HEC MONTRÉAL
École affiliée à l'Université de Montréal

Optimal Asset Allocation and Funding Strategy for a Defined Benefit Pension Fund: An LDI Approach

by
Pierre-Luc Vachon

Thesis presented for the obtention of
a Master of Science (M.Sc.) in Administration
Finance

August 2014

HEC MONTRĒAL

Retrait d'une ou des pages pouvant contenir des renseignements personnels

SOMMAIRE

Ce mémoire applique l’approche d’investissement guidé par le passif pour le plus grand fonds de pension public à prestations déterminées au Canada : le fonds de pension de la fonction publique du Canada. L’objectif de ce mémoire est de trouver la répartition optimale de l'actif et la stratégie optimale de financement afin de minimiser le risque de solvabilité mesuré par la probabilité d'insolvabilité et par la variance de l'excédent à l’échéance. L'insolvabilité doit en effet être évitée dans les fonds de pension à prestations déterminées puisque le répondant du régime est responsable de tout manquement face à ses obligataires. Tout d’abord, nous simulons le passif et l’actif du fonds en fonction de différents facteurs de risque sur une période de 20 ans grâce à une simulation de Monte Carlo. Nous évaluons ensuite le risque de modèle et l'exposition du fonds à chaque facteur de risque. De là, la répartition optimale de l'actif est trouvée et enfin, la stratégie de financement optimale pour combler le déficit est formulée.

Mots clés: gestion actif-passif, investissement guidé par le passif, gestion des risques, prestations déterminées, fonds de pension, allocation d’actifs optimale, taux d’intérêt, taux d’inflation, croissance des salaires, taux de mortalité.

ABSTRACT

This thesis applies the liability-driven investment (LDI) approach to the largest public defined benefit pension fund in Canada: the pension fund of the Public Service of Canada. The objective of this thesis is to find the optimal asset allocation and funding strategy to minimize the solvency risk measured by the probability of insolvency and the variance of the terminal surplus. Insolvency must in fact be avoided in defined benefit pension funds as the plan's sponsor is responsible for any shortcomings. We start by simulating the liability and the asset of the fund as a function of different risk factors over 20 years through a Monte Carlo simulation. We then evaluate the model risk and the exposure of the fund to each risk factor. From there, the optimal asset allocation is found and finally, the optimal funding strategy to fill the funding gap is formulated.

Key words: asset-liability management, liability-driven investment, risk management, defined benefit, pension fund, optimal asset allocation, interest rate, inflation rate, earnings growth, mortality rate.

TABLE OF CONTENT

LIST OF TABLES

LIST OF FIGURES

REMERCIEMENTS

Je tiens premièrement à remercier mon directeur de mémoire, Professeur Martin Boyer de m’avoir suggéré un sujet si intéressant et d’avoir supervisé mon travail. Je remercie également les membres du jury qui ont évalué mon travail.

J’aimerais finalement remercier ma famille et mes amis pour leur support et leurs encouragements tout au long du processus d’élaboration de mon modèle et de la rédaction de mon mémoire.

1. INTRODUCTION

In the last decade, pension funds have drawn a lot of attention as their funding status eroded over time. In fact, market conditions devastated numerous pension funds around North America and Europe. The declining interest rates coupled with the negative to low equity returns simultaneously magnified liabilities and shrunk assets, which resulted in large funding deficits (Martellini, 2006). These deficits are particularly important for defined benefit pension plans since, under these agreements, the fund's sponsor has a contractual obligation to pay its pensioners the promised pensions and is therefore responsible for any shortcomings. "In 2003, the defined benefit pension plans for the companies included in the S&P 500 faced a cumulative deficit estimated at $225 billion […] [while they] were enjoying a total surplus of $239 billion at the end of 1999" (Martellini and Milhau, 2008: 2).

This storm uncovered flaws in pension fund management which was mainly focused on asset return at the time and prompted regulators to develop new accounting rules and regulations. Consequently, from 2006 to 2011, the *Pension Protection Act* was implemented in the United States. Among other things, it requires higher contributions and imposes penalties for actuarial deficits (Leahy, 2011). In Canada, the 1998 International Accounting Standard (IAS) 19 has been amended between 2008 and 2011. It now creates a strong incentive for fund managers to be more prudent in their asset allocation and invest in less volatile asset classes (Andonov *et al.*, 2013). This situation encouraged pension fund managers to use other approaches for investing and managing risk. An approach that gained a lot of popularity is asset-liability management (ALM), which integrates liabilities in the investment decisions. One of the techniques that are widely used under ALM is liability-driven investment (LDI), which is interested in building a liability-hedging portfolio.

In this thesis, we apply the LDI approach to the largest public defined benefit pension fund in Canada: the pension fund of the Public Service of Canada. This fund is segmented in two accounts: the Public Service Pension Fund and the former Superannuation Account. The Public Service Pension Fund was introduced by the Public Sector Pension Investment Board (PSPIB), who was formed to independently invest and manage the contribution payments made on or after April 2000. Both accounts are considered jointly in this thesis.

Our objective is to find the optimal asset allocation and funding strategy to minimize the solvency risk given the fund's liability and its sensitivity to risk factors. This is achieved by modeling the fund's liability and the asset as a function of stochastic risk factors through Monte Carlo simulation over a period of 20 years. It is anticipated that the optimal portfolio will allocate the majority of funds in bonds and that the funding gap will be filled by injecting capital, increasing contribution rates or delaying the age of retirement. Although we apply this study to a particular fund, our results can be generalised to other defined benefit pension funds since most of them share a pension structure and demographics similar to the Public Service Pension Fund. This thesis will hopefully contribute to the knowledge of the risks that affect pension funds and the techniques used to manage them.

This study is relevant because a lot of funds are currently in a precarious financial situation, which makes them likely to fall into insolvency. Furthermore, the financial stability of pension funds is very important to different stakeholders. In fact, these funds protect the standard of living of workers by promising to pay pensions when they retire regardless of market conditions. Also, as we mentioned above, in defined benefit plans, the sponsor is responsible for any shortcomings and must inject new capital to cover deficits. In the case of the Public Service Pension Fund, the sponsor is the government and any additional payments are indirectly borne by the taxpayers. Finally, with their large pool of assets, pension funds are one of the largest players in the financial markets. Their stability is therefore vital not only for the millions of workers covered by these plans and their sponsors but also for the liquidity and stability of financial markets.

The next section reviews the literature on ALM, LDI and various portfolio optimization approaches with liability constraints. In section 3, different theoretical models used to simulate the risk factors are presented. In section 4, the methodology and the different assumptions made by the model are described. In section 5, the results of this study are presented and discussed. We start by assessing the validity of our model. We then perform a valuation of the fund and a sensitivity analysis. Finally we find the optimal asset allocation and funding strategy.

2. REVIEW OF LITERATURE

2.1 ASSET-LIABILITY MANAGEMENT

In the last few years, defined benefit pension funds are under great pressure. On the one side, market conditions crush their surplus and generate large deficits in numerous funds and on the other side, governing bodies impose tighter regulations. Three main stakeholders are affected by this issue. The contributors, who desire to maintain an affordable plan through low contribution rates, the pensioners, who are entitled to pension payments and the pension sponsors, who want to avoid having to inject new capital in case of a deficit. This substantial pressure prompted fund managers to modify their risk management and investment style. A technique that gained a lot of popularity is asset-liability management (ALM) that was already being used in the banking industry since the 1970’s (Choudhry, 2007). Rather than taking an asset-only approach, ALM compares asset to liability and makes sure that the fund remains solvent through time. Selecting a portfolio from the asset-only efficient frontier may result in a non-optimal choice when taking into consideration the surplus of the fund. In fact, this surplus is sensitive to other market factors such as interest rates and inflation (Meder and Staub, 2006).

A variety of strategies are used by managers who apply ALM. In this section, three ALM strategies are discussed: cash-flow matching, duration matching and LDI. Cash-flow matching ensures that every cash outflow (such as a pension payment) is matched with a cash inflow of the same amount (such as the principal repayment on a bond), which creates a perfect hedge. If the cash inflows are coming from government bonds, they are relatively easy to predict. The difficulty however is to predict the timing and the magnitude of the cash outflows to match them with the inflows. Because pensions are indexed, the inflation rate affects the magnitude of the payments. The horizon of the payments is also uncertain because of mortality rates. Leaving mortality rates aside, when a member retires, the only uncertainty pertaining to its pension payments is the indexation (inflation risk). If the portfolio of assets comprises of inflation-linked bonds, the inflows could in theory perfectly match the payments because if inflation would deviate from expectations, both asset and liability would be equally affected. This is however not possible in practice since inflation-linked bonds are available only for some maturities, which would leave some payments uncovered. Also, a lot of those inflation linked bonds pay coupons. The reinvestment of these coupons creates a risk because their return on the markets is uncertain (Martellini, 2006). For these reasons, cash-flow matching is not applied in this thesis.

Because of the impossibility to perfectly match the cash inflows to the outflows, other techniques can be used such as immunization, which hedges the portfolio against interest rate risk. This is relevant since “the pension liability has bond-like features including an inverse relationship to interest rate movements” (Leahy, 2011: 3). Duration matching is a good way to immunize a portfolio against parallel shifts in the yield curve. In fact, duration measures the sensitivity to changes in interest rates. According to Moore (2010), the life annuities offered by defined benefit plans create liabilities with an average duration of 12 years while the typical 40% bond, 60% equity portfolio results in a duration of about 2 years. As a result, liabilities are six times more sensitive than assets to changes in interest rates. If interest rates decrease, liabilities would increase six times as much as assets would. If the magnitude of the interest rate drop is significant, it can rapidly plunge a plan’s funding status into deficit. Whether the interest rates rise or fall and the funding status is improved or worsened, the duration gap creates great volatility in the surplus of the plan. This is why it is important to align the duration of the asset with the duration of the liability. According to Moore (2010), the duration gap is “typically the largest risk to a plan’s funding status”. Duration matching is therefore a good technique to reduce the funding status risk. For this reason, it is applied in this thesis.

An adjustment for convexity can be added to duration matching to immunize more precisely the portfolio to changes in interest rates. Other immunization techniques consist in hedging non-parallel shifts in the yield curve. Regardless, an important drawback of immunization is that it focuses on interest rate risk but disregards other risks such as inflation rate risk and longevity risk. Another drawback is that while hedging the portfolio against a change in interest rates, performance is sacrificed. In fact, immunization often involves investing in different treasury bonds, which are expected to yield lower returns than stocks over time. In order to increase the expected return on the portfolio of assets, other asset classes such as stocks and corporate bonds should be added. It is important to keep a relatively high expected return on the portfolio to keep contribution rates at reasonable levels and still be solvent. (Martellini, 2006).

To reach this combined goal of hedging the portfolio while maintaining a reasonable return, a strategy called liability-driven investment (LDI) is used. Under this strategy, two portfolios are usually formed. The first is the liability-hedging portfolio and the second is the performance portfolio. Because this thesis is centered on LDI, the rest of the review of literature elaborates on this approach.

2.2 LIABILITY-DRIVEN INVESTMENT

In the early 2000's in the UK, new accounting rules and regulations that required to mark-to-market assets and liabilities incited pension plans to adopt LDI and shift assets from equity to long duration bonds as the funding status risk became an important factor. The UK was the precursor to this shift in the U.S. (Moore, 2010). In their case, some funds in the U.S. started implementing LDI after the dot-com bubble burst devastated their funding status (Bragt, 2011). The second wave of adverse market conditions of 2008 encouraged a lot of funds to adopt LDI.

Under the LDI approach, managers must optimize their portfolio while making sure that assets cover liabilities in most possible scenarios. In other words, they must have a positive surplus with a small variance. To reduce the surplus variance, LDI partially closes the duration gap. To do so, the funds are shifted from stocks to bonds and bonds with a longer duration are favored (Leahy, 2011). As we can see, LDI often requires departing from the classical Markowitz framework by which expected return on asset is maximized for each unit of risk measured by the standard deviation of returns (Martellini, 2006). Defined benefit pension funds are faced with a dilemma. They have the incentive to invest in asset classes with low correlation to the liability but with high expected returns like equity in an attempt to decrease contribution rates. However, they also have an incentive to invest in asset classes highly correlated to the liability in order to decrease the volatility of the actuarial surplus (Martellini and Milhau, 2008).

Fund managers who use LDI usually apply the following procedure that leads to the investment decision. First, they compute the value of the liability and the asset. Then they find the objective of the fund such as minimizing the funding gap or funding ratio, minimizing the variance of terminal surplus or minimizing the probability of insolvency. In most cases, the objective is to close the funding gap and reach a surplus that will act as a buffer in case of adverse market conditions or unpredicted changes in the fund's demographics. This buffer inevitably reduces the probability of insolvency. From there, the manager must find the optimal asset allocation. This is done by identifying the factors to which the asset and the liability are sensitive and the level of sensitivity to each factor. From these sensitivities, the liability-hedging portfolio must be created and the manager must decide on how much to hedge and how much to allocate into the performance portfolio, which usually depends on the magnitude of the funding gap. Finally, the manager must formulate a funding strategy to fill the funding gap and reach the fund's objective (Bruder *et al*., 2010).

In their article, Meder and Staub (2006) apply the LDI procedure. They start by identifying the exposure of the asset and the liability to financial and non-financial risk factors. The authors explain that because the indexed pensions of inactive members (retirement pensioners) are only affected by inflation, they should be hedged with a portfolio of inflation-linked bonds. In their case, the pensions of active members (contributors) are exposed to economic growth as well since the pensionable earnings on which the pension payments are calculated are not fixed yet. Active members' pensions generate a negligible risk in the short-term but in 20 years, as contributors retire, they will dominate the liability so it is still important to hedge their exposure. The level of pensionable earnings is a function of two main variables: wage inflation and real wage growth. Wage inflation is strongly linked with general inflation so it can be hedged with inflation-linked bonds. As for the real wage growth, it is linked to productivity increases, which can be proxied by economic growth (GDP). In the long-run, the stock market and the GDP are believed to follow each other. The real wage growth can therefore be hedged with stocks. When combining the wage inflation risk and the real wage growth risk, the resulting liability-hedging portfolio for active members' pensions is comprised of a mixture of inflation-linked bonds and stocks (Meder and Staub, 2006). Meder and Staub (2006) also present the non-financial exposure of the liability, which they call "liability-noise". It includes demographics and mortality rates. According to them, the main factor contributing to liability-noise is the number of participants. Because of their non-market nature, these risks are hard to hedge (Meder and Staub, 2006).

After having identified and analyzed the factors affecting the asset and the liability, Meder and Staub (2006) estimate the corresponding sensitivities. To do so, they create models to represent the value of assets and liabilities. For example, bonds are represented by coupons and principal payments discounted at the risk-free rate plus a bond risk premium. Meder and Staub (2006) neglect the liability-noise and therefore generate the optimal liability-hedging portfolio based only on the sensitivity of the asset and the liability to the financial risk factors and their mutual correlations. The weights that minimize the surplus variance are 80% in nominal bonds, 10% in inflation-linked bonds and 10% in stocks. Because the pension fund tested in their article does not index pensions, a lower weight is allocated to inflation-linked bonds than an indexed fund would. Also, a fund that is less mature would have a higher allocation in stocks and inflation-linked bonds because of its larger wage growth risk. Finally, Meder and Staub (2006) remind us that the liability-hedging portfolio must be supplemented with a performance portfolio to reduce the contributions required to keep a positive balance. This thesis also applies the LDI approach to a fund but, in order to provide a more complete picture, non-financial risk factors are also included.

While Meder and Staub (2006) simply ignore liability-noise, Boyer *et al*. (2011) highlight the importance of managing these risks. More specifically, they found that "to compensate for longevity risk alone, it has been estimated that companies should increase their pension fund provisions by approximately 4%" (Boyer *et al*., 2011: 3). In fact, longevity risk is non-diversifiable because individual mortalities are not independent. They are partially correlated since scientific improvements that lead to increased life expectancy touch most individuals in the population (Boyer *et al*., 2011). Also, longevity bonds and other longevity products are not liquid enough to hedge this risk properly. As a result, fund managers must keep reserves to compensate for longevity risk.

In their article, Boyer *et al*. (2011) apply the CBD model to the Royal Canadian Mounted Police and the Canadian Forces' pension plans in order to measure the exposure of the funds to longevity risk. They start by analyzing the liability ascribable to the retirement pensioners of the Royal Canadian Mounted Police. Firstly, in order to verify the validity of their framework, they model the fund's liability by applying assumptions regarding mortality, demographics, inflation and interest rates the closest to those made by the Office of the Chief Actuary (OCA). The model results in a liability for retirement pensioners only 5% different than the value published in the actuarial report, which shows that the model simulates the fund appropriately. Secondly, they simulate mortality rates with the CBD model and obtain a liability 1.5% larger than what is obtained when using the OCA mortality assumption. This is called the model risk since it measures the impact of using a certain model over another one. With the CBD model, different mortality trajectories are run, which creates a distribution for the plan's liability. To measure the longevity risk, which is associated to the uncertainty around the evolution of mortality rates, the nominal and relative value-at-risk at the 95^{th} and 99^{th} levels are used. Boyer *et al*. (2011) find a relative VaR at the 95^{th} level of about 3%. When combined with the 1.5% model risk, this sums up to an underfunding of about 4.5% of the total liability if we believe the CBD model over the OCA model. According to these results, the fund managers should consider the liability ascribable to retired members to be 4.5% higher if they wish to have a 5% probability of insolvency. The same methodology is applied for the contributors. When combining the two groups (retirement pensioners and contributors) with their corresponding weights, the model risk is found to be 3.2% and the longevity risk at the

95^{th} level is found to be 2.5%, which amounts to an underfunding of about 6% of total liability. When adjusting for longevity risk, the 3.7% actuarial surplus that was published by the fund in 2008 becomes a deficit of about 2.3%. Similar results are obtained for the Canadian Forces. This clearly shows the importance of managing the non-financial exposure of the liability. In this thesis, a similar methodology is used to evaluate the exposure of the fund to risk factors. The model's assumptions are compared to the OCA's assumptions, the model risk is computed and the nominal and relative VaR are found. This thesis however does not focus solely on the longevity risk as it includes several financial and non-financial risk factors. Also, whereas Boyer *et al*. (2011) analyze the impact of the longevity risk on the liability of the pension fund, this thesis models the assets as well and finds the optimal asset allocation to minimize the solvency risk.

The following sections of the review of literature examine a dozen of articles that perform a portfolio optimization based on stochastic liabilities each with different objectives and risk factors. This will help us select the right objectives and risk factors when applying the LDI approach.

2.3 OBJECTIVE OF THE FUND

Whereas the goal of fund managers in defined contribution pension plans is to maximize the expected utility from the accumulated fund, the objective in defined benefit pension plans has to take into account the joint interests of the plan's sponsor, the contributors and the pensioners. As mentioned previously, any shortcomings in the fund are the sponsor's responsibility. Also, contribution rates must remain at acceptable levels so that contributors can afford the plan while the actuarial surplus must remain positive so that pensioners receive their pension payments (Battocchio and Menoncin, 2004).

A recurrent objective is to maximize the funding ratio or equivalently, to maximize the actuarial surplus. This objective is selected by Martellini (2006) and Martellini and Milhau (2008). An important drawback however is that maximizing surplus at all costs may imply taking unnecessary risks. Also, this objective can be attained by increasing contributions, which is not desirable for active members. This is why Josa-Fombellida and Rincon-Zapatero (2008) combined the surplus maximization objective with a contribution minimization objective. This is still problematic since these objectives are contradictory and they still create an incentive for taking excessive risk. In fact, regulators "are only concerned that sufficient assets are present to make payments, not growing a surplus for sponsor comfort" (Leahy, 2011: 6).

To avoid this problem, Delong *et al*. (2008), Josa-Fombellida and Ricon-Zapatero (2010) and Cox *et al*. (2013) used the objective of minimizing the variance of the terminal surplus subject to an expected surplus constraint. This is a risk minimization objective where risk is defined as the variability of the surplus at the end of the projection horizon. As it is discussed in the methodology section, this thesis applies the minimization of the terminal surplus variance as an objective.

Alternate objectives include the minimization of contributions with an expected shortfall constraint, which is applied by Ostaszewski (2011), or the minimization of the probability of insolvency, which is applied by Service and Sun (2004). According to Service and Sun (2004), because asset and liability both have a non-normal distribution, the simple use of the terminal surplus variance as a measure of risk is disputable. Using the probability of insolvency as a risk factor takes into account the different non-normal characteristics of the distribution of asset and liability. For this reason, this thesis also applies the minimization of the probability of insolvency as an objective.

2.4 RISK FACTORS

The choice of factors depends on the risks the manager wants to analyze. It also depends on the level of model complexity desired. The more factors are added, the more complex the model is. Martellini (2006), Delong *et al*. (2008), Martellini and Milhau (2008) and Josa-Fombellida and Ricon-Zapatero (2010) include few factors (two or three) in their model, which allows deriving an analytical solution, while Ostaszewski (2011) includes four factors and uses a Monte Carlo simulation. In his case, Martellini (2006) models the liability as a unique process and does not include factors. This however does not allow analyzing the sensitivity of the liability to different risk factors. Most authors include factors such as the inflation rate, the interest rate, the earnings growth, the mortality rates and the asset returns.

The inflation rate is included by Martellini and Milhau (2008) and is modeled using a geometric Brownian motion. The interest rate is included by Martellini and Milhau (2008), Josa-Fombellida and Ricon-Zapatero (2010) and Ostaszewski (2011) and is modeled using a Vasicek process. The earnings growth is also modeled with a geometric Brownian motion by Josa-Fombellida and Ricon-Zapatero (2010) and Ostaszewski (2011). Mortality rates are modeled by Cox *et al*. (2013) with the Lee-Carter model (1992). Finally, asset returns are modeled with a geometric Brownian motion by Martellini (2006) and Martellini and Milhau (2008) and with a Levy process by Delong *et al*. (2008) and Cox *et al*. (2013). We elaborate on these models as well as alternate models in section 3.

When a model is comprised of numerous factors, most authors include correlation between some of the factors. For example, Ostaszewski (2011) correlates earnings growth to the asset returns and Josa-Fombellida and Ricon-Zapatero (2010) correlate the asset returns to the earnings growth and the earnings growth to the interest rate. All models studied use a linear combination of two Wiener processes where one process is independent and the other process is common to the two correlated factors. For example, in the model of Ostaszewski (2011), the correlated process used is as follows:

$$W_3 = W_1 * \rho + W_2 * \sqrt{1-\rho^2}$$

where W_1 is the common Wiener process, W_2 is the independent Wiener process and ρ is the historical correlation coefficient between the two factors modeled. This is also how correlation is induced between the risk factors in this thesis.

Although some factors are believed to be correlated, other factors are assumed to be independent. This is the case of mortality rates, which are modeled as independent processes by Delong *et al*. (2008) and Ostaszewski (2011).

As we will see in the methodology section, this thesis differs from the current literature by including more factors. In fact, seven risk factors are included compared with three for most models.

