USE OF GENETIC ALGORITHMS FOR OPTIMAL INVESTMENT STRATEGIES

by

Fan Zhang

B.Ec., RenMin University of China, 2010

A Project submitted in partial fulfillment
of the requirements for the degree of

Master of Science

in the
Department of Statistics and Actuarial Science
Faculty of Science

SIMON FRASER UNIVERSITY
Spring 2013

APPROVAL

Name: Fan Zhang

Degree: Master of Science

Title of Project: Use of Genetic Algorithms for Optimal Investment Strategies

Examining Committee: Dr. Derek Bingham
Associate Professor
Chair

____
Dr. Gary Parker
Associate Professor
Senior Supervisor
Simon Fraser University

____
Dr. David Campbell
Assistant Professor
Supervisor
Simon Fraser University

____
Dr. Yi Lu
Associate Professor
Examiner
Simon Fraser University

Date Approved: April 17

Abstract

In this project, a genetic algorithm (GA) is used in the development of investment strategies to decide the optimum asset allocations that back up a portfolio of term insurance contracts and the re-balancing strategy to respond to the changing financial markets, such as change in interest rates and mortality experience. The objective function used as the target to be maximized in GA allows us to accommodate three objectives that should be of interest to the management in insurance companies. The three objectives under consideration are maximizing the total value of wealth at the end of the period, minimizing the variance of the total value of the wealth across the simulated interest rate scenarios and achieving consistent returns on the portfolio from year to year. One objective may be in conflict with another, and GA tries to find a solution, among the large searching space of all the solutions, that favors a particular objective as specified by the user while not worsening other objectives too much.

Duration matching, a popular approach to manage risks underlying the traditional life insurance portfolios, is used as a benchmark to examine the effectiveness of the strategies obtained through the use of genetic algorithms.

Experiments are conducted to compare the performance of the investment strategy proposed by the genetic algorithm to the duration matching strategy in terms of the different objectives under the testing scenarios. The results from the experiments successfully illustrate that with the help of GA, we are able to find a strategy very similar to the strategy from duration matching. We are also able to find other strategies that could outperform duration matching in terms of some of the desired objectives and are robust in the tested changing environment of interest rate and mortality.

Keywords: Genetic algorithms; Investment strategy; Duration matching; ALM; Traditional life insurance

Acknowledgments

I would like to express my sincere gratitude to my senior supervisor Dr. Gary Parker for his excellent guidance and continued interest in the learning process of this project. This report would not have been possible without the support from him. The advice and knowledge from him have been very valuable throughout my studies and I am deeply indebted to him for all that I learned.

I am very grateful to Dr. David Campbell for taking me as his student and for all his patience and engagement in the last few months. I am truly thankful for his help to prepare me step by step for the defence and keep me motivated along the way.

Besides my supervisor, I would like to thank the rest of my thesis committee, Dr. Derek Bingham and Dr. Yi Lu, for their useful comments and insightful questions.

Last but not least, I would like to say thank you to Dr. Tim Swartz for his kind support and encouragement.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Background and Motivation

Genetic algorithm (GA) is a searching technique that could explore the large solutions space of a particular problem and quickly move towards the subspace which contains the best solution according to user-defined objectives. One advantage of GA is its flexibility. It is very easy to implement and can be suitable for many optimization problems, in particular, those that are complicated and multi-dimensional. For a problem with no direct mathematical solution, GA offers an alternative route to solve it.

GA has been widely used in many areas, such as engineering, medicine and finance. In finance, it is particularly helpful in dealing with a complicated optimization problem in asset management. For a fund manager, the objective is selecting assets, such as bonds or stocks, among a wealth of assets available, that allows the portfolio to achieve excellent return while minimizing risk. A common measure for the fund managers' performance would be the Sharpe ratio, which calculates additional return gained for each unit of volatility (Sharpe, 1994). The benchmarks for investment managers are usually the market index and peers. This type of research has been carried out by various researchers and many have shown that GA-aided strategy is able to successfully select a portfolio that could beat the market index consistently in the period studied.

The investment management for insurance companies faces a more complex situation than the pure asset management of fund managers. The reason is that insurance companies need to consider both the asset side and the liability side of the portfolio. In a simplified situation, insurers need to make sure future premiums to be collected and reserves are enough to cover future claims and expenses while having enough capital left to satisfy the capital requirements from regulators (Cooper et al., 2010).

Given the difference in nature of the assets and liabilities facing an insurer, the focus of Asset and Liability Management (ALM) of insurance contracts has been on managing the various risks associated with the mismatch between the asset and liability side (Cooper et al., 2010).

The type of product we are interested in studying is a term life insurance product, which is a type of traditional life products. Traditional life product has been shown to be sensitive to interest risk and therefore managing interest risk has been a top consideration in the ALM practice (Mathis, 1993). Duration has been introduced as a risk measure for interest rate risk and a history of the development of the duration matching strategy can be found in Reitano (1991). Due to its ease of implementation, duration matching has become a common and forefront method in ALM practice for traditional life insurance products and even for companies who do not hedge interest risk actively for this line of business, they have targets in place for the mismatch allowed between the assets and liabilities (Reynolds and Wang, 2007).

In this project we adopt the dollar duration matching approach as introduced in Reynolds and Wang (2007) for term life insurance product. With this method, we are able to find a portfolio of assets that match the dollar duration of liabilities. We re-balance the portfolio at the end of each year to enable the dollar duration from asset and liability sides to be matched again. We are able to calculate return and risk measures for the portfolio by adopting the duration matching strategy, and the performance is used as a benchmark to evaluate the performance under the GA-based strategies.

We use GA to determine a portfolio of assets to back up the insurance portfolio too. However, instead of attempting to find assets that match the dollar duration of liabilities, we aim to find assets that could optimize our objectives under consideration. We are interested in studying if with the help of GA, we are able to find alternative strategies that could outperform the duration matching strategy, by some of the objectives, such as achieving a higher return or a lower semi-standard deviation. We will show that the matched position of dollar duration can also be considered as an objective and GA offers us the flexibility to relax that position to a certain level while accommodating to other objectives a bit more. For example, in some occasions, it may be worthwhile to allow for a duration mismatch for higher profitability.

Duration matching is a method that requires constant re-balancing. As market changes, assets may be bought or sold to maintain the duration matched position. We introduce re-balancing in GA program to let GA benefit from that feature too. We further use GA to help us re-balance actively by selecting strategies based on the known information about interest rate and mortality experience.

From this study, it is shown that GA is able to produce portfolios that are better than the portfolios that duration matching yields, at least in some sense. GA helps find the best allocation for different risk appetites, by allowing for a trade-off between different objectives. It is shown that the re-balancing strategy can be optimized by means of GA. GA is also able to produce a strategy that mimics the performance under the duration matching strategy closely. In addition, it can be observed that duration matching offers a very balanced strategy that GA can learn from as well.

1.2 Literature Review

Genetic algorithm has existed for the past twenty years and has been widely used as an optimization tool. Extensive research has been carried in the field of portfolio optimization of stocks portfolios to either relax the constraints in the portfolio theory or to improve the performance of GA. Some recent contributions include Aranha (2007) who considers transaction costs and introduces the Euclidean distance as the measure for the transaction costs in re-balancing the portfolios. Soam et al. (2012) introduces an active re-balancing strategy to improve the performance of GA-strategy especially in periods of market crashes. Another contribution from that paper is that the trading volumes of stocks are considered in the re-balancing stage as a signal for the future price movements. Chang et al. (2009) introduces different risk measures for evaluating the portfolios and incorporates them into a GA framework. Chan et al. (2002) studies a multi-stage optimization problem for stocks portfolios. One contribution from this paper is the use of the tree representation for scenarios that span multi-periods. It considers a simple scenario that only consists of upward and downward movements in each period, and finds the investment strategy for each scenario. In these previous researches, the performance from the GA-based strategy in stocks portfolio optimization problems has been shown to be efficient and able to outperform other strategies, such as a passive market index tracking approach, by achieving a higher return and a lower risk.

GA has only been recently introduced to the actuarial profession and very few papers have been published by actuaries. One of the papers is Jackson (1997), which compares GA and Newton's method in an asset allocation problem and shows that GA is more efficient in finding the optimal solution as the problem becomes more complex. Tan (1997) uses GA to find asset allocations of various bonds and mortgages to back up a portfolio of single premium deferred annuity contracts in a framework that incorporates price, risk and competitiveness as the objectives.

For the duration matching strategy, which is used as a benchmark for the GA-based strategy, the steps proposed by Reynolds and Wang (2007), who also study the ALM problem for traditional life insurance product and use a dollar duration matching strategy to manage interest risk, are followed. Another strategy to manage bond portfolios to back up liability cash flows is cash flow matching. Details of these two strategies can be found in Fabozzi et al. (2007). Bierwag et al. (1983) discusses the conditions to satisfy in order to achieve immunization for a portfolio that involves several liabilities and shows that duration matching is not the only criteria needed to achieve immunization. Reitano (1991) shows the limitations of the duration matching strategy in terms of non-parallel shifts on the yield curve and develops a multi-variate model for the shifts on yield curve at different points.

1.3 Outline

Inspired by the previous research, such as Tan (1997), we explore whether GA offers a promising alternative as a tool for portfolio optimization in the insurance setting.

Chapter 2 gives an introduction into genetic algorithm. Details of the GA are discussed.

Chapter 3 gives details of the problem and introduces the assumptions used. It also introduces interest rate and mortality scenarios used in the training and testing of the experiments done in Chapter 6.

Chapter 4 discusses the duration matching technique and presents some preliminary results from the duration matching strategy.

Chapter 5 discusses the GA optimization framework. It gives details about the parameters in GA and about the objectives and expected results of the experiments. It also introduces the active re-balancing program in the experiments.

Chapter 6 presents results from the experiments and comments on the various strategies obtained through different methods.

Chapter 7 reviews the advantages of using GA to obtain investment strategies and the advantages of the duration matching strategy, and their limitations. Chapter 8 discusses possible future research areas.

Chapter 2

Genetic Algorithm: What Is It and the Theory Behind It

Genetic algorithm (GA) is an optimization technique that is developed in Holland (1975). GA mimics the phenomenon of natural selection in the evolution of human beings to select optimal solutions by the use of a computer. Based on the concept of survival of the fittest, GA starts with randomly generated solutions to the problem that are analogized to a generation of individuals. By evaluating these solutions using the objective function specified by the user, GA assigns the solutions (individuals) with higher scores a higher chance of being selected as the parents to produce the offspring that form the next generation of solutions. Therefore, the fittest individuals are most likely to pass their genes to the next generation and the children are more likely to fit the environment even better. GA mimics this phenomenon in nature by letting the binary chromosomes which represent the individuals to recombine into new individuals. Although GA is a simplified version of the evolution of life, it is expected that the solutions will become better generation by generation, in the same way the human population evolves through time. GA has become a state-of-the-art technology that is used to solve many real world optimization problems. The details of GA procedures are explained in the following section.

2.1 GA Procedures

In this section, the procedures in GA are introduced. The work flow of GA is displayed in Figure 2.1 and more details are given following this.

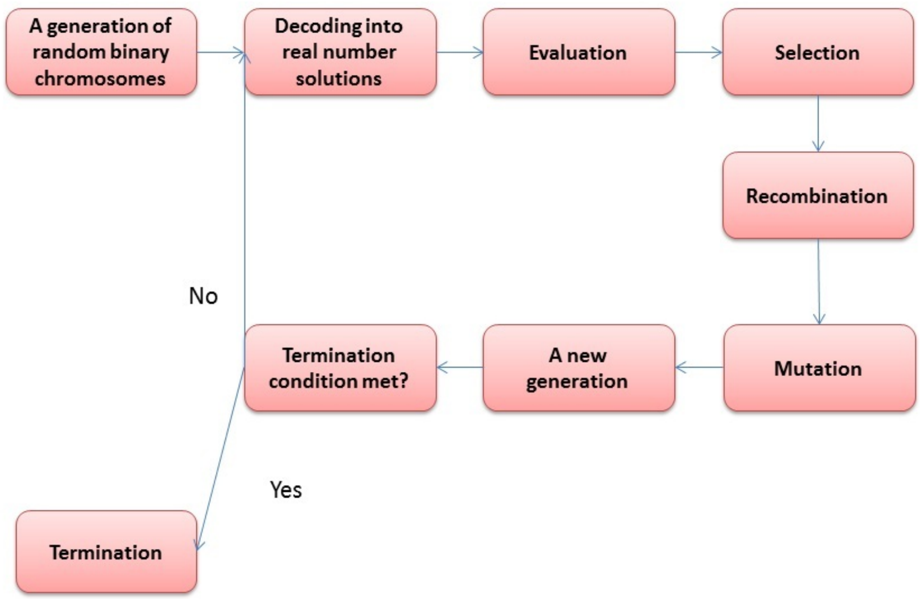

Figure 2.1: GA Procedures

Chromosome Coding: The first step of GA is to start with a number of random solutions to the problem as the initial generation. There are various coding techniques to represent a solution, such as binary, real numbers and permutations (Haupt and Haupt, 2004). For example,

- Binary strings (0101110...0010);
- Real numbers $(2.3, 45.4, \ldots, 56.2, 21.3)$;
- Permutations $(1234, 1243, 2134, \ldots, 3412)$.

The binary representation is the way we will represent our variables. More coding techniques are developed to suit a particular problem. Ideally the coding should be as close as possible to the natural representation of the solution. For example, for the famous

Traveling Salesman Problem (TSP) that attempts to find the cheapest way to travel a finite number of cities, the best coding choice is permutation. For asset allocation problems, two candidates of coding are feasible, binary strings and real numbers coding. The advantage of real numbers coding is that it allows full machine precision while with binary coding, the precision is limited by the number of digits allowed for each variable. In Herrera et al. (1998), a discussion about the advantages and disadvantages of both coding methods is given. According to Gaffney et al. (2010), no coding method outperforms the other in all situations and the performance largely depends on the algorithms for the GA operators other than on the coding method alone. We adopt the binary representation because it is the most original way for representation in a simple GA (SGA). Further analysis could be done to implement the real numbers coding method to decide which suits this problem better.

Therefore, a number of randomly generated binary chromosomes are used as the initial generation, with each one representing an individual solution. We could also start with a set of solutions provided by methods other than random generation. For example, to generate stocks portfolios, we could start with a portfolio of stocks recommended by experts and use GA to aid in selecting a few stocks from them.

The total length of the binary chromosome is determined by the number of variables, and the level of precision allowed for each variable. Let the variable be on the interval $[a, b]$, where a and b are real numbers, let the level of precision required be c decimals, and let the ith variable of interest denoted by x_i, then the number of binary digits d needed for x_i is calculated by

$$2^{d-1} < (b-a) \times 10^c \leq 2^d - 1. \tag{2.1}$$

For example, if a variable is on the interval [0,1] and the level of precision required is 4 decimals. The number of binary digits for each variable can be found by solving

$$2^{d-1} < (1-0) \times 10^4 \leq 2^d - 1. \tag{2.2}$$

Since

$$2^{14-1} < (1-0) \times 10^4 \leq 2^{14} - 1, \tag{2.3}$$

the number of digits required for this variable is 14.

Decoding: In the second step, the binary chromosomes are decoded into real numbers using the Gray decoding method (Rowe et al., 2004). The Gray coding system is considered as a special case of the binary coding system. It also uses the numbers 0 and 1 to represent a real number. The reason for choosing the Gray decoding system instead of binary decoding system is that Gray codes offer smaller *Hamming Distance* than binary codes do. *Hamming Distance* is the number of substitutions required to change one code to another code of equal length. For example, the *Hamming Distance* between [001011] and [101101] is 3.

A brief example is given to illustrate why Gray decoding is more appropriate than binary decoding. Suppose there are two parents, and we denote the first parent as *parent1* and the second parent as *parent2*. As detailed later in this chapter, during reproduction there is a crossover point where the chromosomes of the parents are cut and then joined together to produce two new chromosomes as their offspring. A crossover point randomly falls somewhere in the middle of the chromosomes of the parents and results in two offspring, *offspring1* and *offspring2*. The position of the cut is indicated using "|" and the parents used in this example are 13 and 16. In the first example, regular binary coding system is used. The following equations represent the real numbers converted by binary decoding:

$$\mathit{parent1} = [01|101] = [13],$$
$$\mathit{parent2} = [10|000] = [16].$$

The difference between the real numbers is 3 while the *Hamming Distance* between the binary codes is 4. The crossover results in offspring that diverge from the parents as shown in the following. The two offspring produced are 8 and 21, i.e.,

$$\mathit{offspring1} = [01|000] = [8],$$
$$\mathit{offspring2} = [10|101] = [21].$$

In a second example, Gray decoding system is used to represent the same real numbers 13 and 16 as in the first example, that is,

$$\mathit{parent1} = [01|011] = [13],$$
$$\mathit{parent2} = [11|000] = [16].$$

It can be observed that the *Hamming Distance* of the Gray codes of the two parents is 3, same as the difference between the real numbers.

With the same crossover point, the offspring are given in the following equations:

$$\mathit{offspring1} = [01|000] = [15],$$
$$\mathit{offspring2} = [11|011] = [18].$$

The resulting offspring are more similar to their parents than the offspring in the first example, who are outside the range of the two parents. Therefore these two examples show that with the Gray decoding method the children produced are more aligned with their parents.

Evaluation: Each individual is a solution and a potential candidate that solves the optimization problem of interest. After they are decoded into real numbers in step 2, the objective function takes a solution as input and produces a "score" for the solution.

A multi-objective function can be created by taking a weighted average of several objective functions. The importance of each objective function $f_i(x)$ for i = 1, 2, ..., n, is represented by the weight u_i and the multi-objective function f_{multi} is

$$f_{multi}(x) = \sum_{i=1}^{n} u_i \times f_i(x). \tag{2.4}$$

There are other ways to accommodate multiple objectives, such as the use of "objective sharing" as used in Soam et al. (2012). Objective sharing is a method that assigns a probability to each of the objectives of interest, and in each generation, GA is performed on one of the objectives or an objective function that incorporates a few objectives, that is randomly selected based on its probability.

Selection: As introduced earlier in the Evaluation step, each individual solution is assigned a score. Individual solutions are ranked by their scores and a higher rank results in a higher probability of being selected. The same number of solutions as in the initial generation are selected for reproduction. The order of these selected solutions is shuffled and every two of them become a pair of parents. Stochastic universal sampling (SUS) is used in this study and more details can be found in Baker (1987). There are also other popular selection methods, such as the roulette wheel selection and the tournament selection, and more details can be found in Haupt and Haupt (2004).

Recombination: In this step, single point crossover is used on each parent indicating a cut point. The crossover point is random. The two chromosomes of parents are cut at that crossover position and they recombine to form two offspring. By repeating this for each pair of parents, a new generation of solutions is created.

Mutation: Mutation is an operator used to avoid the problem that the solutions converges too quickly to a sub searching space that is suboptimal and to add a level of heterogeneity into the population. Each element in the binary chromosome is mutated from 0 to 1 or from 1 to 0 with a given probability. A common choice for the number of digits to be mutated is 1 per individual candidate and that is what is used in this study. For a comprehensive investigation on the optimal mutation rates, please refer to Bäck (1993).

Termination Condition: A new generation of solutions are created and the loop is repeated from the second step, which is decoding into real numbers. From there, the rest of the procedures are followed. The loop ends when the termination criteria is met. More details can be found in Haupt and Haupt (2004). The common termination criteria of the iteration are listed below:

- A specified number of generations have elapsed.
- A satisfactory solution is found (for problems that have a definite solution).
- No improvement in solutions for a specified number of generations.

Since for the problem being studied it may not be possible to find a globally optimal solution, only the first and the third type of termination criteria are considered.

Extra Considerations for GA: The work of Holland (1975) was inspired by the evolution of nature. GA operators could be modified to better imitate the phenomenon in nature, which is supposed to improve the performance of GA, such as the speed of convergence.

- Elitism: In nature the parents often co-exist with the children instead of dying off immediately after the children are produced. Elitism is letting a proportion of parents with top performance to be kept in the next generation instead of altering them in the crossover, recombination and mutation processes.
- Mutation rate: Mutation rate is used in GA to introduce the heterogeneity in the population; however, this type of move could also result in slow convergence. While a lot of research has been focused on finding the optimal mutation rates, some have introduced a technique to make it more flexible, which is to change the mutation rate depending on the results from generation to generation. For example, Fogarty

(1989) shows that a mutation rate that decreases exponentially over the generations has superior performance.

- Starting generation: Although it is common to use an arbitrary starting generation, it may save a lot of time to use a generation from a last run or from a generation of solutions that are selected with the knowledge of the particular problem.

2.2 Quantitative Explanation for GA: Schema Theorem and Building Block Hypothesis

In this section, an in-depth discussion is given about why GA works. First of all, let us take a look at the GA jargons. For more details, please see Haupt and Haupt (2004).

- Schema (pl. schemata): a schema S is a pattern in a chromosome. Let (0, 1, $*$) be the symbol alphabet, where "$*$" is a wild card symbol that can represent either 0 or 1, then, for example, the schema [00$*$101$*$] matches strings [0001010], [0001011], [0011010] and [0011011].

- Length of Schema: the length of schema l is the number of bits in a string. For schema [00$*$101$*$], the length l is 7.

- Order of Schema: the order of schema S, denoted as $o(S)$ is the number of fixed positions in a schema. In schema [00$*$101$*$], the order of schema is 5. The number of strings that match schema S is $2^{l-o(S)}$.

- Definition Length of Schema: the definition length of schema S, denoted as $\delta(S)$ is the distance between the first and the last fixed position in it. For example, for schema [00$*$101$*$] the position of the first fixed position is 1 and the position of the last fixed position is 6, and therefore $\delta(S)$ is 5.

- Building Block: short schemata that give a chromosome a high fitness value and increase in number as the GA progresses.

- Building Block Hypothesis: a genetic algorithm seeks near-optimal performance through the juxtaposition of short, low-order, high-performance schemata, called the building blocks.

2.2.1 Schema Theorem

The foundation of the theoretical explanation for the effectiveness of GA is first given by means of the Schema Theorem in Holland (1975). Schema Theorem states that short, low-order, above-average schemata receive exponentially increasing trials in subsequent generations of a genetic algorithm. Inside the black box of GA, the calculations going on are in essence the calculations of schemata. In this subsection, the Schema Theorem that explains the power of the GA operators is presented.

Selection Let $m(S,t)$ denote the number of chromosomes (solution candidates) in generation t that has schema S. Let $f_S(t)$ be the average fitness score evaluated from the objective function of chromosomes belonging to schema S in generation t, and let $f(t)$ denote the average fitness score of all chromosomes in the population in generation t.

The expected number of chromosomes belonging to schema S at time $(t+1)$, $m(S,t+1)$, can be calculated as follows

$$m(S,t+1) = m(S,t)\frac{f_S(t)}{f(t)}, \qquad t = 0, 1, \ldots, \text{last generation} - 1. \tag{2.5}$$

Assume that schema S has a fitness score that is greater than average by $c \geq 0$, then the fitness score of schema S is given by

$$f_S(t) = f(t) + cf(t). \tag{2.6}$$

The expected number of chromosomes with schema S at t+1 is given by

$$m(S,t+1) = m(S,t)\frac{f(t) + cf(t)}{f(t)} = m(S,t)(1+c). \tag{2.7}$$

If c is constant and t starts from 0, the expression for $m(S,t)$ can be obtained from $m(S,0)$, which is the number of chromosomes with schema S at time 0, as

$$m(S,t) = m(S,0)(1+c)^t, \qquad t = 1, 2, \ldots, \text{last generation}. \tag{2.8}$$

This explains why the selection procedure in GA works. Note that this proof is based on the most simple roulette wheel selection method, which specifies that the probability for each candidate solution being selected is proportional to the fitness score that it receives. This implies that the schema with higher fitness value will have a higher chance of appearing in future generations and the growth is exponential when only the effect of selection is considered.

