that allows investment in AA-rated corporate bonds, BBB-rated corporate bonds, large-scale public equity and real estate investment trusts (REITs) with no constraints on allocation changes, except short selling.

For each calibration, both fully connected neural networks and LSTM models are used. Sample open-source codes are hosted at GitHub - Society-of-actuaries-research-institute/FP198-Deep-Learning-for-Liability-Driven-Investment.

for educational purposes that are runnable using either CPUs or GPUs. The training, however, is unlikely to be done on CPUs in practice, given time constraints. All the model training described in this section was performed on GPUs (NVIDIA Tesla K80).

4.4.1 Two Asset Classes with Rebalance Constraint

In the simple calibration example, the asset portfolio starts with an even split between an AA-rated U.S. bond fund and large-scale public equity indexed to the S&P 500. At the end of each quarter, the model needs to determine whether to increase the equity investment allocation by 2%, rebalance it to the asset allocation plan at the beginning of the quarter, or decrease the equity investment allocation by 2%. Table 5 shows the ratio of MSE to average reward value at different optimization steps.

Making comparisons to the benchmark derived through supervised learning using a discounted asset return as the reward function, the model convergence under both FCNN and LSTM is satisfactory.

Table 5
MSE/Average Reward Value for Two Asset Classes with Rebalance Constraint

Optimization Step	FCNN	LSTM
Step 20,000	39.06%	15.60%
Step 40,000	25.43%	11.43%
Step 6,400,000*	7.49%	7.37%
Benchmark	23.90%	23.60%

Note: *The last step in model training.

With the calibrated model, suggested dynamic asset allocation plans can be compared to efficient frontiers based on optimal static investment strategies. In this subsection, FCNN and LSTM generate the same investment strategy. Figure 25 shows that the dynamic investment strategies can be worse than implied by static investment strategies. This is somewhat expected given that more than an 80% bond investment is optimal for a time horizon longer than two years in this example, but the action plan allows only 2% change each quarter, starting from a 50% bond investment.

Figure 25
Dynamic Asset Allocation: Two Asset Classes with Rebalance Constraint

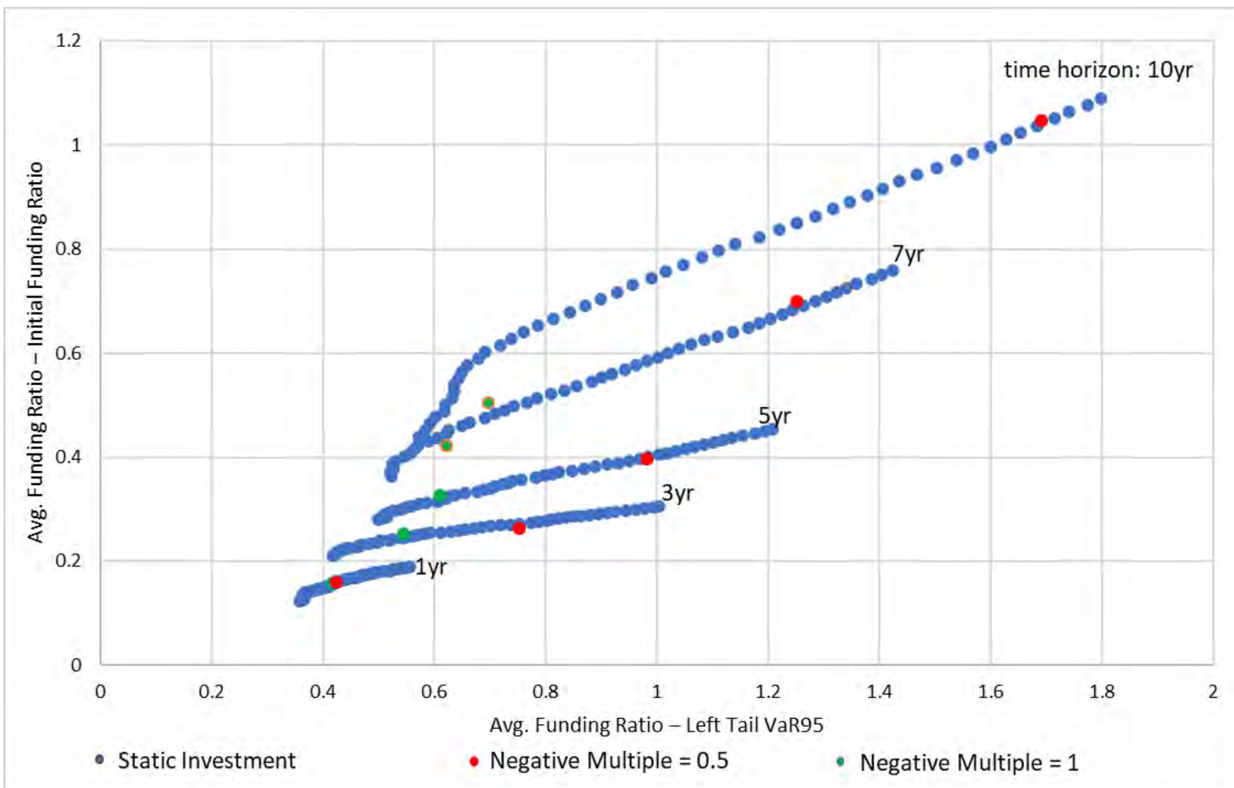

Blue dots stand for efficient frontiers at different time horizons based on static investment strategies, that is, a constant asset mix. Green dots stand for the tradeoff between return and risk measures at different time horizons for the RL dynamic strategy when equal weights are given to positive and negative changes. Red dots represent those for the RL dynamic strategy when half weight is applied to negative changes. Unlike static investment strategies where multiple asset allocations plans can be evaluated, normally only one dynamic strategy is available given one specification of the reward function. However, multiple sets of dynamic investment strategies can be derived by adjusting the reward function to reflect different degrees of risk aversion through RL. In this example, the discounted periodic change in funding ratio is used as the reward function by applying a multiple that can be different from one to negative changes:

$$Reward_t^i = \sum_{j=t}^{N} DF_{j-t} I_j^i \mathrm{R}_j^{fr}$$

where

$Reward_t^i$: total reward at time *t* for simulation *i*

DF_j: discount factor for reward for period *j* given current time *t*

N: Maximum projection period; in this example, it equals 40, which is the quarterly projection for 10 years

R_j^{fr}: change in funding ratio in period *j*

$$I_j^i = \begin{matrix} \text{negative multiple} & \text{if } \mathrm{R}_j^{fr} < 0 \\ 1 & \text{otherwise} \end{matrix}.$$

For some time horizons, RL improves the existing efficient frontier based on optimal static strategies. For a negative multiple greater than 1, RL will generate the same dynamic strategy, because when a negative multiple equals 1, as in this action space, it is already the most conservative strategy.

Figure 26 shows the portion of bond investment in different strategies. With the negative multiple equal to 1, the dynamic investment strategy in this defined action space is increasing the bond investment by 2% each quarter until all investment is in bonds. When the negative multiple equals 0.5, it generates the most aggressive strategy in the allowed action space: increasing the equity investment by 2% each quarter until all investment is in equity.

Figure 26
Bond Investment Portion: Dynamic Strategy with Rebalance Constraint

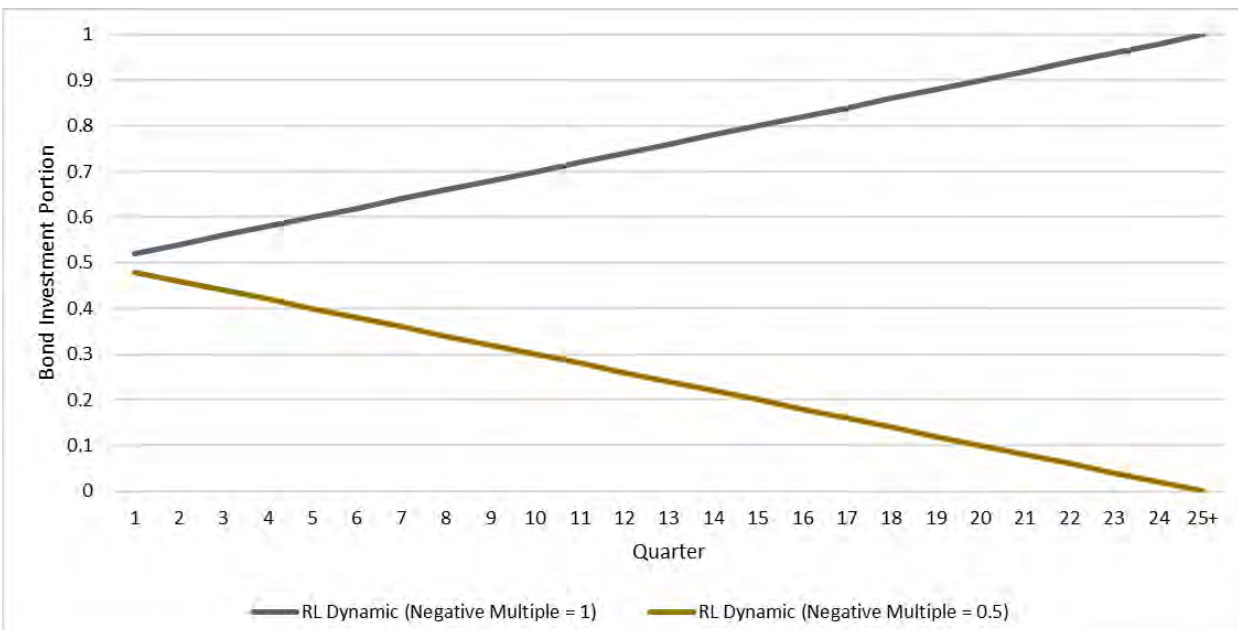

To evaluate the impact of the funding ratio on a dynamic investment strategy, an initial underfunded status (80% funding ratio) is tested as well with a negative multiple of 1. The generated dynamic investment strategy is decreasing the bond investment by 2% each quarter till it is fully invested in equity. This is understandable given the higher probability of closing the funding gap with an equity investment.