2.5 OPTIMAL ASSET ALLOCATION

After having selected and modeled the risk factors, the asset allocation that optimizes the objective function has to be found.

Under its objective of contribution minimization, Ostaszewski (2011) finds that the optimal asset allocation assigns a large portion of assets to stocks relative to bonds so that the expected surplus rises and fewer contributions are needed. The model also highlights an important finding which is that contributions could alternatively be reduced by increasing the retirement age. Overall, although this thesis applies a similar methodology as Ostaszewski (2011) (except that it includes more risk factors), we do not expect to obtain the same optimal asset allocation because we use a different objective, which focuses on risk minimization rather than return maximization.

While Ostaszewski (2011) finds a static allocation, most authors find a contingent strategy where more risk is taken if a fund is in deficit or has a large surplus. This is supported by Zenios and

Ziemba (2007) and Josa-Fombellida and Ricon-Zapatero (2010). Josa-Fombellida and Ricon-Zapatero (2010) find that, for an underfunded plan, the optimal strategy in the first years involves taking a lot of risk with a high allocation in risky assets and as time goes by and the deficit gets smaller, funds are transferred to the savings account. Zenios and Ziemba (2007) add that the larger the surplus is, the greater the allocation in equity. The large surplus acts as a buffer and allows taking more risk while keeping a low probability of falling into a deficit.

Similarly, other authors such as Fabozzi (1997), Martellini (2006), Zenios and Ziemba (2007), Martellini and Milhau (2008), Bruder *et al*. (2010) and Moore (2010) support a contingent strategy but also confirm the LDI approach by which two portfolios are created. In fact, their results show that the optimal asset allocation consists of a liability-hedging portfolio and a performance portfolio.

Fabozzi (1997) explains that a fund that has a funding ratio of 1 has an incentive to hedge its liability in order to keep a positive balance whereas a fund that has a large surplus is able to take more risk with a larger allocation to the performance portfolio and still have a positive surplus if things go bad. As the funding ratio gets lower and close to 1, the fund can gradually go back to the liability-hedging portfolio to reduce risk (Fabozzi, 1997). This is in accordance with the results of Zenios and Ziemba (2007) that were presented above.

In their case, Bruder *et al*. (2010) assess the situation where a fund is in deficit. If the fund is underfunded, a large portion of the asset should be put in the performance portfolio. By increasing the exposure to risky assets, the manager has a chance of closing the funding gap. If the fund is fully funded but has a small surplus, the fund manager should reduce the exposure to risky assets and shift to the liability-hedging portfolio in order to secure assets and make sure the gap does not change (Bruder *et al*., 2010). But on average, according to Moore (2010), most funds should invest 70% of total funds in the liability-hedging portfolio because of the diminishing marginal benefit of hedging the liability.

2.6 SWAPS

Although swaps and swaptions are outside the scope of this thesis, it is important to note that they are often used by fund managers in their investment strategy as a way to reduce the duration gap and reduce the risk of insolvency while keeping funds available for performance generation (Engel *et al.,* 2005; Moore, 2010). As Bruder *et al*. (2010) explain, the liability-hedging portfolio can be achieved by combining an interest rate and inflation rate swap to a fixed income portfolio. While this allows receiving payments that cover the fluctuations in the value of the liability due to the interest rate and the inflation rate, it does not require locking in vast amounts of funds in asset classes such as inflation-linked bonds. Although collateral agreements are required (Bragt, 2011), there is no initial investment, which allows investing funds in the Markowitz efficient portfolio (Bruder *et al*., 2010).

In this thesis, the optimal asset allocation depends on the liability resulting from our model and its exposure to the risk factors. In turn, this exposure depends on our choice of models to describe the risk factors. The following section elaborates on these models.

3. THEORETICAL MODELS

3.1 THE WIENER PROCESS

The Wiener process (also called standard Brownian motion) is a continuous-time stochastic process that is widely used in modeling to describe the path followed by random variables. Three important properties of the process are that it starts at zero, it is almost surely continuous and it has independent increments with a normal distribution. In this thesis, Wiener processes are used as building blocks for other processes.

3.2 INFLATION RATE

The consumer price index (CPI), which is a measure of price inflation, is often modeled with a geometric Brownian motion (Matellini and Milhau, 2008), which is a combination of a Wiener process W and a drift μ. Inflation can therefore be modeled as the growth rate of the CPI and follow the following differential equation (Adam, 2007):

$$\frac{dX_t}{X_t} = \mu dt + \sigma dW_t$$

where X_t is the CPI and $\frac{dX_t}{X_t}$ is the inflation rate.

This model is widely used in finance due to its simplicity. Because increments are normally distributed, the resulting models often allow deriving semi-closed form solutions. A drawback of this model is that volatility (measured by the standard deviation σ) is constant. Empirical evidence however shows that high volatility tends to cluster. In other words, we observe autocorrelation in the volatility processes. This is true for inflation rates as well as stock returns (Cont, 2007).

To take into account this characteristic, generalized autoregressive conditional heteroskelasticity (GARCH) models are used:

$$\sigma_t^2 = a_0 + a\sigma_{t-1}^2 + b\varepsilon_t^2$$

where the squared volatility is a function of its value in the previous period (Cont, 2007). This stochastic volatility can then be used instead of the constant volatility in the Brownian motion above.

Other models for inflation are based on macroeconomic variables and are simple functions of variables such as GDP, population and interest rates (Adam, 2007). For example, there is a strand of models based on the Phillips Curve, which represents the relationship between the unemployment rate and the inflation rate. One of those models is the Accelerationist Phillips Curve developed by Milton Friedman in 1968:

$$\pi_t = \gamma(U_t - U^*) + \pi_t^e$$

where the inflation rate (π_t) is given by the expected inflation rate (π_t^e) plus a fraction of the difference between the actual unemployment rate (U_t) and its natural rate (U^*) across time. The problem with this model is that it requires simulating the unemployment rate. This unemployment rate, in turn, can be modeled from economic growth (measured by the GDP) but this would require modeling more variables and would result in a more complex framework with probably greater model risk (Rudd and Whelan, 2007). For this reason, the geometric Brownian motion with constant volatility is used in this thesis.

3.3 INTEREST RATE

In the literature, two main types of interest rate models are found: short-term interest rate models and market models based on market rates like the Libor or the swap rate (Adam, 2007).

The Vasicek model (introduced in 1977) is a one factor Ornstein-Uhlenbeck stochastic process that represents the evolution of short-term interest rates (Adam, 2007). The model was innovative as it introduced mean reversion to interest rate modeling. In fact, historical interest rates showed a mean reversion behavior over the long-term, which created the need for such model (Elen, 2010). The Vasicek model is used to simulate short-rates, which are instantaneous spot rates. In practice, however, the instantaneous rates are often replaced by the 3-month bond rates. The model is known for its simplicity and ease of implementation. Some limitations of this model are that it allows negative interest rates (with a small probability) and the increments for different maturities are perfectly correlated, which is not realistic. This problem is non-existent however when modeling the rates on bonds with a single maturity, which is what we perform in this thesis. Also, although this slightly changes its properties, the process can be reset at zero when it becomes negative.

The following outlines the model and its parameterization:

$$dr_t = \theta[\mu - r_t]dt + \sigma dW_t$$

where W_t is a Wiener process (or standard Brownian motion), σ determines the amount of volatility, μ is the long-term average rate to which the process reverts and θ determines the speed of mean reversion.

An ordinary least squares (OLS) regression of the interest rate over its lag value (previous period's interest rate) is performed to obtain the maximum likelihood estimators. As presented by Elen (2010), these estimators also have a closed form solution:

$$\hat{\theta} = -\log\left[\frac{n\sum_{i=1}^{n} r_i r_{i-1} - \sum_{i=1}^{n} r_i \sum_{i=1}^{n} r_{i-1}}{n\sum_{i-1}^{n} r_{i-1}^2 - (\sum_{i=1}^{n} r_{i-1})^2}\right] / \Delta t$$

$$\hat{\mu} = \frac{\sum_{i=1}^{n}\left[r_i - e^{-\Delta t\hat{\theta}} r_{i-1}\right]}{n\left(1 - e^{-\Delta t\hat{\theta}}\right)}$$

$$\hat{\sigma} = \sqrt{\frac{\frac{1}{n}\sum_{i=1}^{n}\left[r_i - e^{-\Delta t\hat{\theta}} r_{i-1} - \hat{\mu}\left(1 - e^{-\Delta t\hat{\theta}}\right)\right]^2}{\left(e^{-2\Delta t\hat{\theta}} - 1\right)/2\hat{\theta}}}$$

Although the Vasicek model is designed for short-rate modeling, it is sometimes used to model rates on bonds of longer maturities. In his case, Elen (2010) uses bonds with a 2-year maturity. Because it still describes the right general dynamics, it is used in this thesis to generate the interest rate process on the 20-year Government of Canada bond.

Other short-term interest rate models include Hull & White (1990) and Cox-Ingersoll-Ross (CIR) (1985). Hull & White is a one-factor model in which the short-term interest rate follows a Vasicek process under the risk neutral probability (Adam, 2007). Again, the process can produce negative interest rates. To remediate this problem, the CIR model modifies the Vasicek model slightly by multiplying the Wiener process by the square root of the interest rate (Adam, 2007):

$$dr_t = \theta[\mu - r_t]dt + \sigma\sqrt{r_t}dW_t$$

For the process to be defined, the interest rate cannot be negative.

Although we will not elaborate on these models it is important to note that market models are available and used in the industry. These models use option market prices to find the corresponding interest rates (Adam, 2007).

3.4 EARNINGS GROWTH

Because employees' earnings growth is a risk factor specific to pension funds and it differs significantly from one plan to another, few models are proposed in the literature to describe it. Most authors like Josa-Fombellida and Ricon-Zapatero (2010) and Ostaszewski (2011) use a geometric Brownian motion. This is also applied in this thesis, where the earnings (E_t) are described by the following process with a mean μ and a standard deviation σ:

$$\frac{dE_t}{E_t} = \mu dt + \sigma dW_t$$

Like inflation, a process using other economic variables like GDP or unemployment rates could be used but would increase the model complexity.

3.5 MORTALITY RATES

The Lee-Carter model (1992) is widely used in the academic world as well as in the industry to forecast mortality rates. Although it has been developed to explain all-cause mortality rates in the U.S., the model has also been used for cause-specific mortality in different countries. (Girosi and King, 2007). Here is the equation of the model:

$$\ln(m_{xt}) = a_x + b_x k_t + e_{xt}$$

Lee and Carter explain log-mortality rates with two parameters that varies across age (a_x and b_x) and one parameter that varies across time (k_t) in addition to a random disturbance factor that is normally distributed (e_{xt}). The *k*'s represent mortality changes over time whereas the *b*'s denote the sensitivity of each age group to these changes. In order to fit the model to historical data, maximum likelihood is used. A simple OLS regression cannot be used in this case because only the *m*'s are observable and there is no regressor. The product of two vectors (*b*'s and *k*'s) must be found. Without further constraints, an infinite amount of solutions could hold and maximize the likelihood function. Lee and Carter therefore add that the *b*'s must sum up to one and the *k*'s must sum up to zero. This last constraint makes the *a*'s equal to the average log-mortality rates for each age group, which can be estimated with mortality rates over time (Lee-Carter, 1992).

To estimate the *b*'s and the *k*'s, singular value decomposition (SVD) is used on the matrix of the log-mortality rates after having subtracted the corresponding average mortality rate (*a*'s). The first column of the resulting least square solution represents the maximum likelihood estimators of the *b*'s. The *k*'s are simply the product of the *b*'s and the average log-mortality across age groups for each year (Girosi and King, 2007).

It is important to note that in this model, the *b*'s and *k*'s are independent. Although it simplifies the model, it is also a drawback because evidence in different data sets shows that the *b*'s change over time and are dependent on the *k*'s (Girosi and King, 2007). From there, Girosi and King (2007) explain how to forecast future mortality rates.

Lee and Carter assume the *k*'s follow a random walk with drift:

$$k_t = k_{t-1} + d + e_t$$

where e_t is normally distributed with mean zero and variance found by MLE:
$\hat{\sigma}^2 = \frac{1}{T-1}\sum_{t=1}^{T-1}\left(\hat{k}_{t+1} - \hat{k}_t - \hat{d}\right)^2$ and d (the drift parameter) is $\hat{d} = \left(\hat{k}_T - \hat{k}_1\right)/(T-1)$.

To generate the *k*'s, the following simulation is used:

$$\hat{k}_{T+\Delta t} = \hat{k}_T + (\Delta t)\hat{d} + \sqrt{(\Delta t)e_{T+\Delta t}}$$

Finally, from these *k*'s, we estimate the log-mortality rates as

$$\mu_{T+\Delta t} = \hat{a}_x + \hat{b}_x\hat{k}_{T+\Delta t}$$

Because Lee-Carter is a one factor model, it assumes mortality rates are perfectly correlated across age groups. Alternatively, the Cairns, Blake and Dowd (CBD) model (2006) has a second factor that relaxes this assumption. It allows different evolution in mortality rates for different age groups, which is supported by historical evidence (Boyer *et al.*, 2011).

Both the Lee-Carter and CBD models may overestimate mortality rates because they assume continuous improvements. If mortality improvements evolve in jumps, mortality rates could be lower than predicted by continuous models (Boyer *et al.*, 2011). Also, some theorists expect mortality improvements to decrease over time because they assume humans have a biological limit to their life span. Others adopt a Darwinian perspective and believe that the force of evolution will lead the dominant specie towards immortality (Goldsmith, 2011). Regardless, there are a lot of conflicting theories regarding the future of mortality, which makes the choice of model difficult and uncertain. In this thesis, the Lee-Carter model is selected for its simplicity and fairly good fit of the data since we use a relatively homogeneous demographic group.

3.6 ASSET RETURNS

A simple way to model stock price is to use a geometric Brownian motion (Adam, 2007). As mentioned above, this process allows describing the general diffusion behavior of stock prices and its simplicity allows deriving analytical solutions.

Asset returns from various financial markets however consistently showed more extreme values over time than the normal distribution predicts. Because of this "fat-tail" property, some researchers like Delong *et al*. (2008) use the more realistic Levy process with jumps to simulate future asset prices. Levy processes can create a distribution with heavy tails and skewness by producing jumps in the path of the asset prices as it does not require normally distributed increments (Delong *et al*., 2008). Jump diffusion models are Levy processes, which often combine a Wiener process (the diffusion part) with a Poisson process (the jump part). One of the simplest forms of this model is the Merton jump diffusion model (1976) by which stock prices follow a Brownian motion to which we add a jump process:

$$\frac{dS_t}{S_t}(\alpha - \lambda k)dt + \sigma dB_t + (y_t - 1)dN_t$$

where α is the expected return, σ is the volatility of the stock return, B_t is a standard Brownian motion, N_t is a Poisson process with intensity λ and $(y_t - 1)$ is the size of the jump. The occurrence of the jumps follows a Poisson distribution whereas y_t follows a lognormal distribution (Tankov and Voltchkova, 2009).

Although these processes result in acceptable asset returns, they do not preserve the original distribution and correlations of the historical asset returns. These are especially important because we want to hedge the higher moments of the distribution of the liability with the portfolio of assets. In order to obtain a process with a better distribution, we randomly select returns from the historical sample and adjust them based on our expectation of future returns. This method is further explained in the methodology section.

4. METHODOLOGY

4.1 THE MODEL

In this study, the LDI approach is used to analyze the liability and the asset of the pension fund of the Public Service of Canada and to find the optimal asset allocation and funding strategy. Although the fund is separated into the Superannuation Account and the Public Service Pension Fund, this study is interested in analyzing the effect of risk factors on the two accounts. We also want to build a single optimal asset allocation and funding strategy. For this reason, the model does not differentiate between the accounts and considers the sum of the two.

The first step of LDI is to model and estimate the liability, the asset and the surplus. The fund is therefore modeled for a period of 20 years starting from 2011. Because the 2011 actuarial report was the last to be published, the projections start at this time.

In order to evaluate the exposure of the fund to risk factors, the asset and liability must be modeled as a function of financial and non-financial variables. The Public Sector Pension Investment Board (PSP Investments) discloses in its Statement of Investment Policies (2012: 4) that the funding status is affected by: “the rate of inflation, the rate of increase of members’ pensionable earnings, the rate of return on assets, investment risks (market risk, credit and counterparty risk, concentration risk, liquidity and financing risk and leverage risk), demographic factors (mortality rates, rates of employee terminations and age for retirements) and the actuarial funding cost method used by the Chief Actuary”. Like it is mentioned by the Office of the Chief Actuary (OCA), currency variations are expected to offset each other over time. As a result, the currency risk is neglected. Overall, seven risk factors are selected. On the one side, the liability is modeled as a function of the inflation rate, the interest rate, the growth of employees’ earnings, mortality rates and the number of pensioners. On the other side, the asset is modeled as a function of contribution rates and the return on the portfolio of assets.

4.2 LIABILITY

In this thesis, retired pensioners and contributors from the “main group” are included. Like Ostaszewski (2011) and Boyer *et al*. (2011) did, we exclude disability pensioners and surviving spouses from our model. Contributors on leave without pay and contributors from the “operational group” are also excluded. The reason for this exclusion is that they have different demographics, pay different contributions and receive different benefits than the groups included. In the Public Service Pension Plan, there are 114,125 male retirement pensioners and 78,452 female retirement pensioners, which make up 69% of total pensioners. In their case, the main group contributors make up 89% of total contributors. When combining retirement pensioners and contributors, this thesis includes 80% of total members of the plan. In terms of liability, the retirement pensioners and the active contributors that are in the main group make up respectively 44% and 48% of the total liability, which sums up to 92% of the total plan’s liability. The division of the total liability per group is depicted in Figure 1. As we can see, these groups represent the vast majority of the liability so they engender the majority of the fund’s financial and non-financial risks.

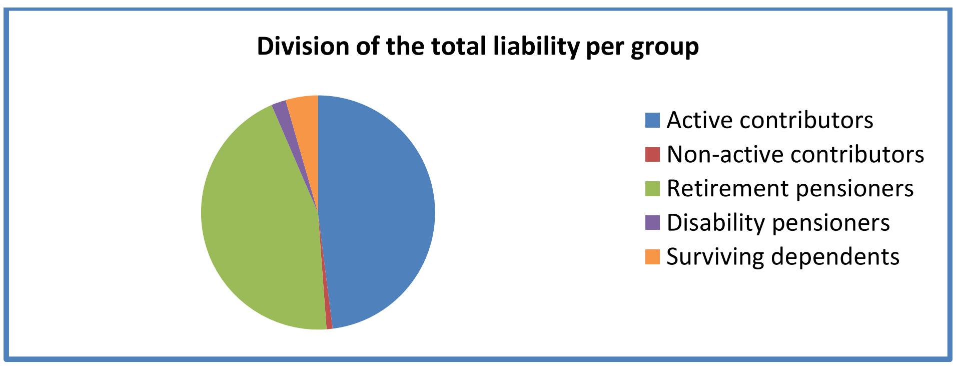

Figure 1

This thesis uses two types of liability valuation. The first is the solvency valuation, which assumes the plan stops operating and has to repay its retirement pensioners. This valuation does not include the benefits of active contributors and therefore focuses solely on the retirement pensioners’ liability, which as mentioned above, makes up 44% of the total liability. Because of this, we also only use 44% of the fund’s total asset in our solvency calculations and in the model in general (the rest being assigned to other retirement groups). The second method of liability valuation is the going concern valuation that assumes the fund continues operating for the projection horizon (20 years) and then goes into perpetuity. A portion of the current active contributors are included because when they retire they become part of the retirement pensioners group. The contributors that are not included in these valuation methods can be assumed to fund their own pension with their contribution payments so they do not participate to the solvency risk.

4.2.1 Inflation rate

The liability of the fund is influenced by the general state of the economy. This is why the level of inflation and the interest rates are being modeled. Firstly, the pension payments are indexed annually with the level of inflation measured by the CPI. In our model, the level of inflation is assumed to follow a geometric Brownian motion. Historical inflation rates fluctuated mostly around 3% so the mean is assumed to be constant and the CPI is assumed to follow random and independent increments. Annual Canadian CPI levels from Statistics Canada ranging from 1991 to 2011 are used to parameterize the stochastic process since future inflation levels are assumed to be similar to the last two decades and their volatility is relatively stable during this period. Just like the OCA does, if the CPI drops, the pension payments are not reduced but the deflation rate is subtracted from the next CPI increase. To test the impact of the data set used to parameterize the Brownian motion, CPI levels from 1914 to 2011 are also used. The reason for using levels from the past century is to include data from an era where the Bank of Canada did not control inflation rates between 1 and 3%. This modification of the behavior of inflation rates can be seen in Figure 2. This data set results in a higher standard deviation for the Brownian motion and therefore allows for some high inflation scenarios.

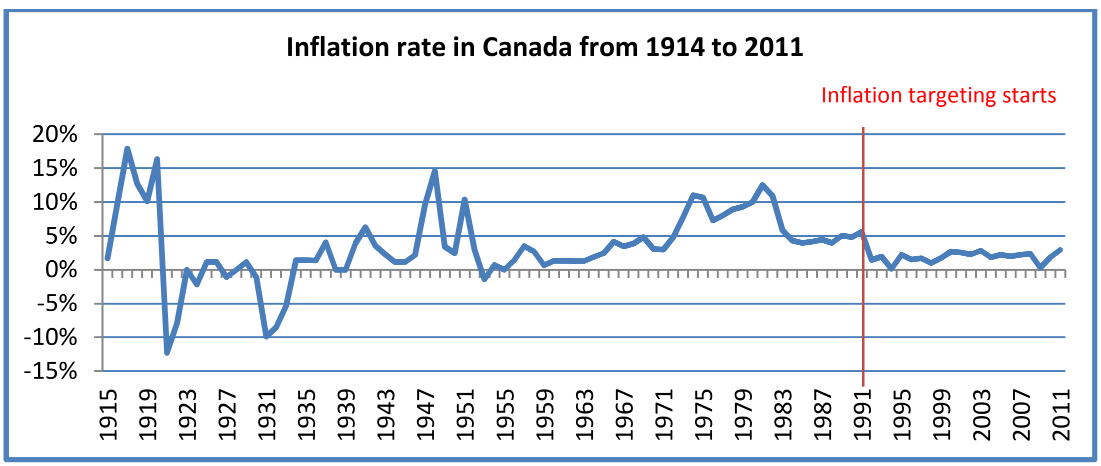

Figure 2

4.2.2 Interest rate

To compute the present value of the liability, the payments made to pensioners must be discounted. The discount rate used by the OCA for the pension fund is a weighted average of the expected return on the different asset classes in which the fund is invested according to their respective weights in the portfolio plus a transaction cost and management fee of 0.3% in total.

The OCA adds risk premiums (as displayed in Table 1) to the long-term government bond return to obtain the expected return on the other asset classes.