Crossover A schema can be kept only if the cut point in crossover falls beyond the definition length. In chromosome A=[0011010] there is schema S_1, [$****01*$], and there is schema S_2, [$*0***1*$]. The definition length for these two schemata is given by

$$\delta(S_1) = 6 - 5 = 1,$$
$$\delta(S_2) = 6 - 2 = 4.$$

There are $l-1$ points that the crossover can fall in. For schema S_1, the crossover point can fall anywhere outside the two numbers, and the schema will survive in the crossover. For schema S_2, the crossover point needs to fall before the first star or after the last star so that the schema will not be destroyed in crossover. The longer the definition length is, the more likely it is destroyed in the crossover. The probability of schema S_1 and S_2 to be destroyed is given by

$$P_d^c(S_1) = \frac{\delta(S_1)}{l-1} = \frac{1}{6},$$
$$P_d^c(S_2) = \frac{\delta(S_2)}{l-1} = \frac{4}{6}.$$

The survival probability from crossover of these two schemata is give by

$$P_s^c(S_1) = 1 - P_d^c(S_1) = 1 - \frac{1}{6} = \frac{5}{6},$$
$$P_s^c(S_2) = 1 - P_d^c(S_2) = 1 - \frac{4}{6} = \frac{2}{6}.$$

Let P_c denote the probability of crossover happening to a chromosome. The probability that schema S can survive from a crossover in this example is given by

$$P_s^c(S) = 1 - P_c P_d^c(S) = 1 - \frac{P_c \delta(S)}{l-1}. \tag{2.9}$$

However, even if the crossover point falls within the definition length, it is possible that the schema will survive. The next example illustrates that for schema S_2, if chromosome A recombines with chromosome B=[1001000], it will survive even when the crossover point falls within the definition length. Suppose that the crossover point falls after the third star, the two parents and the two children after the recombination are given below

$$\textit{parent1} = A = [001|1010]$$
$$\textit{parent2} = B = [100|1000]$$
$$\textit{offspring1} = [001|1000]$$
$$\textit{offspring2} = [100|1010].$$

Schema S_2 is preserved in *offspring2* even though the crossover point falls within the definition length in *parent1*. This is because chromosome B has at least one digit the same as schema S_2. In this example, the second digit of chromosome B is the same as the second digit of schema S_2. Therefore, the survival probability of a schema is

$$P_s^c(S) \geq 1 - \frac{P_c \delta(S)}{l-1}. \tag{2.10}$$

The combined effect of selection and crossover is achieved through multiplying both sides of (2.5) by (2.10), as shown below

$$m(S, t+1) \geq m(S,t) \frac{f_S(t)}{f(t)} \left(1 - \frac{P_c \delta(S)}{l-1}\right). \tag{2.11}$$

Recall that (2.10) gives the expected number of chromosomes belonging to schema S at time $t+1$ only considering the effect from selection. (2.5) gives the survival probability of a schema from crossover. Therefore, the expected number of chromosomes belonging to schema S at time $t+1$ after surviving selection and crossover is obtained by (2.11).

This implies that the schema with short definition length will have a higher chance of appearing more in chromosomes in the future generations.

Mutation Mutation is a GA procedure that mimics the phenomenon that a gene can be perturbed in the reproduction process. Let us denote the mutation probability by P_m. The probability of a schema S surviving a mutation is given by

$$P_s^m(S) = (1 - P_m)^{o(S)}. \tag{2.12}$$

Since $P_m << 1$, this probability can be approximated by

$$P_s^m(S) \approx 1 - P_m o(S). \tag{2.13}$$

The combined effect of selection, crossover and mutation is given by the following equations, which demonstrates the *Schema Theorem*,

$$m(S, t+1) \geq m(S,t) \frac{f_S(t)}{f(t)} \left(1 - \frac{P_c \delta(S)}{l-1} - P_m o(S)\right), \tag{2.14}$$

or

$$m(S, t+1) \geq m(S,t) \frac{f_S(t)}{f(t)} \left(1 - \frac{P_c \delta(S)}{l-1}\right) (1 - P_m)^{o(S)}. \tag{2.15}$$

2.2.2 Building Block Hypothesis (BBH)

The previous section presents mathematical proof of Schema Theorem that states that schemata with above average fitness value, short definition length and low-order will receive exponential growth rates in subsequent generations. These types of schemata are called “building blocks”.

GA seeks the optimal solutions by the juxtaposition of these building blocks. From these short, low order and above average building blocks, it hopes to eventually find the global near-optimal solution, which is long (approaching l), high-order and above average. This theory is first introduced in Holland (1975) and Goldberg (1989). In the following, a simple example is used to illustrate how building block hypothesis works.

Suppose a type of light is [red red yellow yellow green green] and we start with a generation of size 5 that have chromosomes with the building blocks of this solution. Star represents a blank bit that can be any color.

$$\begin{aligned} \textit{chromosome1} &= [rr****] \\ \textit{chromosome2} &= [******] \\ \textit{chromosome3} &= [**yy**] \\ \textit{chromosome4} &= [******] \\ \textit{chromosome5} &= [****gg] \end{aligned}$$

Note that in the first generation, *chromosome1*, *chromosome3* and *chromosome5* have the building blocks for the optimal solution. GA can be applied to solve this simple optimization problem and the optimum solution appears in generation 23.

Chromosomes from generations 1, 2, 10 and 23 are displayed in Figure 2.2. Each row represents a different chromosome. Figure 2.2(a) displays the initial generation of chromosomes, with the first, the third, and the fifth row having the building blocks as introduced earlier. The other chromosomes are kept as blank. The optimal solution would be a row with two red lights, two yellow lights and two green lights. Figure 2.2(b) shows the chromosomes in the second generation. It can be observed that every chromosome has the building block for the optimal solution. In generation 10, which is displayed in Figure 2.2(c), the building blocks in the chromosomes become longer. Finally, in Figure 2.2(d) the buildings blocks in each chromosomes are longer than the building blocks in previous generations and in the last row we are able to obtain the optimal solution. This illustrates how the building blocks can form the optimal solution quickly by the GA procedures.

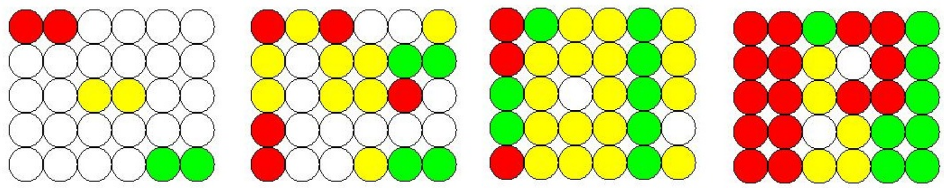

(a) Chromosomes in Generation 1 (b) Chromosomes in Generation 2 (c) Chromosomes in Generation 10 (d) Chromosomes in Generation 23: Optimal Obtained (last row)

Figure 2.2: Traffic Lights: Convergence towards Optimal Solution from Building Blocks. Each row represents a chromosome and each circle represents a traffic light.

The results in Figure 2.2 are obtained by using GA procedures discussed in Section 2.1 . Mutation is used to intentionally introduce random fluctuations in the reproduction process to avoid early convergence into suboptimal solutions. To illustrate the effect of the mutation operator, we start with the same initial generation of chromosomes and disable the mutation procedure. The results from generations 1, 2, 5 and 9 are displayed in Figure 2.3. It is shown that in generation 2 and generation 5, more chromosomes have building blocks and the building blocks become longer. In the last plot of Figure 2.3, all chromosomes are the same and we are not able to further improve them without the mutation operator, and therefore we end up with suboptimal solutions.

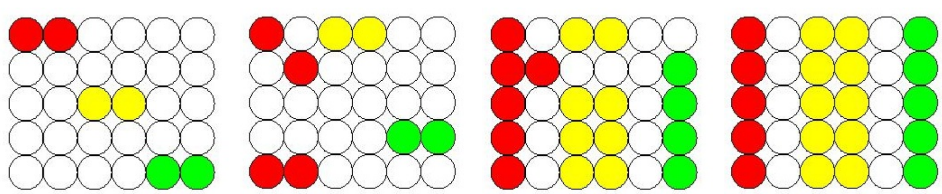

(a) Chromosomes in Generation 1 (b) Chromosomes in Generation 2 (c) Chromosomes in Generation 5 (d) Chromosomes in Generation 9: Suboptimal Solutions

Figure 2.3: Traffic Lights: Early Convergence towards Suboptimal Solution without Mutation Operator. The chromosomes will not improve further from Generation 9, resulting in suboptimal solutions

It is worth noting that there is a lack of comprehensive theoretical basis for GA and the building block hypothesis offers one explanation for the success of GA. It is possible that these building blocks can lead GA to converge to suboptimal areas. The way to deal with this problem is usually through adjusting the objective function. For a detailed investigation on how building block hypothesis works during GA operations, readers are encouraged to read Forrest and Mitchell (1991). While GA has been achieving success in practical applications, the precise theory to explain it has not been found. Altenberg (1995) gives a discussion of the problem and contradiction of the building block hypothesis. For an alternative hypothesis that also attempts to explain why GA works, please refer to Beyer (1997).

2.3 GA Toolbox

In this project, all the calculations are based on the Genetic Algorithm Toolbox for MATLAB (GATbx) that is developed by the Department of Automatic Control and Systems Engineering of The University of Sheffield, UK. The version used is GA Toolbox v1.2 and is published under the GNU General Public License.

Chapter 3

Description of the Problem

In Chapter 2, we introduced the framework of GA and attempted to explain the theory behind GA. In this project, GA is used to solve the portfolio optimization problem in an insurance setting. For a portfolio of term insurance contracts, the decision needs to be made on the asset allocations to back up this portfolio. As the number of choices for the assets gets large, the problem becomes complicated enough that no direct method can be applied to solve it. The problem spans on multi-periods and the re-balancing strategies need to be determined too. The many variables involved in this problem make the solutions space so large that it becomes a challenge to find the optimum solution quickly. In addition, the GA system must tailor to the specific characteristics of liabilities of insurance contracts, instead of focusing only on the assets side, with the latter being the approach used in the stocks portfolios optimization problems.

In this chapter, the details of the problem studied are given. The interest rate and mortality scenarios to be used for the simulation in training and testing of the experiments are also discussed.

3.1 The Problem Studied

The characteristics of the product studied are described in Table 3.1.

In pricing of this product, an interest rate of i=4% is used. A net level premium P^n payable annually at the beginning of the year is assumed from the policyholders. The net

Table 3.1: Description of the Product

Product Characteristics	**Assumptions**
Product Type	Term Life
Sum Assured	$100,000
Term	n=10
Size of Pool	1000
Mortality	1997-2004 CIA Basic Male, Combined
Sex Distribution	100% male
Issue Age	50

premium P^{10} for each policyholder is calculated from the formula below:

$$P^{10} = 100,000 \frac{A_{50:\overline{10}|}}{\ddot{a}_{50:\overline{10}|}}. \tag{3.1}$$

The number of policyholders of age 50 surviving to the beginning of year $t+1$ under scenario s is denoted as $l_{50+t}(s)$, the probability of policyholders of age 50 surviving to the beginning of year $t+1$ under scenario s is denoted as ${}_tp_{50}(s)$, and the number of policyholders of age 50 dying in year t under scenario s is denoted as $d_{50+t}(s)$, $t = 1, \ldots, 10$. No expense is assumed. Policy lapse is not taken into consideration. The premiums collected at the beginning of year t in scenario s, $p(t,s)$, and the claims to be paid at the end of year t in scenario s, $c(t,s)$, can be calculated from

$$p(t,s) = P^{10} \times l_{50+t}(s), \tag{3.2}$$

$$c(t,s) = 100,000 d_{50+t}(s), \tag{3.3}$$

$$l_{50+t}(s) = l_{50}(s) {}_tp_{50}(s), \tag{3.4}$$

$$d_{50+t}(s) = l_{50+t}(s) - l_{50+t+1}(s). \tag{3.5}$$

3.2 Interest Rate Scenarios

A subset of 200 interest rate scenarios are generated from the Cox-Ingersoll-Ross (CIR) model in the training stage of the experiments. CIR model is an important one-factor model for short rate (Zeytun and Gupta, 2007).

For deterministic interest rate scenarios testing, we would test the New York Seven interest rate scenarios, which are prescribed by New York State Regulation 126. For stochastic interest rate testing, a subset of 500 interest rate scenarios are generated from the same CIR model as the one used in training. In this section, the interest rate scenarios to be used in the experiments are discussed.

For the purpose of generating interest rate scenarios for the training and testing purposes of the experiments, the real world parameterization should be used; in pricing, the appropriate model to be used is the risk neutral model. The formulae for the CIR model in this section are taken from Zeytun and Gupta (2007).

Under the risk neutral parameterization of the CIR model, the short rate is assumed to satisfy the stochastic differential equation

$$dr(t) = \kappa(\theta - r(t))dt + \sigma\sqrt{r(t)}dW(t), \qquad r(0) = r_0, \tag{3.6}$$

where κ represents the mean reversion speed, θ represents the long term mean, $\sigma\sqrt{r(t)}$ is the volatility term and $r(t)$ represents the current interest rate. These are all positive constants. $W(t)$ is a standard Brownian motion under the risk neutral measure (Zeytun and Gupta, 2007).

The CIR model assumes a mean-reverting process for the short rate θ and the speed for reversion is equal to κ; with the volatility term $\sigma\sqrt{r(t)}$, the interest rate is always positive and the volatility depends on the interest rate level (Zeytun and Gupta, 2007).

The random variable $r(t)$ has a non-central chi-square distribution, so that the future short rates can be simulated from the following closed form formula:

$$r(t+T) = r(t) + cY, \tag{3.7}$$

where $c = \frac{(1-exp^{-\kappa T})\sigma^2}{2\kappa}$ and Y is a non-central chi-square distribution with $\frac{4\kappa\theta}{\sigma^2}$ degrees of freedom and non-centrality parameter $2cr(t)exp^{-\kappa T}$.

Under the risk neutral measure, the conditional expectation and variance of the short rate are given by

$$\begin{aligned} E[r(t)|\mathcal{F}(u)] =& r(u)e^{-\kappa(t-u)} + \theta\left(1 - e^{-\kappa(t-u)}\right) \\ Var[r(t)|\mathcal{F}(u)] =& r(u)\frac{\sigma^2}{\kappa}\left(e^{-\kappa(t-u)} - e^{-2\kappa(t-u)}\right) + \frac{\theta\sigma^2}{2\kappa}\left(1 - e^{-\kappa(t-u)}\right)^2, \end{aligned}$$

where $r(u)$ is the short rate at time u and $\mathcal{F}(u)$ is the available information as of time u.

Under the real world parameterization, the corresponding stochastic differential equation that the short rate is assumed to satisfy is given by

$$dr(t) = (\kappa\theta - (\kappa + \lambda\sigma)r(t))\, dt + \sigma\sqrt{r(t)}dW^0(t), \tag{3.8}$$

where λ is a constant and W^0 is a Brownian motion under the real world measure.

The risk premium at time t is given by

$$\lambda(t) = \lambda\sqrt{r(t)}, \qquad t = 0, 1, \ldots, 10. \tag{3.9}$$

Let $r(t, s)$ represent the current interest rate for scenario s. The price of a zero-coupon bond under scenario s of maturity T at time t is given by

$$P(t, T, s) = A(t, T)e^{-r(t,s)B(t,T)}, \tag{3.10}$$

where

$$\begin{aligned} A(t,T) =& \Big(\frac{2he^{(h+\kappa+\lambda)(T-t)/2}}{2h + (h+\kappa+\lambda)(e^{h(T-t)} - 1)}\Big)^{2\kappa\theta/\sigma^2}, \\ B(t,T) =& \frac{2(e^{h(T-t)} - 1)}{2h + (h+\kappa+\lambda)(e^{h(T-t)} - 1)}, \\ h =& \sqrt{(\kappa+\lambda)^2 + 2\sigma^2}. \end{aligned}$$

Deterministic Interest Rate Scenarios The New York Seven interest scenarios are first stated in New York State Regulation 126. Details of them can be found in Ramsey (1990). The seven scenarios are:

1. Level interest rate.
2. Increase by 0.5% each year for the next ten years and then remain at that high level.
3. Increase by 1% each year for the next five years and then decrease by 1% each year for the next five years, then remain level.
4. Pop up 3% in the first year and then remain level.
5. Decrease by 0.5% each year for the next ten years and then remain at the low level.
6. Decrease by 1% each year for the next five years and then increase by 1% each year for the next five years, then remain level.
7. Pop down 3% in the first year and then remain level.

According to American Academy of Actuaries (1995), the shocks in the New York Seven scenarios are applied to the 5-year Treasury bond yield. One common practice is to apply the same shocks to the entire yield curve. When in a low interest rate environment, in some scenarios under the New York Seven, it could result in negative interest rates. According to American Academy of Actuaries (1995), a floor of the 50% of the 5-year Treasury bond yield at the beginning of projection is used. A common practice for the ceiling of interest rate is 25% (American Academy of Actuaries, 1995). These approaches are followed for the New York Seven interest rate scenarios testing.

Stochastic Interest Rate Scenarios The CIR model with the same parameters as used in the training stage is used to simulate a sample of 500 interest rate scenarios for testing.

In the training stage of the experiment, the parameters in the CIR model are given below:

- The set of parameters used is $\kappa = 0.2$, $\theta = 0.08$, $\lambda = 0.01$, $\sigma = 0.05$ and $r_0 = 0.04$. t=0,1...,10 and T is from 1 to 10. κ, which represents the speed to revert to the long term mean, is 0.2, the long term mean θ is assumed to be 0.08, the risk premium λ is assumed to be 0.01 and the volatility σ is assumed to be 0.05. t represents the time and T represents the term to maturity of bonds. $r_0 = 0.04$ is chosen since it is the instantaneous bond yield used in pricing. The CIR model does not produce a flat yield curve as used in pricing.

Figure 3.1 displays the yields of bonds of different maturities during the term of the contract in a simulation of 200 interest rate paths. Each plot shows the yields of bond of maturity from 1 to 10 respectively with the x-axis representing time and the y-axis representing the yield from the bond. From the last subplot of this figure it can be seen that for the 10-year bond, the yield has much less change, from t=0 to t=10, than the change from the 1-year bond whose yields are shown on the first subplot of this figure. This is consistent with our expectation, because under a CIR model, interest rates are supposed to approach a long term mean at a speed specified in the model. The yields of the 10-year bond start at a high level and in the long term, there will not be as much rise in the yields as the bonds of shorter maturities.

Figure 3.2(a) presents the yield curve at the start of the first year with x-axis showing the different maturities of bonds and y-axis showing the yield for a particular bond. From this plot, one can roughly check whether the simulated interest rate scenarios are consistent with the parameters in the CIR model. As stated in Kan (1992), the shape of the yield

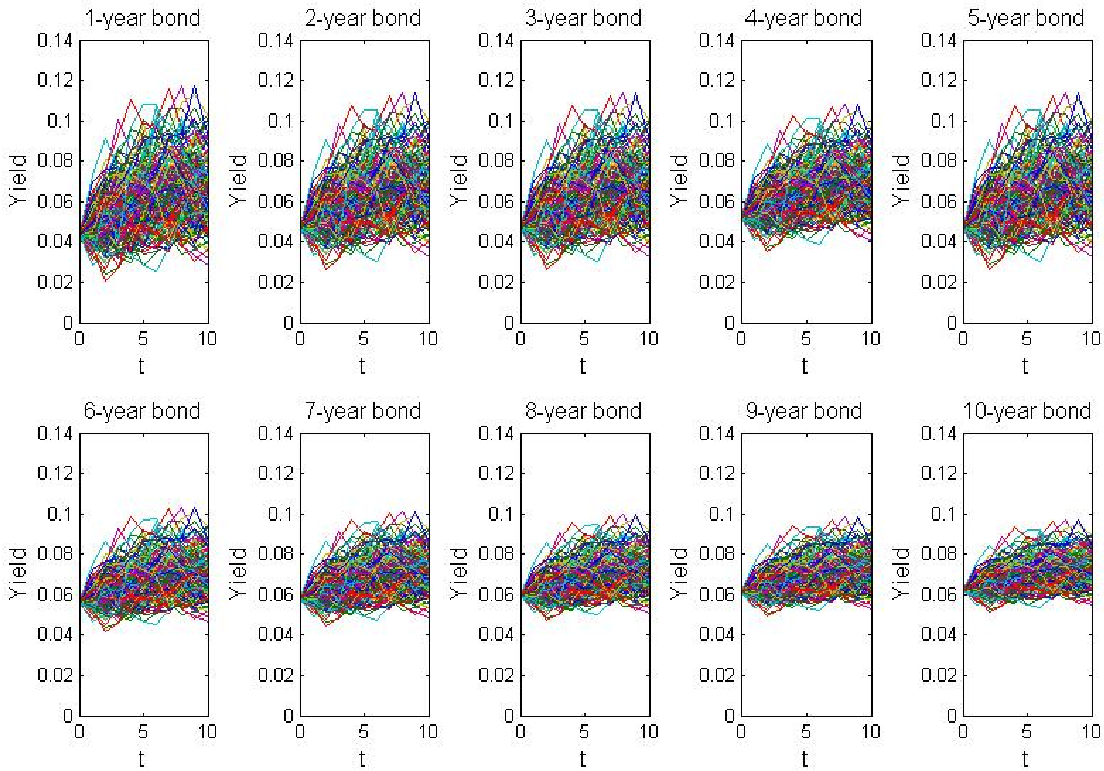

Figure 3.1: Yields of Bonds at t. Each plot shows the yields of a particular bond, with each line representing a scenario. x-axis is the time and y-axis is the value of the yield.

curve from the single factor CIR model should be uniformly increasing when the short rate satisfies the following condition

$$0 \leq r \leq \frac{\kappa\theta}{h}. \tag{3.11}$$

In our case the short rate at time 0 satisfies this condition and the yield curve should be uniformly increasing. The plot of the yield curve as shown in Figure 3.2(a) is consistent with this characteristic of the yield curve.

Figure 3.2(b) presents $E[r(t)]$, the expected value of the short rate at different time. As expected, in the long term it approaches the long term mean parameter θ in the CIR interest rate model.

For deterministic interest rate testing, the New York Seven scenarios are applied to the yield curve at the start of the first year. The yields for the 5-year bond during the term of the contract under the New York Seven scenarios are presented in Figure 3.3. As explained

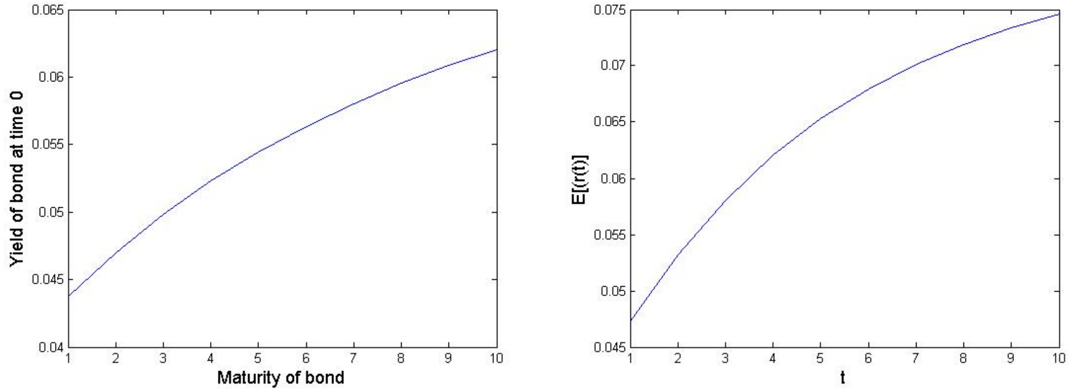

(a) Yield Curve at the Start of First Year

(b) Plot of Expectation of *r(t)*

Figure 3.2: Stochastic Interest Rate from CIR Model

earlier, the yields for other bonds follow the same level of shift to the yield as the 5-year bond.

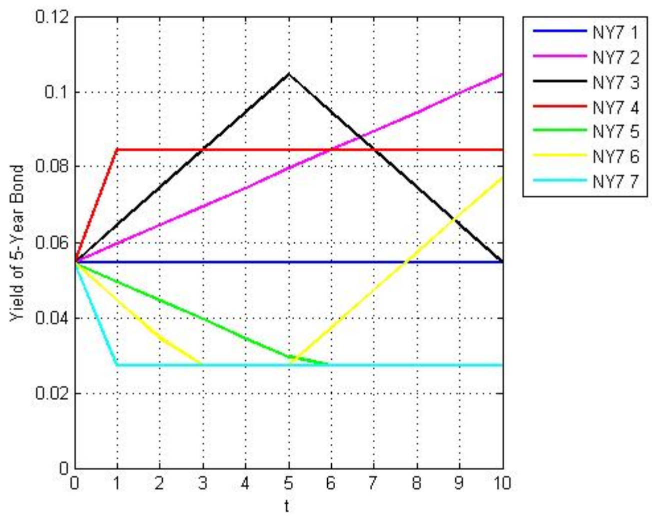

Figure 3.3: New York Seven Scenarios for the Yield of 5-Year Bond

3.3 Mortality Scenarios

The claims incurred are stochastic in nature, with the number of claims and the amount of each claim being random. It is assumed the number of claims to be paid follows a binomial distribution with probability of each claim to be paid out equal to the probability of a person dying during that year.