Based on the analysis till this point, it looks like that given the allowed actions, RL with deep learning models will generate strategies that converge to optimal static investment strategies. However, a few findings are worth noting:

- The dynamic strategy does not try to win in each scenario but rather makes decisions based on the distribution as a whole, no matter what the past experience is. As shown in Section 4.3, it is unlikely to gain by market timing based on the ESG used in this research. Rather, the expected return and volatility are important determinants of selected strategies.

- The generated dynamic strategy is a nearly "static" nonconstant investment strategy given initial funding status. In most cases, when bad experience kicks in, it does not try to react in a tactic way, which means it believes the economic cycles in the ESG will bring the ultimate results to the expected level. This makes us wonder why we do not use static nonconstant investment strategies in the first place with RL. By running through different asset allocation paths, improved efficient frontiers can be directly drawn. However, this is a very expense exercise even in a two-asset-class space. With 40 decision times (a quarterly rebalance for 10 years), using a step of 2% when determining the asset mix (51 possible asset mixes, starting from 0%, 2%, …, till 100% bond investment), it requires 51^{40} asset mix paths to get the full picture. This seems to be an impossible task given the complexity of liability details even with enormous computing power. However, with RL, we can spend much less time to find a few points on the improved efficient frontiers.

- Given the way of defining reward function in this analysis, the dynamic strategy tries to find a balanced strategy across all time horizons in the allowed action space. The choice of balancing is determined by the discount rate. This also means a longer time horizon will have a lower degree of risk aversion.

- The RL-generated dynamic investment strategies are reasonable given allowed actions, although they may not be optimal. Given the limited paths explored by RL in the entire space, there is no guarantee that RL will generate the optimal strategy.

With these findings in mind, the rest of Section 4.4 focuses on the same analysis in different settings to see if improvements over efficient frontiers attainable by static investment strategies can be made.

4.4.2 Two Asset Classes without Rebalance Constraint

Given the limited actions, further improvements may be achieved by relaxing the rebalance constraint. As before, model convergence is assessed with satisfactory results, shown in Table 6.

Table 6

MSE/Average Reward Value for Two Asset Classes without Rebalance Constraint

Optimization Step	FCNN	LSTM
Step 20,000	7.16%	44.81%
Step 40,000	7.94%	10.20%
Step 6,400,000*	3.67%	8.18%
Benchmark	23.90%	23.60%

Note: *The last step in model training.

Figure 27 shows the dynamic asset allocation plans that may improve the efficient frontiers. In this example, a negative multiple of 1 is used for both FCNN and LSTM. FCNN outperforms existing efficient frontiers at almost every time horizon. On the other hand, LSTM generates a constant strategy balancing the influence of different time horizons and therefore stays on the existing efficient frontiers.

Figure 27

Dynamic Asset Allocation: Two Asset Classes without Rebalance Constraint

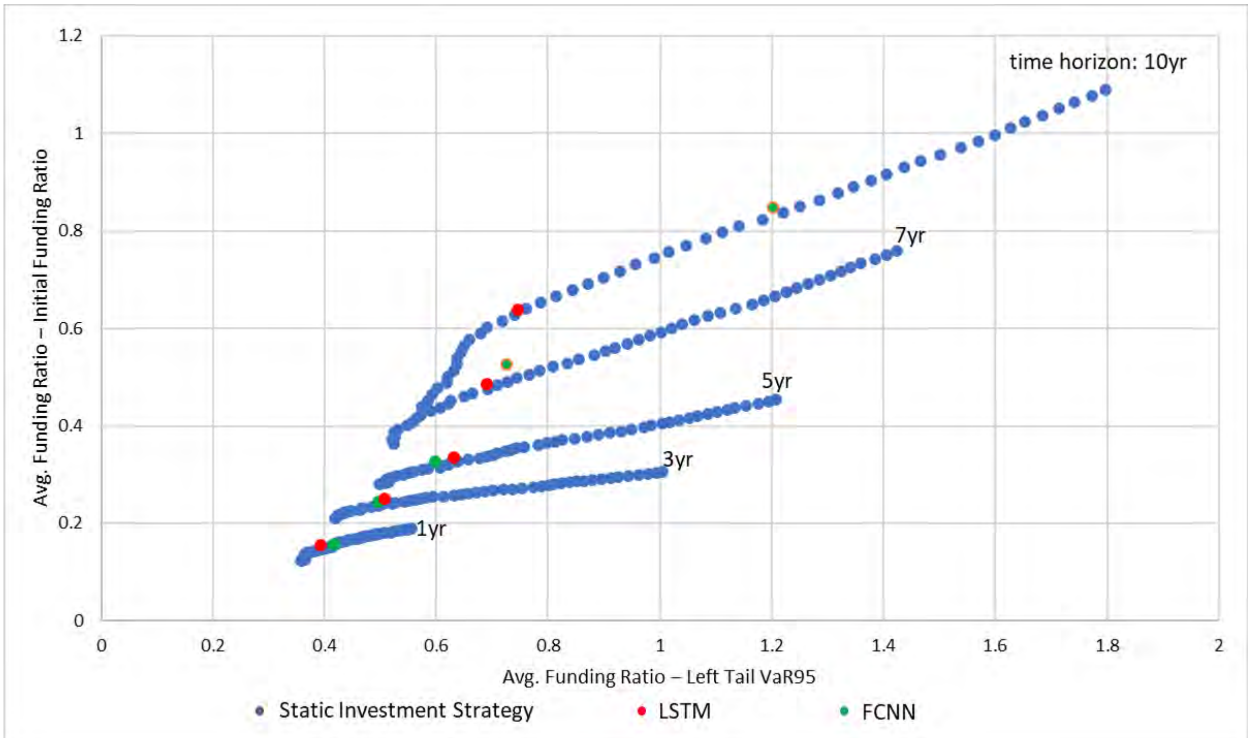

Figure 28 shows the average dynamic asset mix for both FCNN and LSTM. FCNN generates a more promising strategy with a lower bond investment initially, recognizing a lower bond fund return initially, and gaining higher expected equity return for longer return when the degree of risk aversion is implicitly lower.

Figure 28

Bond Investment Portion: Dynamic Strategy without Rebalance Constraint

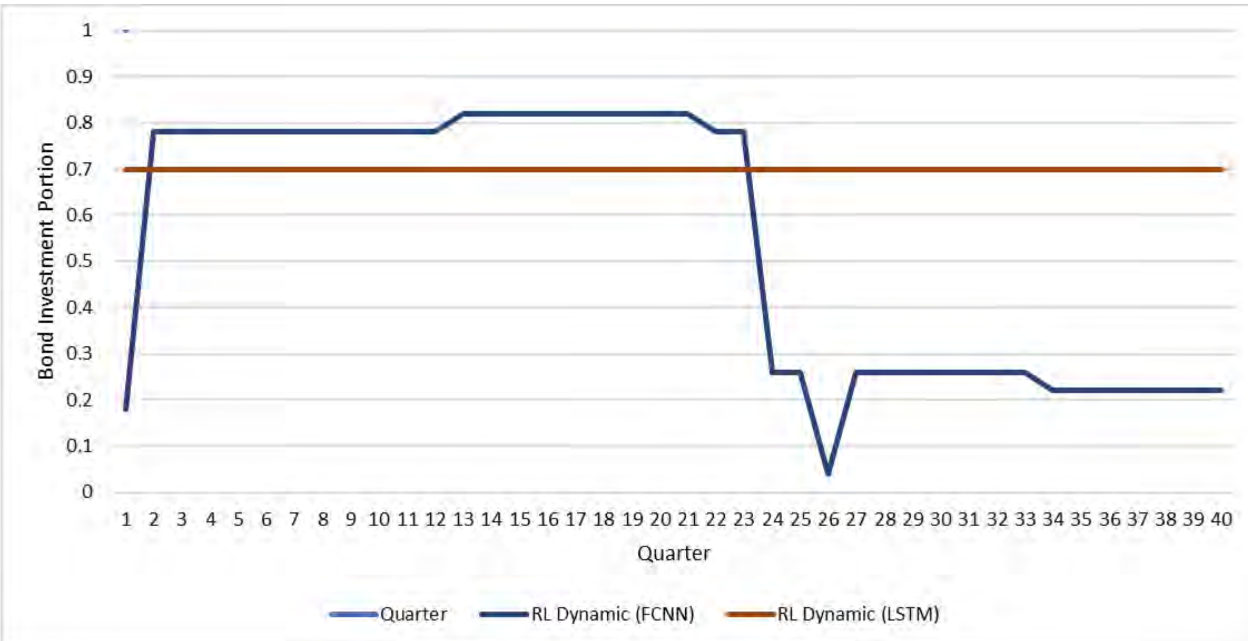

4.4.3 Four Asset Classes without Rebalance Constraint

In addition to AA-rated corporate bonds and large-scale public equity, two additional asset classes are added into the plan: BBB-rated corporate bonds and equity REITs. Table 7 shows worse convergence than in a two-asset-class case but still lower than the benchmark.

Table 7
MSE/Average Reward Value for Four Asset Classes without Rebalance Constraint

Optimization Step	FCNN	LSTM
Step 20,000	52.33%	53.56%
Step 40,000	34.99%	39.86%
Step 6,400,000*	20.32%	22.35%
Benchmark	23.90%	23.60%

Note: *The last step in model training.

In a four-asset-class space, the construction of an efficient frontier can be challenging. The number of possible asset mixes can be calculated as combinations with replacement. With 51 possible dividers within the range between 0% and 100% with an increment of 2%, three dividers need to be chosen to determine the portion assigned to the four asset classes. The chosen dividers can be the same, which means one or more asset classes have zero allocation:

$$C^R(n,r) = \frac{(n+r-1)!}{r!(n-1)!}$$

where

n: the number of possible dividers

r: the number of chosen dividers, which can be repetitive.