	Risk premiums per asset class per year					
	2012	**2013**	**2014**	**2015**	**2016**	**2017+**
Canadian Stocks	5.0%	4.8%	4.6%	4.4%	3.2%	2.0%
US Stocks	5.0%	4.8%	4.6%	4.4%	3.2%	2.0%
Global Equity	6.0%	5.8%	5.6%	5.4%	4.2%	3.0%
Can. Inflation-Linked Bonds	-0.5%	-0.4%	-0.4%	-0.4%	-0.4%	-0.4%
Can. Government Bonds	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
Canadian Corporate Bonds	1.5%	1.3%	1.1%	0.9%	0.9%	0.9%
Global Real Estate	2.5%	2.5%	2.1%	2.1%	1.9%	1.4%
Global Infrastructure	2.5%	2.5%	2.1%	2.1%	1.9%	1.4%

Table 1

The same methodology and the same risk premiums are used in our model. The only difference is that our model allows the long-term government bond return to follow a Vasicek process. Because most payments are due only after decades, using the returns on a 20-year maturity bond is logical. In fact, because the discount rate must represent the risk of the cash flows it discounts, it is more appropriate to use the return on bonds with a similar duration as the cash flows. Different portfolio weights, including the original allocation and different optimal portfolio allocations, are used to calculate the weighted average discount rate. It is important to note that although the discount rate corresponds to the expected portfolio return, the simulated portfolio return can differ from this expectation. This is what the model represents by simulating the discount rate and the portfolio return processes separately.

Because pension payments are made monthly, monthly interest rates are simulated. The Vasicek process is parameterized using monthly 20-year Government of Canada bond returns from 2002 to 2011 (taken from Bloomberg) and the Wiener process inside the Vasicek process is correlated to the inflation rate process. To correlate the interest rate (a monthly process) to the inflation rate (an annual process) is not straightforward because the two vectors have different lengths. To create such a process, we first create an annual process, which is a linear combination of an independent Wiener process and the Wiener process associated to the inflation rate. This process is assumed to be an annual shock to the interest rate. Then, a second process is created by generating twelve standard normal random numbers around each point of the annual process, which is considered to be the mean of each dozen.

Although the OCA uses the expected return on the portfolio of assets to discount the liability, other discount rates are suggested in the literature. Because pension payments attributable to retirement pensioners are known in advance, it could be argued that they can be replicated by long-term bonds and should accordingly be discounted at the long-term Treasury yield. This assumes a lower risk premium associated to those low risk future claims. As a result, in addition to discounting at the expected portfolio return, a second methodology is used in which future payments are discounted at the 20-year Government of Canada bond return.

Finally, as suggested by Wouters (2008), pension payments should be split in two when valuing the liability. The first group of payments are associated to inactive members (the retirement pensioners) and are discounted at the long-term bond yield because of their predictability, while the second group of payments are associated to active members (the contributors) and future members, whose pensionable earnings are uncertain and are discounted at the "plan's expected long-term investment return" (Wouters, 2008: 23). This method is in fact a combination of the two methods mentioned above. All three methods are used and compared.

4.2.3 Earnings growth

Although employees' earnings are affected by the general economy, notably by inflation levels, they also depend on internal decisions made by the employer. In turn, the level of these earnings impacts the magnitude of pension payments because pensions are calculated based on the earnings at retirement. To replicate the nature of earnings in a simple manner, a geometric Brownian motion correlated to the inflation process is used. As mentioned earlier, the correlation is induced by combining an independent and a common Wiener process. The correlated portion represents the impact of inflation and the non-correlated portion represents the impact of internal decisions on earnings.

Because only the earnings of contributors are accessible in the actuarial reports of the Public Service, pensionable earnings are inferred from those of contributors. Although contributors' earnings are segmented in groups based on age and seniority, the volatility of the geometric Brownian motion that expresses earnings growth is parameterized using the average contributor earnings for males and females separately at the pensioners' average age of 56 from 1980 to 2011 for all groups. This simplification should not have a large impact since earnings improvements are relatively homogeneous across age and seniority groups. As for the mean of the geometric Brownian motion, it is set at 3.5%, which is the long-term growth in earnings expected by the OCA.

The average contributor earnings of each group are then assumed to follow the parameterized Brownian motion. We notice that, from age 65 onward, earnings and pensions stay relatively constant representing a cap in earnings progression. The model therefore assumes constant pensionable earnings from age 65 onward. New pensioners each year are assumed to have their pension calculated on the average contributor earnings pertaining to their age and seniority.

The challenge is to approximate the pensionable earnings of inactive members (those who retired on or before 2011) because the actuarial reports only provide us with the detailed contributors' earnings and the average pension received by the pensioners. The 2011 contributors' earnings are therefore used to approximate the pensionable earnings of the inactive members. First we have to find the distribution of pensionable earnings per seniority group, so we simply assume it is the same as the distribution of contributors' earnings. Then, because the pensioners did not retire on 2011 earnings levels, we have to lower our estimates while taking into account that pensions are indexed to follow inflation. From 1983 to 2011, earnings grew 4.8% on average. In the same time period, the CPI grew 3.4% on average. Because earnings grew more than the CPI, the indexation of pensions is not enough to offset the growth in earnings. As a result, by assuming a perfect offset of earnings growth by inflation and using 2011 contributors' earnings, we are overestimating the pensionable earnings of the current pensioners. This can be accounted for by reducing the average pensionable earnings by the exceeding growth in earnings but it would be nearly impossible to find the exact reduction factor for each age and seniority group since pensioners retired on different years and not enough information on pensionable earnings is available in the actuarial reports. As a simplification, we only consider reducing earnings for the pensioners older than 56 (the average age of retirement) to acknowledge that they retired on lower earnings levels. This leaves us with the last three age groups that are included in the contributors' earnings data: 55-59, 60-64 and 65+. To find the appropriate reduction factor, we find the factor that minimizes the difference between the calculated pensions and the pension amounts disclosed in the 2011 actuarial report. The result of this optimization is that no reduction should be included for pensioners in the age groups 55-59 and 60-64 and a reduction of 9.68% should be applied for the age group 65+. This can be explained by the fact that the age group 65+ probably contains pensioners that retired further in the past on lower earnings levels whereas pensioners in the two other age groups may have retired on earnings levels similar to 2011 or at least for which indexation mostly covers earnings growth.

4.2.4 Mortality rates

Besides economic variables, the liability is affected by another exogenous variable: mortality rates. The fund needs to pay their pensioners as long as they are alive so mortality rates affect the size of the fund's liability. The Lee-Carter model is used to simulate mortality rates. Like Delong *et al*. (2008) and Ostaszewski (2011) did, the mortality process is modeled independently from the other risk factors. The model is parameterized using annual Canadian mortality rates (taken from Statistics Canada) from 2000 to 2011 for males and females separately. The Canadian population is assumed to be representative of the pensioners of the Public Service of Canada. Because the mortality data stop with the age group 90+ and mortality rates for the age groups 90-94, 95-99 and 100-104 are needed, these last mortality rates are estimated by applying the same mortality improvement as in the last age group. A linear interpolation is then used to obtain mortality rates for every age within each age group. The rates are then assumed to stay constant from age 104, which is supported by empirical evidence (Goldsmith, 2011). The simulated mortality rates are finally adjusted so that the initial rates are equal to the actual mortality rates published in the 2011 actuarial report. Each year, the number of pensioners of each age is reduced by the corresponding mortality rate.

4.2.5 Number of pensioners

Another factor that affects the value of the liability is the number of pensioners over time. This is affected by mortality rates but also by the number of employees that retire. The OCA does not make any assumptions regarding the growth in the number of pensioners but it assumes the number of contributors grows by 6.4% over 20 years. Because contributors will eventually become pensioners, we assume the same growth rate for pensioners. In the case of our model, a larger growth rate is needed to cover the morality rate and still achieve a long-term growth. Adding approximately 3% of the number of pensioners of each age group at the beginning of each year allows us to reach a similar overall growth of 6.4% for males and females combined. Figure 3 depicts the impact of this growth rate on the number of male and female pensioners separately.

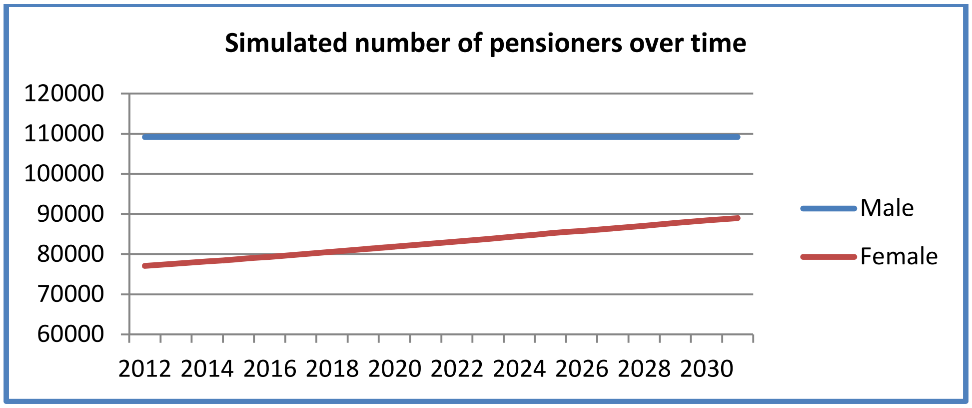

Figure 3

Because of their lower average mortality rate, the number of female members grows while the number of male members remains constant. As a result, the difference between the number of males and females is reduced from 17% to 10% of total pensioners by 2031.

After having determined the number of males and females of each age group, we have to find the corresponding seniority of these pensioners. Because the detailed distribution of pensioners by seniority group is not accessible, the proportion of contributors in each seniority group is applied to the number of pensioners to create a more detailed distribution. This proportion is assumed to remain constant over time. This assumption may underestimate liability if we believe pensioners have more seniority than contributors on average. Knowing that pension payments are made monthly, to add new pensioners at the beginning of each year rather than gradually during the year may slightly overestimate the liability because the payments are discounted on fewer months. Also, the model assumes that in 2011, the age of pensioners and contributors is equal to the median of their age group and their number of years of service is equal to the median of their seniority group. Each year, members gain one year of age and one year of seniority.

4.2.6 Pension plan terms

The liability of the fund also depends largely on the pension plan terms. In this study, for the most part, the same terms are applied in the modeling of the liability. In fact, pensioners are paid a pension that is 2% of their highest average pensionable earnings multiplied by the number of years of service (capped at 35 years). In this model, the average pensionable earnings are replaced by the last earnings a pensioner made before retirement. Doing so overestimates earnings by approximately 6.5%. The retirement pensioners can receive three types of pensions: an immediate annuity, an annual allowance and a deferred annuity. Pensioners with at least 25 years of service receive an immediate annuity. Pensioners with 20 to 24 years of service receive an annual allowance payable from age 50 or 55 if they retired after 2012. This pension is reduced by 5% per year that separates the pensioners from age 60 or 65 if they retired after 2012. Finally, pensioners with less than 20 years of service receive a deferred annuity, which starts being payable at age 60 or 65 if they retired after 2012. When pensioners reach age 65, their pension payments are reduced by 0.625% of the average maximum pensionable earnings (AMPE) multiplied by the number of years of service to account for the payments received by the Canada Pension Plan or the Québec Pension Plan. The AMPE, which was equal to $46,080 in 2011, is assumed to increase annually by its historical average of 2.8%. All pensioners are assumed to retire voluntarily and therefore be entitled to their pension.

4.2.7 Perpetuity

Finally, because the model simulates for 20 years but payments are due for a longer period of time, payments that are made on year 2031 are assumed to go on in perpetuity. The discount rate of the perpetuity has to take into account indexation and mortality. In fact, indexation increases the payments to be made over time while mortality reduces them. Accordingly, the discount rate (which is the same as the one used to discount payments) must be increased by the average mortality rate and reduced by the average inflation rate. Mortality rates are assumed (for simplicity and conservatism) to stagnate at their 2031 level in the future. We use the average mortality rate from age 56 to age 90 because we assume most 2012 pensioners that are still alive in 2031 will have reached at least age 56. The inflation rate and the interest rate are assumed to revolve around their mean value in the future so we use their average value on the 20 years simulated. Also, because payments are made monthly and the pensions are indexed only once a year, the payments each month are the same and are simply one twelfth of the calculated annual pension payment. Overall, the liability each year is the discounted payments owed to the retired members on that year until their death.

4.3 ASSET

To obtain the actuarial surplus, the liability must be subtracted from the asset. The fund invests in eight asset classes: Canadian stocks, U.S. stocks, foreign equity, Canadian inflation-linked bonds, Canadian government bonds, Canadian corporate bonds, global real estate and global infrastructure. The amount of asset depends not only on contributions and pension payments, but also on the market return on these asset classes.

4.3.1 Contribution rates

Firstly, the fund receives contributions from members on each pay. In our model, these payments are assumed to be made annually. The fund is probably not reinvesting the contribution money each week so our annual assumption should not impact the results too severely. These contributions are determined based on the member's level of earnings and if this one is above or below the maximum pensionable earnings (YMPE) for that year. Like the AMPE, the YMPE is assumed to grow at 2.8% annually starting from $50,100 in 2012. According to the OCA, three contribution rates are used: if the earnings are below the YMPE, the rate goes from 6.2% in 2012 to 6.9% in 2014+, if the earnings are above the YMPE, the rate goes from 8.6% in 2012 to 9.0% in 2014+ and if the contributor has more than 35 years of seniority, he pays 1.0%. The same contribution rates are applied in the model and the contributors' earnings are assumed to be the ones that are modeled to find the pensionable earnings above. The government also makes contribution payments to the fund. Based on historical proportions, the OCA projects the government will pay about 1.76 time the employees' contributions. The same is assumed in this model.

4.3.2 Number of contributors

As for the number of contributors, like the OCA, we assume it will grow by 0.3% annually over the next 20 years. Also, because the distribution of contributors remained relatively constant over the last 30 years, the model assumes no variation in the proportion of contributors by age and seniority groups. Contributors are assumed to change groups over time but to be replaced by new contributors. Keeping the distribution of contributors static over time neglects the changing dynamics of the workplace by which employees tend to switch jobs more often than the previous generation. A more accurate model could gradually increase the proportion of employees with less seniority. When retiring, these employees would have on average less seniority and their pensions would be lower. The homogeneity of growth assumption is therefore believed to result in a higher liability.

4.3.3 Portfolio return

A second factor that impacts the asset of the fund is the portfolio return. This one is a weighted average of the returns on the individual asset classes. These future returns are simulated by randomly selecting from historical daily returns (taken from Yahoo Finance) from 2009 to 2014 for eight different index funds that represent the eight asset classes. This is the longest time period for which data are available for all eight index funds. In order to evaluate the impact of this time window, daily returns from 2007 to 2014 are also selected for six of the eight funds for which data are available. Although it does not account for cyclicality or any other time series properties, this method allows obtaining future returns with a similar distribution and correlations as historical returns. Another method that uses correlated Brownian motions has been tested but this one assumes a normal distribution (with a skewness of 0 and a kurtosis of 3). Because we are interested in the third and fourth moments of the distribution, this characteristic is not desirable. The problem with the historical method however is that it restricts returns to historical levels, which were high during the 2009-2014 time period. This results in an average return of 12%. To assume that returns will remain at this level in the next 20 years is not a desirable assumption. In order to have more conservative expected returns, the historical returns on each index fund are reduced by a constant to obtain the same average risk premiums as the OCA. The resulting expected returns along with the corresponding index fund for each asset class are shown in Table 2.

Expected return and index fund used per asset class		
Asset Class	**Index Fund**	**Expected Return**
Canadian Stocks	XIU.TO	5.29%
US Stocks	SPY	5.29%
Global Equity	VEU	6.29%
Canadian Inflation-linked Bonds	XRB.TO	2.36%
Canadian Government Bonds	XGB.TO	2.76%
Canadian Corporate Bonds	XCB.TO	3.71%
Global Real Estate	CGR.TO	4.34%
Global Infrastructure	CIF.TO	4.34%

Table 2

As we will see in the results, our simulated long-term bond return is relatively low, which makes the expected returns on assets also low. This assumption however is in accordance with the idea that returns on financial markets will be more modest in the future. Also, this will produce more conservative results.

The total portfolio return is then computed using different portfolio weights, including the original allocation and different optimal portfolio allocations. Like the OCA does, 0.3% is subtracted from this return to account for transaction costs and management fees.

The overall asset value each year is the market value of the portfolio plus the contribution minus the pension payment for that year. For simplicity, the model assumes that all payments are credited or debited at the end of each year.

4.4 SURPLUS

The actuarial surplus is simply the asset value minus the present value of pension payments, which can be calculated for each of the 20 years simulated. The surplus on the last year simulated (2031) is called the terminal surplus. From these surpluses, the funding status can be determined. If the surplus is positive, the fund is able to meet its obligations in the long run and is said to be fully funded; if it is negative, the fund is said to be in deficit or to be insolvent and the government must step in and fill in the funding gap. An interesting measure is the number of iterations that result in a deficit divided by the total number of iterations, which is the estimated probability of insolvency.

4.5 OBJECTIVE

After having modeled the liability, the asset and the surplus of the pension fund of the Public Service, the next step of the LDI approach is to select an objective for the fund. Like it was mentioned by many authors including Delong *et al*. (2008), Leahy (2011), Ostaszewski (2011) and Cox *et al*. (2013), to maximize the actuarial surplus or similarly the portfolio return is not necessarily a valid goal because it induces excessive risk taking through a large allocation to risky assets. In addition, carrying a large surplus does not benefit directly the members of the fund (except if it is used to decrease contribution rates). As inspired by Delong *et al*. (2008), the minimization of the terminal surplus variance has been selected as an objective. In fact, this objective is focused on reducing the volatility of the surplus, which goes in accordance with the spirit of LDI. To reach this goal, the risk factors affecting the liability must be hedged by the portfolio of assets. When combined with a positive surplus this objective results in a low probability of insolvency. However, if the expected return is so low that the funding status erodes over time, the stability of the surplus is useless. Also, as Service and Sun (2004) pointed out, the simple use of surplus variance as a measure of risk is disputable as it does not take into account the non-normal features of the distribution of the asset and liability. A better measure is the probability of insolvency, which takes into account the whole distribution. To minimize the probability of insolvency, not only the asset must be on average larger than the liability but the distribution of the asset must dominate the one of the liability. In other words, the higher moments of the distribution of the asset (namely the skewness and the kurtosis) must also dominate the ones of the liability (Service and Sun, 2004). Moreover, the main goal of a pension fund is to pay its pensioners and to do so, it must remain solvent (Delong *et al*. 2008). Solvency should therefore be its main concern, which makes the minimization of the probability of insolvency the prime objective. It is important to note however that obtaining a probability of insolvency of zero is theoretically impossible because of potential extreme events. Also, to persist in reducing the probability of insolvency could mean growing a large surplus, which would surely reduce the risk of insolvency but would be very inefficient and costly to maintain. As a result, a small amount of risk must be tolerated. In this thesis, to obtain a deficit in 5% of the scenarios is considered a bearable risk to preserve efficiency. The overall objective of the fund is therefore to obtain a probability of insolvency equal or less to 5% and to minimize the terminal surplus variance.

4.6 SENSITIVITY TO RISK FACTORS

The third step of the LDI approach is to assess the sensitivity of the liability, the asset and the surplus to each of the risk factors. This can be done by performing multiple linear regressions of the valuation measures on the risk factors or by analyzing their scattergraphs. It can also be done by changing the factors in the model (Meder and Staub, 2006). Because most of the factors in our model have a random component, each iteration produces a different potential scenario. With the approximated density function of the liability and the surplus, the probability of insolvency and the value-at-risk can be found. Like Boyer *et al*. (2011) did, the nominal and relative value-at-risk (VaR) are calculated to assess the risk generated by a factor. As mentioned by Choudhry (2007: 834), "VaR is the expected loss of a portfolio over a specified time period for a set level of probability". It depends on the volatility and the correlations of the underlying variables (Choudhry, 2007). As a result, the VaR allows us to assess tail events and to evaluate the magnitude of a risk. In this thesis, the VaR is calculated for each factor individually and for all factors jointly for a probability of 5% and a time horizon of 20 years.

4.7 OPTIMAL ASSET ALLOCATION AND FUNDING STRATEGY

Finally, the last two steps of the LDI approach are to develop the optimal asset allocation and funding strategy. This starts by building the optimal portfolio so that the asset exceeds the liability in most scenarios while keeping a reasonable level of risk. A portfolio of Canadian government bonds could be built to immunize the fund against small changes in interest rate by matching the duration of the asset with the duration of the liability. Conversely, a portfolio with different asset classes could be created. In order to avoid concentration risk, the maximum allocation in one asset class is set at 40%. To make sure that funds are invested in all of the eight asset classes, the minimum allocation is set at 2%. These higher and lower bounds are selected arbitrarily but are supported by the fact that the fund's portfolio weights fluctuate between 7% and 33% so 2% and 40% allows more flexibility while remaining reasonable. To decide on the allocation of the diversified portfolio, a Markowitz portfolio optimization is performed. A drawback of this technique, however, is that it focuses only on the two first moments of the distribution. A multi-moment optimization technique could be used instead to match the four first moments of the distribution of the asset with the distribution of the liability. In this thesis, polynomial goal programming (PGP) is applied to reach that goal. The first step of this method is to find the three optimal portfolios that respectively maximize expected return, maximize the skewness and minimize the kurtosis per unit of variance. Because these optimal portfolios are mutually exclusive, the second step consists in finding the asset allocation that minimizes the distance to the previously found optimal expected return, skewness and kurtosis (Davies *et al*., 2009). In order to create a portfolio with a return distribution close to the distribution of the liability, this method is altered by minimizing the skewness and the kurtosis only. In fact, because the skewness of the liability is equal to -0.11 (or -0.10 when discounting at the long-term bond return) and the minimum reachable skewness on the portfolio of assets is 0.20, the skewness is minimized instead of being maximized.

As it is demonstrated in the results, the fund is in deficit according to our model. Different strategies to close the funding gap are therefore evaluated. Some solutions that are analyzed include: filling the gap with funds from the government, delaying retirement age or increasing contribution rates. But before applying the LDI approach, the validity of the model's framework is assessed and the model risk is analyzed and quantified.

5. RESULTS

5.1 FRAMEWORK VALIDITY

First of all, to evaluate the validity of the model's framework, we compute the liability ascribable to the retirement pensioners with the same assumptions as the OCA (including the inflation rate, interest rate, earnings growth, mortality rates and portfolio return assumptions). This valuation is then compared to the liability disclosed in the 2011 actuarial report.

The liability resulting from the solvency valuation (which discounts only payments owed to retirement pensioners) is equal to $61,991 million whereas the liability ascribable to the retirement pensioners disclosed in the actuarial report is equal to $62,110 million. This 0.19% difference can be attributed to the fact that the OCA uses the exact distribution of members while we use subgroups of age and seniority and assume median age, median seniority and average pensionable earnings for the 2011 members. Because this difference is very small however, we assume its impact to be negligible. This suggests that the model's framework is reliable.