Because the computing time associated with running many stochastic mortality scenarios is too long, six simple deterministic mortality scenarios are tested under the New York Seven interest rate scenarios and the CIR generated interest rate scenarios.

1. Mortality rate is 10% higher than the mortality table.
2. Mortality rate is 20% higher than the mortality table.
3. Mortality rate is 30% higher than the mortality table.
4. Mortality rate is 10% lower than the mortality table.
5. Mortality rate is 20% lower than the mortality table.
6. Mortality rate is 30% lower than the mortality table.

For example, in the first scenario when the mortality rate is 10% higher, q_x of all ages on the mortality table is multiplied by 1.1 to arrive at the new q_x for this scenario.

Chapter 4

Duration Matching Strategy

The impact of interest rates on traditional life insurance product mainly comes from the changes in the discounting factor of liabilities and the changes of the market value of bonds (Reynolds and Wang, 2007).

Duration is a measure for the sensitivity to interest rate risk for a stream of future cash flows. It is first developed by Macaulay (1938) when the word "duration" is coined. It measures the percentage change in price per unit of interest rate change. It is expected that the higher the duration of a cash flow is, the more sensitive it is to interest rate changes. Generally, duration is defined as

$$Duration = \frac{-\partial Price/\partial r}{Price}. \tag{4.1}$$

Dollar duration can be considered as the numerator of the above equation, which measures the change in dollar value per unit of interest rate change. The dollar duration is given by

$$DollarDuration = \frac{-\partial Price}{\partial r}. \tag{4.2}$$

Therefore, the relationship between duration and dollar duration is given as follows

$$Duration = \frac{DollarDuration}{Price}. \tag{4.3}$$

Note that (4.1) and (4.2) give the duration or the dollar duration of a particular cash flow. A simple example of a 5-year bond is given to illustrate the calculations of duration. Suppose the original yield is r=5% and after an increase of 1%, r'=6%. The prices of the

5-year zero-coupon bond, before and after the interest rate changes, $P(5;r)$ and $P(5;r)'$, are given as follows

$$P(5;r) = (1+0.05)^{-5} = 0.7835,$$
$$P(5;r)' = (1+0.06)^{-5} = 0.7473.$$

Therefore, the duration and dollar duration can be calculated as in (4.1) and in (4.2).

$$\begin{aligned}
P(5;r) &= (1+r)^{-5}, \\
\frac{-\partial P(5;r)}{\partial r} &= -(-5/(1+r)^6) = 5(1+r)^{-6} \\
Duration &= \frac{5/(1+r)^6}{(1+r)^{-5}} = \frac{5}{1+r} = \frac{5}{1.05} = 4.7619 \\
Dollar\ Duration &= P(5;r) \times Duration = .7835 \times 4.7619 = 3.7309.
\end{aligned}$$

Note that the dollar duration can also be approximated by:

$$Dollar\ Duration \approx \frac{-(P(5;r)' - P(5;r))}{r' - r} = \frac{-(0.7473 - 0.7835)}{0.01} = 3.62.$$

In the assets or liabilities portfolios of the problem studied, the cash flows have different maturities and different yields. Therefore, each cash flow has a dollar duration. The dollar duration of a portfolio can be obtained by taking the sum of the dollar duration of the individual cash flows. When considering the change in interest rate, we assume that there is a parallel shift on the yield curve, which means the yields of different bonds are increased or decreased by the same magnitude.

The duration matching strategy as introduced in Reynolds and Wang (2007) is adopted. Note that in Reynolds and Wang (2007), durations are matched at a quarterly frequency. However in this project, we adopt an annual re-balancing strategy. It is expected that the more frequent we re-balance the portfolio, the better hedged it should be against the interest rate risk. An improvement could be to increase the frequency at which we re-balance the portfolio.

Duration matching is used in this study as an evaluation criterion to the strategy produced by GA, which takes into consideration several other objectives. Duration matching strategy is based on the concept of matching dollar duration of both assets and liabilities in order to protect the portfolio against parallel interest rate changes. Some disadvantages with the duration matching strategy is that it may only work well with small changes in interest rate and it only considers parallel shift on the yield curve. Both of these conditions can easily be violated in reality. An alternative strategy would be cash flow matching, however duration matching strategy may allow for more combinations of assets that could possibly provide a higher return. Also, sometimes it is impossible to find an asset with a maturity as long as the liability cash flow, making it impossible to achieve cash flow matching (Fooladi and Roberts, 2000).

The sensitivity to interest rate is measured by $DV01$, or the dollar value of a basis point. One reason for choosing dollar duration is that it is shown to be very close to Macaulay duration in Reynolds and Wang (2007). Another reason is that it is a measure for the dollar value mismatch between the assets and liabilities, which is exactly what we are focusing on for insurance portfolios.

In the rest of the report, $DV01$ will be referred to as dollar duration. It can be calculated from the following equation:

$$DV01 = -(P_u - P_d)/2, \tag{4.4}$$

where P_u represents the price of assets after an upward shock of 0.01% is applied to the yield curve and P_d represents the price of assets after a downward shock of 0.01% is applied to the yield curve.

The shock applied here is uniform across all maturities, which means that the yield curve is increased or decreased by 0.01% for bonds of all maturities. In the case of bonds, the bond values decline when there is an upward shock to the yield curve and the bond values increase when there is a downward shock to the yield curve. Therefore, the difference between P_u and P_d is negative and the negative sign in (4.4) is to allow the value of the dollar duration to be positive. Note that $DV01$ is just an approximation for the dollar duration given in (4.2).

Let $Pn_u(t,s)$ denote the discounted value of future premiums at the beginning of year t after an upward shock of 0.01% is applied to scenario s and $Pn_d(t,s)$ denote the discounted value of future premiums at the beginning of year t after a downward shock of 0.01% is applied to scenario s. They can be calculated by

$$Pn_d(t,s) = \sum_{v=t+1}^{n-1} p(v,s) \times P_d(t, v-t, s)$$
$$Pn_u(t,s) = \sum_{v=t+1}^{n-1} p(v,s) * P_u(t, v-t, s),$$

where $P_d(t,T,s)$ is the bond price at time t with T as the term to maturity given a downward shock for scenario s, $P_u(t,T,s)$ is the bond price (same as the bond price for $P_d(t,T,s)$) given an upward shock for scenario s, $p(v,s)$ is the premiums collected for year v in scenario s.

Let $C_u(t,s)$ denote the discounted value of future claims at the beginning of year t after an upward shock 0.01% is applied to scenario s and let $C_d(t,s)$ denote the discounted value of future claims at the beginning of year t after a downward shock 0.01% is applied to scenario s. They are calculated by

$$C_d(t,s) = \sum_{u=t+1}^{n} c(u,s) \times P_d(t,u-t,s)$$
$$C_u(t,s) = \sum_{u=t+1}^{n} c(u,s) \times P_u(t,u-t,s),$$

where $c(u,s)$ is the claims paid for year u in scenario s.

Let $L_u(t,s)$ denote the discounted value of net liabilities at the beginning of year t after an upward shock 0.01% is applied to scenario s and $L_d(t,s)$ denote the discounted value of net liabilities at the beginning of year t after a downward shock 0.01% is applied to scenario s. They can be obtained through

$$L_u(t,s) = C_u(t,s) - Pn_u(t,s),$$
$$L_d(t,s) = C_d(t,s) - Pn_d(t,s).$$

Let $DV01_{L(t,s)}$ denote the dollar duration of the net liabilities of scenario s and $DV01_{A(t,s)}$ denote the dollar duration of the assets of scenario s. The dollar duration of the net liabilities can be obtained by the same method as given in (4.4):

$$DV01_{L(t,s)} = -\left[L_u(t,s) - L_d(t,s)\right]/2. \tag{4.5}$$

At the beginning of the first year, an initial asset allocation is constructed by matching dollar duration of liabilities $DV01_{L(0,s)}$ and dollar duration of assets $DV01_{A(0,s)}$. At time t, the dollar duration each unit of bond of maturity T from scenario s can contribute to is given in the following equation:

$$-\left[P_u(t,T,s) - P_d(t,T,s)\right]/2. \tag{4.6}$$

Let $u(t,T,s)$ denote the units to purchase at time t for bonds of maturity T of scenario s and the units of bonds to be purchased at the beginning of first year is denoted as $u(0,T,s)$.

If (4.6) is multiplied by $u(t, T, s)$, we obtain at time t the total dollar duration bonds of maturity T can contribute to in future times $t+1$, $t+2$, ..., $t+T$:

$$u(t, T, s) \times -\left[P_u(t, T, s) - P_d(t, T, s)\right]/2. \tag{4.7}$$

At the beginning of the first year, the dollar duration for future times 1, 2, ..., 10 from purchasing $u(0, T, s)$ units of bonds of maturity T bonds is:

$$u(0, T, s) \times -\left[P_u(0, T, s) - P_d(0, T, s)\right]/2. \tag{4.8}$$

The dollar duration at time 0 from future claims at times 1, 2, ..., 10 is given by:

$$-\left[c(T, s) \times P_u(0, T, s) - c(T, s) \times P_d(0, T, s)\right]/2. \tag{4.9}$$

The dollar duration at time 0 from future premiums at times 1, 2, ..., 10 is given by:

$$-\left[p(T, s) \times P_u(0, T, s) - p(T, s) \times P_d(0, T, s)\right]/2. \tag{4.10}$$

To match dollar duration between the assets and the net liabilities portfolios, (4.8) must match the dollar duration from net liabilities at time t for future times 1, 2, ..., 10, as given in the following equation:

$$\begin{aligned} u(0, T, s) \times \left\{-\left[P_u(0, T, s) - P_d(0, T, s)\right]/2\right\} = &-\left[c(T, s) \times P_u(0, T, s) - c(T, s) \times P_d(0, T, s)\right]/2 \\ &-\left\{-\left[p(T, s) \times P_u(0, T, s) - p(T, s) \times P_d(0, T, s)\right]/2\right\}. \end{aligned} \tag{4.11}$$

Rearranging the terms in (4.11), the units of bonds to be purchased at the beginning of first year, $u(0, T, s)$, is given as follows

$$\begin{aligned} u(0, T, s) = &\frac{-\left[c(T, s) \times P_u(0, T, s) - c(T, s) \times P_d(0, T, s)\right]/2}{-\left[P_u(0, T, s) - P_d(0, T, s)\right]/2} \\ &- \frac{-\left[p(T, s) \times P_u(0, T, s) - p(T, s) \times P_d(0, T, s)\right]/2}{-\left[P_u(0, T, s) - P_d(0, T, s)\right]/2}. \end{aligned} \tag{4.12}$$

To purchase these bonds, the capital required for scenario s is given by $CR(0, s)$

$$CR(0, s) = \sum_{T=1}^{10} u(0, T, s) \times P(0, T, s). \tag{4.13}$$

In (4.11), if the terms are rearranged, it can be seen that the units of bonds are purchased to achieve a cash-flow matched portfolio. This is because in each year, the value from the matured bonds is equal to the difference between the claims and premiums cash flows. To achieve this position at the start, additional capital may be raised or some capital may be left after the first capital inflow, which is the first year premiums.

Assuming the capital available at the beginning of first year is the premiums collected $p(0,s)$, capital available $CA(0,s)$ of scenario s, after bonds are purchased for cash flow matching is given by

$$CA(0,s) = p(0,s) - CR(0,s). \tag{4.14}$$

At the beginning of the first year, we adopt an arbitrary strategy by investing capital available $CA(0,s)$ into bonds of maturities 1 and 10. Let $u(0,T,s)'$ denote the additional units of bonds to be invested or short-sold. The additional units of 1-year bond and 10-year bond to be purchased or short-sold are calculated by letting the total contribution of the dollar duration from the additional purchase (or short-selling) of these two types of assets be zero and letting the total capital from purchasing (or short-selling) these two types of assets be equal to the capital available. Then $u(0,T,s)'$ can be obtained by solving the following system of equations:

$$u(0,1,s)'P(0,1,s) + u(0,10,s)'P(0,10,s) = CA(0,s)$$

$$u(0,1,s)'\left\{-\left[P_u(0,1,s) - P_d(0,1,s)\right]/2\right\} + u(0,10,s)'\left\{-\left[P_u(0,10,s) - P_d(0,10,s)\right]/2\right\} = 0.$$

The actual units of 1-year bond, 2-year bond, ..., 10-year bond to purchase at the beginning of the first year are

$$u(0,1,s) + u(0,1,s)', u(0,2,s), \ldots, u(0,10,s) + u(0,10,s)',$$

respectively.

In subsequent years, the capital available is calculated by the following equation

$$CA(t,s) = p(t,s) - c(t,s) + u(t-1,1,s), \tag{4.15}$$

where $u(t-1,1,s)$ represents the 1-year bonds from last year that mature now.

The dollar duration of the bonds portfolio is calculated by the following

$$DV01_{A(t,s)} = \sum_{T=2}^{10} u(t-1,T,s) \times \left\{-\left[P_u(t,T-1,s) - P_d(t,T-1,s)\right]\right\}/2, \tag{4.16}$$

where $u(t-1, T, s)$ represents the number of T year bonds from last year that actually has a maturity of $T-1$ in the current year.

The mismatch between the dollar duration of the liabilities portfolio and the assets portfolio is given as follows

$$DV01_{L(t,s)} - DV01_{A(t,s)}. \tag{4.17}$$

We adopt an arbitrary re-balancing strategy to use bonds of maturity 1 and $10-t+1$. Let $u(t,1,s)'$ and $u(t,10-t+1,s)'$ denote the additional bonds to be invested or short-sold. The number of additional units of bonds to purchase or short-sell at time t in order to match dollar duration can be found by solving the following system of equations.

$$u(t,1,s)'P(t,1,s) + u(t,10+t-1,s)'P(t,10+t-1,s) = CA(t,s) \tag{4.18}$$

$$\begin{aligned} &u(t,1,s)'\left\{-\left[P_u(t,1,s) - P_d(t,1,s)\right]/2\right\} \\ &+ u(t,10+t-1,s)'\left\{-\left[P_u(t,10+t-1,s) - P_d(t,10+t-1,s)\right]/2\right\} \\ &\qquad = DV01_{L(t,s)} - DV01_{A(t,s)}. \end{aligned} \tag{4.19}$$

Next, the bond allocations obtained with the duration matching method are presented. The bond allocation at the beginning of the first year is presented in Table 4.1. This is the allocation of bonds that matches the dollar duration of liabilities for each cash flow. It is equivalent to cash flow matching because there is an asset for every liability cash flow. It shows that the weights in earlier years until year 4 are negative and in later years the weights are positive. This means that in earlier years there should be money left after paying off the claims from the premiums collected and in later years, as people age, the death rates increase and this results in higher claims. Therefore, part of our fund is taken to pay off the claims. The numbers in this table and all subsequent tables for bond allocations in this report represent the percentage of total fund value to be invested in a particular bond.

Table 4.1: Bond Allocation at the Beginning of First Year before Adjustments for Duration Matching. This is to achieve a cash-flow matched position for assets and liabilities and the sum of the weights is 68% of the fund value for all scenarios, with some additional capital left.

1-year	2-year	3-year	4-year	5-year	6-year	7-year	8-year	9-year	10-year
-34%	-26%	-18%	-10%	-2%	6%	14%	21%	29%	88%

After cash flow matching, there is some capital left to be invested somewhere. As shown in Table 4.1, 68% of total fund value collected is enough to achieve cash-flow matched position at the start of the first year. The reason only 68% of total fund value is enough to achieve cash-flow matching is that the yields used in cash-flow matching are higher than the interest rate used in pricing. The dollar duration of future claims is higher than the dollar duration of future premiums, therefore, when yields are increased, the decrease in the present value is higher on the liability side than on the asset side, leaving us with additional capital left. Additional capital available is invested in 1-year and 10-year bonds so that it is not left as cash. After the adjustments, the bond allocation at the beginning of first year is shown in Table 4.2. As displayed in Table 4.2, the weights for 1-year bond and 10-year bond change a lot after this adjustment, while the weights for other bonds do not change. The additional capital is invested to retain the duration matched position of the portfolio.

Table 4.2: Bond Allocation at the Beginning of First Year after Adjustments for Duration Matching. The weights of 1-year and 10-year bonds change from Table 4.1 due to re-balancing. The weights are the same for all scenarios and they sum to 100%.

1-year	2-year	3-year	4-year	5-year	6-year	7-year	8-year	9-year	10-year
2%	-26%	-18%	-10%	-2%	6%	14%	21%	29%	84%

As introduced in Chapter 3, after each year in the projection, the portfolio is re-balanced to enable the dollar duration from the asset side and the liability side to match again. The bond allocations in every year are obtained following this approach and they are presented in Table 4.3. In Table 4.3, each row represents the percentage of total capital invested in each bond in each year. With the duration matching method, the bond allocation for each interest rate scenario is different and the results presented in Table 4.3 are the average weights across all scenarios.

Note that the re-balancing strategy may not always offer the most reasonable solution. For example, in the last row of Table 4.3 which shows the asset allocation at the beginning of year 10, a strategy to short sell some two-year bonds is adopted. A more natural strategy would be investing only in one-year bonds since the only cash flow in the future is in the next year. This strategy is the result of solving the system of equations in re-balancing when there is additional capital to be invested somewhere.

For duration matching experiment there is no “training” stage analogous to experiments that use GA, as for duration matching, the same methodology applies for both the “training” and “testing” stages, which is to achieve a dollar duration matched position. Duration matching only requires the current interest rate and does not involve the optimization

Table 4.3: Bond Allocation at the Beginning of Each Year after Adjustments for Duration Matching. The columns are bonds of different maturities. The rows are different periods. The weights are the average across scenarios.

	1	2	3	4	5	6	7	8	9	10
Yr 1	2%	-26%	-18%	-10%	-2%	6%	14%	21%	29%	84%
Yr 2	8%	-13%	-7%	-1%	4%	10%	16%	21%	63%	0%
Yr 3	11%	-6%	-1%	4%	9%	13%	18%	53%	-1%	0%
Yr 4	15%	-1%	3%	8%	12%	16%	48%	-1%	0%	0%
Yr 5	19%	3%	7%	11%	15%	45%	-1%	0%	0%	0%
Yr 6	23%	7%	11%	15%	45%	-2%	0%	0%	0%	0%
Yr 7	29%	12%	16%	46%	-3%	0%	0%	0%	0%	0%
Yr 8	37%	17%	50%	-5%	0%	0%	0%	0%	0%	0%
Yr 9	52%	57%	-9%	0%	0%	0%	0%	0%	0%	0%
Yr 10	126%	-26%	0%	0%	0%	0%	0%	0%	0%	0%

(training) to obtain a strategy. Further results from duration matching strategy therefore are not shown individually, but are presented together with the results from GA experiments for comparison purposes in Chapter 6.

Chapter 5

Multi-Objective Portfolio Optimization and Active Re-balancing

In the previous chapter, a specific allocation of bonds is obtained by adopting the duration matching strategy. Duration matching is a common practice for insurance companies, because it provides a hedge for interest risk, one of the risks the insurance companies are most concerned about for assets and liabilities management over a long horizon. It is also possible to use GA to integrate other strategies for this line of business. Some research, such as Soam et al. (2012), handles the multi-objective optimization problem of selecting stocks portfolio in multi-periods. In that paper, GA is used in the training stage of the experiment to find investment strategies and in the testing stage, and strategies from the training stage are used with active re-balancing to the portfolio in response to the change in the financial markets.

In this study, a similar approach is followed. 200 stochastic interest rate scenarios from the CIR model as introduced in Chapter 3 are used in the training stage of the experiment. In training, mortality is assumed to follow its expectation exactly in each scenario, which means no mortality risk is assumed. The investment strategies obtained in the training stage are tested under a set of deterministic scenarios like the New York Seven scenarios and stochastic scenarios generated from the same model as the one used for the training stage. The description of the scenarios tested is given in Chapter 3.

In the testing stage, we relax the assumption of mortality risk and let the mortality experience change. We would like to use the same framework for both the duration matching method and the GA method, and evaluate how the GA-based strategy compares to the duration matching strategy. Duration matching can naturally provide us with the initial investment strategy and also the re-balancing strategy by the principle of matching durations. With GA, the criteria is specified in the objective function in order for the solutions to move towards a particular area of the searching space. The advantages of using GA are that it may be able to find a better solution than duration matching and it gives the user the flexibility to specify the objectives. There are a few objectives that should be of interest in managing this insurance portfolio.

1. Achieve a high return on the portfolio at the end of the term, measured by the value of the portfolio or the surplus of the portfolio.

2. Under a set of stochastic interest scenarios or deterministic interest rate scenarios, minimize the semi-standard deviation of the fund value across all scenarios at the end of the term.

3. The plot of the surplus of the portfolio against time fits an upward sloping quadratic line well. It is not desirable that the value of portfolio drops suddenly in a year.

5.1 Details of the GA Framework

The details of the elements in GA are given in Chapter 2 and in this section the assumptions for the GA parameters are given. GA framework can be represented by the following:

$$GA = (C, E, I, S, CO, M, T) \tag{5.1}$$

- C-Coding
- E-Evaluation
- I-Initial
- S-Selection
- CO-Crossover

- M-Mutation
- T-Termination

A population size of 100 is used as a rule of thumb, which means in each generation there are 100 individuals, each of which represents a bond allocation strategy. A larger population size such as 200 has also been tested and it does not offer a better solution.

In each chromosome, the strategies for both the initial allocation and re-balancing are encoded and 19 variables are introduced. The first 10 variables in a chromosome represent the percentage of the portfolio invested in each bond at time 0. To be consistent with the re-balancing strategy in duration matching, at time t when there is additional capital left from paying off the claims or when there is not enough money from the matured bonds and the premiums collected to pay off the claims, a combination of bonds with maturity of 1 and $10+t-1$ are bought or short-sold. The last 9 variables in a chromosome represent the percentage of money to be invested in the bond with maturity of 1 at time t, $t = 1, \ldots, 9$ when re-balancing.

For each variable, $d = 7$ digits is used to allow two decimals for each variable on [0,1]. Therefore, there are in total 19×7 digits in each chromosome. The binary (Gray) numbers are converted into real numbers. For the first 10 variables that represent the initial bond allocation, 0.5 is deducted from each number, to allow for some short selling. For the variables that represent the re-balancing strategy at the end of each year, each number is multiplied by 3 and then deducted by 1, giving an interval of [-1,2]. This is to allow the amount of any purchased asset to be at 200% maximum of the capital available and the amount of any short sold asset be at -100% of the capital available. The sum of amounts of the two re-balancing assets should always be equal to the capital available.

Fund value is defined as the total value of the bonds portfolio. Let $FV(t,s)$ denote the fund value at time t for scenario s. Surplus is defined as the difference between the market value of the portfolio plus the present value of the expected future premiums, and minus the present value of expected future claims. Let $C(t,s)$ denote the present value of future claims at time t under scenario s. Let $Pn(t,s)$ denote present value of future premiums at time t under scenarios s. Let $SP(t,s)$ denote the surplus for a given scenario s at time t. The formula for calculating surplus is given in the following:

$$SP(t,s) = FV(t,s) + Pn(t,s) - C(t,s), \tag{5.2}$$

where

$$C(t,s) = \sum_{u=t+1}^{n} c(u,s) \times P(t,u-t,s),$$
$$Pn(t,s) = \sum_{v=t+1}^{n-1} p(v,s) \times P(t,v-t,s).$$

Let $Sv(t)$ denote the semi-standard deviation of surplus at time t across all scenarios. Let $\overline{SP(t)}$ denote the average of surplus at time t across all scenarios. As previously introduced, the number of interest rate scenarios N=200 is used in training. The average surplus and the semi-standard deviation of the surplus across all scenarios at time t are calculated by the following formula:

$$\overline{SP(t)} = \sum_{s=1}^{N} SP(t,s)/N, \tag{5.3}$$

$$Sv(t) = \sum_{s=1}^{N} \left(max\{SP(t,s) - \overline{SP(t)}, 0\} \right)^2 /N. \tag{5.4}$$

For scenario s, it is desirable that the accumulated surplus at different time t, $SP(t,s)$, can fall on a smooth quadratic line. We introduce a model for the accumulated surplus as given in the following:

$$SP(t,s) = y(t,s) = \boldsymbol{\beta}(s)\boldsymbol{x}(t) \tag{5.5}$$

where

$$\boldsymbol{x}(t) = \begin{bmatrix} t \\ t^2 \end{bmatrix},$$
$$\widehat{\boldsymbol{\beta}(s)} = (\boldsymbol{x}(t)^T \boldsymbol{x}(t))^{-1} \boldsymbol{x}(t)^T y(t,s).$$

The residuals from the regression $\boldsymbol{Re}(t,s)$ are calculated by:

$$\boldsymbol{Re}(t,s) = \widehat{\boldsymbol{\beta}(s)}\boldsymbol{x}(t) - y(t,s). \tag{5.6}$$

The "time" standard deviation for scenarios s, $std(s)$, is given as the sum of squared errors divided by the degree of freedom, which is 11-2=9:

$$std(s) = \sum_{t=0}^{10} \boldsymbol{Re}(t,s)^2/9. \tag{5.7}$$

For a given solution, std is the average "time" standard deviation across all scenarios and is calculated by

$$std = \sum_{s=1}^{N} std(s)/N. \tag{5.8}$$

It is worth noting that other models may represent the accumulated surplus better than the simple quadratic model, and this could be improved in future research.