This requires 23,426 static asset plans to be tested to get efficient frontiers similar to those shown before. If the asset allocation path is considered so that static asset plans can be nonconstant, $23{,}426^{40}$ asset allocation paths need to be tested to determine the best one using the grid-searching approach. It is also a challenge, but a less severe one, for RL because at each decision point one of the 23,426 plans needs to be chosen. In this subsection, efficient frontiers based on two asset classes are still used as a benchmark with the expectation that the RL strategies will further improve the efficient frontier compared to the previous subsection. FCNN is used choosing from 1,771 asset mixes, with the allocation on each asset class ranging from 0% to 100% with an increment of 5%. Figure 29 shows the further improvement by bringing more asset classes, and RL is able to generate a better strategy.

Figure 29

Dynamic Asset Allocation: Four Asset Classes without Rebalance Constraint

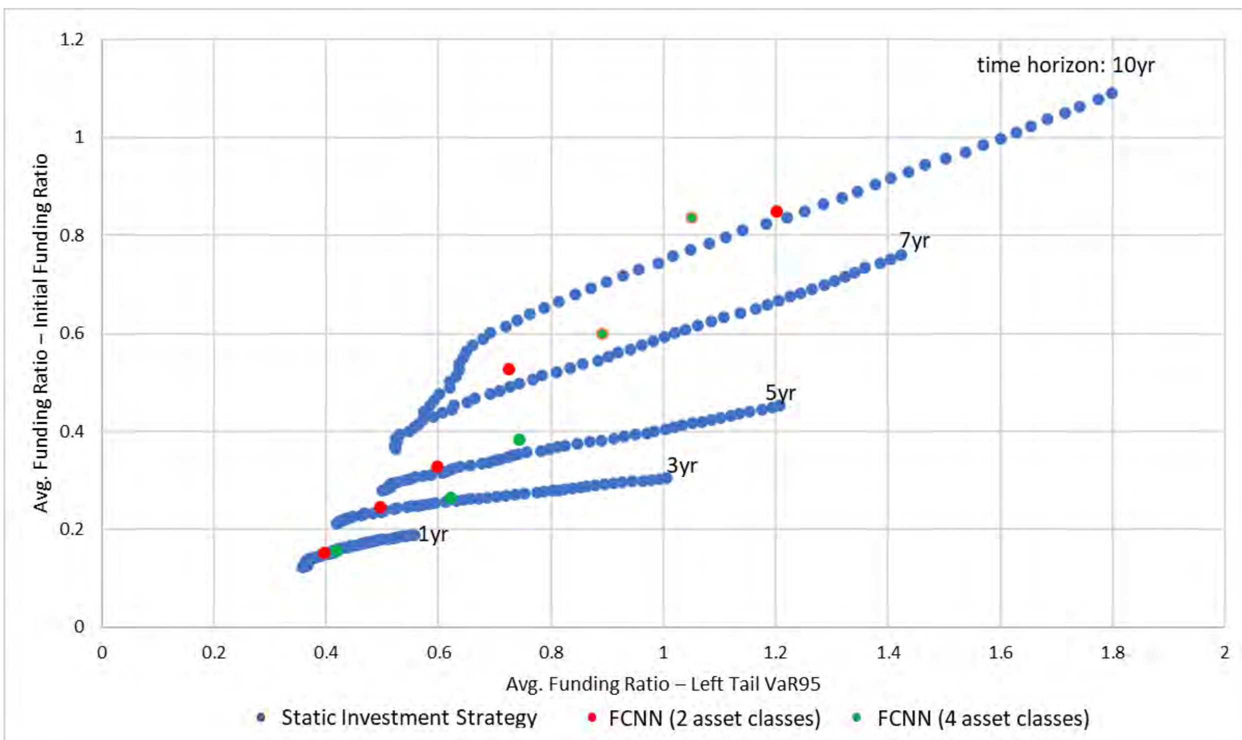

Figure 30 shows the asset mix changing in 10 years.

Figure 30

Asset Mix: Dynamic Strategy without Rebalance Constraint

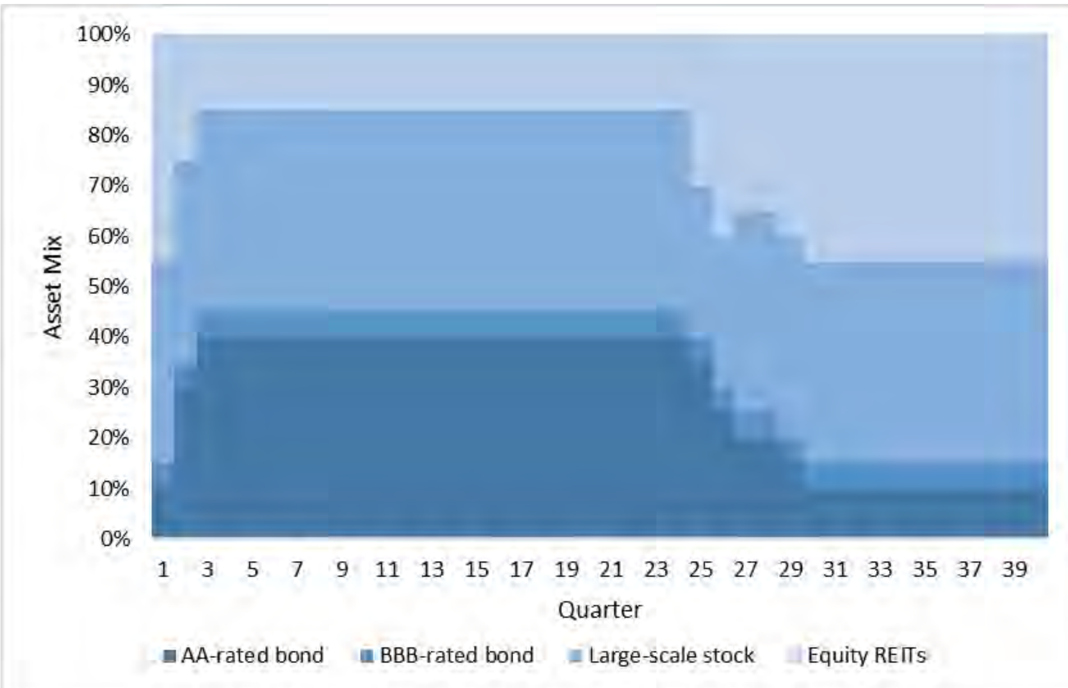

Section 5: Further Developments

This research provides a general framework for using AI models to find an appropriate dynamic LDI strategy. The DB plan example discussed in Section 4 is only a small and simple application. It can be extended to reflect more sophisticated patterns and solve other problems.

1. The ESG is the core of the experiment defining the view of the economic system and asset markets. It should be consistent with the view of the future economy and capital market. The VAR model used in the example is only one possible model to describe the economy. More sophisticated models such as DSGE and agent-based modeling can be used to describe the economy with cause-and-effect relationships and forward-looking views. In the ESG, linear models are used to map asset class returns to fundamental economic factors. Asset markets could be described by other models as well. For asset classes that are not very liquid, a more sophisticated model that considers irrational behaviors may be used to better reflect tail risk. Constraints on asset rebalance and transaction costs may be applied for real applications as well.
2. While at maximum four asset classes are tested in this research, the ESG contains many more asset classes that can be included in searching for an optimal asset mix. It is expected to further improve the RL strategy but at a cost of computing power and training time. However, the cost is not a critical obstacle to incorporating more asset classes. Alternatively, dimensionality reduction may be used to consolidate asset classes with similar distributions and reduce the number of asset classes.
3. The sample DB plan used in this research is an open plan that has a steady liability level and therefore tends to generate an optimal constant asset mix even in a dynamic way. A closed plan may be tested to see if additional expected dynamics can be achieved through RL. Different starting asset mixes, investment goals and views of the future economy and capital markets also may be used to assess their impact on RL.
4. One can include not only a dynamic investment strategy but also a dynamic contribution strategy in the decision space for RL to solve.
5. Unlike static strategies in which one can draw the entire efficient frontier by testing all possible asset allocation plans, RL will generate only the "optimal" asset allocation plan based on the definition of the reward function. By tweaking the weights applied to positive reward and negative reward, the "optimal" strategy will move along the efficient frontier. It is beneficial to adjust the weights and understand their relationship with degree of risk aversion, which is

usually used in a utility function. This will help get the “optimal” strategy in an acceptable range based on the degree of risk aversion.

6. The discount rates used in total reward calculation to discount future rewards are fixed. The rate may be adjusted to depend on the then-current yield curve to reflect the difference in scenarios.
7. When modeling asset selling to meet liability payment requirements, it is assumed that assets are sold proportionally among asset holdings. In practice, different selling orders may be used. Asset managers may also borrow cash for short-term usage to avoid selling asset holdings. These nuances can be modeled and tested to assess the impact on the resulting investment strategy.
8. The framework can be applied to many LDI problems. Insurance companies, DB pension funds or personal financial planning with clearly defined liability can use this framework to find a dynamic investment strategy without compromising liability details and nonlinear relationships. The framework is also not confined to a single economy. For situations with foreign investments, other economies and asset markets can be modeled as well using the same approach.

Section 6: Conclusion

Optimization of a dynamic investment strategy is challenging for LDI in the presence of a dynamic and complex liability portfolio. Dynamic programming that uses a backward induction method to depict all possible outcomes is not a practical option unless the projection of asset and liability is simplified to reduce the dimension. A grid-searching approach that checks all possible investment strategies is not practical given the exponential growth of asset mix paths with more asset classes and a longer time horizon. RL with deep learning models to approximate the reward function is investigated for its potential to optimize a dynamic investment strategy with the level of complexity usually seen in actuarial modeling of assets and liabilities. A few key findings are listed below:

1. With quarterly scenarios and rebalance, investment strategies suggested by RL are driven by average returns rather than periodic fluctuations. This has been tested by using supervised learning with deep learning models to estimate asset returns, which is not effective. RL may be used for tactical asset allocation with higher data frequency but not for strategic asset allocation.