One could argue that the model is biased and estimates a liability that is close to the value obtained by the OCA because pensionable earnings are fitted to the pensions disclosed in the actuarial report (see section 4.2.3). It is important to realise however that only the pensions of the 65+ age group are fitted to account for their lower earnings at retirement and that this adjustment is made before indexation and discounting. Also, it is performed before applying mortality rates and adding the perpetuity. This said, if we remove the earnings adjustment, we obtain a liability that is 5.3% higher, which result in a 5.1% discrepancy with the liability disclosed in the actuarial report. Overall, this shows that the reduction in pensionable earnings is appropriate and improves the solvency valuation.

The going concern valuation, which assumes the fund continues operating for the projection horizon (20 years) and then goes into perpetuity, does not include all active contributors because not all of them will have retired by 2031. The going concern liability of $91,982 million is comprised of $61,991 million for the retirement pensioners and $29,991 million for the contributors. Because the liability ascribable to contributors is equal to $66,772 in the 2011 actuarial report, we can assume that the going concern valuation includes about half of contributors. Nonetheless, this valuation method is useful in assessing the sensitivity of the liability on a longer period of time. More specifically, it allows evaluating the impact of the growth in the number of pensioners or the growth in the pensionable earnings.

Secondly, these valuations are used as a benchmark to evaluate the impact of using stochastic processes to model the risk factors instead of using the OCA assumptions. The risk of using an inadequate process to explain a variable is called the model risk. The stochastic processes used in this thesis along with their estimated parameters are summarized in Table 15 in the appendix.

5.2 MODEL RISK

When using the selected stochastic processes for the inflation rate, interest rate, earnings growth, mortality rates and portfolio return, the liability resulting from the solvency valuation is equal to $63,167 million, which is 1.90% higher than what is obtained with the OCA assumptions. The liability resulting from the going concern valuation is equal to $98,547 million, which is 7.14% higher than what is obtained with the OCA assumptions. As we can see, the going concern valuation is subject to a larger model risk because the earnings growth process takes effect only when including active members in the calculation. Also, the mortality process affects more members and the interest rate process discounts more payments so that the choice of process has a larger impact. While using stochastic processes instead of the OCA assumptions allows us to introduce volatility in the risk factors, it creates a potential modeling error. For this reason it is important to keep in mind the model risk when analyzing the results of this thesis.

The liability values are discussed further in section 5.3. In this section, we first analyze the model risk associated to each risk factor modeled. To do so, the OCA assumptions regarding the risk factor are replaced one at a time by the parameterized stochastic process. We then examine the growth of the asset and liability over the 20 years simulated.

5.2.1 Inflation rate

Because the Bank of Canada committed to keep inflation between 1 and 3%, the OCA assumes the CPI will increase by 2% from 2012 to 2017 and reach 2.3% from 2020 onward. Our model assumes the CPI follows a geometric Brownian motion with a standard deviation of 0.73%. The iterations of the simulated process oscillate mainly between 1% and 3% since it is parameterized on data from the last two decades in which inflation fluctuated within this range. Figure 4 depicts the historical inflation rate and one iteration of the inflation process to demonstrate the fit of the forecast to past data.

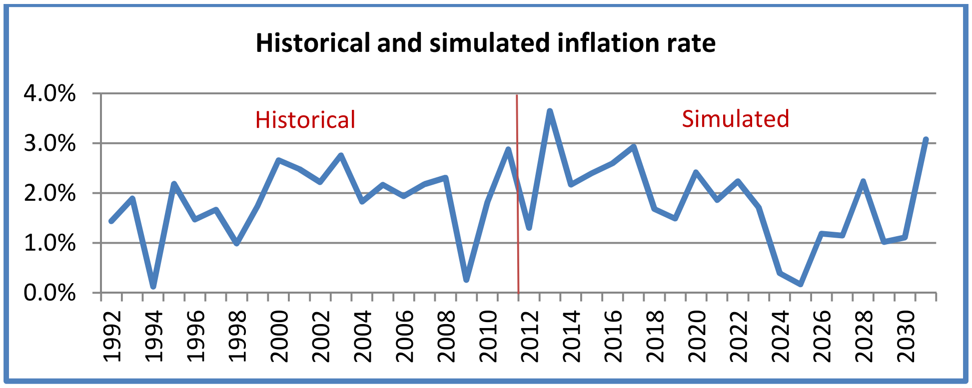

Figure 4

The average process revolves around a mean of 1.85%, which is lower than what the OCA assumes. This is shown in Figure 5.

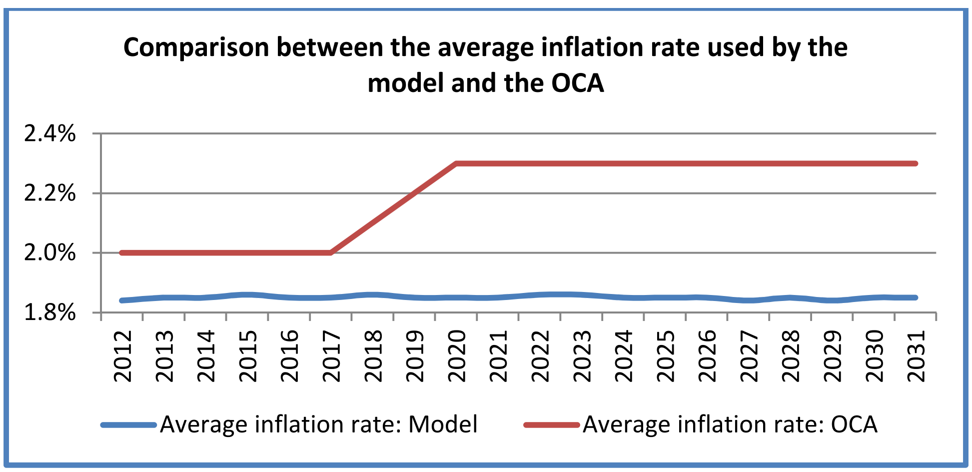

Figure 5

The lower inflation rate results in lower indexed pensions and a lower average liability. In fact, when relaxing the OCA assumption and allowing the inflation rate to follow a geometric Brownian motion, the liability from the solvency valuation and the going concern valuation drops by 2.15% and 2.16% respectively.

In the original model, the inflation process is parameterized on the historical Canadian CPI level from 1991 to 2011. To use this data set implies that the Bank of Canada will continue to control inflation rates in the future. If this is not the case, we should expect higher inflation rates. To evaluate the impact of such scenario, the inflation process is parameterized on CPI levels from 1914 to 2011 (which includes years where there was no inflation targeting). Figure 6 displays an example of one iteration of each process. As it is shown, the extended data set results in a more volatile inflation process.

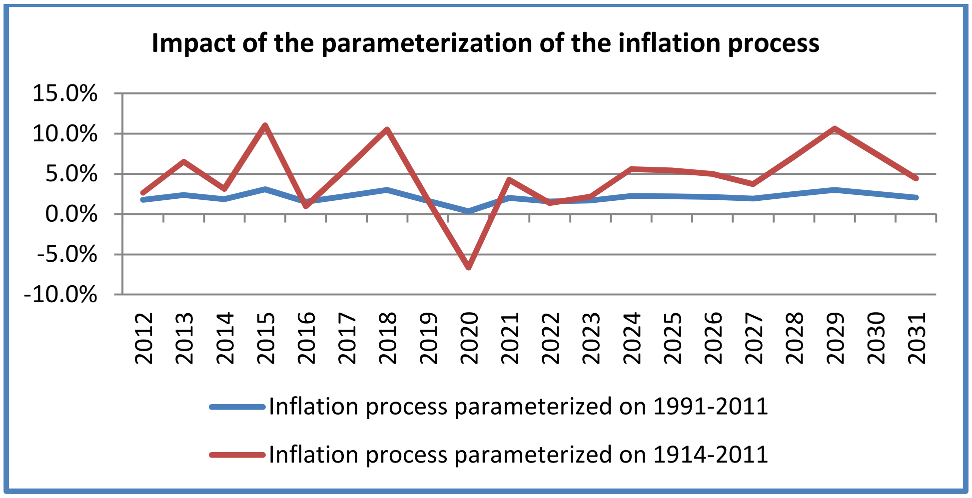

Figure 6

With this higher inflation rate process, the solvency valuation produces a liability 15.15% larger than when the 1991-2011 data set is used. As for the going concern valuation, it produces a liability 13.43% larger. This clearly shows the sensitivity of the valuation to the parameters of the processes. It also shows that if the Bank of Canada decides to stop controlling inflation, the resulting liability could be significantly larger and jeopardize the funding status of the fund.

5.2.2 Interest rate

The interest rate is the yield on the long-term Government of Canada bond. The OCA assumes this one starts at 1.8% and gradually increases to 2.7% by 2017, whereas the model assumes it follows a Vasicek process with a long-term average monthly rate of 0.1%.

Figure 7 shows the fit between the annualized historical interest rate and one iteration of the annualized Vasicek process. We can see the process is capped at zero because of our non-negative constraint. Also, because θ is relatively small (0.0083) and σ is relatively large (0.0001) the process does not revert to its mean rapidly.

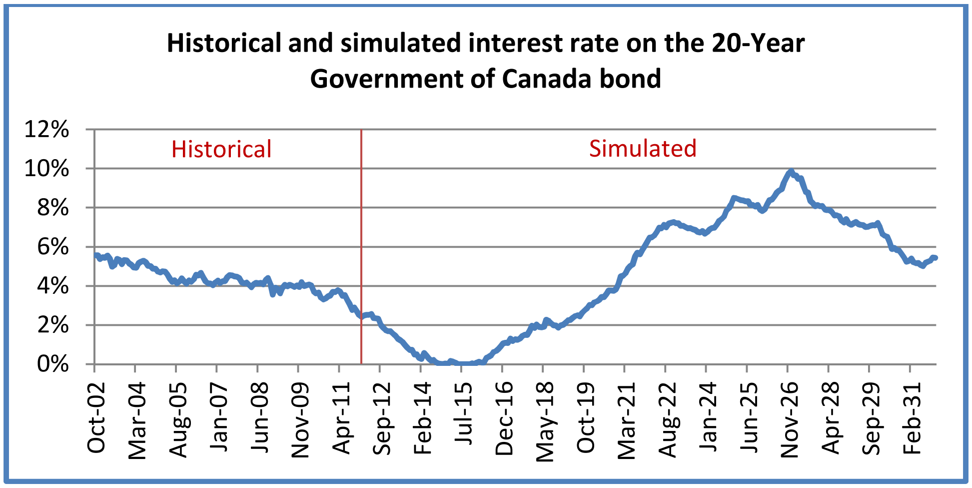

Figure 7

The average process starts at 2.5% and settles around 2.6%. Again, because θ is small, the process does not revert to its annualized mean of 1.2% within the simulation time frame. This is displayed in Figure 8 along with the OCA assumption.

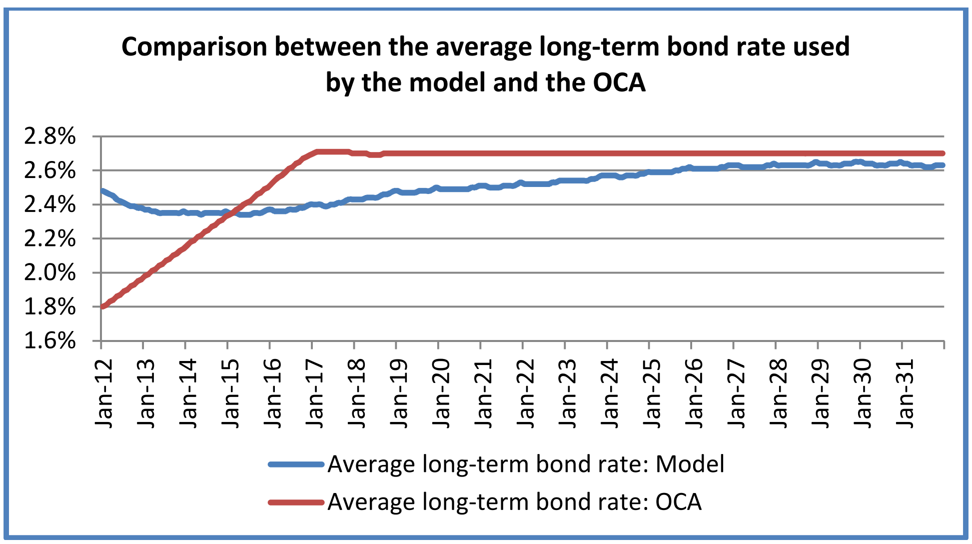

Figure 8

The interest rate is then used to compute the discount rate that the fund uses to discount pension payments and obtain the liability. Both the model and the OCA assume the discount rate is equal to the expected portfolio return, which is a weighted average of the expected returns on risky assets based on the portfolio weights. These expected returns are found by adding risk premiums to the long-term government bond yield. The model uses the same risk premiums, transaction costs and management fees as the OCA, but it allows the long-term government bond yield to follow a Vasicek process. When adding the risk premiums to the long-term bond yield assumed by the OCA, the discount rate goes from 5.47% in 2012 to 4.16% by 2031. The model's discount rate starts at 6.15% in 2012, drops to about 4% by 2017 and remains around this level for the rest of the simulation horizon. This is depicted in Figure 9.

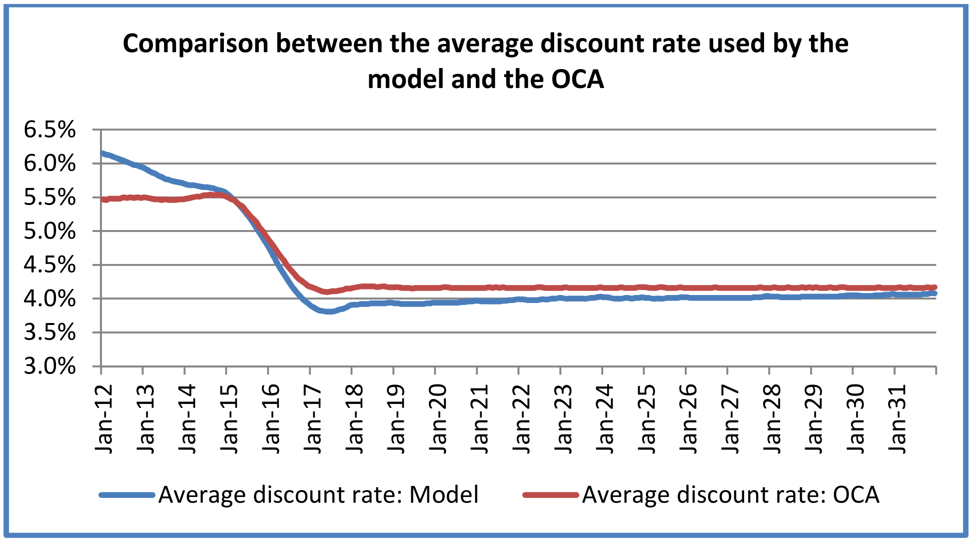

Figure 9

The significant fall in discount rate over time is due to the OCA's expected risk premiums that are projected to shrink by about 50% by 2017. This fall is also in accordance with other Canadian public funds, who saw their average discount rate decrease from 8% in 1994 to 6% in 2010 (Andonov *et al*., 2013) and from 5% in 2011 to 4% in 2012 (Morneau Shepell, 2013). This decrease is however far less important than the drop in the long-term Treasury yield, which went from about 8% to about 3% between 1994 and 2010. This implies that funds should perhaps use an even lower discount rate.

Because our discount rate is on average lower than the OCA's, the discounted payments and the resulting liability are larger. In fact, the liability values from the solvency valuation and the going concern valuation are respectively 1.21% and 1.88% higher. As mentioned earlier, the going concern valuation is affected more severely because more pensions are discounted. The choice of using a Vasicek process instead of the OCA assumption results in a relatively low model risk compared to other variables but a risk is still present. The Vasicek process is indeed a short-term interest rate model but we use it to simulate long-term interest rates. Although it corresponds to the right mean-reverting dynamics more appropriate models could be used instead.

As mentioned in the methodology section, although the discount rate corresponds to the expected return on the portfolio of assets, the simulated portfolio return can differ from this expectation. For this reason, the two processes are modeled separately. As we will see in section 5.2.5, the simulated portfolio return is more stable over time and is on average slightly higher than the expected portfolio return.

Overall, the choice of discount rate is extremely important in the valuation of a fund. As it is explained by Andonov *et al*. (2013), this choice is often guided by accounting standards and regulations and differs from one country to another. Like the pension fund of the Public Service of Canada, U.S. public pension funds are allowed to use the expected rate of return on their portfolio of assets to discount pension payments. On average, they use a discount rate of 8%. Inkmann and Blake (2004) add that discount rates determined as an exogenous variable do not take into account the risk of defaulting on pension payments, which is influenced by the allocation of assets. The higher the allocation in risky assets, the higher the discount rate should be to reflect the increased volatility of the assets and the funding gap. A major problem arises from this method however. As Inkmann and Blake (2004) and Andonov *et al*. (2013) point it out, pension funds that use this discounting method have a strong incentive to invest in riskier asset classes like equity to increase the expected portfolio return and the resulting discount rate, which in turn reduces the value of the liability and allow disclosing a better funding status. On average, these funds invest 10% more in risky assets compared to funds in Canada and Europe, who are mostly under different rules. Accordingly, our results show this relationship. Different portfolio allocations result in different discount rates and the allocations that include more equity result in higher discount rates. Figure 10 shows the discount rate with the original portfolio weights and the minimum variance portfolio weights that allocate far less funds to equity. We elaborate on the matter in section 5.5.3.

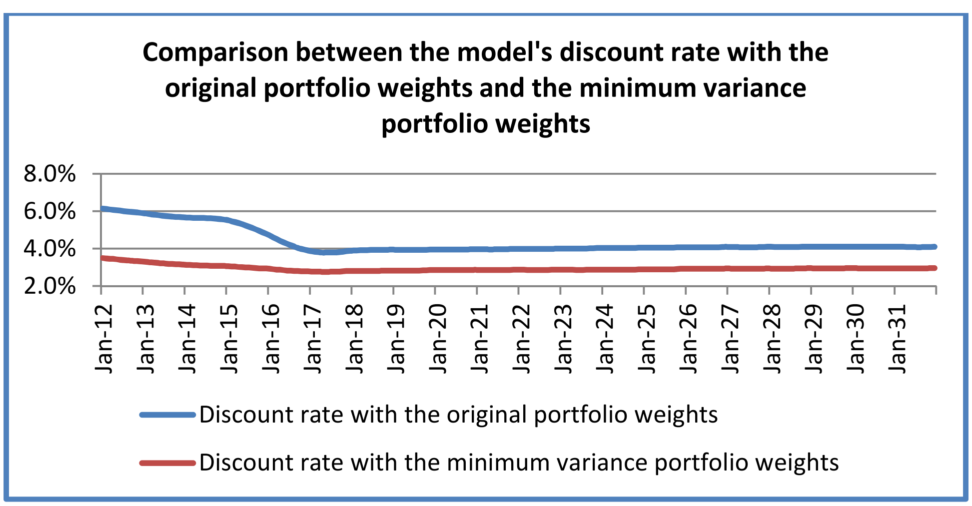

Figure 10

As opposed to the U.S., European pension funds are required to use long-term bond yields to discount pension payments, which supresses the relationship between the asset allocation and the discount rate. On average, they use a discount rate of 4% (Andonov *et al*., 2013). According to Andonov *et al*. (2013: 11), “in Canada, pension fund accounting standards generally require that the discount rate be selected based on market yields of high-quality corporate debt instruments with cash flows that match the timing and amount of the expected benefit payments.” In fact, the International Accounting Standard (IAS) 19 and the International Financial Reporting Standards (IFRS) provide general guidance for the selection of the discount rate assumption but the exact methodology remains at the discretion of the fund. Whereas the Public Service Pension Fund uses the expected portfolio return, other Canadian funds use the rate of return on long-term bonds. This is why the return on the 20-year Government of Canada bond is also applied as the discount rate in this thesis. When using this discount rate, the liability from the solvency valuation is equal to $78,393 million, which is 24.10% higher than when discounting at the expected portfolio return. The going concern valuation results in a liability that equals to $127,352 million, which is 29.23% higher than when using the other discount rate. As we can see, the choice of discount rate has a large impact on the magnitude of the liability.

According to Meder and Staub (2006) however, to discount at the long-term bond return disregards the uncertainty in the earnings growth and in the mortality rates. In fact, these risks are systematic and should therefore be accounted for by a risk premium. This brings us to the third method used to discount the pension payments, which differentiate between active and inactive members. As proposed by Wouters (2008), inactive members’ payments are discounted at the long-term bond rate, whereas active members’ payments are discounted at the expected portfolio return to account for the greater risk due to the uncertain seniority and pensionable earnings. Because the solvency valuation does not include active members, only the long-term bond discount rate is applied, which results in the same liability as when we use this discount rate only. The going concern valuation results in a middle-of-the-road liability equal to $112,153 million, which is 13.81% higher than when discounting at the expected portfolio return but still lower than when discounting at the long-term bond return only. Because the going concern valuation includes the second discount rate whereas the solvency valuation does not, the two valuations are harder to compare and we cannot isolate the impact of changing the risk factors from one valuation measure to the next. For this reason, this method is rejected and only the expected portfolio return and the long-term bond return are used to discount payments in the following sections.

5.2.3 Earnings

Whereas the OCA assumes an annual increase in earnings that starts at 1.75% in 2012 and reaches 3.5% by 2025, the model assumes earnings follow a geometric Brownian motion with a mean equal to 3.5% and a volatility based on historical growth, which results in a standard deviation of 4.19% for males and 4.49% for females. Figure 11 displays the historical male earnings growth and one iteration of the earnings growth process.

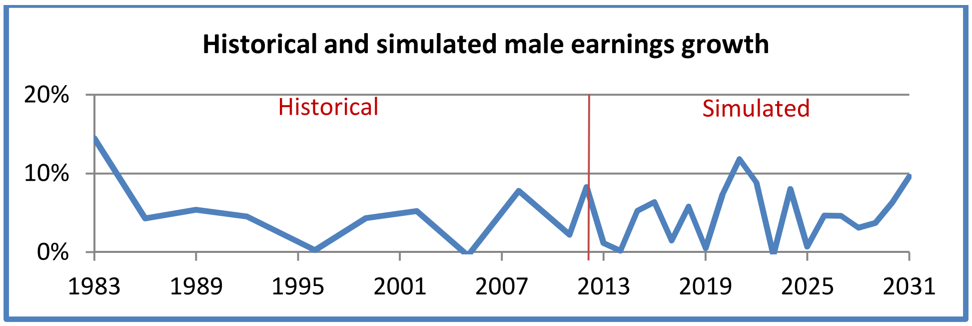

Figure 11

Figure 12 displays the average earnings growth processes for males and females.