Selection is based on the score each individual achieves in the objective function, denoted by $ObjV$, which is a weighted average of the three objectives achieved at the end of the contract. The negative signs in the objective function are used because the second and the third terms are to be minimized. The value of the following objective function is to be maximized:

$$\begin{aligned} &ObjV = a \times \overline{SP(10)} - b \times Sv(10) - c \times std, \\ &a + b + c = 1. \end{aligned}$$

Since it is unrealistic to short sell assets more than what you own, a penalty is given to candidate solutions with the weights invested in any bond that are less than -100%, which indicates a short selling that is more than the total fund value. The way to control this is to assign a smaller score to those with too much short-selling, for example, to those candidate solutions that involve short-selling more than 100% of the total fund value. Therefore, after the score of each individual solution is obtained, it is further examined if too much short-selling is involved during the projection years and the following steps are performed for the scores of the solutions.

1. If given a candidate solution, the weight invested in any bond during the projection years are always bigger than -100%, the original score of this solution is accepted.

2. If given a candidate solution, during the projection years, it ever occurs that the short-selling is more than 100% of the portfolio, the original score achieved by this solution is not accepted and the smallest score from this generation of solutions is assigned to this individual solution.

By doing this, we let the solutions that involve too much short-selling receive the lowest rank.

In order to test that the solution from duration matching is just one of the solutions among the solutions space of GA, we specify another two objective functions.

1.

$$ObjV = a \times \overline{SP(10)} - b \times \sum_{t=0}^{10} \left(max\left\{Sv(t) - Sv(t)', 0\right\}\right)^2; \tag{5.9}$$

2.

$$ObjV = a \times \overline{SP(10)} - b \times \sum_{t=0}^{10} \left(Sv(t)' - Sv(t)'\right)^2, \tag{5.10}$$

where $Sv(t)'$ represents the semi-standard deviation of surplus across scenarios at time t with the duration matching strategy. In the first objective function, we penalize the solutions when they give a higher semi-standard deviation than the duration matching strategy. In the second objective function, we penalize the solutions when they give a different semi-standard deviation than the duration matching strategy. Negative signs are used for the second terms in the two objective functions, because a high value from these terms will be penalized.

Three experiments that use the three objective functions are conducted. They are respectively:

- Experiment 1: $ObjV = 0.5 \times \overline{SP(10)} - 0.25 \times Sv(10) - 0.25 \times std$
- Experiment 2: $ObjV = 0.01 \times \overline{SP(10)} - 0.99 \times \sum_{t=0}^{10} \left(max\left\{Sv(t) - Sv(t)', 0\right\}\right)^2$
- Experiment 3: $ObjV = 0.01 \times \overline{SP(10)} - 0.99 \times \sum_{t=0}^{10} \left(Sv(t)' - Sv(t)'\right)^2$

Experiment 1 is specified to put a heavy weight on the return item in the objective function, which is determined by $\overline{SP(10)}$, the average surplus at the end of the contract across all scenarios. Having experimented with different combinations of weights, we have observed that the weight on return should not be too high. If the weight is too high, the increase in return will be accompanied by a much higher increase in both semi-standard deviation and "time" standard deviation. If the increase in surplus at time t to the initial surplus at time 0 is divided by the semi-standard deviation of the surplus at time t, it can be seen that there is a cut-off point when the increase in surplus cannot justify the increase in semi-standard deviation. The weight 0.5 is chosen at a level close to that cut-off point that could yield satisfactory results and avoid a semi-standard deviation that is too high. Experiment 1 aims to demonstrate that by the specified objective function, the strategy obtained with GA is able to outperform the duration matching strategy, by some of the objectives.

Experiment 2 and Experiment 3 are designed to mimic the behavior of the duration matching strategy. In Experiment 2, a heavy weight is added on the second term to prioritize on semi-standard deviation. In the second term, a penalty is given when the semi-standard deviation from the strategy obtained with GA, $Sv(t)$, is higher than the semi-standard deviation obtained with duration matching, $Sv(t)'$. By doing this, we hope to obtain a strategy that allows us to have lower semi-standard deviation than duration matching in each year during the term of the contract. Experiment 2 aims to show that by specifying this objective function, the use of GA enables us to achieve lower semi-standard deviation than duration matching while maintaining a similar return on surplus during the term of the contact.

In Experiment 3, a heavy weight is also imposed on the semi-standard deviation term. This time, a penalty is given to the solutions when the semi-standard deviation from GA-based strategy is different from the semi-standard deviation from duration matching strategy. By doing this, we are tuning the GA to give a solution that should match the performance from the duration matching strategy very closely. If GA is effective in doing that, we illustrate that duration matching offers one solution among the searching space of all the possible solutions. We also illustrate the ability of GA in moving towards a specific subspace of the searching space by the specified objective function.

5.2 Active Re-balancing

In the training stage of the experiment, GA is used on 200 stochastic interest rate paths generated from the CIR model to select an ideal solution for the initial asset allocation and the re-balancing strategy. In the testing stage, the asset allocations are inherited from the training stage and are tested on the deterministic and stochastic interest rate scenarios.

If the interest rate scenarios simulated by the CIR model are sufficient in training and GA has the power to select the optimal strategy, then it is expected that this strategy should also work well under the interest rate scenarios testing.

A few deterministic mortality scenarios are introduced, with some of them involving extreme shocks to mortality rates. When the mortality experience is adverse, claims incurred can be much higher than expected, and this can greatly alter the dynamics of the portfolio. For example, when the actual mortality experience is higher than expected, we might experience negative cash inflow early.

The re-balancing strategy from training is evolved by assuming no risk from mortality. It is expected that when mortality risk is assumed, this may not be the optimum strategy and can face the risk of making wrong decisions. For example, if the capital available is positive in training for a given year, a negative sign from the re-balancing strategy for a particular bond indicates that we should sell that bond and use the money raised together with the capital available to buy more of another bond. While in testing, if adverse claims experience occurs, the capital available can be negative for the same year, if we simply adopt the re-balancing strategy from training, then it indicates we should buy more of that bond by short selling other bonds. As can be seen from this example, the change of patterns (and signs) of the cash flow may defeat the original purpose of the re-balancing strategy that is optimized for the particular mortality scenario in training.

With the current interest rate becoming available as the input, GA is used to train the active re-balancing strategies specific to different scenarios. We investigate whether this approach can deal with adverse mortality scenarios better than simply using the re-balancing strategies from training (referred to as passive re-balancing). Without active re-balancing, GA produces one strategy and uses that on all the scenarios. As explained earlier this strategy may not even be suitable for some scenarios. Duration matching offers the flexibility to produce one strategy for each scenario. With active re-balancing, GA allows us to have a strategy tailored for each scenario, like duration matching does.

To be more specific, with the active re-balancing program, we start with the bond allocations obtained from the training stage. After each year, given the prevailing interest rate that becomes available, stochastic interest rate scenarios are generated based on this current rate from the CIR model. GA is used to find the optimal re-balancing strategy based on the simulated interest rate paths with the interest rate at time t as the current interest rate. A similar objective function as the one in the training stage is used. The specification of the weights in the objective function largely depends on the expectation of the future mortality experience and also the objectives. It is expected that with the suitable objective function, GA is able to select the active re-balancing strategy that offers better results than the re-balancing strategy obtained from training. However, the current re-balancing program is not able to capture all the available information that allows GA to optimize for this strategy. For example, the interest rate model is not refitted each time new experience becomes available and the mortality is always assumed to be as expected regardless of the actual experience.

In the next chapter, we compare the results given by the active re-balancing strategies with those given by the passive re-balancing strategies under the adverse mortality scenarios.

Chapter 6

Experiment Framework and Results

6.1 Introduction

In this section, we introduce the training stage and the testing stage of the experiments. First, let us look at the framework of experiments in the existing literature that solves the multi-period asset allocation problem for stocks portfolios. Historical data of stock prices, which are readily available, are used as the training data. GA is used on this training data to find the optimal investment strategy. Then, the investment strategy obtained is tested on historical data in a later period for its robustness and efficiency. In some experiments, active re-balancing is considered in testing. The active re-balancing scheme has been introduced mainly because during the financial turmoil, static GA-based strategy may not be able to consistently outperform other commonly used methods, such as tracking the market index for stock portfolios, and can make wrong investment decisions (Aranha and Iba, 2008). The main reason for this is that the strategy is trained for a period without stock market shocks but in the testing period, there are stock market shocks.

A possible design for the re-balancing program is to make use of additional information available. For example, in Soam et al. (2012) daily trading volumes of stocks are used as indicators for the future movement of stocks' prices. Historical data are appropriate for this type of study because training and testing periods are relatively short and data are readily available and representative.

In our research, the span of a term insurance product can be as long as ten, 20 or even 25 years. It is difficult to find a period so long with interest rates that are representative of future interest rates. In Figure 6.1, the historical rates for one-year, five-year and ten-year zero-coupon bonds from 1986 to 2012 are plotted. For example, if the period 1986-1996 is used as the training data, it can be observed that in the next ten years 1996-2006 (testing data) there are large changes to the bond rates. This is one example which shows that the training period cannot predict what will happen in the testing period.

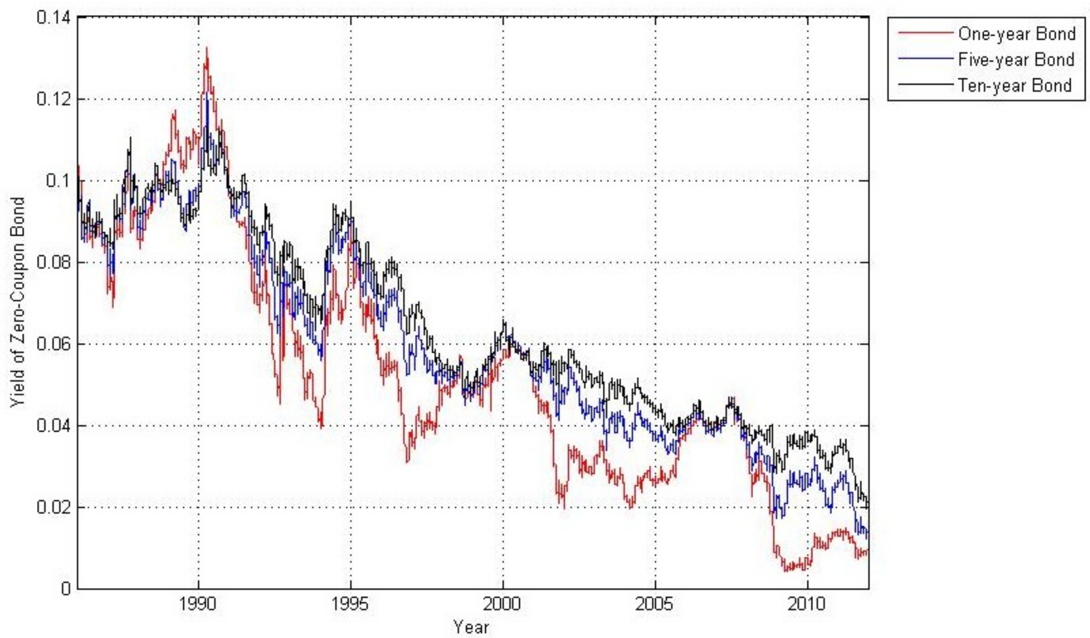

Figure 6.1: Historic Rates of Zero-Coupon Bonds (Bank of Canada)

We choose to simulate interest rate scenarios from a CIR model as the training data. GA procedures are then used to find the investment strategy that is best suited for the simulated interest rate scenarios in the sense of the specified objective function. The process followed is shown in Figure 6.2.

There are a few goals we hope to achieve in the training stage:

- Demonstrate the effectiveness of GA in finding different investment strategies suitable for different risk appetites.
- Illustrate that duration matching is a strategy that belongs to the solutions space of

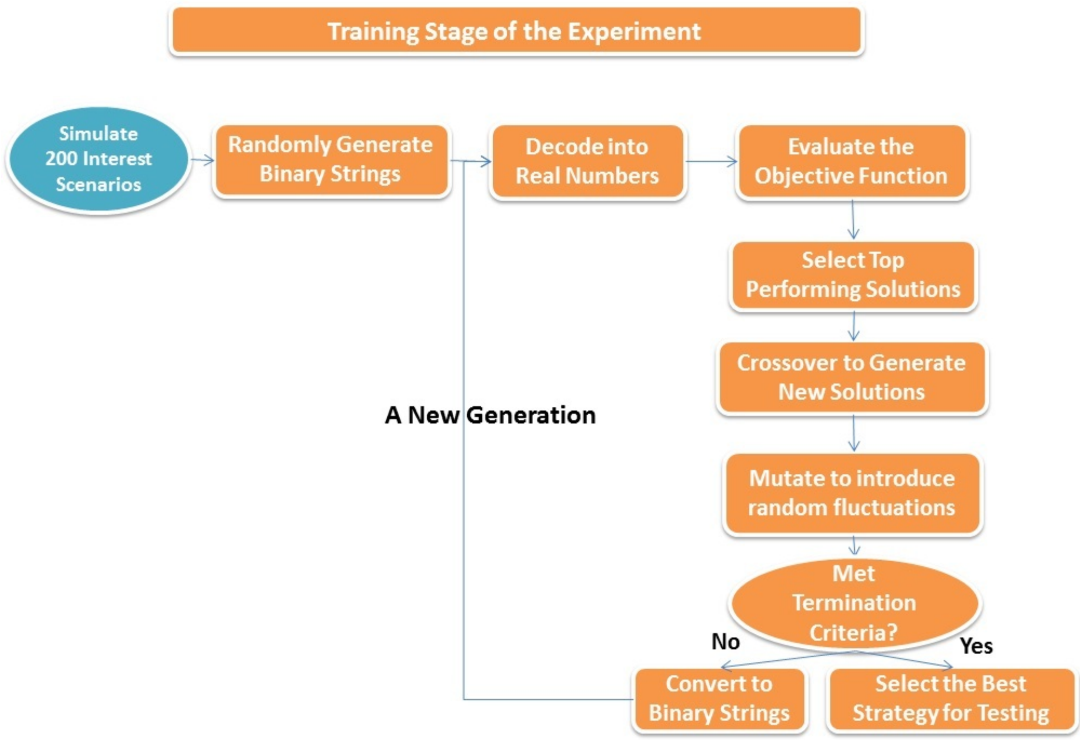

Figure 6.2: Framework of the Training Stage of the Experiment

GA.

In the testing stage, two types of risks facing an insurer are considered, the interest rate risk and the mortality risk. Interest rate scenarios are simulated from the CIR model as the testing data. We adopt the investment strategies obtained in the training stage and test how the portfolio would perform under the simulated stochastic interest rate scenarios. This tests if the scenarios from the CIR model used in the training stage are sufficient or not. In addition, we test our investment strategies on a set of deterministic interest rate scenarios as recommended by insurance regulators. Note that for the interest rate testing we assume the mortality experience to be exactly like expected.

We introduce volatility into mortality experience in the testing stage. We make a few simple deterministic mortality scenarios and test if our investment strategies are still able to achieve excellent results. Each of the mortality scenarios assumes either an increase or a decrease applied to the mortality rates in all the future years, which may represent anti-selection mortality from policyholders or improved mortality experience.

For the purpose of mortality testing, we test each of the scenario under both deterministic interest rate scenarios (New York Seven) and stochastic interest rate scenarios that are from the same CIR model as used in training.

As mentioned in Chapter 5, one advantage of duration matching method is its ability to produce a different strategy for each scenario and this can be a huge advantage in an environment with much variation from expectation. We introduce active re-balancing to GA in the testing stage to allow GA-based strategies to tailor to different scenarios during re-balancing as well. With active re-balancing, we use the information about interest rate and mortality experience that becomes available as an input to the GA and specify the re-balancing strategies at the end of each year as the solutions candidates recommended by GA. Figure 6.3 outlines the flow of the testing stage of an experiment.

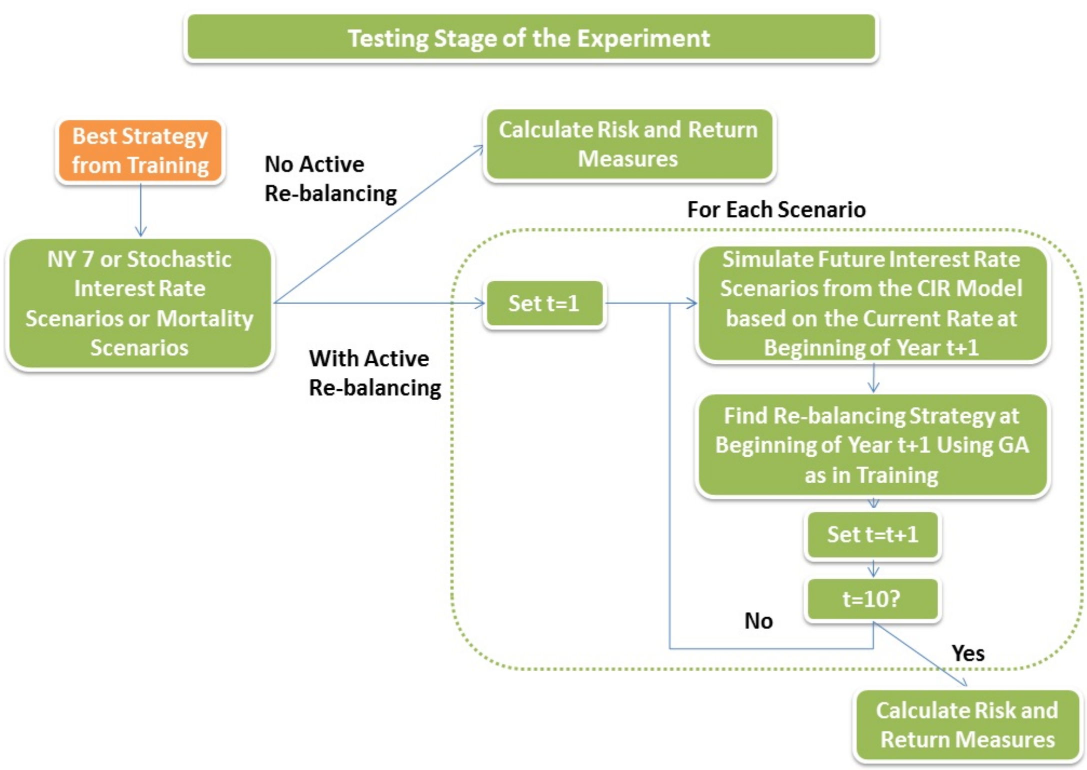

Figure 6.3: Framework of the Testing Stage of the Experiment

There are a few goals we hope to achieve in the testing stage:

- Test the robustness of the GA-based strategies obtained in the training stage and the re-balancing strategies as the interest rate environment or mortality experience changes.

- Test if interest rate simulation from CIR model in training is sufficient.

6.2 Training Stage of Experiments

In this section, we present results from Experiments 1, 2 and 3, as introduced in Chapter 5.

6.2.1 Convergence of GA

First, we present some preliminary results from Experiment 1 in Figure 6.4 and Figure 6.5 to illustrate the convergence of GA.

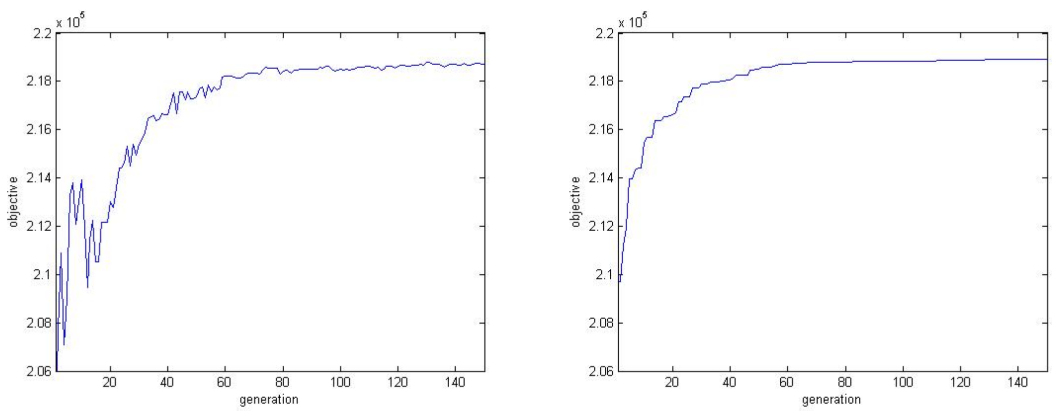

(a) Without Elitism (b) With Elitism

Figure 6.4: Score of the Best Solution in Each Generation. As generations elapse, the best solution improves.

One way to study the convergence of GA is to examine the objective the best solution in each generation achieves. After all, the best candidate available is selected as the proposed investment strategy. In Figure 6.4(a), we see that the objective of the best solution starts at a very low level, which is expected given that the first generation consists of random solutions. As the generations elapse, the best candidate gradually improves and in the last twenty generations, it remains at a high level and does not improve significantly at around generation 120. Note that in some generations, especially in the first few generations, the

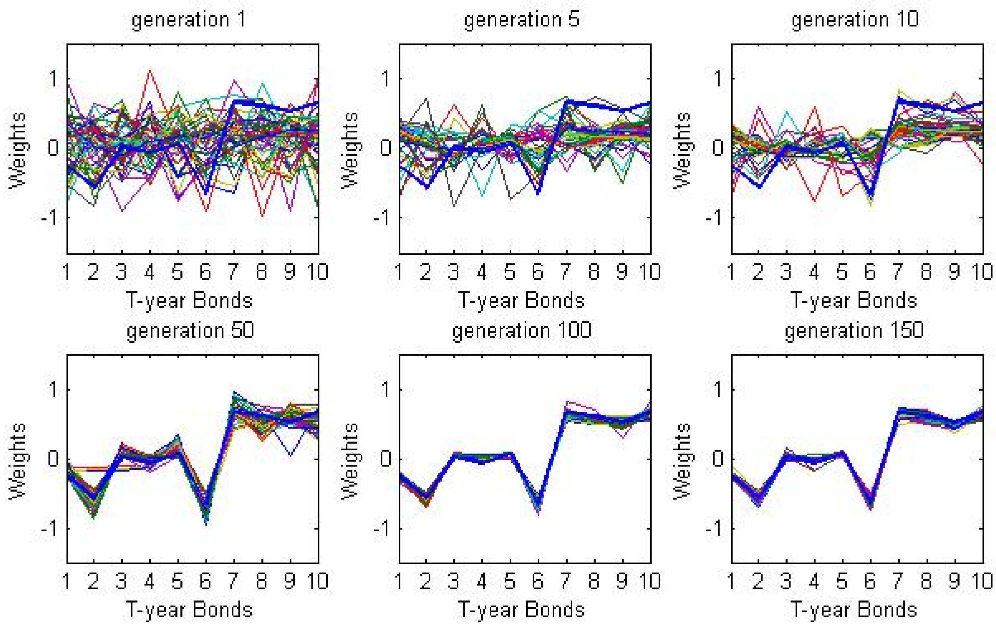

(a) Without Elitism

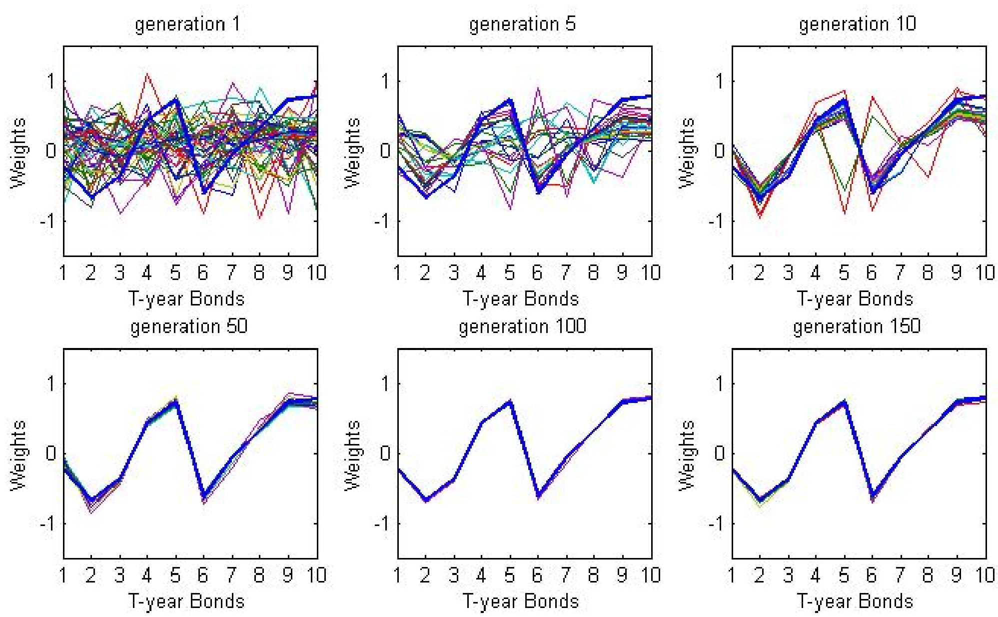

(b) With Elitism

Figure 6.5: Bond Allocations in Different Generations. Each line represents the strategy from an individual solution and all individual solutions are plotted on each small plot.

best solutions drop suddenly to a level much lower than the previous generation. This may be because the best solutions from each generation are not retained, and therefore, causing the convergence to slow down. We can also see that GA indeed involves a certain level of randomness and the exact time for convergence may not be guaranteed, which is considered a weakness of GA.