2. The drivers of the dynamic strategies by RL are liability development and funding status. Compared to static strategies, these dynamic strategies are reasonable but not necessarily the best choices and may achieve a better risk return tradeoff in some cases.

3. RL can reflect risk appetite by adjusting the reward function. By increasing the weight on negative reward (penalty) in total reward calculation, RL moves from aggressive to conservative strategies. However, the relationship between level of risk aversion used in the utility function and the weight on negative reward needs to be explored further.

4. RL can incorporate liability complexity, which is challenging for dynamic programming. RL can handle multiple asset classes with reduced training time, which is challenging for a full-blown asset allocation grid search approach.

The flexibility of AI models allows the inclusion of detailed liability information, refined modeling of the relationship between asset and liability, and other information that is important for decision-making. The new approach opens the door to deriving a dynamic strategy for complicated LDI problems.

References

Ang, Andrew, Bingxu Chen and Suresh Sundaresan. 2013. Liability-Driven Investment with Downside Risk. *Journal of Portfolio Management* 40, no. 1:71–87.

Brandt, Michael W., Amit Goyal, Pedro Santa-Clara and Jonathan R. Stroud. 2005. A Simulation Approach to Dynamic Portfolio Choice with an Application to Learning about Return Predictability. *Review of Financial Studies* 18, no. 3:831–873.

Geyer, Alois, and William T. Ziemba. 2008. The Innovest Austrian Pension Fund Financial Planning Model InnoALM. *Operations Research* 56, no. 4:797–810.

Hinton, Geoffrey, Nitish Srivastava, and Kevin Swersky. Neural Networks for Machine Learning. Lecture 6. http://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf.

Martellini, Lionel. 2006. Managing Pension Assets: From Surplus Optimization to Liability-Driven Investment. EDHEC Risk and Asset Management Research Centre, Nice, France. http://www.pfasoftware.com/PFA%20Articles/Managing%20Pension%20Assets.pdf.

Shang, Kailan, and Zakir Hossen. 2019. Liability-Driven Investment: Benchmark Model. https://www.soa.org/resources/research-reports/2019/liability-driven-investment/.

Sharpe, William F., and Lawrence G. Tint. 1990. Liabilities—A New Approach. *Journal of Portfolio Management* 16, no. 2:5–10.

Silver, David, et al. 2017. Mastering the Game of Go without Human Knowledge. *Nature* 550:354–359.

Appendix A: Reinforcement Learning Model Choice

Reinforcement learning has many variations, although a specific type is used in this research: deep Q-learning. Compared to some other RL models, deep Q-learning is different in the following three areas:

- *Model:* Deep Q-learning is a model-free method, which means the transition probabilities among states are not known explicitly given the complexity of the states, including not only economic data but also funding status and asset allocation. Instead, state transitions are realized through simulation and in this case including economic scenario generation and dynamic asset and liability projections. If transition probabilities can be specified explicitly, dynamic programming may be sufficient to solve the problem, and the need for RL is minimal.

- *Learning method:* Temporal difference is used to improve the reward function $Q(s, a)$ as follows:

  $$Q(s,a) + lr(r + \gamma \max_{a'} Q(s', a') - Q(s,a))$$

  where

  lr: learning rate

  At each state for each action, one simulation gives only one scenario-dependent result. Clearly, many simulations are needed to learn the expected value of the reward function. On the other hand, the reward function can be estimated using a full-blown stochastic calculation, running through possible states and actions in the future. However, with more dimensions than a typical dynamic programing exercise, a full-blown stochastic calculation is unlikely to be finished within a reasonable time framework at a reasonable cost.

- *Policy:* Using temporal difference, an off-policy approach is used for training purposes. In the RL field, on-policy means that only the action that generates the highest value will be chosen in model runs. Off-policy means that the action used in model training is random for a certain portion and certain periods to allow the model to learn the value of different policies. This is beneficial for deep Q-learning to avoid converging to local maxima.

Although deep Q-learning is considered an appropriate approach for LDI analysis with liability details and complex ESG, it has some disadvantages. When the state is defined with many variables, especially continuous variables, model training can be challenging because the deep learning model is supposed to be able to reflect the impact of these variables all together. To achieve a satisfactory accuracy level, it requires a large volume of simulated data and a long training time, which are not always practical. Furthermore, the chosen action is the one with a maximum value of the reward function. Even though off-policy is used, it may still lead to a nonoptimal strategy if the majority of the chosen actions are the ones with the largest reward. Last but not the least, each action has its own reward function to train. Therefore, the number of possible actions is limited. This can be an issue when multiple asset classes are included in the analysis and the number of possible asset mixes explodes. However, given other constraints on investment strategy, the limit on the number of actions may not be a bottleneck in practice.

Two alternatives to deep Q-learning are available in the RL field that may address these potential disadvantages of deep Q-learning but at the same time have their own shortcomings.

Policy gradient methods avoid the issue of selecting the best action based on a comprehensive reward function that takes care of all states and actions. Instead, they directly optimize the policy used to choose actions. This somehow avoids the issue of undesired model convergence performance and can deal with continuous states and actions. But unlike deep Q-learning where simulations can be repeatedly fed into the training process, they can be used only once for policy gradient methods. Policy gradient methods are on-policy methods, which means they will always use the best action estimated. This can be an issue for a complex LDI analysis in which simulation takes many computing resources. However, policy gradient methods may be helpful when the ESG and dynamic liability projection are simplified to reduce simulation time.

Actor-Critic methods learn both the reward function and the policy to make sure the model converges to some meaningful solutions. However, they are also on-policy methods, and simulations can be used only once during training.

Although other model choices are available and may outperform deep Q-learning in some cases, deep Q-learning is still considered to be a practical choice given its capability for repeatedly using training data.

Appendix B: Economic Scenario Generation

The appendix is an excerpt of the appendix in Shang and Hossen (2019). It provides technical notes of the economic scenario generation model used in experiments of deep learning model training. Details of the reasonableness of individual scenarios can be found in that paper.

B.1 Economic Scenario Generation for Fundamental Economic Factors

In this report, a vector autoregressive (VAR) model is used to generate scenarios for fundamental economic factors, including real GDP growth rate, inflation rate, unemployment rate, short-term interest rate, long-term interest rate, credit spread, personal consumption growth rate and investment growth rate. Asset return scenarios are determined based on their relationship with these fundamental economic factors. Table B.1 describes the historical data of fundamental risk factors used in this example.

Table B.1
Fundamental Economic Factor Historical Data

Fundamental Economic Factor	Indicator	Notation	Data Source
Real GDP growth rate	U.S. quarterly notional GDP growth rate—inflation rate	gdpgr	Bureau of Economic Analysis (seasonally adjusted at annual rates)
Inflation rate	CPI—all urban consumers (current series)	cpi	Bureau of Labor Statistics
Unemployment rate	U-3 rate	unemploy	Bureau of Labor Statistics
Short-term interest rate	U.S. three-month Treasury bill rate	m3tb	Bloomberg (USGG3M Index)
Long-term interest rate	U.S. 10-year Treasury bond yield	tb10y	Bloomberg (USGG10Y Index)
Credit spread	U.S. AA-rated finance corporate bond 10-year credit spread	aa10y	Bloomberg (c02310Y Index)
Consumption growth rate	U.S. personal consumption expenditures	pconsump	Bureau of Economic Analysis (seasonally adjusted at annual rates)
Investment growth rate	U.S. gross private domestic investment	gpdinv	Bureau of Economic Analysis (seasonally adjusted at annual rates)

Quarterly historical data from 1991Q1 to 2016Q4 are used. This analysis covers three full economic cycles. A VAR model is used to describe the relationship of the fundamental economic factors based on these historical data. By incorporating lagging variables into the analysis through VAR, relationships among leading, coincident and lagging economic factors can be better reflected. For example, the short-term interest rate is largely controlled by the Fed, after reviewing economic growth, unemployment and other economic conditions. Time is needed before making rate decisions. For simplicity, VAR(1) is used so that the evolution of fundamental economic factors is affected by their values in the previous quarter. A quarter is likely to be enough for the interaction among fundamental economic factors. Having a higher order of VAR model can improve the results only marginally in this example:

$$\mathbf{F}_t = \mathbf{c} + \mathbb{A}_1\mathbf{F}_{t-1} + \mathbf{e}_t$$

where

$\mathbf{F}_t =$ $(\mathrm{gdpgr}_t, \mathrm{cpi}_t, \mathrm{unemploy}_t, \mathrm{m3tb}_t, \mathrm{tb10y}_t, \mathrm{aa10y}_t, \mathrm{pconsump}_t, \mathrm{gpdinv}_t)^T$, a column vector with eight elements as the value of fundamental economic factors at time t or during period t

$\mathbf{c} =$ a column vector with eight elements to represent the constant terms of the eight fundamental economic factors

$\mathbb{A}_1 =$ an 8×8 matrix containing the model parameters describing the linear dependence of fundamental economic factors and

$\mathbf{e}_t =$ a column vector with eight elements to store the error terms that cannot be explained by linear models.

Table B.2 shows the fitted model parameters ($\mathbb{A}_1$ and $\mathbf{c}$) based on the historical data. It also shows σ, the standard deviation of error vector $\mathbf{e}_t$.