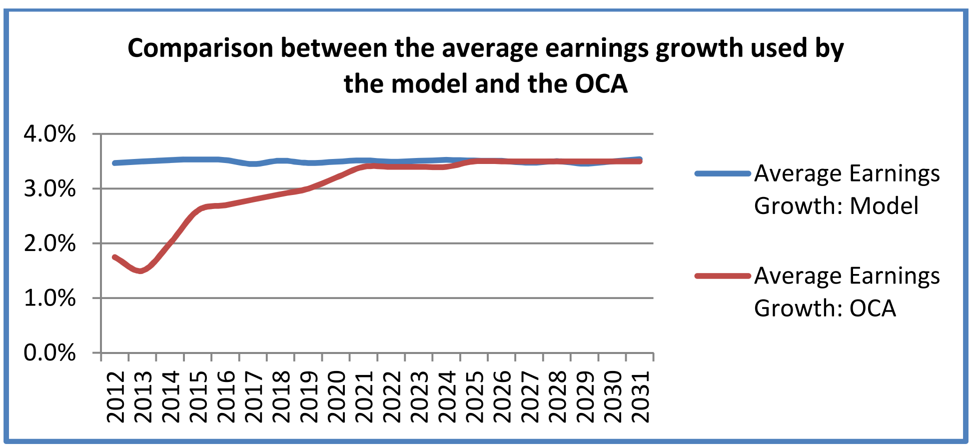

Figure 12

Because the model assumes higher earnings growth than the OCA until 2025, the pensionable earnings of active members are expected to be larger in the future, which makes the going concern liability 3.81% higher than when using the OCA assumption.

Other assumptions regarding the earnings include the YMPE, the AMPE and the calculation of the pensionable earnings. The OCA assumes the year's maximum pensionable earnings (YMPE) will increase from 2.5% in 2013 to 3.6% in 2021. It also assumes that the average maximum pensionable earnings (AMPE) will go from 1.5% in 2013 to 3.5% in 2025. The YMPE is used to determine the contribution rate and the AMPE is used for the CPP adjustment. The model assumes the YMPE and the AMPE will increase annually by 2.8%. This assumption has a negligible impact on the magnitude of the liability. At worst, it underestimates the solvency valuation liability by 0.0003% and the going concern valuation liability by 0.0165%. Finally, when calculating pension amounts, the OCA uses the highest 5-year average earnings as the pensionable earnings, whereas the model uses the last earnings before retirement. In the model, average earnings are always growing, which makes the last earnings before retirement the largest. By using the highest earnings, the model overestimates pensionable earnings by an average of 6.5%. It has no effect on the solvency valuation as this one does not include new pensioners but it overestimates the going concern valuation by 3%.

5.2.4 Mortality rates

The OCA assumes mortality rates will decline at a decreasing rate for 20 years and then level off for the following years, whereas we assume mortality rates will follow the Lee-Carter model and decrease at a random rate, which is stationary over time. Figure 13 shows the average mortality rate over time.

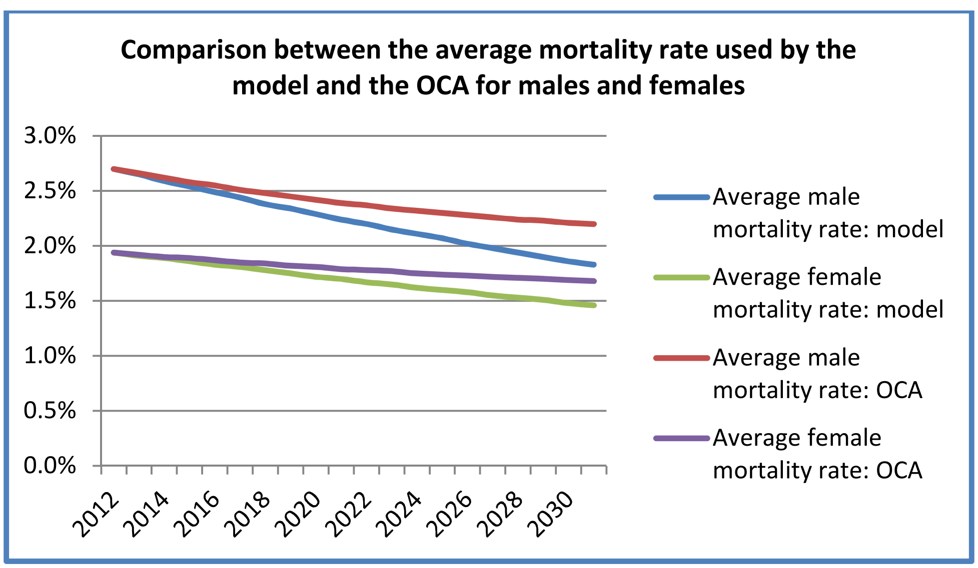

Figure 13

On average, the calibrated process generates better improvements, which results in mortality rates 0.14% lower than the OCA assumes. This is due to the fact that the OCA assumes decreasing improvements, which implies a cap in mortality rate improvements, whereas the Lee-Carter model does not make such assumption, which implies unlimited improvements in mortality rates. As we can also see in Figure 13, the OCA and the model both assume lower mortality rates and improvements for females. This is because male mortality rates are assumed to “catch up” female rates in the future.

Whereas the OCA assumes less important improvements for older pensioners, the Lee-Carter model assumes that the expected improvements will follow their historical average for each age group. This results in average improvements ranging from 0.98% for people aged 25-29 to 2.96% for people aged 70-74. Figure 14 displays one iteration of the male mortality improvements. It shows that the improvements do not seem to follow any logical order across age group. It also shows that the improvements are stationary over time and perfectly correlated across age groups. These are all limitations of the Lee-Carter model.

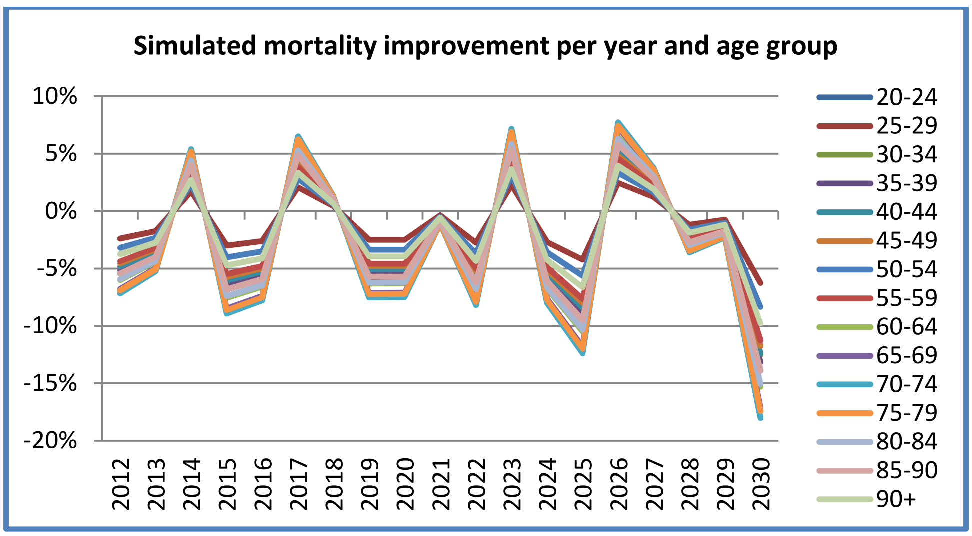

Figure 14

Overall, because our model forecasts better improvements, the simulation of mortality rates with the Lee-Carter model instead of using the OCA assumption results in a liability that is 5.43% higher for the solvency valuation and 7.30% higher for the going concern valuation. Because the going concern valuation includes more members, the choice of mortality model has a larger impact, which translates to a larger model risk. Also, to assume constant improvements instead of diminishing improvements makes a substantial difference in the long run. For this reason, the model risk associated to the mortality rate process is relatively larger than the other risk factors.

5.2.5 Portfolio return

As mentioned above, the simulated portfolio return differs from the expected portfolio return used to discount pension payments. In fact, the two processes are simulated separately. Although the same risk premiums are used in the two processes, they are included differently. The risk premiums for each asset class and each year are added to the long-term bond returns to obtain the expected asset returns each year, which are combined to produce the expected portfolio return that is used as a discount rate. On the other hand, the average expected returns on the different asset classes are used to adjust the returns from the random historical sample that simulates future portfolio returns. Because of this adjustment, the time window used to select the asset returns have a small impact, which produces almost identical portfolio returns and optimal allocations. The resulting portfolio return process has a slightly smaller long-term average than the expected portfolio return (3.69% instead of 4.06%). This can be seen in Figure 15.

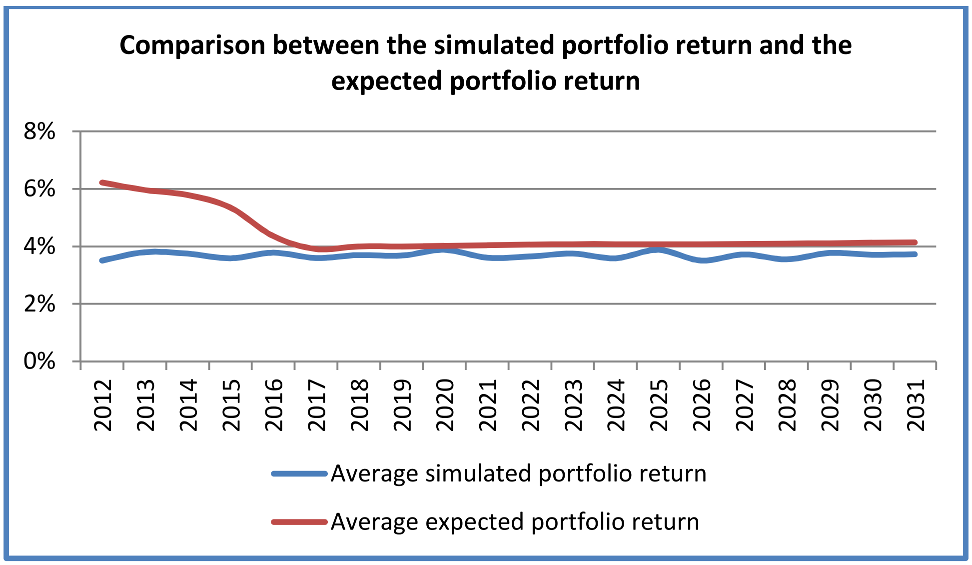

Figure 15

The model risk created by the addition of the random portfolio return to the model can be estimated by looking at its impact on the asset value. This one decreases by only 1.24% after 20 years of compounding.

Also, because the random selection of historical returns does not take into account time series properties, we verify if our method is appropriate by testing some of those properties. First, we test for autocorrelation by regressing historical asset returns on their lagged value. Statistically significant autocorrelation is found in U.S. and foreign stocks, corporate bonds, global real estate and global infrastructure. Then, different moving averages are used to evaluate if there is cyclicality in the mean, variance and covariance of historical returns. The tested variables are all relatively stationary except for the variance of historical returns on stocks and global real estate, which are higher in 2009 (after the financial crisis) and decreased over time. Figure 16 shows the relative stationarity of historical returns and Figure 17 shows the volatility bursts that affected stock returns and real estate.

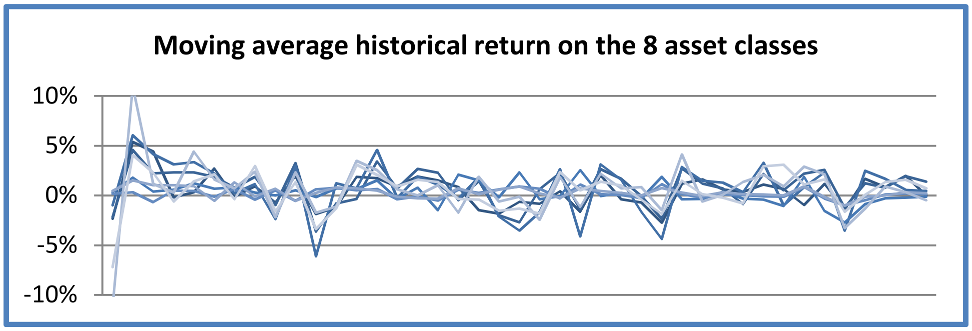

Figure 16

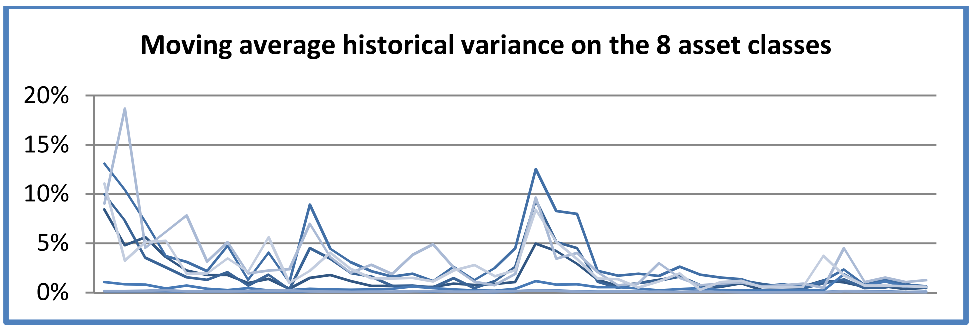

Figure 17

As we mentioned earlier, these time series properties cannot be replicated by our simulated asset returns, which contributes to the model risk. Other models such as GARCH models may simulate these returns better, but this is outside the scope of this thesis.

5.2.6 Correlations

An important feature of our model is the inclusion of linear correlation between the inflation rate, interest rate and revenue growth processes. Table 3 shows the historical and modeled coefficients of correlation between the different variables.

Comparison of the coefficients of correlation of the historical and modeled risk factors

	Historical coefficients of correlation			
	Inflation rate	Interest rate	Earnings growth	Mortality rates
Inflation rate		-0.67	0.60	0
Interest rate	-0.67		-0.06	0
Earnings growth	0.60	-0.39		0
Mortality rates	0	0	0	
	Modeled coefficients of correlation			

Table 3

As we can see, the only modeled coefficient of correlation that is different than its historical value is the coefficient for the interest rate with the earnings growth, which is still relatively small and negative but not equal. This is not surprising because these two processes are not directly correlated to one another. They display an indirect correlation because they are both correlated to the inflation rate. As we expected, because the mortality process is modeled independently we obtain a correlation of zero with the other risk factors.

5.2.7 Asset and liability value

Overall, when combining the different stochastic processes that simulate the risk factors, we obtain a very similar value for the asset and a different value for the liability than the OCA. This is depicted in Figure 18 and Figure 19 for asset and liability respectively.

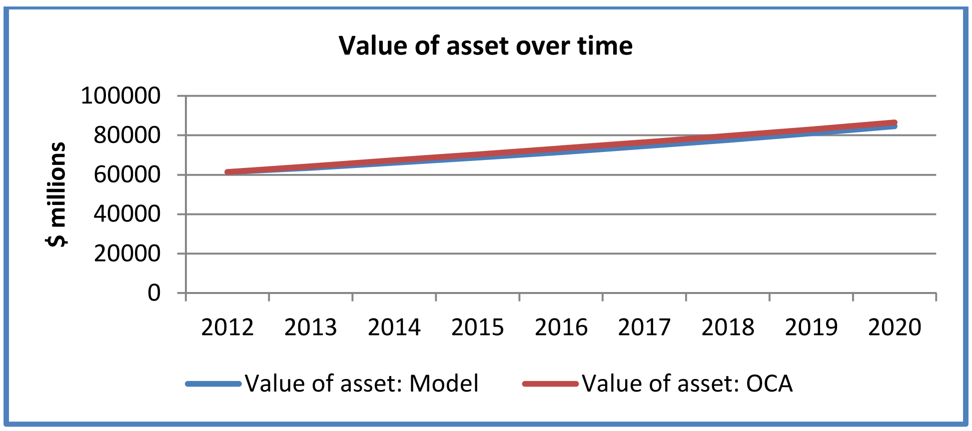

Figure 18

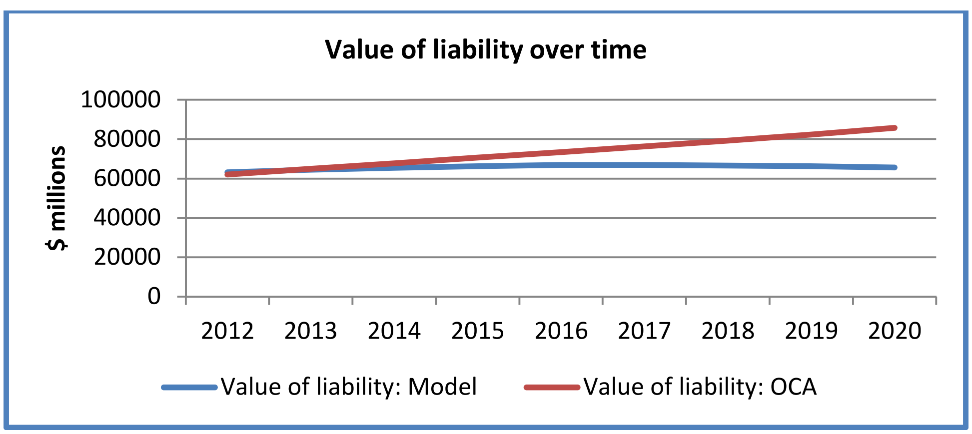

Figure 19

The asset starts at the same value as the OCA because we use their 2011 asset value as our starting point. The two processes are considered to be relatively similar since our model predicts an annual growth in asset that is only 0.27% smaller than what the OCA expects. As for the annual growth in liability however, it is 3.63% smaller than what the OCA predicts. In order to test the impact of our assumptions, the liability is increased by 3.63% each year. This results in a terminal surplus that is 42.58% smaller than without the liability adjustment. Again, as we can see, our choice of model has a large impact on the valuation of the liability and the resulting actuarial surplus. In order to take into consideration the model risk, we evaluate the impact of the liability adjustment on the optimal solution in section 5.6.5.

After having assessed the model risk related to the different simulated factors, the next section exposes the detailed results of the valuation of the fund. This corresponds to the first step of the LDI approach.

5.3 VALUATION OF THE FUND

As mentioned in the above section 5.2, the liability resulting from the solvency valuation is equal to $63,167 million. As for the initial asset value, it is equal to $61,414 million. This creates an initial deficit of $1,753 million. If we discount pension payments at the long-term bond return instead of the expected portfolio return, the solvency valuation liability is equal to $78,393 million, which creates an initial deficit of $16,979 million. The distribution of the asset and liability is shown in Figure 20 and the distribution of the initial surplus is shown in Figure 21.

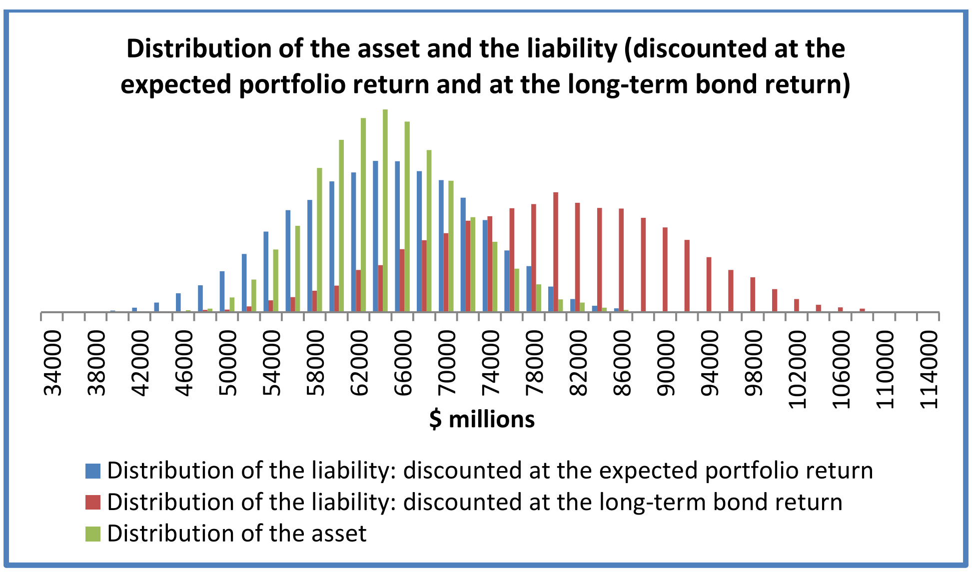

Figure 20

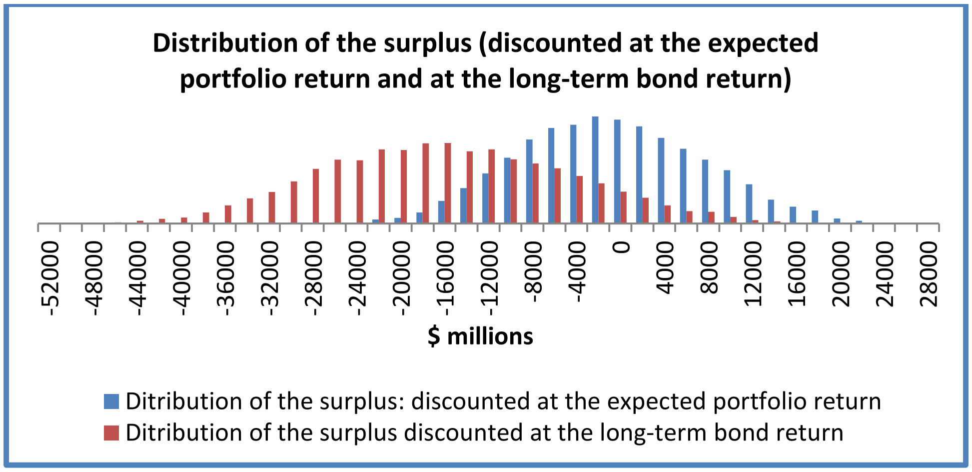

Figure 21

As we can see in Figure 20, the distribution of the value of the liability with both discount rates partially dominates the value of the asset. The mismatch is even greater with the more conservative long-term bond discount rate. This results in a larger deficit on average, which can be observed in Figure 21. From these distributions, we can get the probability of insolvency, which is estimated by the number of iterations that result in a negative surplus divided by the total number of iterations. Because the average liability is far greater than the average asset, our model finds an initial probability of insolvency of 59% when discounting at the expected portfolio return and 93% when discounting at the long-term bond return.

Alternatively, the going concern valuation results in a liability of $98,548 million if we discount at the expected portfolio return and $127,352 million if we discount at the long-term bond return.

Whereas the fund is initially in deficit, its financial situation improves over time. In fact, the terminal surplus has a mean of $97,877 million and a standard deviation of $61,021 million when discounting at the expected portfolio return and it has a mean of $94,324 million and a standard deviation of $60,997 million when discounting at the long-term bond return. The growing surplus can be partly explained by the fact that contributions become larger than pension payments over time. This is depicted in Figure 22.