In Chapter 2 we introduce a way to improve GA that is called “elitism” and it involves retaining the top parent solutions from previous generation and letting them be part of the solutions in the next generation. An experiment using GA with “elitism” has also been carried out. In this experiment, the top 10% of the parent solutions are kept in the next generation. The score of the best solution in each generation is shown in Figure 6.4(b). It can be observed that from generation to generation, the best candidate improves consistently and the convergence occurs at around generation 60, greatly reducing the generations needed for convergence. Also, we check the converged level of objective for both experiments, and it shows that with “elitism”, the objective achieved from the best solution is 0.24% higher than the best solution from the experiment that does not implement “elitism”. It indicates that “elitism” is an effective method to improve the performance of GA.

From these results, we are confident that with GA the performance of the best solutions displays an upward trend and gradually improves over generations. With the implementation of “elitism”, the improvement is more consistent and the convergence is reached more quickly. It can also be observed that in the last few generations, the objective of the best solution remains a flat line, with no further improvement on the objective. This indicates that the 150 generations used are sufficient for convergence. In order to reduce the running time, it may be a good idea to implement “elitism” to the GA program and specify that if no improvement happens in five consecutive generations, GA program will terminate, instead of running 150 generations every time.

Another way to examine the convergence of GA solutions is through checking the bond allocations in each generation. In Figure 6.5(a), each small plot represents the bond allocations of all candidates in generations 1, 5, 10, 50, 100 and 150 for the GA without “elitism”. Each point represents what percentage of the portfolio is invested in bonds of different maturities. The points are connected together to form lines to reveal allocation patterns. It is obvious that the plots look chaotic in generations 1, 5 and 10 because solutions in the first generation are randomly selected. Candidates converge to a clear pattern from generation 50. In generation 150, almost all the existing solutions converge to a single solution, which is ideal. In Figure 6.5(b), we present the same results for GA with "elitism" implemented. It is obvious that the convergence is much faster.This is another way to show the successful convergence of GA. It can be observed that the converged asset allocations are different under the GA with and without "elitism" despite using the same initial generation. This may indicate that there are multiple solutions to arrive at the same objective.

In addition, we can check from Figure 6.5(a) and Figure 6.5(b) that the short selling of any bond in these two portfolios is both less than 100% of the total value of the fund at the beginning of the first year. This is consistent with our expectation because a heavy penalty is imposed on solutions with excessive short-selling of a single bond. However, the total amount of short-selling of all bonds is more than 100% of the fund value. This may be impossible to achieve in practice and a further improvement could be limiting the total amount of short-selling to be less than 100% of the fund value or to be less than or equal to the maximum amount of short-selling allowed under the duration matching strategy. An experiment forcing the total short-selling to be reduced is conducted too and by the same objective function as the one used for Experiment 1, GA is able to find a solution that produces very similar results in both training and testing as the strategy in Experiment 1, with total short selling smaller than 100% of the total fund value. When a stricter short selling constraint is imposed, say, less than 50% of the total fund value, the results become worse than the results when a higher short selling amount is allowed.

6.2.2 Results from the Training Stage of Experiments

Having examined convergence of GA from Experiment 1, we present the results in the training stage for Experiments 1, 2 and 3 in this subsection.

Figure 6.6 shows the results from Experiment 1.

As introduced earlier, in Experiment 1, we aim to find a strategy that can outperform the strategy from duration matching by the risk and return measures we specify. Figure 6.6(a) shows the average accumulated surplus across all scenarios for each year, with the blue line representing the strategy proposed by Experiment 1 and the red line representing the strategy from duration matching. The x-axis is time and the y-axis is the accumulated surplus at each time. Both strategies start at a surplus of the same level. The initial surplus is not zero because we use a flat interest rate in pricing and we use that interest rate as the short rate in the CIR model for the interest rate paths simulation, which means that the yield curve at the beginning of the first year is not flat as assumed in pricing. The initial

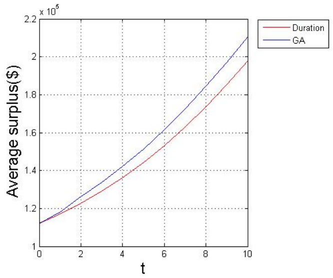

(a) Average Surplus

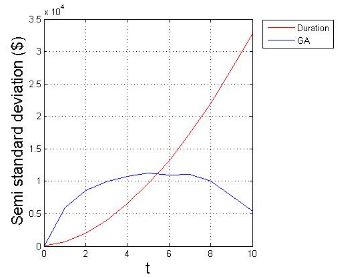

(b) Semi-standard Deviation

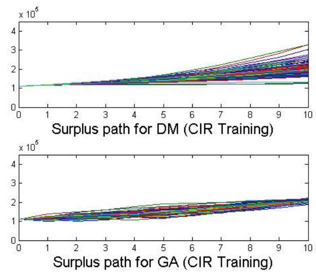

(c) Surplus Path

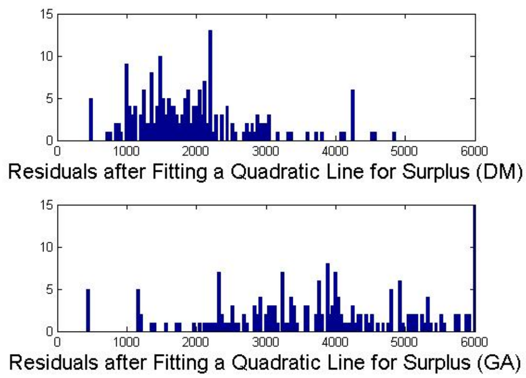

(d) Residuals after Fitting Surplus Path to Quadratic Line

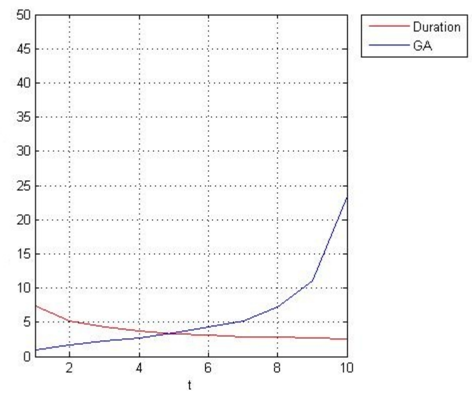

(e) $\frac{\overline{Sp(t)} - Sp(0)}{Sv(t)}$

Figure 6.6: Experiment 1 Training. Aims to let GA strategy outperform DM strategy by certain objectives.

surplus is due to the fact that the yields of bonds generated by the CIR model are higher than the interest rate used in pricing. The calculations for both the GA-based strategy and the duration matching strategy use the same interest rate paths simulated by the CIR model. The GA-based strategy is able to produce a slightly higher accumulated surplus in each year and results in higher surplus than the duration matching strategy at the end of the contract, which is ideal.

Figure 6.6(b) shows the semi-standard deviation of surplus across scenarios for both strategies in each year. GA has a higher semi-standard deviation in early years and has a lower semi-standard deviation in later years. It can be observed that the semi-standard deviation at the end of the contract from duration matching is a lot higher than that of GA-based strategy. From this, we can tell that GA selects a strategy that allows us to reduce the semi-standard deviation across scenarios at the end of the term, as we have specified in the objective function. However, the reduction in semi-standard deviation may be a result of letting the GA program be too greedy for that objective. This may indicate that the weight we put on semi-standard deviation could be too high and we sacrifice extra return on surplus to achieve this.

Figure 6.6(c) displays the paths that surplus follows for all scenarios during the term of the contract. The surplus paths from the duration matching strategy become more and more widely dispersed in later years while with the GA-based strategy, they become a bit more widely dispersed until in the middle when they start to become narrower approaching the end of the term. The reason for this particular shape of the surplus paths for the GA-based strategy is that we put a priority on the semi-standard deviation at the end of the contract. As we can see from this figure, GA can help us achieve this goal. However, it loses to the duration matching strategy by the semi-standard deviation measure in the early years.

Figure 6.6(d) displays the distribution of the sum of squared residuals for different scenarios after fitting the surplus paths to quadratic lines. As discussed before in Chapter 5, the reason for making this an objective is that an upward sloping shape of the surplus indicates consistent growth in all the years, which is considered desirable. It is shown that with the duration matching strategy, the residuals are smaller, indicating a better fit to quadratic lines, which is a desirable feature. This is also consistent with the results in Figure 6.6(c). The surplus paths from the GA-based strategy are not as smooth as the surplus paths from the duration matching strategy, and hence the larger residuals. We notice that in Experiment 1, we put more weight on risk and return while putting less weight on the measure for consistent growth.

Figure 6.6(e) shows the ratio between the difference of the average surplus at time t and the initial surplus at time 0 to the semi-standard deviation at time t. It is a measure of whether the additional volatility can be justified by the additional return, which is akin to the Sharpe ratio used to evaluate stocks portfolios. GA-based strategy yields worse result than the duration matching strategy in the first few years, however in later years, it results in a much better performance according to this measure.

From the training results of Experiment 1, we observe that the GA-based strategy offers a trade-off between the objectives. It offers better results for the return and semi-standard deviation at the end of the contract while suffering from higher volatility on surplus in the early years and loses the consistent growth from year to year on surplus.

Figure 6.7 presents the results from Experiment 2.

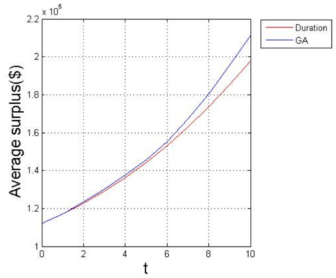

(a) Average Surplus

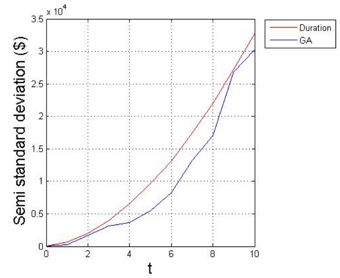

(b) Semi-standard Deviation

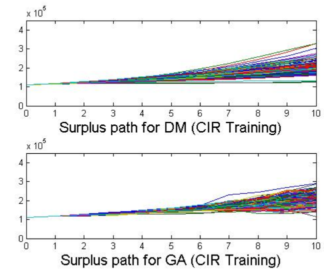

(c) Surplus Path

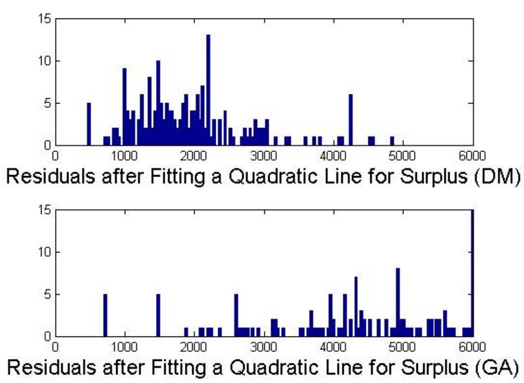

(d) Residuals after Fitting Surplus Path to Quadratic Line

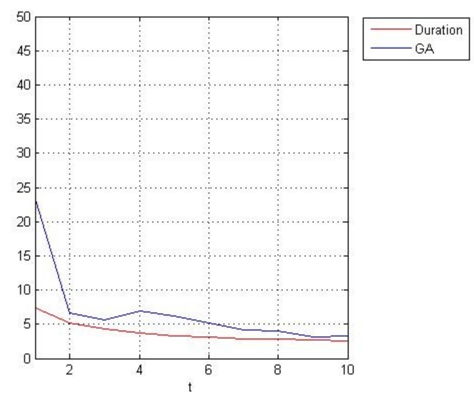

(e) $\frac{\overline{Sp(t)} - Sp(0)}{Sv(t)}$

Figure 6.7: Experiment 2 Training. Aims to let GA strategy achieve a lower semi-standard deviation in all years than DM strategy.

As introduced earlier, in Experiment 2, we aim to let GA find a solution that mimics the performance of duration matching while achieving an even lower semi-standard deviation when possible.

Figure 6.7(a) shows that the GA-based strategy is still able to produce a slightly higher average accumulated surplus in each year than the duration matching strategy. The results look very much like the results from Experiment 1 as shown in Figure 6.6.

However, the graph for semi-standard deviation looks very different from that in Experiment 1. Figure 6.7(b) shows that with the GA-based strategy we are able to achieve a lower semi-standard deviation in all years than the duration matching strategy, although as the term approaches maturity, the semi-standard deviation under both strategies become very close to each other. From this, we can see that GA is indeed very sensitive to the objective function and is very effective in finding the solution that is perfectly aligned with the objective function.

In Figure 6.7(c), the surplus paths of GA-based strategy appear to be narrower than the surplus paths of the duration matching strategy in the early years, and in the last few years, they begin to look like that of duration matching. This is consistent with the results shown in Figure 6.7(b), with GA-based strategy having a much lower semi-standard deviation in earlier years and becoming more like duration matching approaching the end of the term.

Figure 6.7(d) displays the distribution of the sum of squared residuals after fitting the surplus paths to quadratic lines for all scenarios. We could see for the duration strategy, it is much more to the left, indicating a better fit to quadratic lines than the GA-based strategy. Similar to Experiment 1, we put less weight on this measure in the objective function of Experiment 2, and hence the larger residuals of the GA-based strategy.

Figure 6.7(e) shows the ratio between the difference of surplus at time t and the initial surplus at time 0 to the semi-standard deviation at time t. With the objective function in Experiment 2, we are able to produce a higher ratio in all years although as we approach the maturity of the term, this measure under the GA-based strategy becomes very close to that from the duration matching strategy. This result is consistent with the results in other figures of this experiment.

Figure 6.8 displays the results from Experiment 3.

In Experiment 3, we aim to let GA find a solution that exactly mimics the performance of duration matching.

In Figure 6.8(a) we could see that the average accumulated surplus in all years from the GA-based strategy is very closely aligned to that of the duration matching strategy, resulting in a slightly higher end of term surplus. Figure 6.8(b) shows that the semi-standard deviation from both strategies are very close to each other in all years. Figure 6.8(c) displays the surplus paths for all scenarios during the term of contact. Each line represents the surplus path of a scenario. As expected the shapes of both strategies look very identical. Figure 6.8(d) displays the distribution of the sum of squared residuals after we fit the surplus path to a quadratic lines for all scenarios. We could see that this time the GA-based strategy is able to achieve a slightly better result than duration matching. Figure 6.8(e) shows that the ratio of the increase in surplus to the semi-standard deviation at time t. Again, both achieve very similar results by this measure.

From the results of the three experiments we have carried out, we see that GA is very

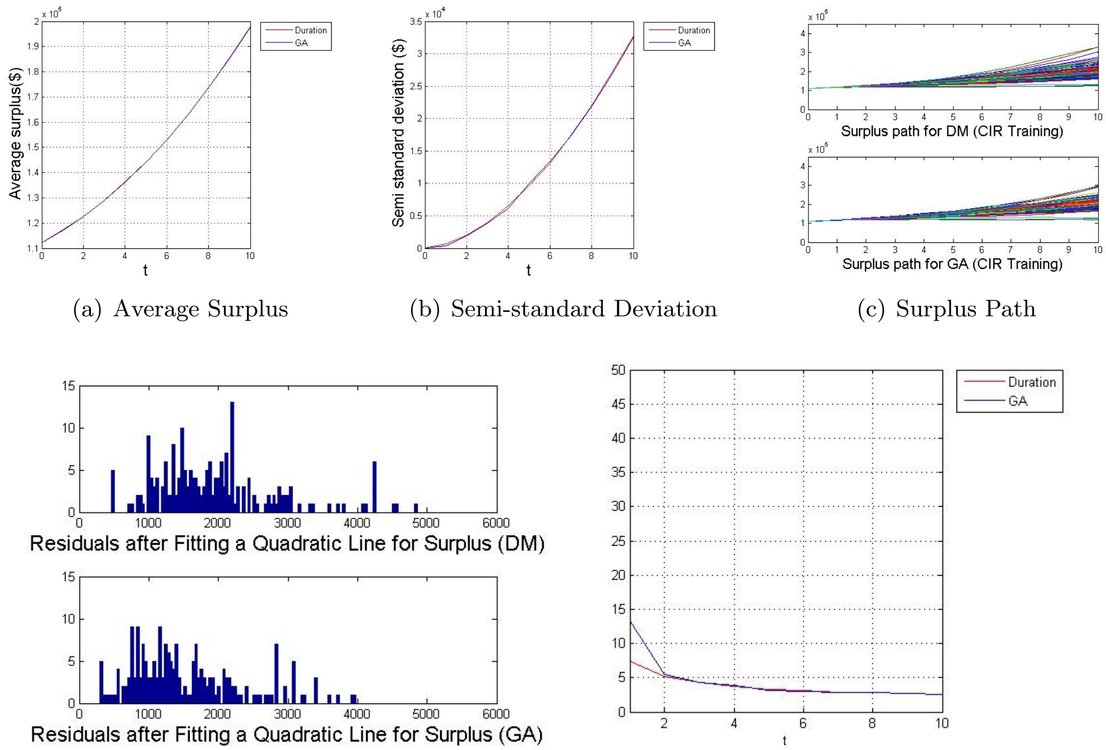

(a) Average Surplus (b) Semi-standard Deviation (c) Surplus Path

(d) Residuals after Fitting Surplus Path to Quadratic Line (e) $\frac{\overline{Sp(t)} - Sp(0)}{Sv(t)}$

Figure 6.8: Experiment 3 Training. Aims to let GA strategy mimic DM strategy.

effective in tuning the solutions so that they are perfectly aligned with the objective functions we specify. From the results in Experiment 3, we also see that GA is capable of generating a strategy that mimics the duration matching strategy very closely. In addition to the duration matching solution, GA is capable of generating other solutions that meet different preferences from the user. From Experiment 3, we see that by specifying semi-standard deviation alone, we can let GA find solutions that are perfectly aligned with that of duration matching in terms of the three objectives we specify.

Last but not least, we present the average dollar duration of the bonds portfolio across scenarios for the three experiments in Figure 6.9. Consistent with our expectation, the dollar duration of GA-strategy follows duration matching most closely in Experiment 3, as presented in the third plot. In Experiment 1, GA-based strategy is able to produce a lower level of long term volatility while having a higher dollar duration in every year. Experiment 2 aims to offer a lower volatility in all years and results in the dollar duration being very high in some years.

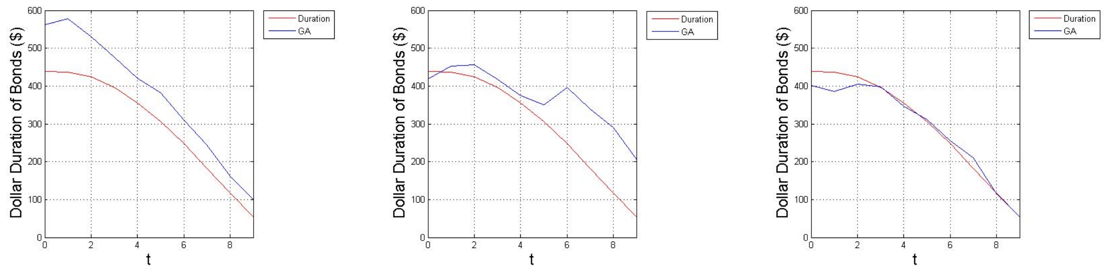

(a) Experiment 1 (b) Experiment 2 (c) Experiment 3

Figure 6.9: Average Dollar Duration of Bonds Portfolios: GA Compared to DM in Experiments 1, 2 and 3

While Experiments 2 and 3 are designed to illustrate one aspect of GA, we choose the asset allocation strategy obtained from Experiment 1 to conduct the testing, because it offers the lowest long term volatility among the three and has the highest potential so far to outperform the duration matching strategy. Note that with the duration matching method we obtain different weights for each scenario and that is where GA might have a disadvantage in the testing stage. We expect that if we simply adopt a very similar strategy as duration matching in the testing stage, such as the strategy from Experiment 3, we will end up with worse results than the strategy from duration matching. This is because the net cash flows will be very different when we introduce adverse mortality scenarios and although the strategy from Experiment 3 works well in the training stage, it may not be able to adapt as well as the duration matching strategy which has the flexibility to tailor to a particular scenario in testing.

6.3 Testing Stage of Experiments

In the testing stage of the experiments, we test how the investment strategy obtained from the training stage would perform under both deterministic and stochastic interest rate scenarios. The claims experience is set to be the same as expected for both types of interest rate testing and this is relaxed to be deterministic scenarios later. The deterministic interest scenarios tested are as follows:

- New York Seven scenarios with a current interest rate of 4%. The 4% used is the short rate used in the CIR model.

The stochastic interest scenarios tested are:

- 500 stochastic interest rate scenarios from CIR model with the same parameters as the ones used in training.

6.3.1 Deterministic Interest Scenarios

We present the testing results under deterministic interest rate scenarios with the strategy obtained in Experiment 1 in this subsection. Again, note that this is assuming a deterministic mortality experience as well. As introduced earlier, for duration matching, the portfolio is always actively re-balanced. For the GA-based strategy, no active re-balancing means the asset allocation obtained in the training stage, including both the initial weights and the re-balancing weights, are used directly in the testing stage without considering additional information about interest rate or mortality experience as the projection progresses. For GA-based strategy, if active re-balancing is used, only the initial asset allocation, but not the re-balancing weights, obtained from the training stage is kept. Each year, the re-balancing weights are obtained through training a new GA-based strategy with the additional information about interest rates and claims experience as they become available. We do not consider active re-balancing in this section and will discuss active re-balancing strategies in Section 6.5. More details about active re-balancing are given in Chapter 5.

First, we present detailed results in Table 6.1 and Table 6.2 for the surplus at the end of the contract under the New York Seven scenarios for both GA-based strategy and the duration matching strategy. We compare them to the surplus obtained from the training stage as well. In these two tables, the seven columns in the middle show the end surplus under the New York Seven scenarios. The last column is the average surplus at the end of the contract for both strategies in the training stage. The last row shows the percentage change of surplus from the New York Seven testing to the average value in training.

Comparing these two tables, we discover that the percentage changes from GA-strategy are a lot more modest than the percentage changes from the duration matching strategy. Duration matching strategy performs better under two scenarios, NY2 and NY4. Both of these scenarios present an increase in interest rates and then a subsequent flattening of interest rates. GA-strategy sacrifices the excess return from these scenarios but achieves much more stable results under the scenarios with severe drops in interest rates, such as scenarios NY5, NY6 and NY7. For the flat scenario, NY1, the GA-based strategy is also able to outperform the duration matching strategy.

Table 6.1: Experiment 1 New York Seven Testing End of Term Surplus (in thousands of dollars). The first row is the end of term surplus achieved through the Experiment 1 strategy under the New York Seven scenarios. The second row is the change to the average end of term surplus in training.

	NY1	NY2	NY3	NY4	NY5	NY6	NY7	Ave in Training
EoT Surplus, Testing	214	210	215	208	209	212	214	210
% Change from Training	2%	0%	2%	-1%	0%	1%	2%	N/A

Table 6.2: Duration Matching New York Seven Testing End of Term Surplus (in thousands of dollars). The first row is the end of term surplus achieved through the duration matching strategy under the New York Seven scenarios. The second row is the change to the average end of term surplus in training.