Table B.2
VAR(1) Model Parameters

Variable	$\mathbb{A}_1$								c	σ
	unemploy	gpdinv	pconsump	gdpgr	m3tb	aa10y	tb10y	cpi		
gdpgr	−0.04	0.09	0.49	−0.48	−0.07	−0.25	0.06	−0.45	0.83	0.59
cpi	0.09	−0.06	0.05	0.14	0.13	−0.22	−0.09	−0.03	0.27	0.59
unemploy	0.94	−0.01	−0.13	−0.02	−0.03	0.29	0.09	0.02	−0.12	0.20
m3tb	−0.01	0.01	0.04	0.08	0.92	−0.15	0.03	−0.09	0.22	0.43
tb10y	0.03	−0.02	0.00	0.02	0.07	−0.14	0.86	−0.06	0.41	0.52
aa10y	−0.07	0.04	0.09	−0.09	−0.03	1.02	0.02	−0.03	0.30	0.31
pconsump	0.01	0.00	−0.03	0.08	0.06	−0.54	−0.02	0.10	1.58	0.50
gpdinv	0.71	0.15	3.67	−1.24	0.06	−1.62	−0.26	−1.79	−3.08	1.95

Based on the fitted VAR(1), the stable values of fundamental risk factors $\bar{\mathbf{F}}$ can be derived:

$$\bar{\mathbf{F}} = \mathbf{c} + \mathbb{A}_1\bar{\mathbf{F}}$$

Table B.3 lists the stable values based on VAR(1) along with the historical mean and standard deviation. The VAR(1) suggests a lower future economic growth rate, inflation rate, interest rates, consumption and investment growth rate than in the past 26 years. Credit spread is expected to be a little bit higher, and the unemployment rate is expected to stay at the same level. These model-implied expectations are good checkpoints to assess the model's reasonableness against a model user's view on future economic development.

Table B.3
VAR(1) Stable Values

Variable (Quarterly)	VAR(1)	Historical Data	
	Stable Value (%)	Mean (%)	Standard Deviation (%)
gdpgr (quarterly)	0.48	0.54	0.68
cpi (quarterly)	0.44	0.57	0.59
unemploy	6.06	6.06	1.58
m3tb	0.94	2.62	2.18
tb10y	3.17	4.51	1.75
aa10y	1.29	1.09	0.67
pconsump (quarterly)	0.99	1.18	0.63
gpdinv (quarterly)	0.69	1.23	3.05

The error terms of fundamental economic factors are expected to be zero. However, the error terms are not independent of each other. It is important for the economic scenario generator to capture the correlation when developing future scenarios. Table B.4 shows the correlation matrix of the error vector $\mathbf{e}_t$.

Table B.4
VAR(1) Error Term Correlation Matrix

	gdpgr	cpi	unemploy	m3tb	tb10y	aa10y	pconsump	gpdinv
gdpgr	1.00	−0.68	−0.07	0.15	0.03	−0.01	−0.14	0.45
cpi	−0.68	1.00	−0.17	0.18	0.33	−0.24	0.62	0.04
unemploy	−0.07	−0.17	1.00	−0.33	-0.12	0.20	−0.22	−0.26
m3tb	0.15	0.18	−0.33	1.00	0.48	−0.31	0.33	0.23
tb10y	0.03	0.33	−0.12	0.48	1.00	−0.39	0.37	0.24
aa10y	−0.01	−0.24	0.20	−0.31	−0.39	1.00	−0.34	−0.14
pconsump	−0.14	0.62	−0.22	0.33	0.37	−0.34	1.00	−0.01
gpdinv	0.45	0.04	−0.26	0.23	0.24	−0.14	−0.01	1.00

Based on VAR(1), stochastic scenarios of fundamental economic factors can be constructed as

$$\mathbf{F}_t = \mathbf{c} + \mathbb{A}_1\mathbf{F}_{t-1} + \boldsymbol{\sigma} \cdot \mathbb{L}\boldsymbol{\varepsilon}_t$$

where

$\boldsymbol{\sigma}$ is a column vector containing the standard deviation of the error terms of eight fundamental economic factors
$\boldsymbol{\varepsilon}_t$ is a column vector containing eight independent random variables following the standard normal distribution and
$\mathbb{L}$ is an 8×8 lower triangular matrix so that the error term correlation matrix can be decomposed as $\mathbb{L} \times \mathbb{L}^{\boldsymbol{T}}$.

Using Cholesky decomposition, a correlation matrix $\mathbf{CM}$ such as that in Table B.4 can be decomposed as the product of a lower triangular matrix $\mathbb{L}$ and its transpose $\mathbb{L}^{\boldsymbol{T}}$, given that the correlation matrix is positive definite. $\mathbb{L}\boldsymbol{\varepsilon}_t$ has the same correlation matrix from which $\mathbb{L}$ is derived, as shown below:

$$\mathrm{Cov}(\mathbb{L}\boldsymbol{\varepsilon}_t, \mathbb{L}\boldsymbol{\varepsilon}_t) = \mathbb{E}(\mathbb{L}\boldsymbol{\varepsilon}_t(\mathbb{L}\boldsymbol{\varepsilon}_t)^{\boldsymbol{T}}) = \mathbb{E}(\mathbb{L}\boldsymbol{\varepsilon}_t(\boldsymbol{\varepsilon}_t)^{\boldsymbol{T}}\mathbb{L}^{\boldsymbol{T}}) = \mathbb{L}\mathbb{E}(\boldsymbol{\varepsilon}_t(\boldsymbol{\varepsilon}_t)^{\boldsymbol{T}})\mathbb{L}^{\boldsymbol{T}} = \mathbb{L} \times \mathbb{I} \times \mathbb{L}^{\boldsymbol{T}} = \mathbb{L} \times \mathbb{L}^{\boldsymbol{T}} = \mathbf{CM}$$

where $\mathbb{I}$ is an 8×8 identity matrix, because $\boldsymbol{\varepsilon}_t$ contains independent random variables, all of which have an expected value of zero.

The economic status is projected for each quarter under each scenario. This not only is helpful for understanding the scenarios but also is critical for asset return generation because it can differentiate between economic recession and expansion. During an economic recession, higher volatility and correlation are usually observed and need to be reflected in stochastic asset returns. Real GDP growth rate, unemployment rate, consumption growth rate and investment growth rate are used in the following logistic model to predict whether the economy is in recession:

$$R_t = \frac{1}{1 + e^{-\boldsymbol{\beta}\mathbf{F}_t}}$$

where

R_t = the probability that the economy is in recession during period t
$\mathbf{F}_t = \begin{pmatrix} 1, \mathrm{gdpgr}_t, \mathrm{gdpgr}_{t-1}, \mathrm{gdpgr}_{t-2}, \mathrm{unemploy}_t, \mathrm{unemploy}_{t-1}, \mathrm{unemploy}_{t-2}, \\ \mathrm{pconsump}_t, \mathrm{pconsump}_{t-1}, \mathrm{pconsump}_{t-2}, \mathrm{gpdinv}_t, \mathrm{gpdinv}_{t-1}, \mathrm{gpdinv}_{t-2} \end{pmatrix}^T$, a column vector with 13 elements containing the constant term and fundamental economic factors during periods t, t − 1 and t − 2 and
$\boldsymbol{\beta}$ is a row vector with 13 elements containing the model parameters for variables in $\mathbf{F}_t$.

Historical data for U.S. economic cycles (1999Q1 to 2016Q4) from the National Bureau of Economic Research (NBER) and fundamental economic factors are used to calibrate the logistic model, with parameters shown in Table B.5. Current and previous two quarters of values of fundamental economic factors are used to determine the current status of the economy. Changes from quarter to quarter are more important than the absolute value of economic factors for predicting the economic status. The calibrated logistic model has 100% accuracy in matching the history of the past 26 years.

Table B.5
Economic Recession Prediction: Logistic Model Parameter

Variable	Period	Parameter (β)
Intercept		78.0
Unemployment rate	Current quarter *t*	66.3
	Previous quarter *t* − 1	−51.3
	Previous quarter *t* − 2	−32.2
Investment growth rate	Current quarter *t*	0.3
	Previous quarter *t* − 1	−3.4
	Previous quarter *t* − 2	4.3
Consumption growth rate	Current quarter *t*	−1.2
	Previous quarter *t* − 1	−9.9
	Previous quarter *t* − 2	0.1
Real GDP growth rate	Current quarter *t*	−25.3
	Previous quarter *t* − 1	−8.4
	Previous quarter *t* − 2	−17.0

B.2 Economic Scenario Generation for Asset Return

With the fundamental stochastic scenarios generated, return scenarios of each asset classes can be constructed based on their relationships with fundamental economic factors. Linear models are used to describe the relationships based on historical data:

$$\mathbf{y}_t = \boldsymbol{\alpha} + \boldsymbol{\phi}_1\mathbf{y}_{t-1} + \boldsymbol{\phi}_2\mathbf{y}_{t-2} + \mathbb{B}_0\mathbf{F}_t + \mathbb{B}_1\mathbf{F}_{t-1} + \mathbb{B}_2\mathbf{F}_{t-2} + e_t$$

where

$\mathbf{y}_t$ is a column vector containing the returns of all asset classes during period t
$\boldsymbol{\alpha}$ is a column vector containing the constant terms of all asset classes
$\boldsymbol{\phi}_i$ is a column vector containing parameters to govern the relationship between current return and return for all asset classes during period $t-i$
$\mathbb{B}_i$ is a matrix with eight columns that contains all asset classes' model parameters for the eight fundamental economic factors during period $t-i$ and
$\mathbf{F}_t = (\mathrm{gdpgr}_t, \mathrm{cpi}_t, \mathrm{unemploy}_t, \mathrm{m3tb}_t, \mathrm{tb10y}_t, \mathrm{aa10y}_t, \mathrm{pconsump}_t, \mathrm{gpdinv}_t)^T$, a column vector with eight elements as the value of fundamental economic factors at time t or during period t.

Table B.6 lists the asset classes and historical data used to calibrate the linear models.