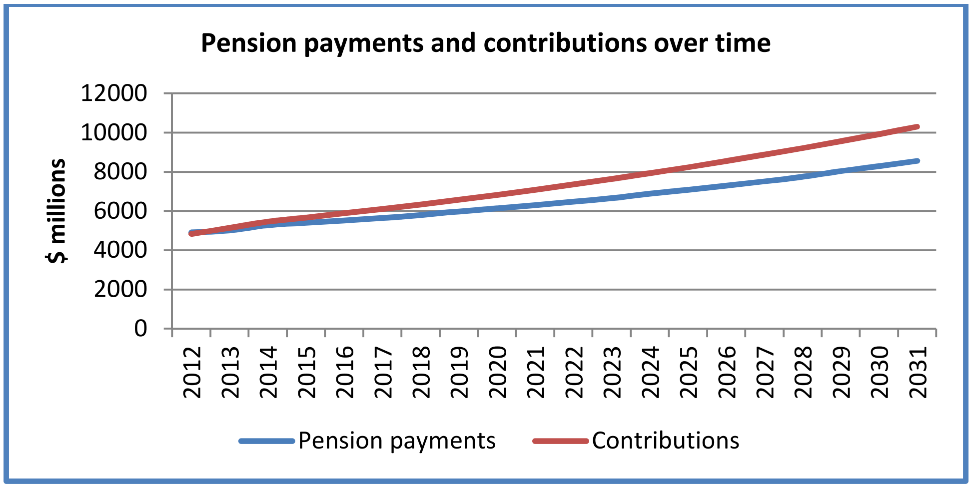

Figure 22

Whereas the pension payments and the contributions are initially similar, the contributions grow 1.1% more than the pension payments annually.

Also, because we assume a relatively low inflation rate, the pension payments are indexed moderately whereas the contribution payments, which are a function of earnings, increase at a greater rate.

The overall financial situation of the fund over time is displayed in Figure 23 when discounting at the expected portfolio return and in Figure 24 when discounting at the long-term bond return.

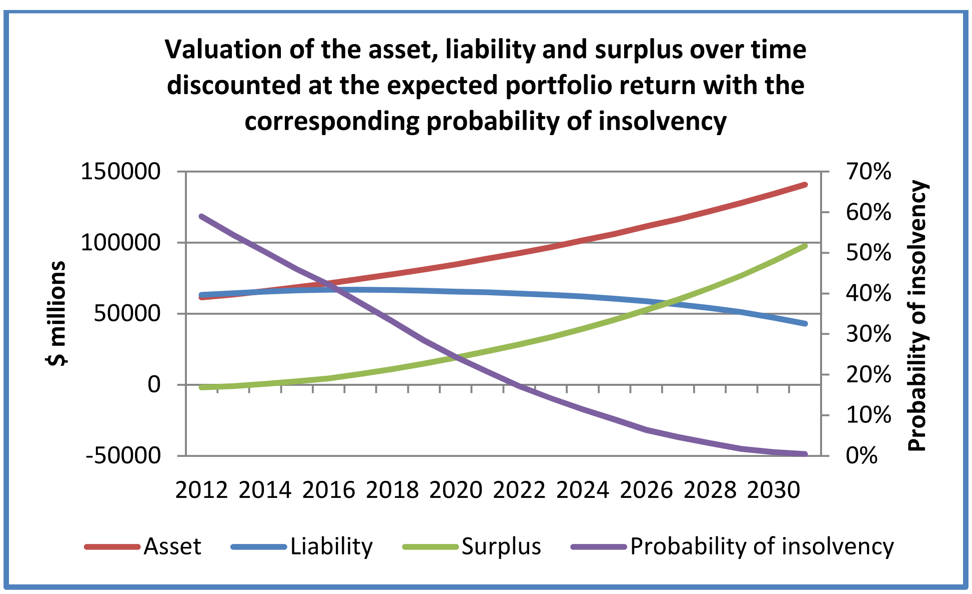

Figure 23

As we can see, when using the expected portfolio return to discount pension payments, the surplus becomes positive within 3 years and the probability of insolvency becomes lower than 5% by the year 2027.

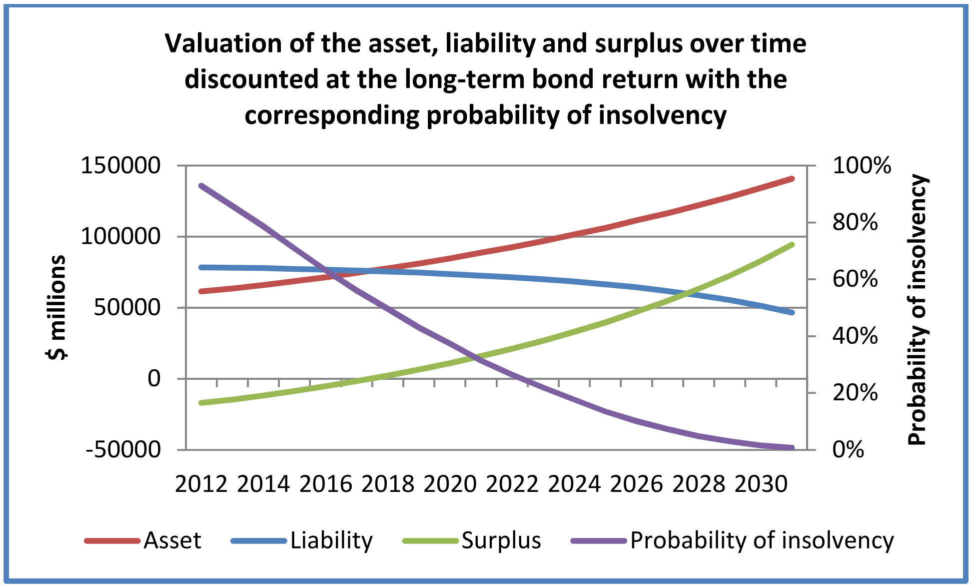

Figure 24

When using the long-term bond return to discount payments instead, the surplus becomes positive within 7 years and the probability of insolvency becomes lower than 5% by the year 2028.

Keeping in mind that the chosen objective of the fund is to maintain a probability of insolvency of at most 5% and to minimize the terminal surplus variance, the fund has to not only find the portfolio of asset that hedges its liability the best but it also has to find a strategy to fill the funding gap. As we saw in this section, according to our model and with the current asset, the fund has a deficit, which is even greater when discounting at the long-term bond return. Before building the optimal asset allocation and funding strategy however, we must evaluate the sensitivity of the asset, the liability and the surplus to the chosen risk factors. This corresponds to the next step of the LDI approach, which will help us to find the optimal asset allocation.

5.4 SENSITIVITY TO RISK FACTORS

5.4.1 Regressions

We start by performing a multiple linear regression of the terminal surplus on the standardized risk factors. The factors are standardized to allow comparing their impact on the surplus. In fact, a drop of 1% in mortality rates is not comparable to a drop of 1% in earnings growth but the standardization adjusts for their mean and standard deviation to make them comparable. From their regression coefficient, the risk factors can then be ranked based on the extent at which they influence the surplus. Table 4 shows the results of the regression. All the coefficients are statistically significant at a 99% confidence level.

Regressions of the terminal surplus on the standardized risk factors		
Risk factor	**Regression coefficient**	**Rank**
Inflation rate	-0.21×10^{10}	3
Interest rate	0.15×10^{10}	4
Earnings growth	0.92×10^{10}	2
Mortality rates	0.08×10^{10}	5
Portfolio return	5.58×10^{10}	1

Table 4

It is interesting to see that the interest rate ranks second to last whereas the portfolio return is the factor that affects terminal surplus the most. This is not surprising however since the terminal surplus is not discounted to year 0. As a result, the interest rate, which affects directly the discount rate, has little impact on the terminal surplus. Conversely, the asset value that is used to compute the terminal surplus grows at the portfolio rate of return for 20 years so it is influenced significantly by the magnitude of this rate.

In order to verify the impact of the interest rate on the valuation of the fund, we perform a regression of the going concern surplus on the standardized risk factors. The going concern surplus is simply the difference between the asset value at time 0 and the present value of the pension payments made to active and inactive members. The results are shown in Table 5 where all regression coefficients are statistically significant at a confidence level of 99%.

Regressions of the going concern surplus on the standardized risk factors		
Risk factor	**Regression coefficient**	**Rank**
Inflation rate	-0.20×10^{10}	3
Interest rate	1.29×10^{10}	1
Earnings growth	-0.30×10^{10}	2
Mortality rates	0.05×10^{10}	5
Portfolio return	0.15×10^{10}	4

Table 5

As we expected, the interest rate becomes the risk factor that influences the going concern surplus the most. This is because the interest rate is used to discount pension payments that are subtracted to the asset in the going concern surplus. We can conclude that the valuation of the fund is extremely sensitive to changes in interest rates because of the discounting and the asset value is affected significantly by the portfolio return in the long run.

A limitation of these regressions is that it can only measure linear relationships between the dependent variable and its regressors. In order to account for other types of relationships, we draw scattergraphs of the terminal surplus and the going concern liability as a function of the different risk factors.

5.4.2 Scattergraphs

The graphs of the terminal surplus as a function of the inflation rate, the interest rate, the earnings growth and the average mortality rate all result in no clear relationship. Conversely, the graph of the terminal surplus as a function of the portfolio return displays a positive convex relationship. This is shown in Figure 25.

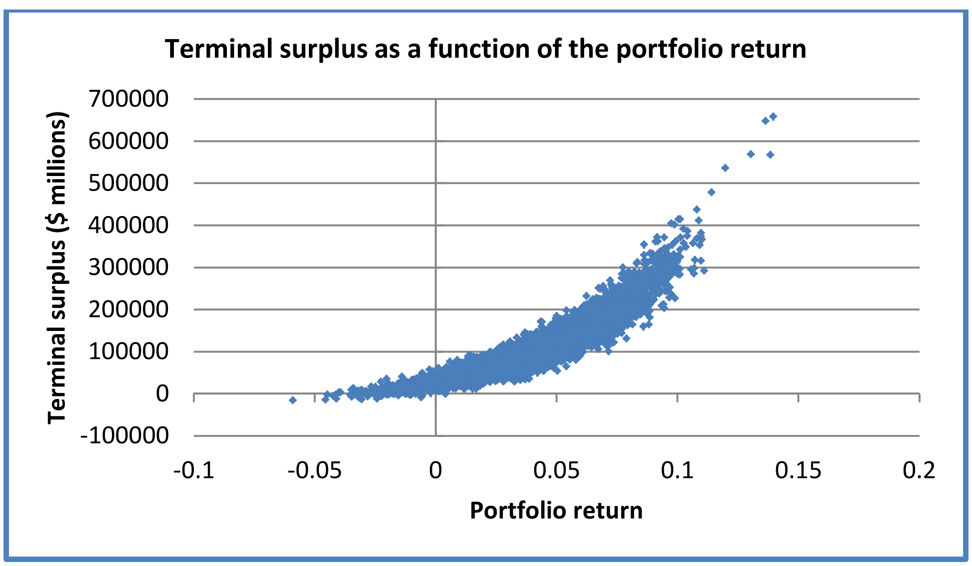

Figure 25

To better capture the influence of the discount rate, we draw the graph of the going concern liability as a function of the interest rate. Like in the regression analysis, the use of the going concern valuation instead of the terminal surplus results in a much clearer relationship for the interest rate. This is shown in Figure 26 where a negative convex relationship can be observed.

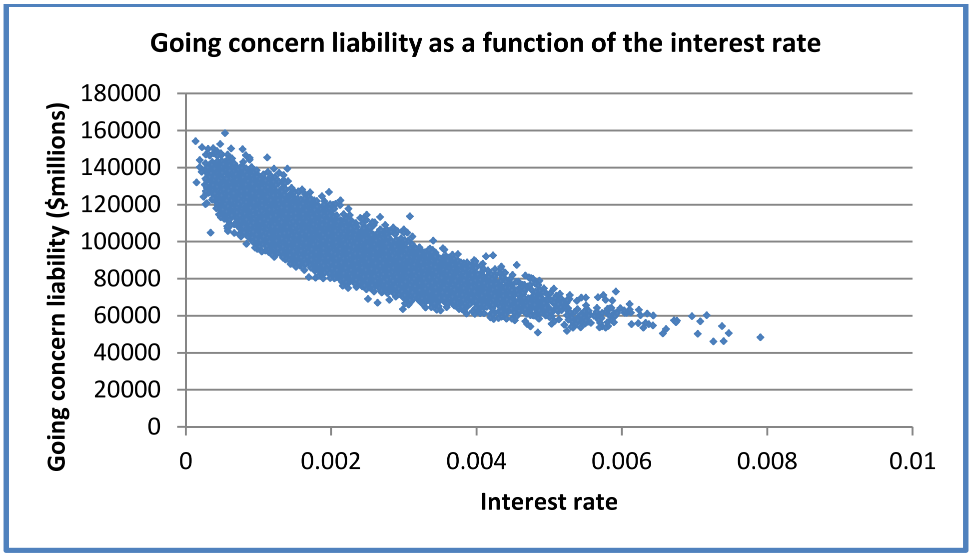

Figure 26

From these graphs, we can conclude that not only the portfolio return displays a positive relationship with the terminal surplus and the interest rate displays a negative relationship with the going concern liability but these relationships are convex.

5.4.3 Variation of the risk factors

The next step of our sensitivity analysis is to change the risk factors in the model to evaluate their impact on the asset, the liability and the surplus. In each simulation, a risk factor is increased or decreased by 1% of its value. For simplicity, we measure the sensitivity by moving the risk factor in the direction that is believed to cause an adverse change in the fund's valuation. Because the interest rate and the portfolio return have a convex relationship with the valuation measures, we first evaluate if the sensitivity is symmetric for an increase and a decrease of 1%. To verify this, we increase and decrease the interest rate by 1%. In both cases, the solvency valuation liability moves by the exact same magnitude, which represents 0.21% of its initial value. It seems that on a small interval, the convexity of the relationship does not affect our results. Table 6 shows the results of our sensitivity test when discounting at the expected portfolio return and Table 7 shows the results when discounting at the long-term bond return. The change is in percentage of the initial value and the ranking based on the sensitivity of the valuation measure to the risk factor is indicated in parenthesis.

Change in the valuation measure for a variation of 1% in the risk factors (discounting at the expected portfolio return)							
Risk factor	**Change**	**Solvency valuation liability**	**Going concern valuation liability**	**Terminal liability**	**Terminal Asset**	**Solvency valuation surplus**	**Terminal surplus**
Pensioners growth	1% inc		0.47% (1)	0.86% (2)	-0.45% (3)		-1.02% (2)
Inflation rate	1% inc	0.17% (3)	0.15% (5)	0.28% (4)	-0.14% (6)	-6.02% (3)	-0.33% (7)
Interest rate	1% dec	0.21% (2)	0.26% (3)	0.10% (5)		-7.56% (2)	-0.49% (5)
Earnings growth	1% inc		0.15% (4)	0.29% (3)	0.33% (4)		0.35% (6)
Mortality rates	1% dec	0.34% (1)	0.37% (2)	1.04% (1)	-0.24% (5)	-12.17% (1)	-0.80% (4)
Contribution rate	1% dec				-1.31% (1)		-1.88% (1)
Portfolio return	1% dec				-0.61% (2)		-0.87% (3)

Table 6

Change in the valuation measure for a variation of 1% in the risk factors (discounting at the long-term bond return)							
Risk factor	**Change**	**Solvency valuation liability**	**Going concern valuation liability**	**Terminal liability**	**Terminal Asset**	**Solvency valuation surplus**	**Terminal surplus**
Pensioners growth	1% inc		0.51% (1)	0.86% (2)	-0.45% (3)		-1.09% (2)
Inflation rate	1% inc	0.18% (3)	0.16% (5)	0.28% (4)	-0.14% (6)	-0.85% (3)	-0.35% (7)
Interest rate	1% dec	0.23% (2)	0.28% (3)	0.11% (5)		-1.06% (2)	-0.52% (5)
Earnings growth	1% inc		0.16% (4)	0.29% (3)	0.33% (4)		0.35% (6)
Mortality rates	1% dec	0.39% (1)	0.43% (2)	1.12% (1)	-0.24% (5)	-1.79% (1)	-0.90% (4)
Contribution rate	1% dec				-1.31% (1)		-1.95% (1)
Portfolio return	1% dec				-0.61% (2)		-0.90% (3)

Table 7

As we can see, while the ranking is the same, the sensitivity of the valuation measures when discounting at the long-term bond return is larger because the variations in value are discounted less. It is interesting to see that when increasing the earnings growth by 1%, the terminal surplus increases by 0.35%. Because the liability is based on the magnitude of the pensionable earnings, we expected the surplus to drop instead. When earnings grow however, the contribution rate is applied to larger earnings and therefore results in larger contributions, which in turn grow the fund's surplus. It seems that the effect on the contributions outweighs the effect on the liability in our model. Regardless, if we would have decreased the earnings growth by 1% instead, the percentage change would simply be the opposite (-0.35%). In this analysis, the growth in the number of pensioners and the contribution rate are included to assess their impact on the valuation measures even if they are not stochastic and are not considered as risk factors. In fact, they are both dependent on internal decisions.

Overall, when considering only the sensitivity of the liability to the factors, both the mortality rates and the growth in the number of pensioners rank first interchangeably. In fact, these two factors determine the number of pensioners over time, which moves the value of the liability twice as much as the other factors. Then comes the interest rate that determines the present value of the liability. Again, the terminal liability is not influenced significantly by the interest rate since its value is not discounted. The next risk factor is the earnings growth that influences the magnitude of pension payments. Finally, the inflation rate ranks last. These results are in accordance with Delong *et al*. (2008), who come to the conclusion that the two main factors affecting the value of the liability are the intensity of mortality and the interest rate.

The value of asset is slightly influenced by the growth in the number of pensioners, the inflation rate, the earnings growth and the mortality rates because these factors affect the pension payments that are debited to the account. The contribution rate and the portfolio return are however far more important in determining the value of the asset. According to our results, the asset is most sensitive to fluctuations in the contribution rate.

When we combine the asset and the liability and evaluate the sensitivity of the surplus, the contribution rate comes in first position with a coefficient about twice as big as the other ones. The growth in the number of pensioners is second. Then the portfolio return and the mortality rates are third and fourth with a coefficient about twice as big as the last three. The interest rate and the earnings growth are in fifth and sixth position. Finally, the inflation rate is last.

5.4.4 Value-at-risk

The last part of our sensitivity analysis is to examine the value-at-risk of the distribution of the liability, the asset and the surplus. In this thesis, we choose a confidence level of 95% for the VaR. For simplicity, we only present the relative VaR, which is the ratio of the VaR to the mean value of the distribution. This measure facilitates the comparison of the volatility generated by the different risk factors since the mean of the distribution is different from one risk factor to the next.

The relative VaR of the solvency valuation liability is equal to 1.2121 when discounting at the expected portfolio return and 1.2295 when discounting at the long-term bond return. This means that the value of the liability is expected to be under 121.21% and 122.95% of its average with a probability of 95%. Alternatively, the going concern liability has a relative VaR of 1.2833 when discounting at the expected portfolio return and 1.3065 when discounting at the long-term bond return.

To illustrate how the model risk affects the relative VaR, let us take for example the more volatile inflation process that is parameterized with data from 1914 to 2011. In this case, the value-at-risk of the solvency valuation liability is equal to 1.3714 instead of 1.2121. Because the volatility of the liability resulting from the inflation process triples, the VaR increases significantly. As we can see, the choice of model has a significant impact on the volatility of the liability and therefore on its VaR.

In the rest of this section, the VaR is used to analyze the fluctuation in asset, liability and surplus caused by each risk factor. To do so, we perform different simulations allowing only one factor to be stochastic while all other factors are kept constant at their average value. Table 8 and Table 9 display the relative VaR for the different valuation measures and the different risk factors when discounting at the expected portfolio return and at the long-term bond return respectively. The two tables also include the ranking based on the magnitude of the relative VaR for each valuation measure. The “Total” row indicates the relative VaR when allowing all risk factors to be stochastic.

Relative value-at-risk per risk factor (discounting at the expected portfolio return)						
Risk factor	**Solvency valuation liability**	**Going concern valuation liability**	**Terminal liability**	**Terminal Asset**	**Solvency valuation surplus**	**Terminal surplus**
Inflation rate	1.0275 (2)	1.0227 (3)	1.0419 (4)	0.9773 (3)	3.0338 (2)	0.9568 (3)
Interest rate	1.1578 (1)	1.1993 (1)	1.0737 (2)		7.3808 (1)	0.9680 (4)
Earnings growth		1.0928 (2)	1.1511 (1)	0.8413 (2)		0.8245 (2)
Mortality rates	1.0077 (3)	1.0112 (4)	1.0543 (3)	0.9960 (4)	1.5698 (3)	0.9787 (5)
Portfolio return				0.4741 (1)		0.2805 (1)
Total	1.2121	1.2833	1.2441	0.4593	8.6480	0.2416

Table 8

Relative value-at-risk per risk factor (discounting at the long-term bond return)						
Risk factor	**Solvency valuation liability**	**Going concern valuation liability**	**Terminal liability**	**Terminal Asset**	**Solvency valuation surplus**	**Terminal surplus**
Inflation rate	1.0296 (2)	1.0243 (3)	1.0428 (4)	0.9773 (3)	1.1454 (2)	0.9540 (3)
Interest rate	1.1740 (1)	1.2183 (1)	1.0796 (2)		1.8136 (1)	0.9612 (4)
Earnings growth		1.0968 (2)	1.1511 (1)	0.8413 (2)		0.8223 (2)
Mortality rates	1.0099 (3)	1.0144 (4)	1.0592 (3)	0.9960 (4)	1.0486 (3)	0.9743 (5)
Portfolio return				0.4741 (1)		0.2582 (1)
Total	1.2295	1.3065	1.2502	0.4593	2.0596	0.2113

Table 9

As we can see, while the ranking is the same, the relative VaR when discounting at the long-term bond return is larger because the deviations from the mean values are discounted less and are therefore less attenuated.

The relative VaR for the terminal asset and the terminal surplus are smaller than 1 because, as opposed to the liability, the value-at-risk is on the left side of the distribution. The smaller the VaR (and the relative VaR), the larger the risk is. The VaR for the solvency valuation surplus is also on the left side of the distribution but because both the VaR and the mean value are negative, the relative VaR becomes larger than 1.

We can see that the interest rate and the earnings growth create a lot of fluctuation in the value of the liability and the surplus as indicated by their high relative VaR. In fact, the interest rate ranks first and the earnings growth ranks second based on their impact on the liability calculated with the solvency and going concern valuations. The inflation rate and the mortality rates in their case are far less impactful as indicated by their small relative VaR. They rank second to last and last in all valuation measures except for the terminal liability where it is the reverse. As for the portfolio return, it has a large impact on the value of the asset and the surplus in which it ranks first.

5.4.5 Overall ranking

By combining the different surplus measures and by considering that the terminal surplus does not allow assessing the risk of the discount rate at its full extent, the rankings of the regression analysis and the VaR analysis agree. Whereas the regression analysis cannot differentiate between the ranks of the first two risk factors, the VaR analysis shows that the portfolio return is the risk factor that generates the largest negative surplus values. Table 10 shows the combined ranking.