	NY1	NY2	NY3	NY4	NY5	NY6	NY7	Ave in Training
EoT Surplus, Testing	165	218	212	223	145	167	143	200
% Change from Training	-17%	9%	6%	12%	-27%	-16%	-28%	N/A

We also present the performance of the portfolio throughout the term of the contract under the New York Seven scenarios in Figure 6.10. As a reminder, the details about the New York Seven scenarios were introduced earlier in Chapter 3.

In Figure 6.10(a), we present the average accumulated surplus across New York Seven scenarios. We could see that the GA-based strategy is able to outperform duration matching in all years. Figure 6.10(b) shows that the semi-standard deviation under the GA-based strategy is higher in the early years and lower in the later years, ending with a much lower semi-standard deviation at the end of the contract. Figure 6.10(c) displays that the GA-based strategy is able to achieve excellent results for the ratio between the increase in surplus to the semi-standard deviation too.

Figure 6.10(d) shows the surplus paths at different times for the seven scenarios tested. It is consistent with the results in the training stage, with the GA-based strategy able to achieve a much narrower looking shape than the duration matching strategy towards the end of the contract. It can be observed that in the lower plot of the figure, which presents the surplus paths under GA, in some scenarios, the surplus keeps growing during the term of the contract while for some others, the surplus paths do not grow as fast. In the upper plot of this figure, which presents the surplus paths under the duration matching strategy,

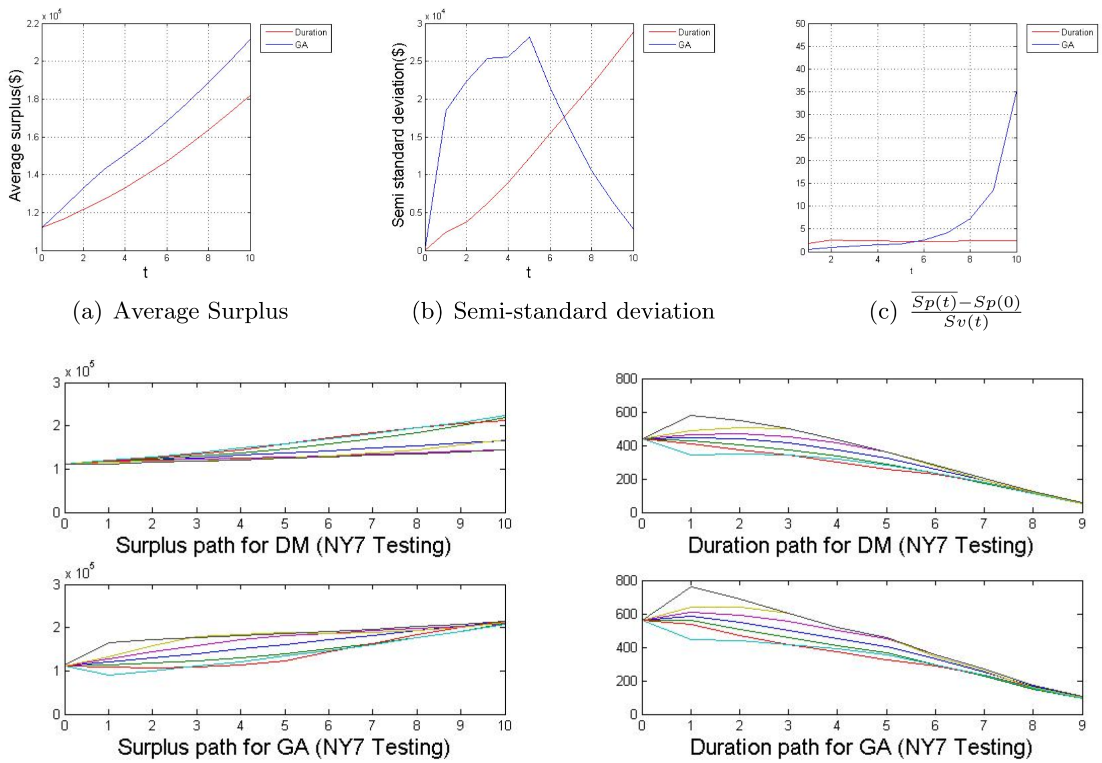

(a) Average Surplus (b) Semi-standard deviation (c) $\frac{\overline{Sp(t)}-Sp(0)}{Sv(t)}$

(d) Surplus Path (e) Duration Path

Figure 6.10: Experiment 1 New York Seven Testing

we could observe that the surplus paths under all scenarios keep growing. At the end of the contract, the variation among the seven scenarios under the GA-based strategy is much smaller than the variation from the duration matching strategy. From this we can see that the GA-based strategy from Experiment 1 handles shocks in interest rates well and allows us to achieve our objective. The CIR scenarios used in training are able to capture the New York Seven scenarios that involve the shocks. However, the particular shape of the surplus paths is one disadvantage of this GA strategy. For example, it can be observed that in the first five years, GA-strategy has a higher semi-standard deviation on surplus among different scenarios, as shown in Figure 6.10(d). GA-based strategy intentionally allows this to happen because following this strategy, the surplus paths under different scenarios can join to achieve very lower semi-standard deviation at the end of the contract.

Finally, in Figure 6.10(e) we present paths for the dollar duration of the bond portfolio. It indicates that the GA-based strategy chooses bonds of longer maturity to invest on average.

6.3.2 Stochastic Interest Scenarios

In this subsection, we present the results from testing the robustness of the investment strategy from Experiment 1 under 500 stochastic interest rate scenarios. The results are presented in Figure 6.11.

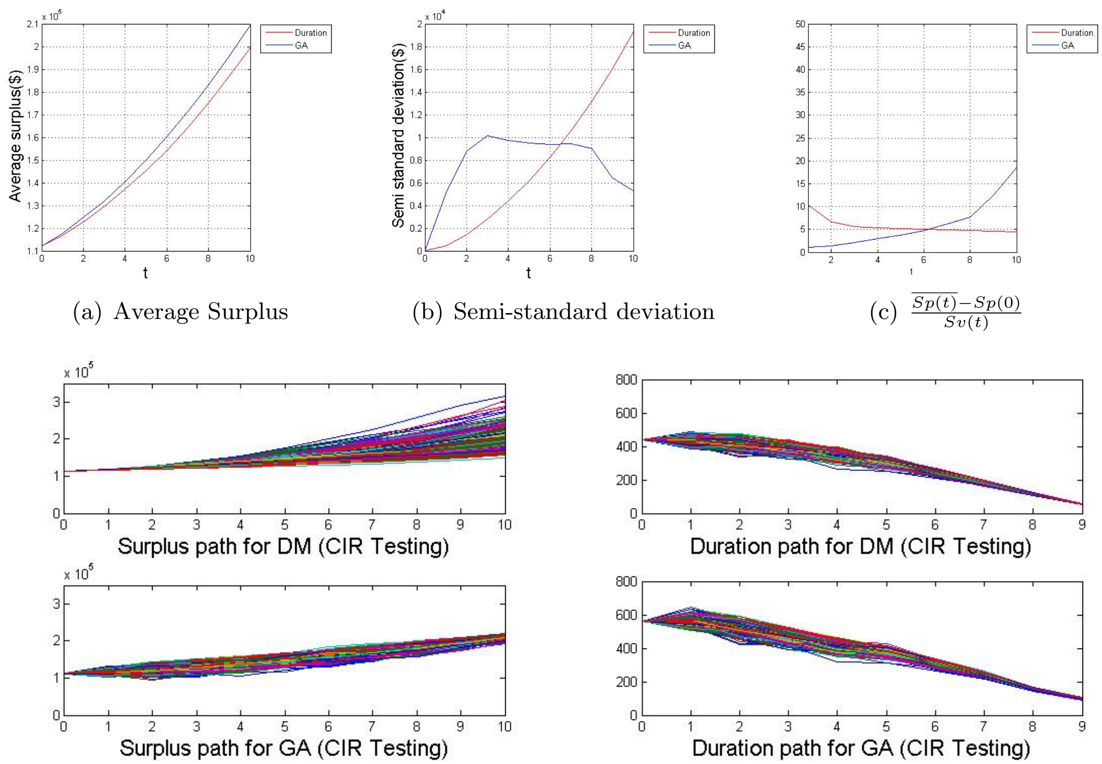

(a) Average Surplus (b) Semi-standard deviation (c) $\frac{\overline{Sp(t)-Sp(0)}}{Sv(t)}$

(d) Surplus Path (e) Duration Path

Figure 6.11: Experiment 1 Stochastic Interest Testing

In Figure 6.11(a) shows the accumulated average surplus across the 500 scenarios tested. It shows that GA-based strategy achieves very similar results as the duration matching strategy, while being slightly better at the end of the contract. In Figure 6.11(b), we present the semi-standard deviation of surplus across the 500 scenarios tested. Compared to the results in training as shown in Figure 6.6(b), again the testing results are consistent, with GA-based strategy achieving a lower level of semi-standard deviation at the end of the term. In Figure 6.11(c), consistent with our expectation, the ratio between the increase in surplus to the semi-standard deviation is very favorable for the GA-based strategy. In Figure 6.11(d), we present the surplus paths across time. Consistent with the results in

Figure 6.11(b), with the GA-based strategy, the plot of surplus paths looks much narrower during the term of the contract than the surplus paths that the duration matching strategy yields, which again looks like a fan. It is worth noting though, that GA sacrifices the high return in some scenarios to achieve this. As we compare the two plots in Figure 6.11(d), we can see that for a few scenarios, duration matching offers much higher end of term surplus. In Figure 6.11(e), we present the paths of dollar duration across time for the bonds portfolio. Again, GA seems to invest in bonds of longer maturities.

From Figure 6.11 we can see that the results are similar to those in the training stage as presented in Figure 6.6 and the strategy from Experiment 1 is very robust under the testing. It indicates that the 200 interest rate scenarios used in the training of the GA-based strategy are sufficient to represent the distribution of the interest rate paths. Further research can be done to use different parameters in the CIR model to generate the interest rate scenarios in the testing stage.

6.3.3 Mortality Scenarios

Last but not least, we introduce mortality scenarios to the testing. We introduce a few simple deterministic mortality scenarios as discussed in Chapter 3. Adverse mortality experience can be caused by anti-selection behavior from policyholders or a pandemic like the Influenza in 1918. Since our mortality scenarios involve an increase or decrease to the mortality rates in all future years, the adverse scenarios are more likely to represent the effects from the anti-selection behavior to mortality and the favorable scenarios are more likely to represent the phenomenon that people are living longer.

The strategy obtained from Experiment 1 is tested on each of the six mortality scenarios with either CIR interest rate scenarios or New York Seven scenarios. The surplus paths of the six mortality scenarios under the testing of New York Seven scenarios are presented in Figure 6.12. Each plot represents one of the six mortality scenarios and each line represents the paths that the surplus follows under one of the New York Seven scenarios.

It can be observed that when mortality experience is better than expected, the GA-based strategy from Experiment 1 is able to generate higher average surplus across the scenarios while maintaining lower semi-standard deviation than the duration matching strategy. The other results not shown confirm this as well, showing that Experiment 1 strategy is able to outperform duration matching in most of the New York Seven scenarios, with a higher average surplus at the end of the term. The semi-standard deviation among the New York

Seven scenarios is lower too.

When mortality experience is worse than expected by 10%, the other results not shown indicate that Experiment 1 strategy is not able to outperform duration matching on both of the objectives. It is able to achieve a slightly higher average surplus however also having a much higher semi-standard deviation at the end.

When mortality experience is worse than expected by 20% or even 30%, the results reveal that Experiment 1 strategy is not suitable. As can be observed from Figure 6.12, Experiment 1 strategy yields surplus paths that have larger variation among scenarios. The detailed results indicate that on average, the end of term surplus is lower too.

The surplus paths under the six mortality scenarios with CIR-generated interest rate scenarios are presented in Figure 6.13. The results are similar to those under the New York Seven scenarios. The GA-based strategy of Experiment 1 performs well when mortality experience is better than expected, achieving higher surplus and lower semi-standard deviation. When mortality experience is adverse, the GA-strategy achieves very similar or slightly worse on the surplus side and suffers from increased volatility across scenarios.

The results of Experiment 1 strategy from mortality testing suggest that this strategy is not optimal for the adverse mortality experience, especially when the actual mortality rates are 20% or 30% higher. From the training stage for Experiment 1 we have seen that the strategy is optimized in a way that favors a few objectives more than others. For example, the strategy obtained is able to achieve a slightly higher return and a much lower semi-standard deviation at the end of the term. However, it falls short on the measure, "time" standard deviation, as evidenced by Figure 6.6(c). In this figure, it can be seen that although we are able to achieve a narrow variation in the surplus at the end among all the scenarios, for some scenarios, the surplus paths flatten approaching the end of the contract, indicating a faster growth of surplus in earlier years followed by a slower growth in later years. This is not most desirable because ideally we would like the surplus to grow consistently and that is why we choose to fit them into a quadratic line in the first place. To summarize, the GA-based strategy of Experiment 1 sacrifices "time" standard deviation for the low semi-standard deviation at the end of the contract. This is caused by an overly heavy weight on the semi-standard deviation relative to the "time" standard deviation in the objective function.

We specify semi-standard deviation as a top priority because it is certainly very desirable to gain comfort that we will achieve less variation on surplus under different scenarios at the

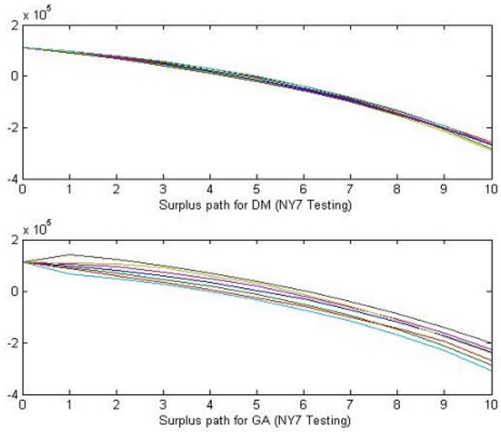

(a) Surplus Path: Mort Increase 10%

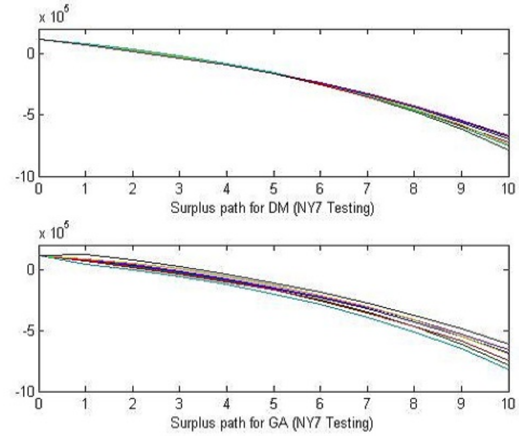

(b) Surplus Path: Mort Increase 20%

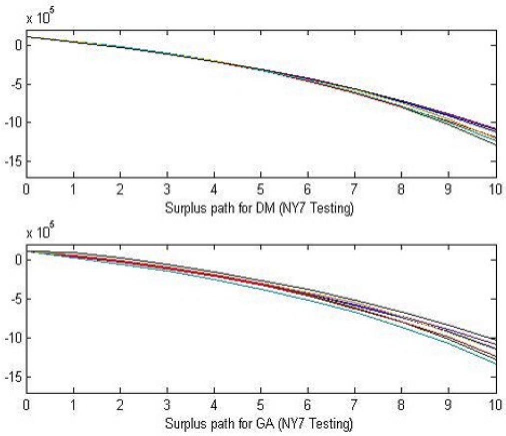

(c) Surplus Path: Mort Increase 30%

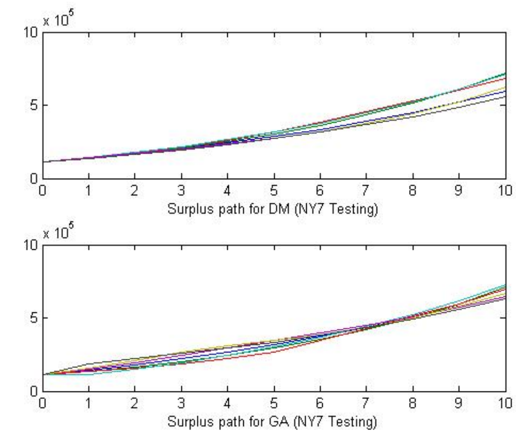

(d) Surplus Path: Mort Decrease 10%

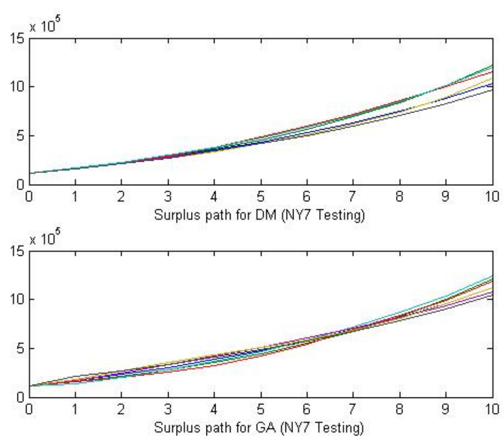

(e) Surplus Path: Mort Decrease 20%

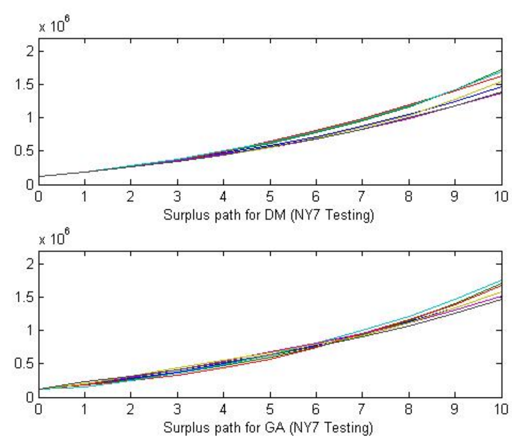

(f) Surplus Path: Mort Decrease 30%

Figure 6.12: Experiment 1 Mortality Testing with New York Seven. In each plot, a shock is applied to the mortality rate and each line represents one of the New York Seven scenarios.

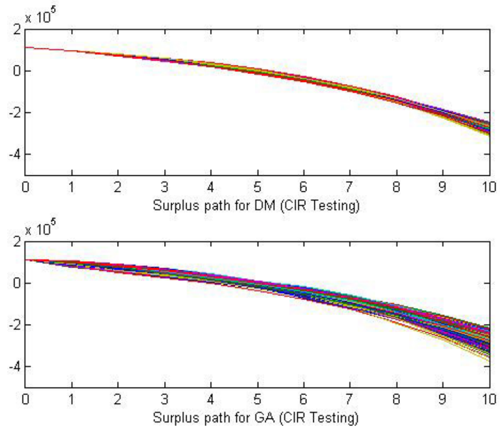

(a) Surplus Path: Mort Increase 10%

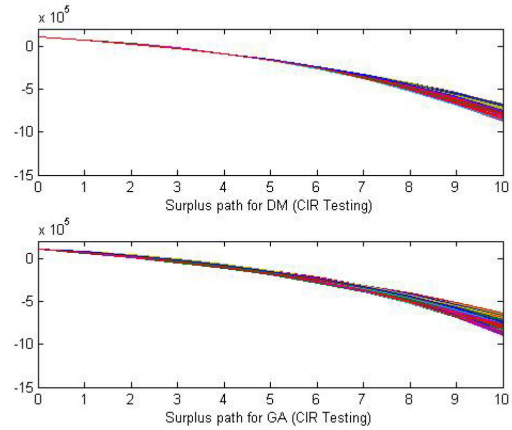

(b) Surplus Path: Mort Increase 20%

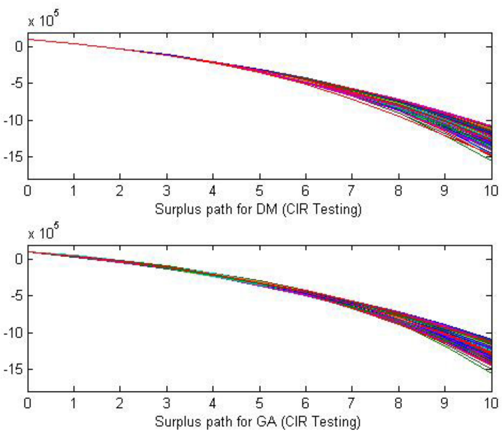

(c) Surplus Path: Mort Increase 30%

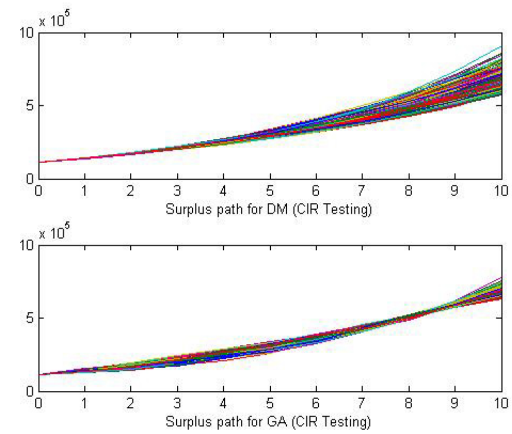

(d) Surplus Path: Mort Decrease 10%

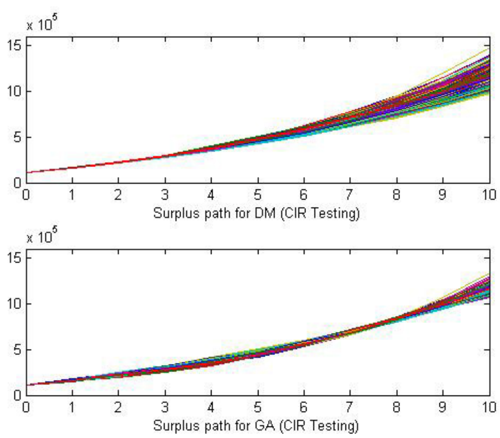

(e) Surplus Path: Mort Decrease 20%

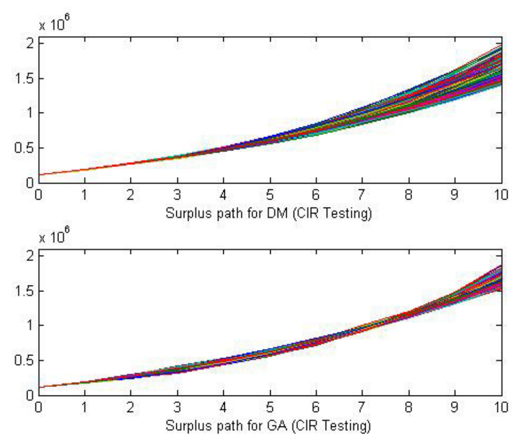

(f) Surplus Path: Mort Decrease 30%

Figure 6.13: Experiment 1 Mortality Testing with Stochastic Interest. In each plot, a shock is applied to the mortality rate and each line represents one of the CIR simulated scenarios.

end of the term of the insurance contract. It is reasonable to assume that for an insurance company, the performance at the end of the term is very important. As we have seen, in order to achieve this, this sacrifices the return on surplus and also the upward quadratic shape for the surplus paths, which indicates a consistent growth in surplus every year, with the latter being an advantage of the duration matching strategy.

There may be a possibility that there exists a solution in training that allows the GA-based strategy to outperform duration matching by all three objectives we specify, instead of being much better in one objective while much worse in another. We hope this strategy is more aligned with the duration matching strategy in terms of the “time” standard deviation objective and we let GA search for solutions conditional on this priority. The strategy obtained from the new experiment may be able to yield more stable results for mortality testing.

We design an experiment called Experiment 4 that aims to outperform duration matching in each of the three objectives. The objective function is given as follows:

- Experiment 4: $ObjV = 0.5 \times \overline{SP(10)} - 0.05 \times Sv(10) - 0.45 \times std$

The results from the training stage are presented in Figure 6.14. Figure 6.14(a) presents the average accumulated surplus across scenarios. It shows that both strategies offer very similar results with GA-based strategy being a bit higher in the end.

In Figure 6.14(b), we present the semi-standard deviation of surplus across scenarios. It is satisfactory that GA-based strategy achieves lower semi-standard deviation in most of the years and also at the end of term.

In Figure 6.14(c), we present the surplus paths which tell us how the surplus of each scenario evolves. The first thing that can be observed is that the shape of both plots look similar. Compared to Experiment 1, this time GA yields a strategy that offers more consistent growth, judging from the more smooth and upward sloping shape for the surplus paths. It also shows that with GA-based strategy, the semi-standard deviation is smaller during the term of the contract, judging from the narrower range the surplus paths appear to be.