Table B.6
Asset Return Historical Data

Asset Class	Return Type	Time Period	Data Source
Treasury bond yield curve (terms: 1, 2, 3, 5, 7, 10, 20 and 30 years)	Yield	1991Q1–2016Q4 except for 20-year bond yield starting from 1993Q4	U.S. Department of the Treasury
AAA-, AA-, A- and BBB-rated corporate bonds	Credit spread	1996Q4–2016Q4	BofA Merrill Lynch US Corporate AAA, AA, A, BBB Effective Yield (BAMLC0A3CAEY, BAMLC0A2CAAEY, BAMLC0A1CAAAEY, BAMLC0A4CBBBEY) *Federal Reserve Economic Data*
	Default rate	1991Q1–2016Q4	*2016 S&P Annual Global Corporate Default Study and Rating Transitions Report*
Public equity, large cap	Dividend yield	1991Q1–2016Q4	S&P 500 Index (^GSPC) *Yahoo Finance*
	Capital return		
Public equity, mid cap	Dividend yield	2000Q3–2016Q4	iShares Core S&P Mid-Cap (IJH) *Yahoo Finance*
	Capital return		
Public equity, small cap	Dividend yield	2000Q3–2016Q4	iShares Core S&P Small Cap (IJR) *Yahoo Finance*
	Capital return		
Public equity, high dividend yield	Dividend yield	2004Q1–2016Q4	iShares Select Dividend ETF (DVY) *Yahoo Finance*
	Capital return		
Equity REITs	Cap rate	1991Q1–2016Q4	FTSE NAREIT US Real Estate Index—All Equity REITs
	Capital return		
Mortgage REITs	Cap rate	1991Q1–2016Q4	FTSE NAREIT US Real Estate Index—Mortgage REITs
	Capital return		
Wage index	Total return	1991Q1–2016Q4	U.S. Compensation/Employed Compensation of Employees, Received: Wage and Salary Disbursements (A576RC1) and All Employees: Total Nonfarm Payrolls (PAYEMS) *Federal Reserve Economic Data*
Plan sponsor equity	Total return	1991Q1–2016Q4	GE stock price (used as an example) *Yahoo Finance*

Table B.7 lists the parameters of linear models for asset return economic scenario generation used in this report. To avoid overfitting, we removed from the final models explanatory variables with a parameter not statistically different from zero (p-value of t-test > 0.3). The last column of Table B.7 shows the adjusted R^2 of the linear models. It indicates the portion of asset return volatility that can be explained by the linear relationships. Certain asset classes, including infrastructure project index, private equity, mortgage REITs, oil and gold, have a low adjusted R^2. They are driven more by idiosyncratic factors than by fundamental economic factors. The second-to-last column of Table B.7 contains the standard deviation of the idiosyncratic factors, which cannot be explained by the linear models. Stochastic scenarios of asset returns can be generated as

$$\mathbf{y}_t = \boldsymbol{\alpha} + \boldsymbol{\phi}_1 \mathbf{y}_{t-1} + \boldsymbol{\phi}_2 \mathbf{y}_{t-2} + \mathbb{B}_0 \mathbf{F}_t + \mathbb{B}_1 \mathbf{F}_{t-1} + \mathbb{B}_2 \mathbf{F}_{t-2} + \boldsymbol{\sigma} \cdot \boldsymbol{\varepsilon}_t$$

where

$\boldsymbol{\sigma}$ is a column vector containing the standard deviation of error terms of all asset return models and

$\boldsymbol{\varepsilon}_t$ is a column vector containing independent random variables following a standard normal distribution for all asset return models.

Table B.7

Asset Return Model Parameters (Returns in Percentage Format)

			Autocorrelation		Fundamental Risk Factors																									Adjusted R^2
Lag (Quarter)			2	1	2	1	0	2	1	0	2	1	0	2	1	0	2	1	0	2	1	0	2	1	0	2	1	0	0	
Asset Class		Intercept	ϕ_2	ϕ_1	gdpgr			cpi			unemploy			m3tb			tb10y			aa10y			pconsump			gpdinv			σ	
Treasury bond zero rate (term)	1	(0.09)	0.00	0.75	(0.07)	0.00	0.00	(0.01)	0.00	0.00	0.00	0.00	0.00	0.00	(0.68)	0.88	0.00	(0.27)	0.34	0.00	0.00	0.00	0.00	0.00	0.00	0.02	(0.01)	0.00	0.14	100%
	2	(0.20)	0.00	0.80	0.00	0.00	0.00	0.00	0.00	(0.08)	0.00	0.00	0.00	0.15	(0.75)	0.72	0.00	(0.53)	0.63	0.00	0.00	0.00	0.00	0.11	0.00	0.00	0.00	0.00	0.16	99%
	3	(0.23)	0.00	0.75	0.00	0.00	0.00	0.00	0.00	(0.08)	0.00	0.00	0.00	0.17	(0.66)	0.61	0.00	(0.62)	0.77	0.00	0.00	0.00	0.00	0.11	0.00	0.00	0.00	0.00	0.15	99%
	5	(0.24)	0.00	0.74	0.00	0.00	0.00	0.00	0.00	(0.06)	0.00	0.00	0.00	0.08	(0.34)	0.31	0.00	(0.73)	0.97	0.00	0.00	0.00	0.00	0.14	0.00	0.00	0.00	(0.02)	0.10	100%
	7	(0.05)	0.00	0.79	0.00	0.00	(0.04)	0.00	0.00	(0.07)	0.00	0.00	0.00	0.00	(0.11)	0.13	0.00	(0.82)	1.02	0.00	0.00	0.00	0.00	0.04	0.00	0.00	0.00	0.00	0.06	100%
	10	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	100%
	20	0.03	0.00	0.85	0.00	0.00	0.00	0.00	(0.06)	0.00	0.00	0.00	0.00	(0.05)	0.23	(0.18)	0.00	(0.76)	0.92	0.00	0.04	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.08	100%
	30	0.22	0.00	0.78	0.00	0.00	(0.07)	0.00	(0.05)	0.00	0.00	0.00	0.00	(0.08)	0.28	(0.22)	0.00	(0.65)	0.88	0.00	0.09	(0.08)	0.05	(0.10)	0.00	0.00	0.00	0.03	0.09	100%
AAA-rated corporate bond	Credit Spread	0.18	0.00	0.87	(0.16)	(0.11)	0.00	(0.16)	(0.06)	0.01	0.00	0.00	0.00	0.00	(0.22)	0.25	0.00	0.14	(0.17)	0.02	(0.77)	0.76	(0.03)	0.14	0.00	0.02	0.00	0.00	0.15	92%
	Default Rate	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	100%
AA-rated corporate bond	Credit Spread	(0.06)	0.00	0.76	(0.17)	0.00	0.00	(0.16)	0.00	0.00	0.00	0.00	0.00	0.00	(0.20)	0.27	0.00	0.16	(0.23)	0.00	(0.65)	0.91	0.00	0.11	0.00	0.02	0.00	0.00	0.17	95%
	Default Rate	0.01	(0.21)	0.92	0.00	0.00	0.00	0.00	0.04	0.00	0.00	0.07	(0.07)	0.00	0.05	(0.05)	0.00	0.00	0.01	(0.03)	0.05	0.00	(0.01)	0.00	0.00	0.00	0.00	0.00	0.04	73%
A-rated corporate bond	Credit Spread	(0.10)	0.00	0.59	(0.14)	(0.13)	(0.02)	(0.15)	(0.18)	0.00	0.00	0.00	0.00	0.00	(0.19)	0.28	0.00	0.18	(0.26)	0.00	(0.69)	1.25	0.00	0.21	0.00	0.01	0.00	0.00	0.16	96%
	Default Rate	(0.09)	0.00	0.62	0.01	0.00	(0.02)	0.02	0.00	(0.03)	0.06	0.00	(0.08)	0.00	0.05	(0.06)	0.00	0.00	0.03	0.00	0.05	0.05	(0.03)	0.00	0.05	0.00	0.00	0.00	0.05	79%
BBB-rated corporate bond	Credit Spread	0.10	0.00	0.52	0.00	(0.13)	(0.11)	(0.03)	(0.26)	(0.28)	0.00	0.00	0.00	0.00	(0.21)	0.25	0.00	0.39	(0.43)	(0.26)	0.00	0.98	0.00	0.31	0.00	0.00	0.00	0.00	0.20	96%
	Default Rate	0.20	(0.10)	0.65	(0.04)	0.00	(0.11)	(0.05)	0.00	(0.10)	0.00	0.00	(0.07)	0.00	0.00	(0.05)	0.08	0.00	0.00	0.15	0.00	0.00	0.00	0.02	0.00	0.00	0.00	0.03	0.11	76%
Public equity, large cap	Dividend Yield	0.26	0.00	0.91	0.00	0.00	0.00	0.00	0.16	0.00	0.00	0.00	(0.02)	0.06	0.00	(0.08)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.18	88%
	Capital Return	(0.22)	0.00	(0.16)	6.73	2.08	1.24	5.91	1.06	0.00	1.13	0.00	0.00	0.00	0.00	0.00	0.00	(4.45)	4.05	0.00	0.00	(5.15)	(3.49)	0.00	0.00	(1.39)	(0.21)	0.00	5.87	32%
Public equity, mid cap	Dividend Yield	2.35	0.00	0.00	0.37	0.00	0.00	0.40	0.00	(0.23)	0.00	0.00	0.00	0.00	0.00	0.09	0.00	(0.33)	0.00	0.00	0.00	(0.01)	(0.24)	0.00	0.00	(0.06)	0.00	0.00	0.33	43%
	Capital Return	(15.06)	(0.14)	(0.34)	5.44	0.00	0.00	4.01	0.00	0.00	0.00	0.00	2.46	0.93	0.00	0.00	(5.78)	(5.04)	8.55	(0.29)	0.00	0.00	0.00	0.00	6.63	(0.88)	0.00	0.00	5.86	52%
Public equity, small cap	Dividend Yield	2.13	(0.13)	0.00	0.53	0.00	0.00	0.37	0.00	0.00	0.00	(0.46)	0.56	0.15	(0.28)	0.25	0.00	(0.41)	(0.01)	0.00	(0.57)	0.36	(0.33)	0.00	0.00	(0.05)	0.00	0.00	0.29	51%
	Capital Return	(19.71)	(0.18)	(0.50)	8.08	3.27	3.61	7.53	0.31	(2.07)	1.07	(2.05)	3.50	1.77	3.66	(4.45)	(6.57)	(6.58)	10.41	7.64	(4.42)	(2.42)	(2.00)	3.11	6.03	(0.91)	(0.49)	0.00	4.85	59%
Public equity, high dividend yield	Dividend Yield	1.42	0.00	0.41	0.00	0.00	(0.16)	0.00	0.00	(0.14)	0.00	0.00	0.00	0.00	0.20	(0.19)	0.00	0.00	0.00	0.00	0.42	0.00	0.00	0.00	0.25	0.00	0.00	0.03	0.19	80%
	Capital Return	3.42	0.00	(0.14)	0.00	0.00	0.00	0.00	0.00	0.74	0.00	0.00	0.00	(6.45)	5.81	0.00	0.00	0.00	0.00	9.05	(10.17)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	4.97	55%
Equity REITs	Cap Rate	0.05	0.42	0.48	0.00	0.00	0.00	0.00	0.00	(0.09)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	(0.24)	0.27	0.00	0.00	0.00	0.08	0.00	0.00	0.00	0.19	79%
	Return	(6.10)	0.00	(0.16)	3.81	2.23	0.00	6.28	4.77	0.00	0.00	0.00	0.00	(3.28)	6.55	(4.11)	0.00	0.00	0.00	16.04	0.00	(15.38)	0.00	0.00	0.00	0.00	0.00	0.00	6.45	49%
Mortgage REITs	Cap Rate	2.56	0.00	0.00	0.00	0.00	0.00	0.00	0.87	0.00	(0.03)	0.12	0.00	0.00	0.00	0.00	0.08	0.00	0.00	0.00	0.00	0.00	0.00	(0.91)	0.00	0.00	0.00	0.04	1.15	11%
	Return	1.85	0.15	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	6.55	0.00	(8.12)	0.00	0.00	0.00	0.00	0.00	(0.74)	9.82	14%