Ranking of the risk factors for the regression analysis and the VaR analysis	
Risk factor	**Rank**
Portfolio return	1
Interest rate	2
Earnings growth	3
Inflation rate	4
Mortality rates	5

Table 10

The analysis by which the risk factors are changed by 1% results in a slightly different ranking when we exclude the non-stochastic factors (contribution rate and pensioners growth). Whereas the mortality rates rank last in the regression and the VaR analysis, they rank second in this case. Table 11 shows the resulting ranking for the risk factors only.

Ranking of the risk factors for the 1% variation analysis	
Risk factor	**Rank**
Portfolio return	1
Mortality rates	2
Interest rate	3
Earnings growth	4
Inflation rate	5

Table 11

To understand the difference in ranking, we have to remember that the VaR depends on the volatility of the underlying random variables which is, in our case, determined by the volatility of the risk factor. Similarly, in the regression analysis, variables with a small standard deviation do not impact the value of the surplus enough to make a linear relationship noticeable and are assigned a small regression coefficient. It turns out that the standard deviation of the mortality rates is about ten times smaller than the standard deviation of the inflation rate, the interest rate and the earnings growth. As a result, the mortality rates cause the surplus to vary ten times less than the other factors and therefore ranks last. Conversely, the standard deviation of the portfolio return is ten times larger so it causes the surplus to vary ten times more and therefore ranks first.

The second analysis however is independent on the volatility of the underlying risk factors since they are all changed by the same percentage (1%). This allows assessing the sensitivity of the surplus to a fixed change in the value of the risk factor. From this, we can conclude that a movement in the mortality rates has a large impact on the value of the surplus but it is not necessarily the greatest risk since its process is relatively stable according to our model. In fact, the risk does not simply depend on the sensitivity to a factor, it also depends on the probability of an adverse change in the factor and the magnitude of such change. In our model, a large fluctuation in mortality rates is far less likely than for the portfolio return.

Overall, this section shows that the fund should mainly manage the interest rate risk and the portfolio return risk. This result is in accordance with the findings of Wouters (2008) and Moore (2010), who show that the duration gap and the volatility of the portfolio returns are the two largest risks to a fund’s funding status. To manage these risks, the fund must design a portfolio of assets that hedges the exposure of the liability to interest rates and have a small variance. More generally, the portfolio of assets must help the fund to minimize the probability of insolvency and the variance of the terminal surplus. This corresponds to the next step of the LDI approach, which is to find the optimal asset allocation.

5.5 ASSET ALLOCATION

According to Morneau Shepell, a consulting company that surveyed 100 Canadian public companies in 2013, the average asset allocation as at December 31, 2012, was 53% in equity, 42% in fixed income and 5% in other assets. In the case of the fund, the Public Sector Pension Investment Board (PSP Investments) is in charge of managing its assets. In 2011, they invested 32% in Canadian stocks, 17% in U.S. stocks, 18% in global equity, 5% in Canadian inflation-linked bonds, 6% in Canadian government bonds, 9% in Canadian corporate bonds, 9% in global real estate and 4% in global infrastructure. This sums up to 67% in equity, 20% in fixed income and 13% in other assets, which is more risky than the average Canadian allocation. In this section, we evaluate other asset allocations and their impact on the different valuation and risk measures.

5.5.1 Bond-only portfolio

The first allocation evaluated includes only bonds. In fact, in order to hedge the fund against changes in interest rates, the LDI approach suggests matching the duration of the asset to the duration of the liability. An efficient way to do so is to invest only in government bonds and select the appropriate maturities to obtain the same duration as the liability. As mentioned in the review of literature, the goal of duration matching is to synchronise the sensitivity of the asset and the liability to the interest rates so that a parallel shift in the yield curve affects both sides of the balance sheet equally. In the case of the fund, the duration of the liability depends on the magnitude and timing of the pension payments. The Macaulay duration of the liability is equal to 10.27 years whereas the modified duration (which is simply the Macaulay duration divided by the average yield) is equal to 9.84 years. The modified duration indicates that the liability would increase by 9.84% if the interest rate would drop by 100 basis points. To build the portfolio of assets, we select zero coupon bonds from the Government of Canada. Their yields were taken from the Bank of Canada database. Because the liability is modeled starting from 2012, we use their yield as at December 30th, 2011. We then find a combination of weights that makes the modified duration of the asset also equal to 9.84 years. One of these combinations consists of investing 28% in the one-year zero coupon bond and about 3% in each of the zero coupon bonds with maturities going from 2 years to 25 years with one year increments. With such a portfolio, a variation of 100 bps in the interest rate results in a mismatch of 2% between the change in value of the asset and the liability. A variation of 1000 bps, however, results in a mismatch of 29%, but it is important to understand that duration matching immunises against small changes in interest rates. To obtain a better fit however, we can make an adjustment that takes into account convexity. The convexity of the liability is equal to 137.93. When adjusting the weights of the portfolio so that the asset also has a convexity of 137.93, a variation of 100 bps in the interest rate results in a mismatch of only 0.3% between the asset and the liability. The weights of this adjusted portfolio consist of 44% in the 5-year zero coupon bond and about 3% in each of the zero coupon bonds with maturities going from 6 years to 23 years. Note that when discounting at the long-term bond return instead of the expected portfolio return, the modified duration goes from 9.84 to 10.84. The portfolio weights would be shifted slightly towards bonds with longer maturities.

As mentioned above, allocating assets based on the duration of the liability allows hedging the interest rate risk, which is the primary risk of the fund after the portfolio return risk. It disregards however other risks such as inflation rate risk. To hedge this risk as well, the duration matching portfolio could use inflation-linked government bonds instead of nominal government bonds. The maturities of the bonds in the portfolio would also be chosen as to match the duration of the asset to the duration of the liability. In practice, however, inflation-linked bonds are available with only some maturities (by increments of 5 years). This makes the duration matching harder and the inflation rate risk hedging even harder. In fact, in order to hedge the inflation rate risk, the maturities of the inflation-linked bonds must fit the timing of the pension payments. We would therefore need bonds with maturities that end each month, which is not possible.

Moreover, the problem with such allocation is that including only government bonds in the portfolio results in a low expected return. The duration matched portfolio with nominal bonds has an expected return of about 1.8%. This return is even lower for the inflation-linked bonds. With these low portfolio returns, it becomes costly for contributors to maintain the actuarial surplus. Investing in asset classes with higher expected returns such as stocks allows reducing the need for contributions. Another problem with a bond-only allocation is that if interest rates are expected to rise (which is likely in the future) the value of the portfolio will shrink. For these reasons, we evaluate more diverse asset allocations that include the eight asset classes.

5.5.2 Diversified portfolio

As mentioned in the methodology section, we set the minimum allocation at 2% and the maximum allocation at 40%. In theory, a way to reduce the probability of insolvency is for the distribution of the asset to dominate the distribution of the liability in most scenarios. To achieve such a distribution, we first study the distribution of the liability.

In our model, the value of the liability is a function of the various risk factors chosen. Although most of these risk factors are normally distributed, the liability follows a complex distribution that is composed of the product of various distributions. Even if we don't know the exact distribution of the liability, the different iterations of our stochastic model create an expected distribution. We evaluate the four first moments of this distribution. The mean of the liability when discounting at the expected portfolio return is equal to $63,167 million, its standard deviation is $8,374 million, its skewness is -0.11 and its kurtosis is 2.67. When discounting at the long-term bond return, the mean is equal to $78,393, the standard deviation is $11,210 million, the skewness is -0.10 and the kurtosis is 2.66. As we can see, the skewness and kurtosis are very similar with the two discount rates and they are close to what the normal distribution predicts (skewness of 0 and a kurtosis of 3).

Secondly, we need to analyze the distribution of the return on the different asset classes that are combined to form the optimal portfolio. Table 12 displays the four first moments of their distributions.

Four first moments of the distribution of the 8 asset classes' simulated returns

Asset Class	Mean	Variance	Skewness	Kurtosis
Canadian Stocks	0.0476	0.0313	0.5481	3.5698
US Stocks	0.0544	0.0363	0.4978	3.2077
Global Equity	0.0611	0.0627	0.7606	3.8311
Canadian Inflation-linked Bonds	0.0138	0.0091	0.3168	3.0074
Canadian Government Bonds	0.0218	0.0019	0.2231	3.1447
Canadian Corporate Bonds	0.0338	0.0022	0.2400	3.1616
Global Real Estate	0.0423	0.0665	0.7631	4.0458
Global Infrastructure	0.0376	0.0445	0.4608	3.5539

Table 12

As we can see, bonds have a lower expected return but a lower variance, skewness and kurtosis. Overall, bonds have the distribution that is the closest to the distribution of the liability so they are expected to constitute the majority of the optimal portfolio.

As we mentioned in the methodology section, to find the combination of assets that fits the distribution of the liability the best, we use a method called polynomial goal programming (PGP). From this optimization, with the minimum and maximum weight constraints, we obtain an allocation of 6% in stocks, 81% in bonds and 13% in other assets, with an expected portfolio return of 2.65% compared to 4.43% for the original allocation. The detailed asset allocation can be found in Table 13.

For comparison, we also use PGP to find the portfolio that maximizes the skewness and minimizes the kurtosis, since in an asset-only framework a large positive skewness is desirable. This results in a drastic shift towards risky assets. The portfolio is in fact composed of 44% in stocks, 16% in bonds and 40% in other assets with an expected portfolio return of 4.03%.

Besides these portfolios, we also find the conventional Markowitz optimal portfolios by minimizing the variance for each level of expected return. We test four portfolios on the efficient frontier including the minimum variance portfolio. This minimum variance portfolio is composed of 12% in stocks, 83% in bonds and 5% in other assets, with an expected return of 3.00%. Again the detailed allocation of this portfolio can be found in Table 13.

Allocation of the different portfolios

Asset Class	Original allocation	Minimum skewness minimum kurtosis	Maximum skewness minimum kurtosis	Minimum variance
Canadian Stocks	32%	2%	2%	8%
US Stocks	17%	2%	40%	2%
Global Equity	18%	2%	2%	2%
Canadian Inflation-linked Bonds	5%	31%	12%	3%
Canadian Government Bonds	6%	40%	2%	40%
Canadian Corporate Bonds	9%	10%	2%	40%
Global Real Estate	9%	11%	11%	2%
Global Infrastructure	4%	2%	29%	3%

Table 13

In addition to these asset allocations, several other combinations have been tested but they produced inferior results in terms of probability of insolvency and terminal surplus variance.

5.5.3 Choice of discount rate and addition of funds

As we mentioned in section 5.2.2, because the discount rate based on the expected portfolio return is an average weighted by the portfolio weights, it changes when we modify the allocation. In turn, the discounted pension payments and the distribution of the liability also change. As a result, the allocation that produces the lowest probability of insolvency is the one that maximizes the expected portfolio return by investing mostly in stocks. This highlights the conflict that arises when a fund uses the expected portfolio return to discount its pension payments. Like Inkmann and Blake (2004) and Andonov *et al*. (2013) mentioned, the manager has an incentive to take excessive risk and invest the majority of the asset in stocks to increase the discount rate and disclose a better funding status.

Keeping in mind that our objective is not to maximize the expected return on the portfolio, we decided to use the long-term bond return as a discount rate. Because this rate is independent from the asset allocation, the distribution of the liability remains constant, which facilitates the comparison between the different portfolios. Moreover, it is more conservative and prudent to use the long-term bond rate or return rather than discounting at a higher rate, which could potentially underestimate the liability. This method is also in accordance with most other Canadian pension funds and complies with European standards.

As we saw in section 5.3, the initial probability of insolvency is equal to 93% when discounting at the long-term bond return. This is far from the 5% goal, which makes it difficult to evaluate which allocation allows reaching our objective the best. In order to facilitate the comparison of the different portfolio allocations, we add extra funds to fill the funding gap and obtain a probability of insolvency of 5% in 2012. As we will discuss in section 5.6.1, the injection of funds to fill the gap in 2012 is one of the strategies that could be applied to reach the objective of the fund.

5.5.4 Comparison of the portfolios

With these additional funds and the long-term bond return as the discount rate, the distribution of the asset with the different allocations can then be evaluated and compared. The first four moments of the distribution of the asset with the different portfolio allocations are displayed in Table 14.

Comparison of the distribution of liability and the distribution of asset with the different allocations

Moment	Liability	Asset					
		Original allocation	Government bond-only	Inflation-linked bond-only	Minimum skewness, minimum kurtosis	Maximum skewness, minimum kurtosis	Minimum variance
Mean	7.84E+10	9.97E+10	9.78E+10	9.75E+10	9.82E+10	9.92E+10	9.85E+10
Variance	1.26E+20	1.08E+20	1.63E+19	8.44E+19	2.38E+19	1.26E+20	1.15E+19
Skewness	-0.0993	0.2768	0.1461	0.2772	0.1618	0.3264	0.1190
Kurtosis	2.6574	3.0995	3.0035	3.1155	3.0606	3.1318	3.0664

Table 14

As we can see, the first three moments (mean, variance and kurtosis) of the distribution of the asset dominate the first three moments of the distribution of the liability. The kurtosis of the liability is not dominated however. It is important to note that the skewness and kurtosis of the asset can differ from those of the portfolio returns because the value of the asset is also affected by the pension payments and the contributions received. In fact, whereas the minimum reachable skewness of the portfolio returns is 0.20, the skewness of the asset with the allocation in the "minimum skewness, minimum kurtosis" portfolio is 0.16, which is even lower. Also, we notice that the allocation in government bonds results in a skewness of 0.15 and a kurtosis of 3.00, which is lower than for the "minimum skewness, minimum kurtosis" portfolio. This is because this allocation does not follow the minimum and maximum allocation constraints of 2% and 40%. As for the minimum variance portfolio, whereas its skewness is lower than the "minimum skewness, minimum kurtosis" portfolio, suggesting a better fit with the liability, its kurtosis is slightly larger.

Regardless which portfolio matches the distribution of the liability the best, we have to find the allocation that achieves our objective the best. In fact, the allocation that minimizes the probability of insolvency is the minimum variance portfolio. Figure 27 shows the probability of insolvency over time for the different asset allocations.

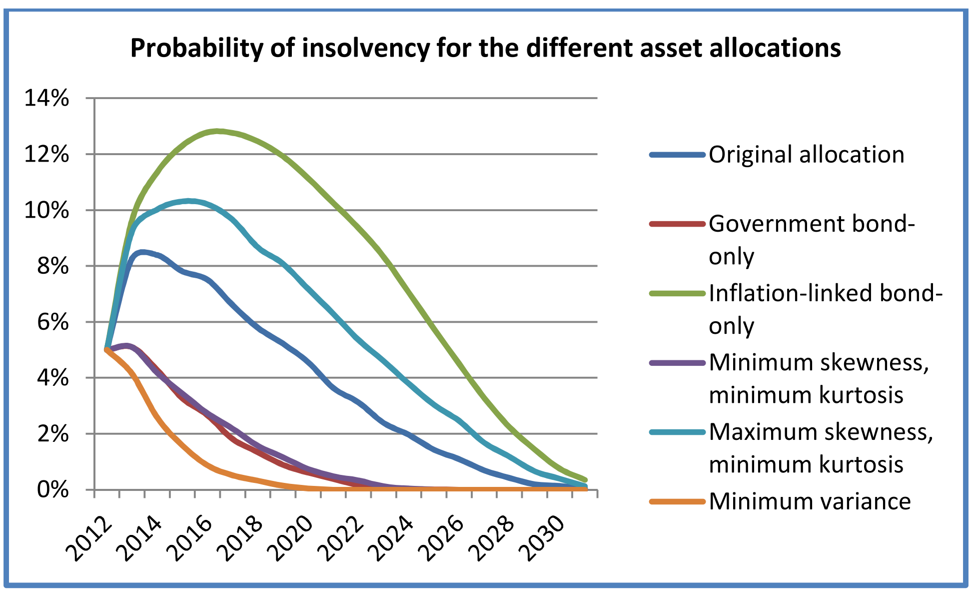

Figure 27

As we can see, only the minimum variance portfolio allows keeping a probability of insolvency under 5% over the 20 years simulated. The allocations in government bonds only and in the "minimum skewness, minimum kurtosis" portfolio result in a significant reduction in the probability of insolvency but not as much as the minimum variance portfolio. Besides the probability of insolvency, our second objective is to minimize the terminal surplus variance, which is shown for the different allocations in Figure 28.

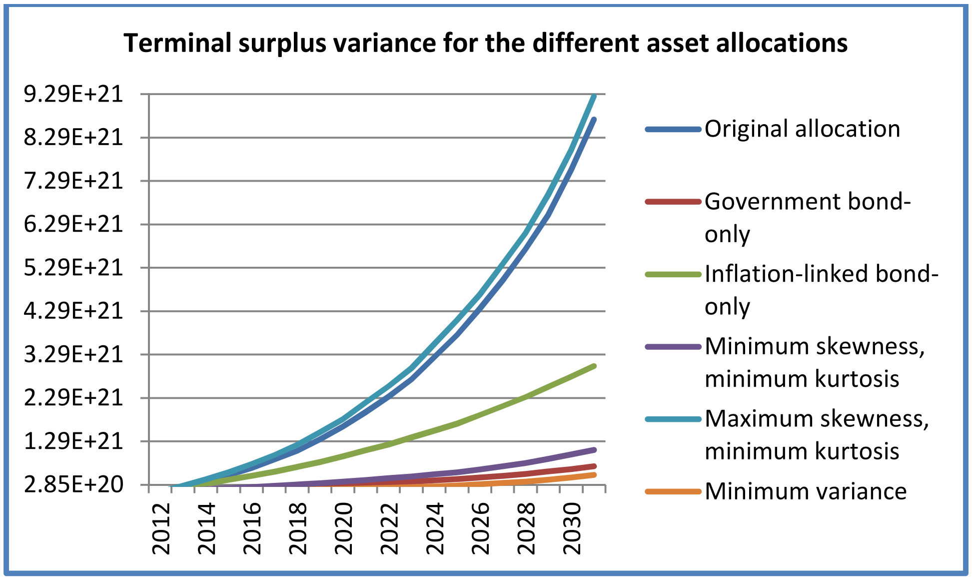

Figure 28

In all cases, the terminal surplus variance increases with time because our forecast becomes less precise and more pension payments are indexed and discounted over a longer period of time, which creates more uncertainty. The allocations that include a lot of stocks, like the original allocation and the "maximum skewness, minimum kurtosis" portfolio, result in the highest terminal surplus variance. Conversely, the allocations with a lot of government bonds result in the lowest variance. This can be explained by the fact that the returns on Government of Canada bonds have the lowest variance among the eight asset classes included. The allocation that reduces the terminal surplus variance the most is the minimum variance portfolio. Although the government bond-only allocation also results in a low terminal surplus variance, it does not comply with our minimum and maximum allocation constraints of 2% and 40%. Also, it does not minimize the probability of insolvency, which is our primary objective. Nonetheless, because it allows duration matching, it will still be considered in section 5.6.

The minimum variance portfolio allocation improves the risk profile of the fund in various ways. Firstly, by allocating more funds in long-term bonds rather than stocks, the duration of the portfolio goes from approximately 1 year with the original allocation to 4.15 years, which offers a better protection against interest rate risk. Overall, we can see that this allocation allows hedging part of the risk generated by the underlying stochastic factors because the relative value-at-risk of the terminal surplus is increased. Whereas a low relative VaR for the liability is preferable as it denotes smaller extreme liability values, a large relative VaR for the surplus is preferable. In fact, low extreme surplus values must be avoided. For this reason, the increase in relative VaR indicates an improvement in the risk profile. When adding the extra funds to close the initial funding gap, the relative VaR of the terminal surplus goes from 0.2113 to 0.3059. It is further increased to 0.6851 when changing the asset allocation from the original weights to the minimum variance portfolio weights. Because the relative VaR is higher, it indicates that the surplus fluctuates less due to changes in the risk factors. Among all portfolio allocations, the minimum variance allocation results in the highest relative VaR of the terminal surplus.

5.5.5 Optimal asset allocation

Considering that the minimum variance portfolio minimizes the probability of insolvency and minimizes the terminal surplus variance while complying with the asset allocation constraints, it is considered as the optimal asset allocation. It also corresponds to the definition of the liability-hedging portfolio introduced in the LDI approach as it hedges the fund against the underlying risk factors the best.

Because the minimum variance portfolio is also on the Markowitz efficient frontier, it can also be considered as a performance portfolio. In fact, in an asset-only framework, managers often select portfolios on the efficient frontier for their superior performance per unit of risk. Likewise, Martellini (2006) also considers the mean-variance optimal portfolio as the performance portfolio. For this reason, we do not need to build two different portfolios like the LDI approach suggest. In our case, the minimum variance portfolio is the liability-hedging portfolio as well as the performance portfolio.

Our optimal asset allocation is in line with the findings of Moore (2010) and Leahy (2011) who suggest shifting assets from equity to bonds. More precisely, our 83% allocation to bonds is in accordance with the findings of Meder and Staub (2006) and Cox *et al.* (2013) who obtain an optimal portfolio with respectively 90% and 83% of assets allocated to bonds. Conversely, Ostaszewski (2011) finds a balanced optimal portfolio that allocates 40% in equity. This is not surprising because the objective of his model is to minimize the contribution rate, which is achieved by maximizing the expected return on the portfolio.

An important limitation of this thesis is that it results in a static portfolio allocation. If this model were to be applied however, the fund managers could rerun the portfolio optimization as new market and fund-specific data uncover. Also, if the funding status of the fund improves and the objective shifts towards generating a better return on asset, the portfolio of the fund could move up along the Markowitz efficient frontier, which would maintain the lowest asset variance for every level of expected return. The increased expected return would however be at the expense of the surplus variance and the probability of insolvency. If we take for example the portfolio with the highest variance on the efficient frontier and compare it to the minimum variance portfolio, the expected return increases from 3.00% to 4.90% but the terminal surplus variance becomes 50 times larger. As for the probability of insolvency, it becomes 15 times larger. If the fund has a probability of insolvency of 0.1% however, making it 15 times larger is not necessarily a problem since it is still under 5%.

The next three sections evaluate how three of our assumptions influence the optimal allocation. We start by removing the minimum and maximum allocation constraints. We then make adjustments to the asset returns to correlate them with the liability. Finally, we discount the liability with the expected portfolio return.

5.5.6 Removing the minimum and maximum allocation constraints

To verify the impact of the allocation constraints on the optimal allocation, we perform an unconstrained Markowitz portfolio optimization. The resulting minimum variance portfolio shifts funds away from global equity and inflation-linked bonds and reallocates them into the other asset classes. This results in an asset variance that is 9% smaller and a terminal surplus variance that is 12% smaller. Because the expected return is smaller however, the probability of insolvency increases by 1.3% on average, which makes this solution inferior and shows that our constraints do not affect the optimal allocation.