Last, in Figure 6.14(d), we show the residuals after we fit the surplus paths to quadratic lines, which also act as a measure for the consistency in growth. Consistent with the result from Figure 6.14(c), the GA-based strategy offers good results in terms of “time standard deviation”. As we can see from this plot, the distribution of the residuals from GA are more to the left of the graph.

Compared to the dollar duration of the duration matched portfolio, the portfolio in Experiment 4 has a slightly higher dollar duration in all years. However, its dollar duration is lower than the portfolio under the strategy in Experiment 1.

From these results in the training stage, we can see that the strategy from Experiment 4 indeed offers satisfactory results among all the objectives of interest. The dollar duration of the bonds portfolio from following this strategy is examined too and it is found that it falls in the middle between the strategy from Experiment 1 and the duration matching strategy. It offers an improvement to all the objectives studied, however for some objectives, the improvement is smaller than from Experiment 1, as a trade-off for improvement in other objectives. To summarize, the strategy from Experiment 4 is more similar to the duration matching strategy than Experiment 1, evidenced from a lower dollar duration than the portfolio under the strategy in Experiment 1.

Before we move on to mortality testing, we first test the Experiment 4 strategy under stochastic CIR scenarios and deterministic New York Seven scenarios, without assuming any mortality volatility. Figure 6.15 displays the surplus paths from these two tests. Under the New York Seven scenarios, the GA-based strategy offers a higher average surplus across the scenarios while also having a large semi-standard deviation. More detailed results show that GA strategy achieves a higher end of term surplus for the first five scenarios out of New York Seven scenarios and has an average surplus 5.31% higher than that of duration matching.

As for the results under CIR scenarios testing, it can be observed that the semi-standard deviation of the GA-based strategy is much lower during the projection years. Also, the average surplus is higher with the GA-based strategy. This again indicates that the interest rate scenarios used in training are sufficient.

Finally, we present the surplus paths from mortality testing for the strategy from Experiment 4. The mortality testing results under New York Seven scenarios are presented in Figure 6.16 and the mortality testing results under CIR scenarios are presented in Figure 6.17.

From Figure 6.16, we can see that a major improvement from Experiment 4 is that the semi-standard deviation from the adverse mortality scenarios is reduced. The trade-off is that in the more favorable mortality scenarios, the semi-standard deviation is increased. If we compare these results to the results from the strategy in Experiment 1, we discover that

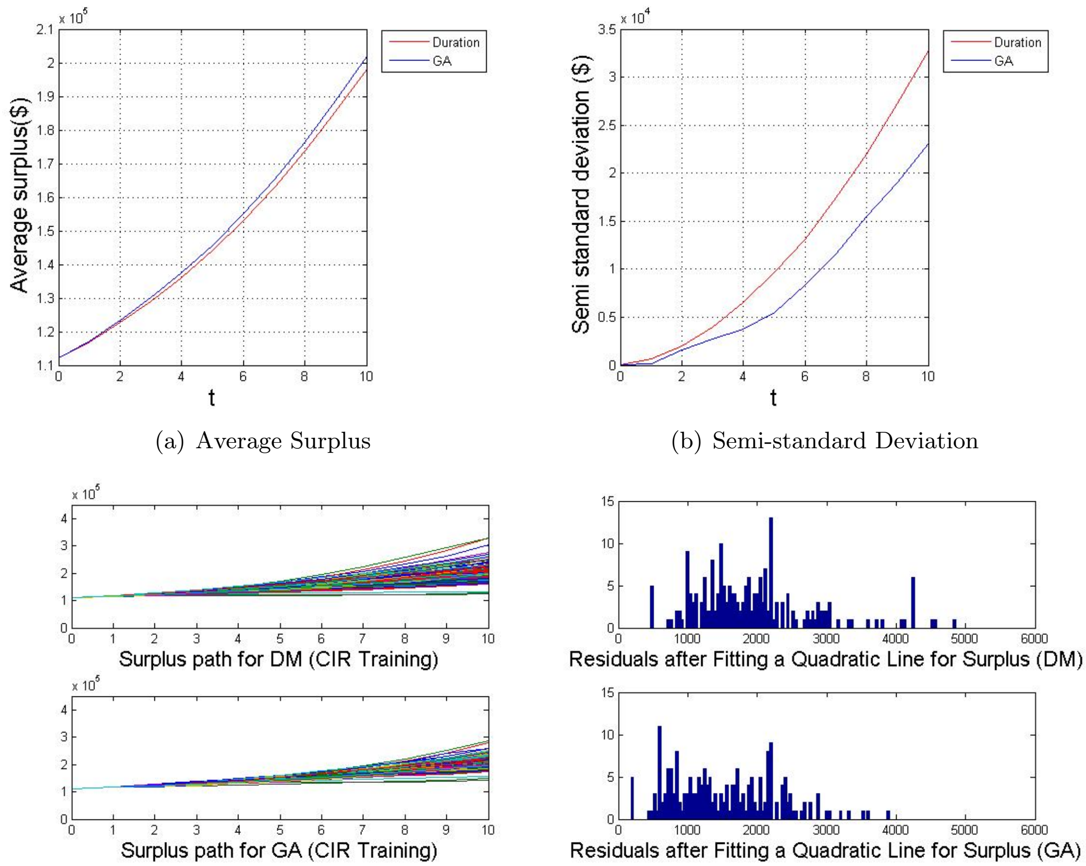

(a) Average Surplus

(b) Semi-standard Deviation

(c) Surplus Path

(d) Residuals after Fitting Surplus Path to Quadratic Line

Figure 6.14: Experiment 4 Training. Aims to outperform duration matching by all three objectives

Experiment 1 may be a more suitable strategy for mortality scenarios that have decreased mortality rates or increased mortality rates up to 10% more than expected, by offering a higher average surplus and lower semi-standard deviation than the results Experiment 4 strategy yields. However when mortality rates are 20% and 30% higher than expected, the strategy from Experiment 4 offers the protection by producing a higher average surplus and lower semi-standard deviation than the results from the Experiment 1 strategy.

From Figure 6.17, again we can see that a major improvement on the semi-standard deviation side compared to Experiment 1. The detailed results indicate that for the favorable mortality scenarios, we are able to achieve a higher average surplus with the new GA-based

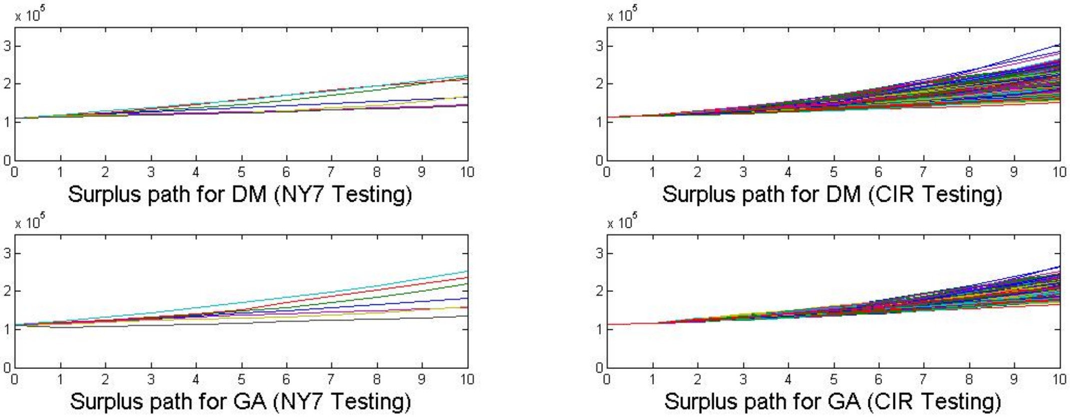

(a) Surplus Path: Experiment 4 New York Seven Testing

(b) Surplus Path: Experiment 4 CIR Testing

Figure 6.15: Experiment 4 New York Seven and Stochastic Interest Testing

strategy while having the semi-standard deviation much reduced. For the adverse mortality scenarios, we have a slightly lower average surplus accompanied by much lower semi-standard deviation. If we compare these results to results from Experiment 1 in Figure 6.13, we can indeed see much improvement on the semi-standard deviation objective for the scenarios with increased mortality. The trade-off is, on the one side, sometimes we are losing a little on surplus and on another side, we have slightly increased semi-standard deviation from the favorable mortality scenarios.

From these results, we show that Experiment 4 allows us to have the trade-off we desire, if we expect the actual mortality experience to be a lot worse than the assumption for mortality rate used in pricing. We lose some of the advantages in the favorable scenarios while gaining more stability in the adverse scenarios. Compared to the duration matching strategy, it offers a trade-off as well, by having a higher average surplus accompanied by a slightly higher semi-standard deviation in the favorable mortality scenarios and having a slightly lower average surplus accompanied by lower semi-standard deviation in the adverse mortality scenarios.

We have experimented with stochastic mortality scenarios. The results show that with a reasonable number of stochastic mortality scenarios, for example, 500, the results from the duration matching strategy are stable. However the results are very volatile for the GA-based strategy if no active re-balancing is used even when we use a large number of mortality

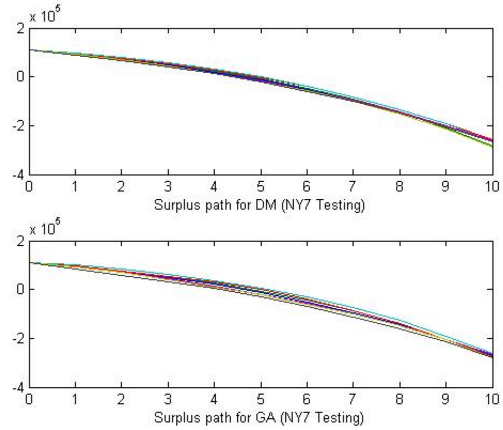

(a) Surplus Path: Mort Increase 10%

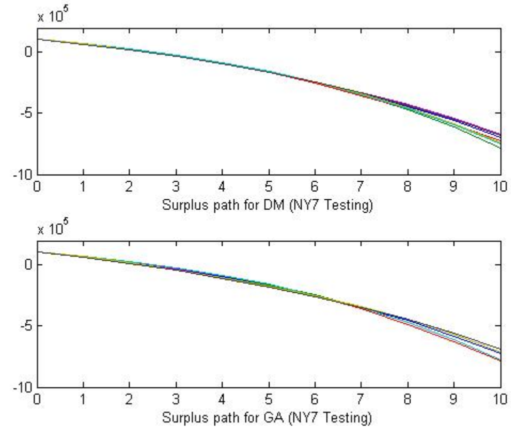

(b) Surplus Path: Mort Increase 20%

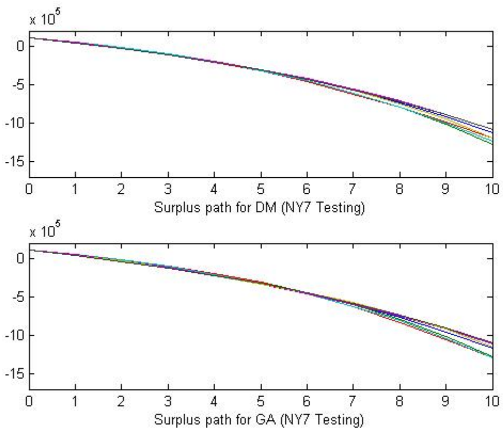

(c) Surplus Path: Mort Increase 30%

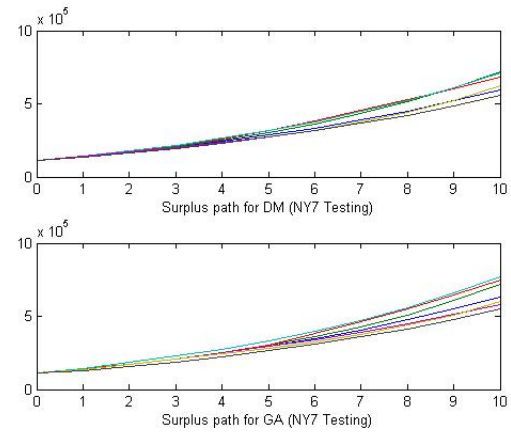

(d) Surplus: Mort Decrease 10%

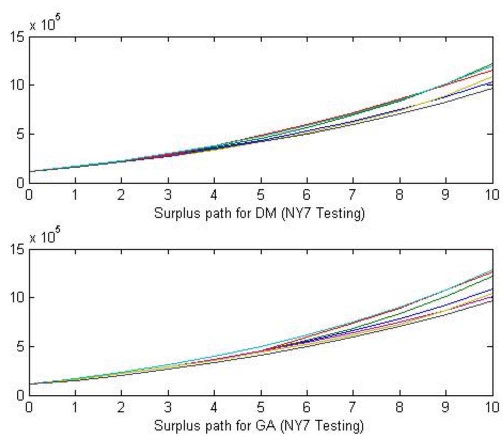

(e) Surplus Path: Mort Decrease 20%

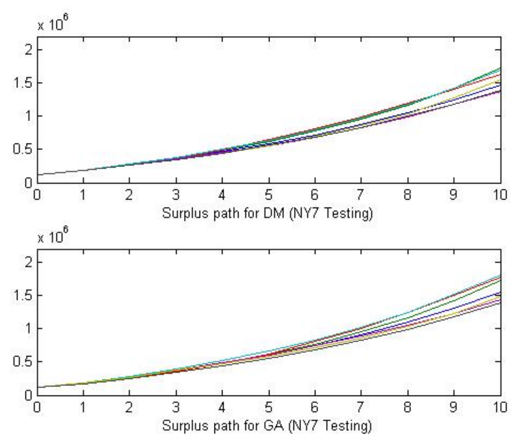

(f) Surplus Path: Mort Decrease 30%

Figure 6.16: Experiment 4 Mortality Testing with New York Seven. In each plot, a shock is applied to the mortality rate and each line represents one of the New York Seven scenarios.

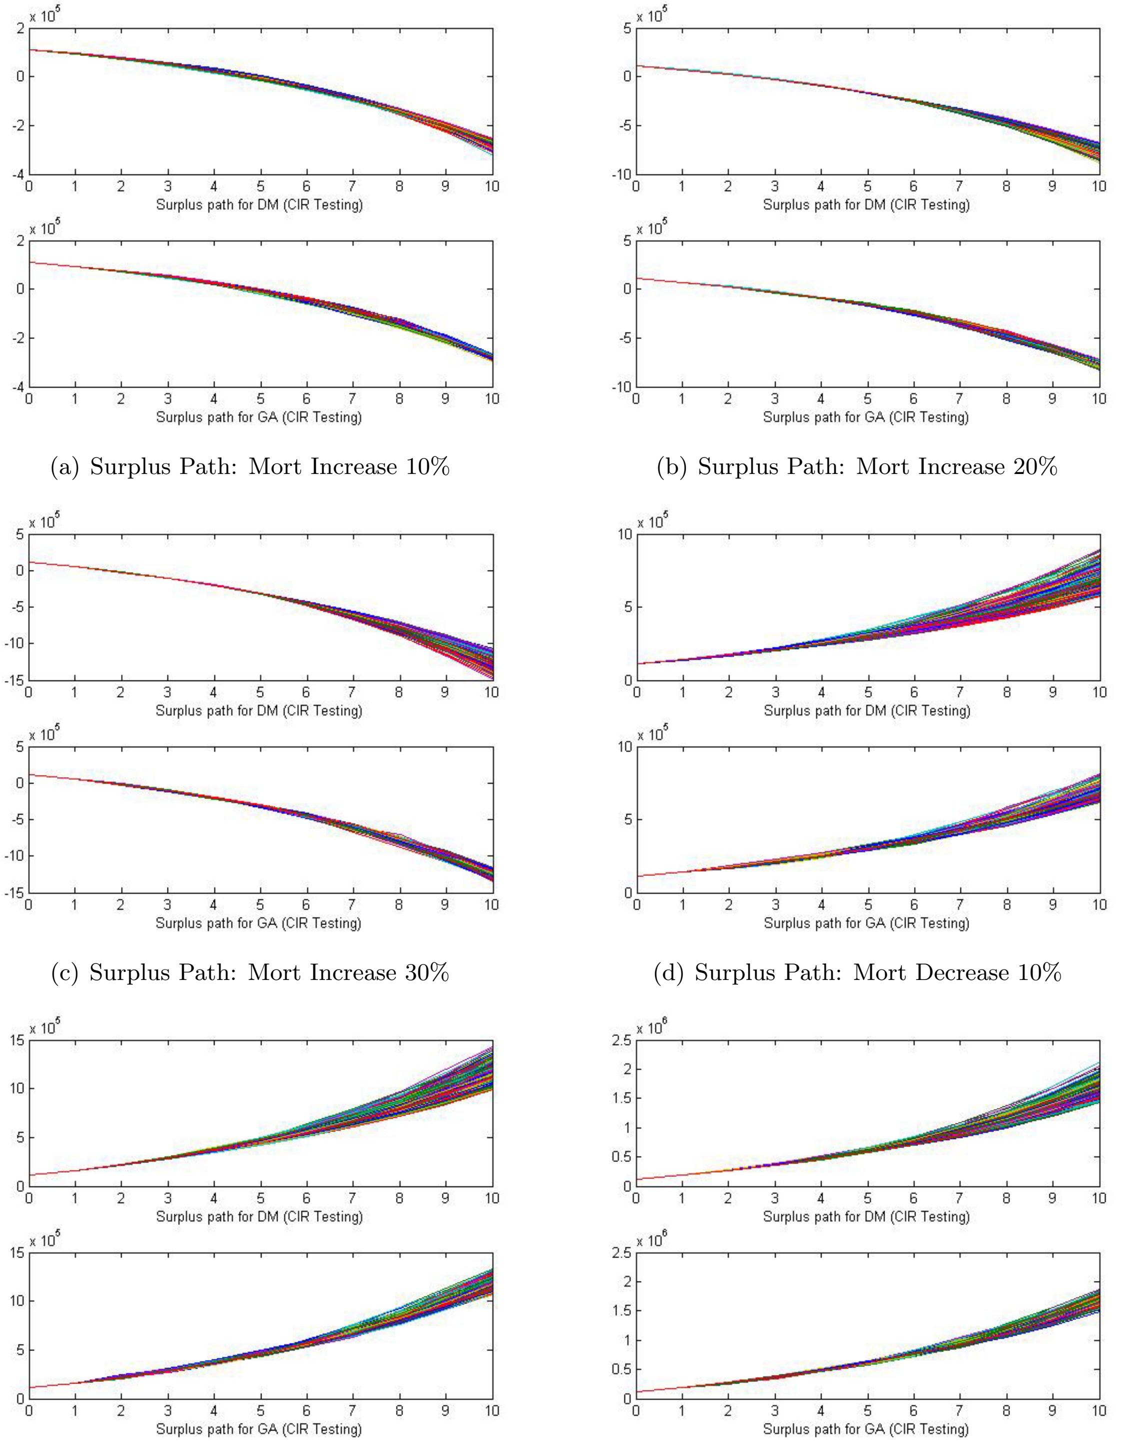

(a) Surplus Path: Mort Increase 10%

(b) Surplus Path: Mort Increase 20%

(c) Surplus Path: Mort Increase 30%

(d) Surplus Path: Mort Decrease 10%

(e) Surplus Path: Mort Decrease 20%

(f) Surplus Path: Mort Decrease 30%

Figure 6.17: Experiment 4 Mortality Testing with Stochastic Interest. In each plot, a shock is applied to the mortality rate and each line represents one of the CIR simulated scenarios.

scenarios. This indicates that the GA-based strategy is very sensitive to the volatility from the mortality experience.

6.4 Analysis

The results from the testing of interest rate scenarios seem to indicate that GA is able to find a strategy that produces better results than the duration matching strategy under the New York Seven scenarios and the 500 stochastic CIR scenarios, at least in some sense. From the testing of a few deterministic mortality scenarios, we observe that the strategy from GA in Experiment 1 is not able to outperform duration matching when the mortality experience is adverse.

In this section, we aim to study what attributes the strategy from GA has that enables it to outperform duration matching in interest rate scenarios testing or lose its advantage in the mortality scenarios testing.

Firstly, we intend to separate the effect of initial allocations and the re-balancing allocations. In duration matching, at the end of each year, assets are bought or sold to enable the dollar duration of the assets and liabilities to match again. On the other hand, for GA, the information for re-balancing is encoded as part of the "genes" of each solution candidate. In the experiments we perform earlier, when we evaluate an individual in GA, we are evaluating both its initial allocation strategy and its re-balancing strategy.

We design an experiment called Experiment 5 that allows us to see the effect from using GA to select the re-balancing strategies alone. In this experiment, we fix the initial weights for individuals in any generation to be the same as the initial weights we obtain from duration matching. In future years, when re-balancing is required, we use GA to select the weights of the re-balancing assets only. Therefore, we have a hybrid of GA and duration matching strategy that starts with a duration matched portfolio and re-balances using GA under an appropriate objective function.

We adopt the assumptions in Experiment 1 to perform Experiment 5 and the results indicate that the accumulated average surplus across all scenarios is very similar to that from Experiment 1. The semi-standard deviation at the end of the projection is about twice of the semi-standard deviation in Experiment 1. The surplus paths look very much like that from duration matching in the first two years, which is expected since they both start from the same initial allocation of bonds. After the first two years, they start to widen.

In the last few years they look much narrower than the surplus paths from the duration matching strategy, although a bit wider than Experiment 1, as evidenced by the variation across scenarios at the end of the term. These results indicate that the re-balancing strategy is effective in reducing the semi-standard deviation for surplus at the end of the term.

It may be helpful to start from examining the surplus paths that the GA-based strategy in Experiment 5 and duration matching strategy each presents to us. The surplus paths of the duration matching strategy look like a fan, very narrow in early years and gradually becoming more and more widely dispersed in later years. For the GA-based strategy, the surplus paths are more widely dispersed in the first few years but in later years the shape of the paths remains like a tube, resulting in a lower level of semi-standard deviation at maturity. One hypothesis is that with GA we are able to take advantage of the information in all projection years while with duration matching we adjust the portfolio year by year only using limited information, i.e., current interest rates.

We also examine the average dollar duration across all scenarios and the duration paths of the portfolio from following the Experiment 5 strategy. The results reveal that at the start, both GA-based and duration matching strategies have the same level of dollar duration. This is expected since the initial allocations are the same. However, it seems that GA chooses to re-balance to a portfolio of much higher dollar duration from the second year, just like it does in Experiment 1.

The bond allocation obtained from Experiment 5 is given in Table 6.3 . Each row represents the percentage of total fund invested in bonds of different maturities for a specific year. We compare this table to Table 4.3 which displays the asset allocations obtained from the duration matching strategy, averaged across scenarios. It is obvious that GA chooses to invest heavier in bonds with longer durations from the second year. For example, at the end of the first year, when both re-balance for the first time, GA-based strategy chooses to invest -30% in 1-year bond and 37% in 10-year bond while the duration matching strategy chooses to invest 8% in 1-year bond and 0% in 10-year bond.

Another method that allows us to gain an insight into GA-based strategy is through checking whether the GA-based or the duration matching strategy matches future cash flows more closely. We compare each asset cash flow at different times and we discover that in Experiment 5, GA-based strategy chooses to invest a lot more heavily in bonds of longer durations instead of matching future liability cash flows closely, like what duration matching does. We take the difference between the future cash flow from net liabilities

Table 6.3: Bond Allocation at the Beginning of Each Year after Adjustments for GA in Experiment 5. The columns are bonds of different maturities. The rows are different periods. The weights are the average across scenarios.

	1	2	3	4	5	6	7	8	9	10
Yr 1	2%	-26%	-18%	-10%	-2%	6%	14%	21%	29%	84%
Yr 2	-30%	-13%	-7%	-1%	4%	10%	16%	21%	63%	37%
Yr 3	-9%	-6%	-1%	4%	9%	13%	18%	53%	20%	0%
Yr 4	-3%	-1%	3%	8%	12%	16%	48%	17%	0%	0%
Yr 5	-2%	3%	7%	11%	15%	45%	20%	0%	0%	0%
Yr 6	1%	7%	11%	15%	45%	20%	0%	0%	0%	0%
Yr 7	6%	12%	16%	46%	20%	0%	0%	0%	0%	0%
Yr 8	11%	17%	50%	22%	0%	0%	0%	0%	0%	0%
Yr 9	16%	57%	27%	0%	0%	0%	0%	0%	0%	0%
Yr 10	64%	36%	0%	0%	0%	0%	0%	0%	0%	0%

(claims less premiums) and assets (bonds that are maturing) for both the GA-based strategy and duration matching strategy, and the results are shown in Table 6.4 and Table 6.5 respectively. In these two tables, each row shows, at the beginning of a particular year in the projection, the difference between the future cash flows on the asset and liability side. It can be observed from the first row that both strategies offer the same results, which is expected given that we let GA-based strategy start at a duration matched position in Experiment 5. From the second year, duration matching has more cash flow from 1-year bond than the liabilities side while GA-based strategy has more cash flow from 10-year bond than the liabilities side. This is consistent with the results in Table 6.3. It indicates that if we fix the first year allocation of the GA-based strategy, it decides to invest very heavily on 10-year bond when re-balancing in the second year in order to achieve the objectives.