Based on the fitted linear models, the stable values of asset return $\bar{y}$ can be derived as follows:

$$\bar{\mathbf{y}} = \boldsymbol{\alpha} + \boldsymbol{\phi}_1\bar{\mathbf{y}} + \boldsymbol{\phi}_2\bar{\mathbf{y}} + \mathbb{B}_0\bar{\mathbf{F}} + \mathbb{B}_1\bar{\mathbf{F}} + \mathbb{B}_2\bar{\mathbf{F}}$$

Table B.8 compares the stable values with average historical returns. As expected, differences are noticeable for asset classes with a low adjusted R^2. For asset classes with a high adjusted R^2, historical trends could also lead to a difference. For example, the stable values of interest rates are much lower than historical averages because of the downward trend in the historical data. These need to be checked to make sure the models are consistent with the model user's view of the future economy.

Table B.8
Asset Return Linear Model Stable Value

Asset Class		Linear Model	Historical Data	
		Stable Return (%)	Mean Return (%)	Standard Deviation (%)
Treasury bonds, 0 rate (for given term)	1 year	1.24	2.93	2.23
	2 years	1.58	3.25	2.22
	3 years	1.84	3.49	2.15
	5 years	2.42	3.93	1.98
	7 years	2.84	4.26	1.85
	10 years	3.17	4.51	1.75
	20 years	3.72	4.82	1.54
	30 years	3.82	5.05	1.50
AAA-rated corporate bonds	Credit spread	0.42	0.47	0.59
	Default rate	0.00	0.00	0.00
AA-rated corporate bonds	Credit spread	0.44	0.58	0.79
	Default rate	0.02	0.02	0.08
A-rated corporate bonds	Credit spread	0.99	1.04	0.94
	Default rate	0.06	0.06	0.11
BBB-rated corporate bonds	Credit spread	1.90	1.83	1.08
	Default rate	0.18	0.18	0.25
Public equity, large cap	Annualized dividend yield	2.05	2.00	0.52
	Capital return (quarterly)	1.76	2.04	7.66
Public equity, mid cap	Annualized dividend yield	1.36	1.29	0.46
	Capital return (quarterly)	2.38	2.32	9.39
Public equity, small cap	Annualized dividend yield	1.13	1.12	0.48
	Capital return (quarterly)	2.78	2.57	9.92
Public equity, high dividend	Annualized dividend yield	3.57	3.60	0.45
	Capital return (quarterly)	1.49	1.27	7.87
Equity REITs	Cap rate (quarterly)	1.29	1.56	0.43
	Capital return (quarterly)	1.52	1.42	9.53

Asset Class		Linear Model	Historical Data	
		Stable Return (%)	Mean Return (%)	Standard Deviation (%)
Mortgage REITs	Cap rate (quarterly)	2.89	2.98	1.26
	Capital return (quarterly)	−0.81	−0.85	10.79

The relationship between asset returns and fundamental economic factors is not always linear. During an economic recession, higher volatility and correlation are often observed. Asset return economic scenarios need to be further adjusted to reflect the nonlinear relationship. Table B.9 shows the volatility of idiosyncratic factors (error terms) and the correlation between systemic factors (prediction by linear models) and idiosyncratic factors, using either all the data or the data in a recession. Volatility and correlation behaved quite differently in a recession.

Table B.9
Asset Return Linear Model Idiosyncratic Factors: Volatility and Correlation

Asset Class		Idiosyncratic Factor Volatility (%)		Correlation with Systemic Factors	
		All Periods	Recession	All Periods	Recession
Treasury bond, 0 rate (for given term)	1 year	0.1	0.1	0.0%	27.2%
	2 years	0.2	0.1	0.0%	−21.9%
	3 years	0.2	0.2	0.0%	−24.9%
	5 years	0.1	0.1	0.0%	3.7%
	7 years	0.1	0.1	0.0%	18.0%
	10 years	0.0	0.0	0.0%	0.0%
	20 years	0.1	0.1	0.0%	−39.9%
	30 years	0.1	0.1	0.0%	−28.8%
AAA-rated corporate bonds	Credit spread	0.1	0.2	0.0%	−1.5%
	Default rate	0.1	0.2	0.0%	−1.5%
AA-rated corporate bonds	Credit spread	0.0	0.0	0.0%	0.0%
	Default rate	0.0	0.1	0.0%	32.2%
A-rated corporate bonds	Credit spread	0.2	0.2	0.0%	−4.8%
	Default rate	0.0	0.1	0.0%	−9.2%
BBB-rated corporate bonds	Credit spread	0.2	0.1	0.0%	8.7%
	Default rate	0.1	0.1	0.0%	−33.1%
Public equity, large cap	Dividend yield	0.2	0.4	0.0%	0.1%
	Capital return	5.9	5.5	0.0%	38.0%
Public equity, mid cap	Dividend yield	0.3	0.4	0.0%	18.6%
	Capital return	5.9	5.6	0.0%	6.9%
Public equity, small cap	Dividend yield	0.3	0.3	0.0%	−16.7%
	Capital return	4.9	4.1	0.0%	4.5%
Public equity, high dividend	Dividend yield	0.2	0.2	0.0%	11.8%
	Capital return	5.0	6.2	0.0%	21.3%
Equity REITs	Cap rate	0.2	0.2	0.0%	−53.0%
	Capital return	6.5	7.3	0.0%	50.1%
Mortgage REITs	Cap rate	1.2	2.4	0.0%	8.9%

Asset Class		Idiosyncratic Factor Volatility (%)		Correlation with Systemic Factors	
		All Periods	Recession	All Periods	Recession
	Capital return	9.8	10.7	0.0%	−10.1%

Therefore, the idiosyncratic part of asset return economic scenarios $\boldsymbol{\sigma} \cdot \boldsymbol{\varepsilon}_t$ needs to be adjusted to reflect nonconstant volatility and nonlinear relationships, with the following steps:

Step 1: Predict whether the economy is in recession for each period under each scenario, based on fundamental economic factors, such as the logistic model described in Section B.1.

Step 2: Construct two correlation matrices from the error terms of asset return models. The first correlation matrix, $\mathbf{CM}_{\mathrm{All}}$, describes the general relationships among error terms of all asset classes, using all historical data. Theoretically, it is better to use data only during economic expansions. In this specific case, using data during economic expansions generates a correlation matrix that is not positive semidefinite, a condition that ensures consistency among correlations. The issue can be addressed using all historical data. The resulting correlation matrix is only slightly more conservative. The second correlation matrix, $\mathrm{CM}_{\mathrm{Res}}$, describes the relationships among error terms only in economic recessions. Cholesky decomposition then can be performed to get the lower triangular matrices $\mathbb{L}_{\mathrm{All}}$ and $\mathbb{L}_{\mathrm{Res}}$ to generate correlated idiosyncratic factors.