5.5.7 Adding correlation between the asset and the liability

One of the characteristics that make an asset class a good investment for the risk management of pension funds is its correlation to the liability and its underlying risk factors. The asset classes with the largest correlations are therefore prioritised in the liability-hedging portfolio. According to Meder and Staub (2006), stocks, real estate and infrastructure funds are correlated to the GDP and to earnings growth, whereas bonds are correlated to interest rates. By definition, inflation-linked bonds are correlated to inflation. These relationships are however not very strong in our historical data set in which we obtain small correlation coefficients. This could be due to the limited data set and the fact that we used monthly data, but as Adam (2007) explained, the correlation between different economic variables is very volatile over time. This is due to what we call “regime changes”. This lack of evidence regarding the correlation of assets to the risk factors supports our model. In fact, the simulated asset returns are randomly selected from historical returns so that they have no connection to the risk factors. Nonetheless, we test the impact of linking the asset returns to the risk factors. This seems especially important for inflation-linked bonds since their main appeal is their inflation protection feature, which is neglected by our method.

In order to create a relationship between the asset returns and the risk factors, some adjustments are made in accordance with Meder and Staub (2006). The returns on stocks, real estate and infrastructure are adjusted by adding a 20% correlation to the earnings growth to simulate a relationship with the GDP. The bond returns are also adjusted to include an inverse relationship with the interest rate. Finally, the returns on the inflation-linked bonds are adjusted upward and downward by the unexpected inflation. When considering the adjustment on the inflation-linked bond-only, we can see that the results do not differ significantly. Whereas the terminal surplus variance is 0.1% larger, the probability of insolvency is on average 0.2% smaller. When considering the adjustments on all asset classes, the terminal surplus variance is 0.4% larger and the probability of insolvency is on average 0.3% smaller. These small differences have a negligible impact on the optimal solution and the minimum variance portfolio remains the optimal asset allocation.

5.5.8 Discounting at the expected portfolio return

Finally, it is important to remember that the optimal portfolio was formulated when discounting at the long-term bond return. Our results show however that the minimum variance portfolio is still the optimal asset allocation when discounting at the expected portfolio return. Whereas the solvency valuation liability and the going concern valuation liability are 17% and 21% larger than with the original allocation because the expected portfolio return is smaller, the minimum variance portfolio generates a much smaller variance of the return on asset. This in turn makes the average probability of insolvency 79% smaller than with the original allocation. Moreover, the surplus variance is also minimized with this allocation. For this reason, we still retain the minimum variance portfolio as the optimal asset allocation.

By investing assets in the optimal portfolio, the fund improves its risk profile but it does not solve the initial deficit of $16,979 million, which creates a probability of insolvency of 93%. This is why the last step of the LDI approach is to elaborate a funding strategy that will help the fund reach its objective. In the previous analysis, we added extra funds in 2012 to facilitate the comparison between the different allocations. We will see in the next section that this could be a good strategy to reach a probability of insolvency of 5%. Other solutions include delaying retirement age or increasing contribution rates.

5.6 FUNDING STRATEGY

When we account for the risk generated by the stochastic factors, not only the expected average liability increases but reserves should be added to avoid insolvency if adverse scenarios uncover. In fact, the liability-hedging portfolio is not perfect. Although some of the risk generated by the stochastic factors is hedged, some remains. As mentioned previously, we cannot cover all of the inflation rate risk with inflation-linked bonds since not enough different maturities are available. Also, most of the time, non-financial risks cannot be hedged with the financial markets. In other words, the market is incomplete. Although some longevity bonds and swaps are available to hedge fluctuations in mortality rates, these markets are not developed enough to be an effective solution against longevity risk (Cox *et al*., 2013). For these reasons, a fund cannot expect to be perfectly hedged and reserves must be kept to account for the remaining risk and avoid defaulting on its pension payments. The reserve that must be kept to maintain a probability of insolvency of 5% is approximated by the VaR with a confidence level of 95%. This measure however is dependent on the validity of our model. To accounts for the possibility that we model the fund or the risk factors with the wrong model, we must increase the reserve even higher by the measured "model risk".

For our model, when discounting at the long-term bond return, the model risk is equal to 3.18% and the relative VaR is equal to 1.23, which calls for additional reserves of 26.91% of the liability or $21,095 million. If we would have discounted at the expected portfolio return instead, the additional reserves would be 23.51% of the liability, which is equal to $14,851 million. By considering the liability 26.91% larger and therefore reserving an additional $21,095 million, we make sure that the fund has enough asset to cover the pension payments in 95% of scenarios while accounting for the model risk. Leaving aside the reserve for the model risk, the amount required to cover the $16,979 million deficit and maintain a probability of insolvency of 5% is equal to $34,968 million. This is the amount that was added in section 5.5.3 to facilitate the comparison of the different allocations.

5.6.1 Injecting capital

This additional capital corresponds to the first strategy that we consider. In fact, in order for the fund to solve its deficit and include precautionary reserves, the government could inject $34,968 million at the beginning of 2012. As we saw in section 5.3, a surplus is developed over time as the asset is projected to grow faster than the liability. This results in an average surplus of $120,591 million in 2031. Considering that only $37,971 million is needed in 2031 to ensure a probability of insolvency of 5%, $82,620 million could be repaid to the government, which would result in a 4.2% annual return on the initial injection of capital.

Another way to return money to the government while gaining autonomy would be to lower the government's contribution ratio. In 2012, the government contributed 1.76 times the contributors' payments. This ratio could be lowered to 0.8 in the first 4 years and then lowered to 0.4 for the following years. This would allow the fund to rely less on government funding and still maintain a probability of insolvency lower than 5%.

Conversely, instead of repaying the government, the fund could use the exceeding surplus to reduce contribution rates. Different reduction strategies were tested but the most efficient is to reduce contribution rates by 2% in the first two years and then reduce them by 4% in the following years. The terminal value of asset and the terminal surplus would be reduced by 42.31% and 58.69% respectively, but this would avoid growing a surplus while maintaining a probability of insolvency of 5%.

5.6.2 Duration matching

If the fund decides to invest only in government bonds to match the duration of the asset to the duration of the liability, the lower expected portfolio return (2.22%) generates a higher probability of insolvency because the asset grows at a lower rate. In fact, even if the government injects funds to make the probability of insolvency equal to 5% in 2012, this one reaches 5.1% in 2013 before gradually decreasing. In order to maintain a probability of insolvency of 5% throughout, the fund can increase the contribution rates by 1% in the first year, decrease them by 2% from year 2014 and decrease them by 4% from year 2018. As we can see, this solution is more costly to contributors so the minimum variance portfolio allocation is preferred.

While injecting $34,968 million may seem drastic, it allows reaching the fund's objective by 2012. If we allow a delay for reaching the solvency goal, other more flexible and less costly strategies could be implemented.

5.6.3 Increasing contribution rates

Instead of injecting capital in 2012, the fund could increase contribution rates. If the fund were to double the contribution rates to about 18%, the terminal value of the asset and the terminal surplus would be 2.5 times and 3.5 times larger respectively. As a result, the probability of insolvency would reach 5% in 5 years. A less drastic 20% increase to about 11% would make the terminal value of the asset and the terminal surplus 1.3 times and 1.5 times larger respectively, which would allow the probability of insolvency to reach 5% in 10 years.

5.6.4 Delaying the age of retirement

Another solution to close the funding gap without injecting money in 2012 is to delay the age of retirement. If the active members collectively decide to retire one year later and all pension payments start being payable when pensioners are one year older, the solvency valuation liability and the going concern liability would drop by 3.37% and 5.68% respectively. The initial surplus and the terminal surplus would increase by 15.31% and 16.83% respectively. As a result, the probability of insolvency would reach the 5% goal in 11 years. If the age of retirement is delayed by 5 years instead, the solvency valuation liability and the going concern liability would drop by 15.57% and 25.77% respectively. The initial surplus and the terminal surplus would increase by 70.79% and 76.37% respectively, which would result in a probability of insolvency lower than 5% in 4 years. Although delaying the retirement age by 10 years would allow reaching the solvency goal in 2 years, this is not a feasible solution as it would require people to work past age 70.

5.6.5 Integrating model risk

Although the "model risk" was left aside in the previous funding strategies, it is important to realise that our predictions could be erroneous. For this reason, we should include an additional reserve of 3.18% or $2,490 million to account for this risk. Again, this could be injected by the government in 2012 or it could be reached gradually by increasing contributions or delaying the retirement age.

To illustrate the model risk, we adjust the liability growth upward by 3.63% to fit the growth rate assumed by the OCA. In fact, as we saw in section 5.2.7, the liability produced by our model grows at a slower pace than the OCA predicts. This adjustment made terminal surplus 42.58% smaller and increased the probability of insolvency to 14% by 2019. If the fund did not keep additional reserves to account for the model risk, it could either increase contribution rates by 40% in the first 4 years or delay retirement age by 1 years.

5.6.6 Final remarks

Although increasing the contribution rate or the age of retirement solves the funding gap and builds precautionary reserves, the "fairness" of these solutions is questionable. In fact, contributors are asked to significantly increase their payments to the fund or delay their retirement so that pensioners can be paid. This is an example of intergenerational transfer problems where one generation is penalised for the benefit of another generation. For this reason, the initial injection of money by the government seems like the best strategy.

While this thesis focuses mainly on the solvency of the fund, its liquidity is also important. On a year-to-year basis, the only way that the fund can fail to make a payment is if liquid assets are smaller than the year's difference between pension and contribution payments. Because pension and contribution payments are highly correlated (80%), this creates a natural hedge. In fact, when the pension payment is high, the contribution payment tends to be high as well. To prevent not having enough money, the fund can keep a cash reserve each year of the amount of the expected pension payment minus the expected contribution payment plus 2 standard deviations (for a 95% level of certainty) or 3 standard deviations (for a 99% level of certainty). This represents about 1% of total assets each year for 2 standard deviations and 1.1% for 3 standard deviations. This liquidity management could be combined with the above solvency management to create a better risk management strategy.

Finally, as we explained in the methodology section, our model starts simulating the liability and the asset of the fund in 2012. This is because the last actuarial report was published in 2011. For this reason, the analysis and solutions of this thesis are tailored to the situation the fund was in at the end of 2011. It is therefore relevant to examine how this situation evolved since then. Because we only have information regarding the asset through the financial statements prepared by PSP Investment, we can only assess the evolution of the asset. Since 2011, stock returns have been positive and interest rates have been relatively stable. This resulted in a 17% annual growth in asset on average in the last three years, which is higher than what is forecasted by the model. As a result, the government could inject less capital to cover the now smaller deficit. Because the allocation of the fund remained similar, except for a shift from Canadian equity to foreign equity, our recommendation regarding the optimal portfolio does not change. Overall, we can assert that the analysis and solutions contained in this thesis are still valid and relevant in 2014.

6. CONCLUSION

The objective of this thesis was to find the optimal asset allocation and funding strategy for the pension fund of the Public Service of Canada by applying the liability-driven investment (LDI) approach. This is relevant because the latest market conditions and regulations call for new investment and risk management strategies to remediate the widespread deficits and the high solvency risk in defined benefit pension funds.

In order to contribute to the knowledge of the risks affecting pension funds and the techniques used to manage them, we built a model that simulates the liability and the asset of the fund over 20 years as a function of different risk factors. Overall, the model predicts a deficit of $16,979 million. This deficit however is dependent on the processes that are chosen to simulate the risk factors. In fact, the processes used in this thesis might not describe the simulated variables properly. This is referred to as the model risk and we estimate that it is equal to 3.18% of the value of the liability.

From the simulated liability and asset, a sensitivity analysis was conducted to see how the risk factors impacted these values. The two most important risk factors are the portfolio return and the interest rate, which is in accordance with the findings of Wouters (2008) and Moore (2010). Conversely, the mortality rates have a large impact but are less likely to fluctuate widely according to our model.

The next step was then to find the optimal asset allocation that minimizes the solvency risk measured by the probability of insolvency and the variance of the terminal surplus. This portfolio was found to be the minimum variance portfolio which best covers the liability against the risk factors. Like we anticipated, more funds are allocated to bonds instead of stocks, which makes the portfolio return less volatile and reduces the sensitivity of assets to interest rates by increasing duration. This high allocation to bonds is in line with the findings of Meder and Staub (2006), Moore (2010), Leahy (2011) and Cox *et al.* (2013). This optimal portfolio however still leaves some risk uncovered. For this reason, reserves must be kept so the fund has enough money to pay pensions if adverse scenarios uncover. In order to maintain a probability of insolvency of 5% while accounting for the model risk, 26.91% of the liability or $21,095 million should be injected by the government in addition to filling the funding gap. Conversely, instead of injecting funds in 2012, a

5% probability of insolvency could be reached in a decade by increasing contribution rates by 20%. Alternatively, we could delay the retirement age by 1 year, which would allow reaching the goal in 11 years. The fairness of these solutions is however questionable since contributors would pay for the pensions of inactive members, which would result in an intergenerational transfer where one generation is penalised for the benefit of another generation. The injection of capital by the government is therefore preferred.

Although, we applied the model to the pension fund of the Public Service of Canada, the average Canadian fund has a similar asset allocation as well as a similar pension structure and demographics so that we can generalize our results to other defined benefit funds.

While few models in the literature include more than three risk factors, this thesis includes seven, which allows learning about the impact of more risk factors. Also, the particularity of this thesis is that the LDI approach is used to find the optimal asset allocation and funding strategy. This provides a structure that can be applied by pension fund managers who wish to obtain a different perspective on their risk exposure and learn about potential techniques to manage this exposure. Also, our analysis of the model risk allows evaluating the impact of the choice of process for the risk factors on the valuation measures. In fact, this thesis does not claim that our assumptions regarding the risk factors are more appropriate than the assumptions made by the Office of the Chief Actuary. It rather intends to provide information on the impact of different process choices.

Whereas the basic processes selected in this thesis give a general idea of the risk factors' dynamics, other more complex processes could be used to obtain a more precise model. Another way to improve this study would be to include interest rate and inflation rate swaps and swaptions in the portfolio of assets and evaluate their impact on the optimal solution. Finally, it would be interesting to apply this model to other funds and compare the results.

REFERENCES

Adam, Alexandre (2007). *Handbook of Asset and Liability Management: From Models to Optimal Return Strategies*, John Wiley & Sons, Chichester, England, 550 p.

Andonov, Aleksandar, Rob Bauer and Martijn Cremers (2013). Pension fund asset allocation and liability discount rates: Camouflage and reckless risk taking by US public plans? *Working Paper,* University of Notre Dame, 1-49.

Battocchio, Paolo and Francesco Menoncin (2004). Optimal pension management in a stochastic framework. *Insurance: Mathematics and Economics,* 34.1, 79-95.

Boyer, M. Martin, Lars Stentoft and Joanna Mejza (2011). Measuring longevity risk for a Canadian pension fund. *CIRANO-Scientific Publication,* 2011s-43, 1-32.

Bragt, David van (2011). The pros and cons of liability-driven investing: Finding the right solution. *AEGON Global Pensions View*, May 2011 issue, 1-6.

Bruder, Benjamin, Guillaume Jamet and Guillaume Lasserre (2010). Liability-driven investment. *Lyxor White Paper*, 2, 1-44.

Cairns, Andrew JG, David Blake and Kevin Dowd (2006). A Two-Factor Model for Stochastic Mortality with Parameter Uncertainty: Theory and Calibration. *Journal of Risk and Insurance,* 73.4, 687-718.

Choudhry, Moorad (2007). *Bank Asset and Liability Management: Strategy, Trading, Analysis,* John Wiley & Sons, Singapore, 1415 p.

Cont, Rama (2007). Volatility clustering in financial markets: Empirical facts and agent-based models. *Long Memory in Economics*, Springer Berlin Heidelberg, 289-309.

Cox, John C., Jonathan E. Ingersoll Jr and Stephen A. Ross (1985). A theory of the term structure of interest rates. *Econometrica: Journal of the Econometric Society*, 385-407.

Cox, Samuel H. *et al.* (2013). Managing capital market and longevity risks in a defined benefit pension plan. *Journal of Risk and Insurance,* 80.3, 585-620.

Davies, Ryan J., Harry M. Kat and Sa Lu (2009). Fund of hedge funds portfolio selection: A multiple-objective approach. *Journal of Derivatives & Hedge Funds,* 15.2, 91-115.

De Jong, Frank (2008). Pension fund investments and the valuation of liabilities under conditional indexation. *Insurance: Mathematics and Economics,* 42.1, 1-13.

Delong, Lukasz, Russell Gerrard and Steven Haberman (2008). Mean–variance optimization problems for an accumulation phase in a defined benefit plan. *Insurance: Mathematics and Economics,* 42.1, 107-118.

Elen, Emile A.L.J. van (2010). Term structure forecasting: Does a good fit imply reasonable simulation results? A thesis submitted to Tilburg School of Economics and Management, Tilburg University, Netspar Theses, 1-40.

Engel, Janwillem P.W., Harry M. Kat and Theo P. Kocken (2005). Strategic interest rate hedges, or how derivatives can help solve the pension fund crisis Part II. *Alternative Investment Research Centre Working Paper*, 0024, 1-22.

Fabozzi, Frank J. (1997). *Pension Fund Investment Management*, Vol. 25, John Wiley & Sons, New Hope, Pennsylvania, 299 p.

Friedman, Milton (1968). The Role of Monetary Policy. *American Economic Review*, 58, 1-17.

Girosi, Federico and Gary King (2007). Understanding the Lee-Carter mortality forecasting method. *Princeton University Press*, 1-35.

Goldsmith, Theodore (2011). *An Introduction to Biological Aging Theory*, second edition, Azinet Press, 1-44.

Hull, John and Alan White (1990). Pricing interest-rate-derivative securities. *Review of financial studies,* 3.4, 573-592.

Inkmann, Joachim and David Blake (2004). Liability valuation and optimal asset allocation. *Working Paper*, London School of Economics, 1-27.

Josa-Fombellida, Ricardo and Juan Pablo Rincón-Zapatero (2008). Funding and investment decisions in a stochastic defined benefit pension plan with several levels of labor-income earnings. *Computers & Operations Research,* 35.1, 47-63.

Josa-Fombellida, Ricardo and Juan Pablo Rincón-Zapatero (2010). Optimal asset allocation for aggregated defined benefit pension funds with stochastic interest rates. *European Journal of Operational Research,* 201.1, 211-221.

Leahy, Matthew P. (2011). Liability driven investing. *Congress Wealth Management White Paper,* June 2011 issue, 1-8.

Lee, Ronald D. and Lawrence R. Carter (1992). Modeling and forecasting US mortality. *Journal of the American Statistical Association,* 87.419, 659-671.

Martellini, Lionel (2006). Managing pension assets: From surplus optimization to liability-driven investment. *Working Paper,* EDHEC Business School, 1-18.

Martellini, Lionel and Vincent Milhau (2008). The benefits of structured products in asset-liability management. *Working Paper,* EDHEC Business School, 1-31.

Meder, Aaron and Renato Staub (2006). Linking pension liabilities to assets. *UBS Global Asset Management Working Paper,* March issue, 1-11.

Merton, Robert C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics,* 3.1, 125-144.

Moore, David W. (2010). Understanding duration risk in pension plans: The case for LDI. *New England Pension Consultants,* January issue, 1-6.

Morneau Shepell. Economic assumptions in accounting for pension and other post-retirement benefits. *2013 Survey,* 1-10.

Ostaszewski, Krzysztof (2011). Determination of optimal contribution rate and optimal investment portfolio of defined benefit pension plan under the expected shortfall constraint. *Working Paper*, Illinois State University, 1-25.

PSP Investment, Public Service Pension Plan account financial statements as at March 2014, 1-26.

PSP Investment, Statement of investment policies, standards and procedures for assets managed by the Public Sector Pension Investment Board, March 2012, 1-16.

Public Service of Canada, Actuarial Report on the Pension Plan for the Public Service of Canada, 1980-2011, Office of the Chief Actuary.

Rudd, Jeremy and Karl Whelan (2007). Modeling inflation dynamics: A critical review of recent research. *Journal of Money, Credit and Banking,* 39.s1, 155-170.

Service, David and Jie Sun (2004). Asset allocation in the light of liability cash flows. *Institute of Actuaries of Australia Financial Service Forum*, Centre for Actuarial Research, Australian National University. 1-23.

Tankov, Peter and Ekaterina Voltchkova (2009). Jump-diffusion models: A practitioner's guide. *Banque et Marchés,* 99, 1-24.

Vasicek, Oldrich (1977). An equilibrium characterization of the term structure. *Journal of financial economics,* 5.2, 177-188.

Wouters, Anton (2008). Implementing asset-liability management. *Risk and Rewards,* 51, Society of Actuaries, 1-7.

Zenios, Stavros A. and William T. Ziemba (2007). *Handbook of Asset and Liability Management Volume 2: Applications and Case Studies,* North-Holland, Amsterdam, 1147 p.

APPENDIX

Stochastic processes used and their estimated parameters

Process	Equation	Parameters
Inflation rate: Brownian motion	$\frac{dX_t}{X_t} = \mu dt + \sigma dW_t$	μ = 0.0185 σ = 0.0073
Interest rate: Vasicek process	$dr_t = \theta[\mu - r_t]dt + \sigma dW_t$	θ = 0.0083 μ = 0.0010 (1.2% annual) σ = 0.0001
Earnings growth: Brownian motion	$\frac{dE_t}{E_t} = \mu dt + \sigma dW_t$	μ = 0.0350 σ (male) = 0.0419 σ (female) = 0.0449
Mortality rates: Lee-Carter model	$\ln(m_{xt}) = a_x + b_x k_t + e_{xt}$	

	a_x		b_x	
	Male	Female	Male	Female
20-24	-7.9637	-8.4082	-0.2656	-0.1537
25-29	-7.6760	-8.1850	-0.1147	0.0663
30-34	-7.4529	-8.0027	-0.2375	-0.3506
35-39	-7.2706	-7.8486	-0.2504	-0.3056
40-44	-7.1164	-7.7150	-0.2345	-0.2407
45-49	-6.4974	-6.9613	-0.2215	-0.1853
50-54	-6.1179	-6.5364	-0.1545	-0.1542
55-59	-5.4380	-5.8055	-0.2121	-0.2414
60-64	-5.0370	-5.3878	-0.2943	-0.3142
65-69	-4.3131	-4.7193	-0.3334	-0.3080
70-74	-3.8976	-4.3222	-0.3522	-0.2952
75-79	-3.1556	-3.5940	-0.3390	-0.3061
80-84	-2.7344	-3.1771	-0.2896	-0.2794
85-90	-2.0767	-2.4325	-0.2655	-0.2911
90+	-1.6833	-2.0104	-0.1822	-0.2100

	k_t	
	Male	Female
2012	0.4208	0.5983
2013	0.4779	0.6838
2014	0.5399	0.7717
2015	0.6010	0.8554
2016	0.6577	0.9392
2017	0.7198	1.0292
2018	0.7770	1.1125
2019	0.8379	1.1996
2020	0.8968	1.2813
2021	0.9588	1.3693
2022	1.0165	1.4521
2023	1.0729	1.5411
2024	1.1351	1.6272
2025	1.2023	1.7085
2026	1.2585	1.7973
2027	1.3159	1.8823
2028	1.3793	1.9673
2029	1.4394	2.0531
2030	1.4960	2.1393
2031	1.5503	2.2228

Table 15