All these investigations allow us to arrive at the conclusion that when we specify semi-standard deviation at the end of maturity as a top objective to optimize, GA tends to invest much more heavily in bonds of longer maturities, resulting in a much higher dollar duration of the portfolio than duration matching during the term of the contract. In addition, as we have seen in Experiment 5, although we force the portfolio to start as a duration matched portfolio, GA quickly adjusts to a longer dollar duration portfolio at the next time it has the chance to re-balance the portfolio.

From this, it is not hard to understand why the GA-based strategy from Experiment 1 is able to outperform the duration matching strategy in the training stage. In order to achieve

Table 6.4: Difference Between Future Asset and Liability Cash Flow for GA in Experiment 5. The table shows the difference between assets and liabilities cash flows (in thousands of dollars) in future years (columns) given we are in year t (rows).

	1	2	3	4	5	6	7	8	9	10
Yr 1	130	0	0	0	0	0	0	0	0	-23
Yr 2	-70	0	0	0	0	0	0	0	-23	369
Yr 3	-12	0	0	0	0	0	0	-23	265	0
Yr 4	13	0	0	0	0	0	-23	225	0	0
Yr 5	15	0	0	0	0	-23	223	0	0	0
Yr 6	-10	0	0	0	-23	259	0	0	0	0
Yr 7	-18	0	0	-23	269	0	0	0	0	0
Yr 8	-34	0	-23	287	0	0	0	0	0	0
Yr 9	-45	-23	297	0	0	0	0	0	0	0
Yr 10	-95	323	0	0	0	0	0	0	0	0

the low semi-standard deviation at the end of the contract as we specify in the objective function, GA favors strategies that invest more heavily in longer-term bonds in order to lock the returns. As the yields change during the term of the contract, we may observe higher volatility of surplus from the GA-based strategy. However, if we have a heavier weight on long term bonds, although we may experience volatility from their values during the term, when they mature, say, at the end of the contract, the return is guaranteed.

When a heavy weight is put on the performance at the end of the term, the “safest” and most obvious way would be to invest in 10-year bonds that will have a definite value at maturity. As we can see from the strategy in Experiment 1, GA selects a strategy that invests heavily in 10-year bond either at the start or when it re-balances for the first time. As the portfolio progresses, although GA re-balances so that the dollar duration of the portfolio keeps decreasing, the dollar duration is always higher than a duration matched portfolio.

From the plot of the surplus paths in Experiment 1, as shown in Figure 6.6, we can also see that the volatility of the surplus under different scenarios is higher in the middle. However, approaching the end of the contract, the volatility greatly reduces. As explained earlier, the trade-off the Experiment 1 strategy makes it sacrifice the volatility during the contract but able to get very low semi-standard deviation in the end.

This GA-based strategy from Experiment 1 works well under the interest rate testing, evidenced by the results from Figure 6.10 for New York Seven scenarios testing and Figure

Table 6.5: Difference Between Future Asset and Liability Cash Flow for Duration Matching in Experiment 5. The table shows the difference between assets and liabilities cash flows (in thousands of dollars) in future years (columns) given we are in year t (rows).

	1	2	3	4	5	6	7	8	9	10
Yr 1	130	0	0	0	0	0	0	0	0	-23
Yr 2	138	0	0	0	0	0	0	0	-23	-3
Yr 3	147	0	0	0	0	0	0	-23	-6	0
Yr 4	158	0	0	0	0	0	-23	-10	0	0
Yr 5	171	0	0	0	0	-23	-16	0	0	0
Yr 6	187	0	0	0	-23	-23	0	0	0	0
Yr 7	209	0	0	-23	-35	0	0	0	0	0
Yr 8	239	0	-23	-55	0	0	0	0	0	0
Yr 9	289	-23	-94	0	0	0	0	0	0	0
Yr 10	397	-213	0	0	0	0	0	0	0	0

6.11 for stochastic interest scenarios testing. For stochastic interest rate scenarios, this is largely because the scenarios we use to train and test are from the same interest rate model, therefore we expect the results to be consistent, as long as the CIR model is well-represented by the scenarios used in the training stage. For New York Seven scenarios, one explanation might be that the CIR model we use for training is able to capture the New York Seven scenarios. In addition, recall that the scenarios GA-based strategy does not perform as well as the duration matching strategy are those that involve a sharp rise in interest rates. This may be because when interest rates drop, investing in bonds of longer maturity can lock in the initial higher interest rates through to end of the contract. On the other hand, when interest rates rise, especially when they rise sharply, having longer term bonds is not desirable because even yields from shorter term bonds can exceed the yields that were locked in previously.

However, there is limit as to how much higher dollar duration is allowed in the portfolio than the duration matched portfolio. One way to understand is by the means of objective function. A higher dollar duration implies a higher return, which means a heavier weight on the return in the objective function. As different weights on the return are tested, the results show that there is a cut-off point when the extra return is accompanied by much higher volatility. This indicates that it is important to find the balance between the several objectives, such as return and volatility, and by the use of GA, we are able to find a strategy that balances the several objectives. Another possible explanation is that, if we invest too heavily in longer term bonds, we may need to short sell a lot of shorter term bonds to raise the capital. In the short term, when claims are incurred, with a short position in the short term bonds, we may need to sell longer term bonds to raise the capital to pay off the claims. As a result, as the projection progresses, we lose the longer term bonds and will not be able to secure the returns.

The GA-based strategy in Experiment 1 works well when the mortality scenarios are favorable. However it does not work as well when the mortality experience is adverse. One possible explanation for this may come from the shape of the surplus paths. Recall that the volatility of surplus in the middle of the surplus paths from the GA-strategy can be much higher than that of the duration matching strategy. The disadvantage of this is that when the mortality experience is worse and more claims are incurred than expected during the term of the contract, the higher volatility of surplus in the middle may make the shape of surplus paths to be more widely dispersed.

When we introduce Experiment 4, we could see from the training and testing results that we obtain a more modest strategy than the strategy in Experiment 1. It puts less weight on the end of term semi-standard deviation and ensures the shape of the surplus paths to be more consistent with the surplus paths of duration matching. The benefit from this is that it allows this new strategy to be more similar to duration matching in a way that it responds to adverse mortality scenarios better while at the same time, it loses some advantages from Experiment 1.

6.5 Re-balancing Strategies

From the experiments we have carried out, we are able to observe the trade-offs between different objectives, for example, between return, which is measured by the surplus, risk, which is measured by the semi-standard deviation, and also the "time standard deviation", which measures the consistency of growth of surplus.

Without active re-balancing, the strategy from Experiment 1 works well under interest rate scenarios and mortality scenarios with 10% increase in mortality rate or with a decrease in mortality rates. If one expects that the actual experience is more likely to be one of these scenarios, then the strategy from Experiment 1 may be a desirable choice.

In order to deal with the adverse mortality scenarios, we introduce a more balanced strategy, Experiment 4 strategy, that behaves more similar to the duration matching strategy. The disadvantage of this strategy is that it loses some of the advantages in the more favorable scenarios in order to offer the protection for the adverse mortality scenarios.

Ideally we hope to adopt the Experiment 1 strategy and let it take advantage of the active re-balancing program to reduce the losses from the adverse mortality scenarios.

We examine the re-balancing strategy obtained in Experiment 1 first. We discover that it may not be a suitable strategy under adverse mortality scenarios. For example, the re-balancing weights it selects may be tailored made for the mortality scenario used in training, which assumes no risk in mortality. Note that the surplus keeps growing in training, while the surplus keeps dropping when the actual claims are much higher than expected in testing. This means that we may have extra cash for re-balancing in training, while we may need to short sell bonds to raise money in testing. Therefore the re-balancing strategy from training may suggest buying and short selling the wrong bonds.

On further examination of the dollar duration position of the portfolio from Experiment 1, we confirm this hypothesis. For example, in the +30% mortality scenario, there is a time when the dollar duration of the portfolio goes from positive to negative when the re-balancing weights signal a large short-sell of long term bonds. This may not be an intention of the strategy because when the strategy is trained, mortality and cash flows are as expected and perhaps the re-balancing weight is intended to be a purchase of long-term bonds to obtain a slightly higher duration but not to drastically lower duration and turn it into negative. It can be observed that in later years the dollar duration of the portfolio remains very low and that easily explains why the returns from this testing are not satisfactory. This happens because as explained earlier the signs and patterns of the cash flows change when adverse mortality is introduced. The dollar duration of the portfolio from the Experiment 1 strategy starts at a higher dollar duration than the duration matched portfolio and under expected mortality experience, it should be able to maintain that position. One re-balancing strategy we find that works well to improve the average return on the Experiment 1 strategy is to let GA find the re-balancing strategy that allows the portfolio to maintain a higher dollar duration position.

We also examine the re-balancing strategy and the dollar duration of the portfolio in the +30% mortality scenario for the duration matching strategy and the GA-based strategy from Experiment 4. By comparison, the duration matching strategy chooses to short sell a lot more 1-year bonds in order to buy longer term bonds in the testing than the strategy from Experiment 4. This may be the reason why the duration matching strategy achieves a higher average return than the strategy from Experiment 4 yields. The dollar duration of the two portfolios confirms this belief, showing that the re-balancing strategy from Experiment 4 lets the portfolio have a lower dollar duration than the duration matched portfolio during the projection. One possible route to improve the performance of Experiment 4 in testing may be to let GA actively select the strategy that raises the dollar duration of the portfolio by a certain level.

The advantage of the duration matching strategy is that it responds to the adverse mortality experience well. Therefore, we make a re-balancing program that combines the initial allocation from Experiment 1 and incorporates an active re-balancing program that may have objectives that are very different from Experiment 1. For example, if we expect the future mortality to be much worse than expected, then it may not be desirable to put too much weight on the semi-standard deviation at the end of the contract.

One way to specify the objective function in the re-balancing program is to select a re-balancing strategy that allows the semi-standard deviation of surplus at the end of next year to be optimized. This is similar to the premise of duration matching where it looks at current rates and tries to minimize the year-to-year change in surplus. However, in the re-balancing program itself, there can be multiple objectives as well. Other than year-to-year tracking of semi-standard deviation, another way is to look at a longer term, say the semi-standard deviation at the end of the contract. The latter is what Experiment 1 puts the heaviest emphasis on.

The specification on the objective function reflects the user's preference. For example, if the user would like to mimic the behavior of the duration matching strategy, then he may need to put a heavy weight on the volatility of surplus in next year; if the user would like to deviate from the duration matching strategy, then he may need to put a heavy weight on the volatility at the end of the term. The decision also requires the user's judgment on the future trend for mortality experience, as different expectations of future experience may result in different preferences of objectives.

We add this simple active re-balancing program to the initial allocation of the GA-based strategy from Experiment 1 and test it under the New York Seven scenarios for the mortality scenario that has 30% higher mortality rates than expected. Note that to be consistent with the duration matching strategy, we use the current interest rate available as the starting short rate in the CIR model, and we update the future expected premiums and claims after taking into account of the actual claims incurred. However, we do not adjust the expected mortality rate in the future, which means for the mortality experience in the future, we still use the same assumption as originally used in pricing.

The difference in the objective function for the active re-balancing program from the objective function used in training is that we do not only need to maintain the balance between return and volatility, but we also need to maintain a balance between short term volatility and long term volatility of the surplus across scenarios. The need for the latter objective is in response to the adverse mortality experience.

The results for before and after the implementation of active re-balancing are presented in Table 6.6. It is shown that the results from NY1, NY4, NY5 NY6 and NY7 are improved while for the other three scenarios, the results are worse. On average, the results improve as well and are higher than the results from the duration matching strategy. Further examining the active re-balancing strategies, we discover that they avoid some of the mistakes the passive re-balancing strategies make, however, they are are not optimal. One obvious weakness is that the interest rate model used to simulate the interest rate scenarios in the re-balancing program. The effectiveness of the re-balancing program relies very much on the interest rate model and the current model still use the old parameters regardless of the new experience available. This could lead the GA to search for strategies optimized for the simulated interest rate paths but not for the actual experience.

Table 6.6: Experiment 1 New York Seven Testing Active Re-balancing. The first row is the end of term surplus (in millions of dollars) achieved through the Experiment 1 strategy under the New York Seven scenarios without using active re-balancing. The second row are similar results but with the use of active re-balancing. Mortality rates are 30% higher than expected.

	NY1	NY2	NY3	NY4	NY5	NY6	NY7
Active Re-bal	-1.11	-1.28	-1.27	-1.28	-1.06	-1.11	-1.02
Passive Re-bal	-1.14	-1.28	-1.23	-1.33	-1.09	-1.15	-1.03
DM	-1.12	-1.28	-1.19	-1.23	-1.08	-1.20	-1.09

Although the current re-balancing program can help improve the performance by the specification of the suitable objectives, it is worth noting that in the re-balancing stage, GA needs additional information as input compared to the duration matching strategy. For example, it relies on a good interest rate model that can represent the future interest rate well. Ideally the interest rate model is refitted each time new interest rates become available.

Also, the objective function in the active re-balancing program is flexible and requires more tuning before a suitable objective function can be found.

Chapter 7

Conclusion

In this project, we build a framework with GA to handle an ALM problem for a term life insurance product. We have demonstrated that GA is definitely able to find a solution to this type of problem effectively, evidenced by the reasonable running time for each program investigated and the solution, which consists of 19 variables, representing both the initial allocation and the re-balancing strategy, and is robust in both stochastic and deterministic testing for interest rate and deterministic testing for mortality.

To be more specific, from Experiment 1, we are able to obtain a strategy that offers better results than the duration matching strategy in two of the three objectives specified in the training stage of the experiments. This strategy achieves excellent results in New York Seven and CIR interest rate testing. It is able to sacrifice the return in some scenarios while boosting returns in the more adverse interest rate scenarios. The average return is higher than that of duration matching strategy and the semi-standard deviation across scenarios is much lower. This strategy performs well with favorable mortality scenarios or a mortality scenario with 10% increase to mortality rates. However, it may not be optimal when we introduce adverse mortality scenarios. The major reason for this is that the static re-balancing strategy may be tailored for the mortality scenario used in training and may not give the best results when mortality experience is a lot worse than expected.

We introduce a strategy from Experiment 4 that offers better results than the duration matching strategy in all three objectives in the training stage of the experiments, indicating that it is a more balanced strategy that the strategy from Experiment 1. The results from testing show that it sacrifices some of the advantages from Experiment 1, while offering a better protection from the adverse mortality scenarios.

The choice of the strategy to be used depends largely on one's risk preference and the forecast for the future interest rate and mortality experience. If one expects the mortality rates to only drop or increase slightly, say by 10%, with high likelihood, then the strategy from Experiment 1 may be a good choice. If one wants to be on the safe side from the adverse mortality experience and expects the mortality rates to be a lot higher than expected with high likelihood, then duration matching strategy should be adopted. If one would like to benefit from the favorable scenarios and would like to have some protection for the adverse scenarios too, then the strategy from Experiment 4, which falls in the middle between the strategies from Experiment 1 and duration matching, should be considered. Note that this is assuming no active re-balancing in GA strategies.

It is expected that with an active re-balancing program, the performance of GA strategies could be further improved. For example, for the Experiment 1 strategy, we have observed that the re-balancing strategies obtained in the training stage works well for the expected mortality scenarios. However, under the adverse mortality scenarios, when the signs and patterns of the cash flows change, the static re-balancing strategies could lead us to make wrong decisions. We make a simple active re-balancing program that takes into account of some of information available, such as the current interest rate and the actual claims occurred, and this has improved the performance of the GA-based strategy from Experiment 1 in the adverse mortality scenarios. The weaknesses of our re-balancing program is that the interest rate model is not refitted and expected mortality rates are not updated based on the new information that becomes available. This could lead us to suboptimal decisions and further improvement to the re-balancing program is recommended.

One discovery we make from the experiments is that for a solution to perform well in all the testings we conduct in the study, it cannot have too much weight on any one of the objectives, but rather, a balance needs to be maintained. The strategy we find that is able to perform well in all the testing is the strategy that is able to outperform duration matching in the training by all the specified objectives, while having a slightly higher dollar duration than the duration matched portfolio. It shows that the duration matching is indeed a method that can balance different objectives well.

From Experiment 5 we also discover that, the initial allocation and the re-balancing strategy are both very important, since only when they both work together, are we able to achieve the lowest semi-standard deviation, indicating that they both contribute to the reduction in the semi-standard deviation.

The merit of GA we recognize through our research is that GA is a very flexible tool that allows us to find strategies that meets the needs of different risk appetites. Duration matching method offers us one strategy. It is shown that GA is able to find a strategy that mimics duration matching strategy closely. In addition, with GA we are able to achieve strategies that are different from duration matching but also have their own merits, such as, being able to outperform duration matching in some sense.

Another advantage of using GA is that it may allow the users to directly monitor the real objectives, such as return and risk. Note that for ALM practice, duration matching is an effective tool however it is not the end goal. After all, return and risk are the true goals.

GA allows the user the freedom to specify the objective function and will quickly adjust the solutions to tailor to the specific objective function. This flexibility is a double-edged sword. On the down side, this means it lacks a good theory like the duration matching theory behind it, and relies much on the user's judgment on what objectives are important. The specific weights to put on the objectives require tuning too. Ideally, GA is used as a component of an ALM program, and other theories or knowledge are in place to assist in making the framework more comprehensive.

Chapter 8

Future Work

There are a few areas where further research could be explored. First of all, more elements in the natural evolution can be introduced into the GA used in this project to make the GA program more complicated. For example, as introduced in Chapter 2, the mutation rate can be made more adaptive, for example, declining over the generations, instead of being a fixed value, to increase the speed of convergence. Also, a model for the adaptive crossover and mutation rates can be introduced into the GA program, based on the concept that each pair of chromosomes should receive a crossover or mutation probability based on its individual performance, see Lin et al. (2003).

Another possible area to improve the performance of the current GA is to use real-valued or continuous representation for the solutions, which means that each variable is not represented by a few digits of binary numbers but simply by a real number. For one thing, this is probably the most natural way of representation for this problem. For another, there are many variables in this problem (19 variables) and the implementation of the continuous GA can allow us to save running time and storage. A good reference on how to implement the real-valued GA is Adewuya (1996).

Further research is also recommended to make GA deal with adverse mortality experience better. Ideally, the interest rate model or mortality model could be refitted as the experience becomes available. Or, other information available can be used as an indicator for the interest rate or mortality outlook. Also, for the re-balancing, transactions costs could also be incorporated.

Last but not least, more objectives could be explored. For example, the deviation from the duration matched portfolio could be specified as an objective and this may allow us to study specifically the impact on return and risk measures due to these deviations.

Bibliography

Adewuya, A. A. (1996). *New methods in genetic search with real-valued chromosomes*. PhD thesis, Massachusetts Institute of Technology.

Altenberg, L. (1995). The schema theorem and prices theorem. *Foundations of genetic algorithms*, 3:23–49.

American Academy of Actuaries (1995). Life practice note 1995-4: Interest rate models.

Aranha, C. (2007). *Portfolio management with cost model using multi objective genetic algorithms*. PhD thesis, The University of Tokyo.

Aranha, C. and Iba, H. (2008). A tree-based ga representation for the portfolio optimization problem. In *Proceedings of the 10th annual conference on Genetic and evolutionary computation*, pages 873–880. ACM.

Bäck, T. (1993). Optimal mutation rates in genetic search. In *Proceedings of the 5th International Conference on Genetic Algorithms*, pages 2–8, San Francisco, CA. Morgan Kaufmann Publishers Inc.

Baker, J. E. (1987). Reducing bias and inefficiency in the selection algorithm. *Proceedings of the Second International Conference on Genetic Algorithms and their Application*, pages 14–21.

Beyer, H. (1997). An alternative explanation for the manner in which genetic algorithms operate. *BioSystems*, 41(1):1–15.

Bierwag, G., Kaufman, G. G., and Toevs, A. (1983). Immunization strategies for funding multiple liabilities. *Journal of Financial and Quantitative Analysis*, 18(1):113–123.

Chan, M.-C., Wong, C.-C., Cheung, B. K., and Tang, G. Y. (2002). Genetic algorithms in multi-stage portfolio optimization system. In *proceedings of the eighth international conference of the Society for Computational Economics, Computing in Economics and Finance, Aix-en-Provence, France*.

Chang, T.-J., Yang, S.-C., and Chang, K.-J. (2009). Portfolio optimization problems in different risk measures using genetic algorithm. *Expert Systems with Applications*, 36(7):10529–10537.

Cooper, E., André, A., Gabovich, S., and Pomerantz, A. (2010). Revisiting the role of insurance company alm within a risk management framework. *Goldman Sachs Asset Management*.

Fabozzi, F. J. et al. (2007). *Fixed income analysis*. Wiley, Hoboken, NJ.

Fogarty, T. C. (1989). Varying the probability of mutation in the genetic algorithm. In *Proceedings of the third international conference on Genetic algorithms*, pages 104–109. Morgan Kaufmann Publishers Inc.

Fooladi, I. J. and Roberts, G. S. (2000). Risk management with duration analysis. *Managerial Finance*, 26(3):18–28.

Forrest, S. and Mitchell, M. (1991). *Towards a stronger building-blocks hypothesis: Effects of relative building-block fitness on GA performance*. Morgan Kaufmann Publishers Inc., San Mateo, CA.

Gaffney, J., Pearce, C., and Green, D. (2010). Binary versus real coding for genetic algorithms: A false dichotomy? *ANZIAM Journal*, 51:C347–C359.

Goldberg, D. (1989). *Genetic algorithms in search, optimization, and machine learning*. Addison-Wesley Professional, Reading, MA.

Haupt, R. L. and Haupt, S. E. (2004). *Practical Genetic Algorithms*. John Wiley & Sons, Inc., New Jersey.

Herrera, F., Lozano, M., and Verdegay, J. L. (1998). Tackling real-coded genetic algorithms: Operators and tools for behavioural analysis. *Artificial intelligence review*, 12(4):265–319.

Holland, J. (1975). *Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence*. University of Michigan Press, Ann Arbor, MI.

Jackson, A. (1997). Genetic algorithms for use in financial problems. IAA AFIR Colloquium Cairns.

Kan, R. (1992). Shape of the yield curve under cir single factor model: A note. *Working Paper, University of Chicago*.

Lin, W.-Y., Lee, W.-Y., and Hong, T.-P. (2003). Adapting crossover and mutation rates in genetic algorithms. *Journal of Information Science and Engineering*, 19(5):889–904.

Macaulay, F. R. (1938). Some theoretical problems suggested by the movements of interest rates, bond yields and stock prices in the United States since 1856. *NBER Books*.

Mathis, R. H. C. (1993). Duration of life insurance liabilities and asset liability management. IAA AFIR Colloquium Rome.

Ramsey, C. A. (1990). New york regulation 126. In *1990 Valuation Actuary Symposium Proceedings*. Society of Actuaries.

Reitano, R. R. (1991). Multivariate duration analysis. *Transactions of the Society of Actuaries*, 43:335–376.

Reynolds, C. W. and Wang, D. W. (2007). Interest rate hedging on traditional life and health. Society of Actuaries.

Rowe, J., Whitley, D., Barbulescu, L., and Watson, J.-P. (2004). Properties of gray and binary representations. *Evolutionary Computation*, 12(1):47–76.

Salomon, R. (1998). Short notes on the schema theorem and the building block hypothesis in genetic algorithms. In *Evolutionary Programming VII*, pages 113–122.

Sharpe, W. F. (1994). The sharpe ratio. *Journal of portfolio management*, 21:49–58.

Soam, V., Palafox, L., and Iba, H. (2012). Multi-objective portfolio optimization and rebalancing using genetic algorithms with local search. In *Evolutionary Computation (CEC), 2012 IEEE Congress*, pages 1–7. IEEE.

Tan, R. (1997). Seeking the profitability-risk-competitiveness frontier using a genetic algorithm. *Journal of Actuarial Practice*, 5(1):49.

Zeytun, S. and Gupta, A. (2007). *A Comparative Study of the Vasicek and the CIR Model of the Short Rate*. ITWM, Kaiserslautern, Germany.