Step 3: Generate correlated idiosyncratic factors $\mathbf{I}_t^i$ for all asset returns during period t under scenario i:

$$\mathbf{I}_t^i = \begin{cases} \boldsymbol{\sigma}_{\mathbf{All}} \cdot \mathbb{L}_{\mathbf{All}} \boldsymbol{\varepsilon}_t^i & \text{if the economy is not in recession during period } t \text{ under scenario } i \\ \boldsymbol{\sigma}_{\mathbf{Res}} \cdot \mathbb{L}_{\mathbf{Res}} \boldsymbol{\varepsilon}_t^i & \text{otherwise} \end{cases}$$

where

$\boldsymbol{\sigma}_{\mathbf{All}}$ is a column vector containing the standard deviation of error terms of all asset return models in normal periods, as shown in column 3 of Table B.9

$\boldsymbol{\sigma}_{\mathbf{Res}}$ is a column vector containing the standard deviation of error terms of all asset return models in normal periods, as shown in column 4 of Table B.9

$\boldsymbol{\varepsilon}_t^i$ is a column vector containing independent random variables following a standard normal distribution for all asset return models

$\mathbb{L}_{\mathbf{All}}$ is a lower triangular matrix so that the error term correlation matrix $\mathbf{CM}_{\mathbf{All}}$ can be decomposed as $\mathbb{L}_{\mathbf{All}} \times \mathbb{L}_{\mathbf{All}}{}^T$ and

$\mathbb{L}_{\mathbf{Res}}$ is a lower triangular matrix so that the error term correlation matrix $\mathbf{CM}_{\mathbf{Res}}$ can be decomposed as $\mathbb{L}_{\mathbf{Res}} \times \mathbb{L}_{\mathbf{Res}}{}^T$.

At this step, nonconstant volatility and nonlinear relationships among idiosyncratic factors of asset return models have been taken care of. The economic scenario generation formula becomes

$$\mathbf{y}_t = \boldsymbol{\alpha} + \boldsymbol{\phi}_1 \mathbf{y}_{t-1} + \boldsymbol{\phi}_2 \mathbf{y}_{t-2} + \mathbb{B}_0 \mathbf{F}_t + \mathbb{B}_1 \mathbf{F}_{t-1} + \mathbb{B}_2 \mathbf{F}_{t-2} + \mathbf{I}_t^i$$

$$\text{Let } \mathbf{y}_t^p = \boldsymbol{\alpha} + \boldsymbol{\phi}_1 \mathbf{y}_{t-1} + \boldsymbol{\phi}_2 \mathbf{y}_{t-2} + \mathbb{B}_0 \mathbf{F}_t + \mathbb{B}_1 \mathbf{F}_{t-1} + \mathbb{B}_2 \mathbf{F}_{t-2}$$

$$\mathbf{y}_t = \mathbf{y}_t^p + \mathbf{I}_t^i$$

Step 4: Adjust $\mathbf{I}_t^i$ to reflect nonzero correlation between idiosyncratic factors and systemic factors during recessions:

$$\mathbf{J}_t^i = \begin{cases} \left(\boldsymbol{\rho}_r \cdot \mathbf{y}_t^p + \sqrt{\mathbf{1} - \boldsymbol{\rho}_r{}^2} \cdot \mathbf{I}_t^i\right) \dfrac{\sigma_{\mathbf{Res}}}{\sqrt{\boldsymbol{\rho}_r{}^2 \left(\boldsymbol{\sigma}_{\mathbf{Res}}^p\right)^2 + (\mathbf{1} - \boldsymbol{\rho}_r{}^2)(\sigma_{\mathbf{Res}})^2}} & \text{if in recession during period } t \text{ under scenario } i \\ \mathbf{I}_t^i & \text{otherwise} \end{cases}$$

where $\boldsymbol{\rho}_r$ is a column vector containing the nonzero correlation between idiosyncratic factors and systemic factors, as shown in column 6 of Table B.9.

The economic scenario generation formula becomes

$$\mathbf{y}_t = \boldsymbol{\alpha} + \boldsymbol{\phi}_1\mathbf{y}_{t-1} + \boldsymbol{\phi}_2\mathbf{y}_{t-2} + \mathbb{B}_0\mathbf{F}_t + \mathbb{B}_1\mathbf{F}_{t-1} + \mathbb{B}_2\mathbf{F}_{t-2} + \mathbf{J}_t^i$$

During economic recessions, Treasury bond yields go down, and credit spread and the default rate go down. Public equities (large cap, mid cap, small cap and high dividend class) behave similarly but with different volatility during recessions. Both dividend yields and capital returns go down. Equity REITs have a much stronger correlation with the macroeconomy than mortgage REITs. Cap rate and capital returns drop drastically for equity REITs during recessions. However, mortgage REITs are less correlated with the general economic cycles and could move in the opposite direction of equity REITs. These patterns are consistent with the correlation numbers shown in Table B.9. Wage inflation has a strong correlation with the economic cycle and with low wage inflation during recessions. Wage inflation is a factor that may affect pension liabilities if the benefit is linked to future wages.

B.3 Economic Scenario Generation for Bond Fund Return

With generated Treasury bond yield curve scenarios, credit spread and default rate, the model permits calculation of bond fund income rates and capital returns, based on bond fund investment strategy. Bond fund income rates can be used to project the positive cash flows (coupon payments) from bond investments. These cash flows can be used to meet benefit payout requirements or be reinvested. Bond fund capital returns are useful for projecting the bond fund values. This report uses five bond funds constructed to invest in Treasury bonds and AAA-, AA-, A- and BBB-rated corporate bonds, respectively. Their returns are generated using the following steps:

Step 1: Set the term mix of each bond fund. To match the long-term nature of the liability, all bond funds contain the mix of maturities shown in Table B.10. The bonds have an average maturity of 19 years and will be rebalanced each year to maintain the term mix.

Table B.10
Bond Fund Term Mix Assumption

Bond Maturities (Years)	Mix (% of Fund)
1	2.5%
3	2.5%
5	2.5%
7	2.5%
10	15%
20	50%
30	25%

Step 2: Using the assumption that the bonds have an annual coupon rate that is the same as the bond yield at the same term, calculate the coupon payments and income rates for each projection period. Coupon payments and redemptions are then reinvested in the same term mix.

Step 3: Evaluate the existing bonds each year with the new yield curve. In addition to the new yield curve and shorter maturity that could change the bond value from the beginning of the period, default and credit rating migration also are reflected in the revaluation. As part of the economic scenario generator, default rates are scenario dependent. Rating migration probabilities and recovery rates are kept fixed, as shown in Table B.11. If a bond is in default, its remaining value will be determined by the product of the recovery rate and the bond value. If a bond changes its rating, the value will be changed based on the yield curve of bonds with the new credit rating. The capital return is then calculated as the percentage change of total bond value.

Table B.11
Rating Migration and Recovery Rate Assumptions

	Rating Migration (%)				Junk[1]	Default Rate (Scenario Dependent)	Recovery Rate (%)
	AAA	AA	A	BBB			
AAA	87.05	9.03	0.53	0.05	3.34—Default Rate$_{AAA}$	Default Rate$_{AAA}$	49
AA	0.52	86.82	8.00	0.51	4.15—Default Rate$_{AA}$	Default Rate$_{AA}$	49
A	0.03	1.77	87.79	5.33	5.08—Default Rate$_{A}$	Default Rate$_{A}$	37
BBB	0.01	0.10	3.51	85.56	10.82—Default Rate$_{BBB}$	Default Rate$_{BBB}$	25

[1]Bonds downgraded below rating BBB (junk bonds) are assumed to be reinvested in investment-grade bonds. The recovered value from defaulted bonds also is assumed to be reinvested in investment-grade bonds.
Sources: The rating migration assumption is set according to the *2016 S&P Annual Global Corporate Default Study and Rating Transitions Report.* The recovery rate assumption is set according to *Moody's Corporate Default and Recovery Rates, 1920–2010* (Moody's Investors Service, February 28, 2011).

Step 4: After the valuation, rebalance the bond fund to the target term mix and credit rating.

Table B.12 lists the average simulated bond fund return for each bond fund. The bond fund return generation approach described here is only one possible approach. Many other bond fund investment strategies exist and can be applied. In practice, if the bond fund has enough historical data, linear models can be used to simulate future income rates and capital returns, as was done for other asset classes in Section B.2.

Table B.12
Average Simulated Bond Fund Return

Fund Type	Average Return (%)
Government bond fund	4.22
AAA-rated corporate bonds	3.60
AA-rated corporate bonds	4.59
A-rated corporate bonds	4.31
BBB-rated corporate bond	5.13

About The Society of Actuaries

The Society of Actuaries (SOA), formed in 1949, is one of the largest actuarial professional organizations in the world dedicated to serving more than 32,000 actuarial members and the public in the United States, Canada and worldwide. In line with the SOA Vision Statement, actuaries act as business leaders who develop and use mathematical models to measure and manage risk in support of financial security for individuals, organizations and the public.

The SOA supports actuaries and advances knowledge through research and education. As part of its work, the SOA seeks to inform public policy development and public understanding through research. The SOA aspires to be a trusted source of objective, data-driven research and analysis with an actuarial perspective for its members, industry, policymakers and the public. This distinct perspective comes from the SOA as an association of actuaries, who have a rigorous formal education and direct experience as practitioners as they perform applied research. The SOA also welcomes the opportunity to partner with other organizations in our work where appropriate.

The SOA has a history of working with public policymakers and regulators in developing historical experience studies and projection techniques as well as individual reports on health care, retirement and other topics. The SOA's research is intended to aid the work of policymakers and regulators and follow certain core principles:

Objectivity: The SOA's research informs and provides analysis that can be relied upon by other individuals or organizations involved in public policy discussions. The SOA does not take advocacy positions or lobby specific policy proposals.

Quality: The SOA aspires to the highest ethical and quality standards in all of its research and analysis. Our research process is overseen by experienced actuaries and nonactuaries from a range of industry sectors and organizations. A rigorous peer-review process ensures the quality and integrity of our work.

Relevance: The SOA provides timely research on public policy issues. Our research advances actuarial knowledge while providing critical insights on key policy issues, and thereby provides value to stakeholders and decision makers.

Quantification: The SOA leverages the diverse skill sets of actuaries to provide research and findings that are driven by the best available data and methods. Actuaries use detailed modeling to analyze financial risk and provide distinct insight and quantification. Further, actuarial standards require transparency and the disclosure of the assumptions and analytic approach underlying the work.

Society of Actuaries
475 N. Martingale Road, Suite 600
Schaumburg, Illinois 60173
www.SOA.org