



Henley Business School

ICMA Centre

Essays on Robust Portfolio Selection and Pension Finance

Thesis submitted in partial fulfilment of the requirements for
the degree of Doctor of Philosophy

by

Emmanouil Platanakis

Reading, October 2015

To my parents, George and Kalliopi

"μή μου τοὺς κύκλους τάραττε!" - Αρχιμήδης

Declaration

I confirm that this is my own work and the use of all material from other sources has been properly and fully acknowledged.

Date: _____

Emmanouil Platanakis

Acknowledgements

First and foremost, I would like to express my deepest gratitude to my supervisor, Charles Sutcliffe (ICMA Centre – University of Reading) for his guidance, support and faith that he showed to me. Charles has been very encouraging, helpful and thoughtful, while his comments and feedback have been direct and insightful. Also, Charles has made me think harder about the motivations, potential contributions, importance and positioning of my work.

Furthermore, I would like to thank the ICMA Centre for the doctoral scholarship that it has granted me during my doctoral studies as well as their generous additional funding in order to cover my travel expenses for various conference presentations. Many faculty members of the ICMA Centre have also assisted me in a number of ways. In particular, discussions with Ioannis Oikonomou, Chris Brooks, Sotiris Tsolacos, Nicholas Chen and Alfonso Dufour have been very inspiring.

I would also like to express my appreciation for my fellow doctoral candidates Andrei Stancu and Chris Godfrey for helping me towards the data collection process. In addition, I must thank the participants of the 4th International Conference of the Financial Engineering and Banking Society (Surrey, UK), 8th Portuguese Finance Network International Conference (Vilamoura, Portugal), 12th Annual International Conference on Finance (Athens, Greece) and the Annual Workshop of the Dauphine-Amundi Chair in Asset Management (Paris, France) for their comments on the paper presented in Chapter 3, as well as two anonymous referees for the European Journal of Finance where the study presented in Chapter

3 has been published after 3 rounds of revision. Furthermore, I must thank the participants of the 5th International Conference of the Financial Engineering and Banking Society (Nantes, France), 4th European Business Research Conference (London, UK) and the ICMA Centre Internal Research Seminar (Reading, UK) for their valuable comments on the study presented in Chapter 4. Finally, I must also thank the participants of the ICMA Centre Internal Research Seminar (Reading, UK) for their comments as well as an anonymous referee for the Insurance: Mathematics and Economics journal where the study presented in Chapter 5 has been accepted for revise and resubmit.

Finally, my special gratitude goes to my parents, George and Kalliopi, and to my brother, Konstantinos, who have always incentivized me to pursue education and achieve my targets. They have always been a source of inspiration and I think that I would not be able to pursue this doctoral degree without their support and motivation.

Abstract

This thesis examines three different, but related problems in the broad area of portfolio management for long-term institutional investors, and focuses mainly on the case of pension funds. The first idea (Chapter 3) is the application of a novel numerical technique – robust optimization – to a real-world pension scheme (the Universities Superannuation Scheme, USS) for first time. The corresponding empirical results are supported by many robustness checks and several benchmarks such as the Bayes-Stein and Black-Litterman models that are also applied for first time in a pension ALM framework, the Sharpe and Tint model and the actual USS asset allocations. The second idea presented in Chapter 4 is the investigation of whether the selection of the portfolio construction strategy matters in the SRI industry, an issue of great importance for long term investors. This study applies a variety of optimal and naïve portfolio diversification techniques to the same SRI-screened universe, and gives some answers to the question of which portfolio strategies tend to create superior SRI portfolios. Finally, the third idea (Chapter 5) compares the performance of a real-world pension scheme (USS) before and after the recent major changes in the pension rules under different dynamic asset allocation strategies and the fixed-mix portfolio approach and quantifies the redistributive effects between various stakeholders. Although this study deals with a specific pension scheme, the methodology can be applied by other major pension schemes in countries such as the UK and USA that have changed their rules.

Table of Contents

Acknowledgements.....	iv
Abstract.....	vi
List of Figures	x
List of Tables.....	xi
1 Introduction and Overview	1
2 Literature Review	9
2.1 Introduction	9
2.2 Markowitz Portfolio Theory	10
2.3 Modern Portfolio Theory Criticisms.....	10
2.4 Portfolio Approaches Dealing with Estimation Risk.....	12
2.4.1 Bayes' Estimators	13
2.4.2 Constraints on Portfolio Weights.....	13
2.4.3 'Resampled Efficiency'	14
2.4.4 Optimal Mixture Portfolios	15
2.4.5 Portfolios with Moments Restrictions	16
2.4.6 Robust (Worst-Case) Optimization	16
2.5 Performance Measures and Portfolio Evaluation.....	19
2.5.1 Performance Measures based on Lower Partial Moments	20
2.5.2 Performance Measures based on Drawdown.....	21
2.5.3 Performance Measures based on Value-at-Risk.....	22
2.6 Asset Liability Management (ALM)	22
2.7 Operational Research and ALM Techniques	25
2.7.1 Stochastic Programming ALM Models	27
2.7.2 Single Period ALM Models (Portfolio Theory).....	28
2.7.3 Stochastic Simulation ALM Models.....	29
2.7.4 Dynamic Programming and Stochastic Control ALM Models	30
2.8 Techniques for Generating Scenarios	31
2.8.1 Sampling Approach	31
2.8.2 Statistical Approach	32
2.8.3 Alternative Techniques	32
2.9 Pension Schemes Design and Intergenerational Transfers.....	33
2.9.1 Defined Benefit (DB) Pension Schemes	34

2.9.2 Defined Contribution (DC) Schemes	36
2.9.3 Intergenerational Effects	38
2.10 Conclusions	41
3 Asset Liability Modelling and Pension Schemes: The Application of Robust Optimization to USS	44
3.1 Introduction	44
3.2 Robust Optimization ALM Model.....	50
3.3 Application to USS.....	62
3.4 Data and Analysis	64
3.5 Results	75
3.6 Robustness Checks.....	83
3.7 Conclusions	87
Appendix 3.A – Projected Contribution Rate.....	89
Appendix 3.B - Uncertainty Sets	89
Appendix 3.C – Conversion to a Second Order Cone Problem	91
Appendix 3.D - Actuarial Liability Models.....	96
4 Socially Responsible Investment Portfolios: Does the Optimization Process Matter?.....	99
4.1 Introduction	99
4.2 Related literature and motivation of the study	103
4.3 Model and dataset	115
4.3.1 Portfolio evaluation metrics.....	115
4.3.2 Dataset	119
4.4 Results	122
4.4.1 Main results.....	122
4.4.2 Robustness tests	136
4.4.3 Additional analyses	140
4.5 Conclusions	144
Appendix 4.A: Mathematical definition of implemented optimization approaches	147
4.A.1 Markowitz portfolio optimization.....	147
4.A.2 Robust estimation approach.....	147
4.A.3 Black-Litterman approach.....	149
4.A.4 Risk-parity portfolio construction.....	150
4.A.5 Reward-to-risk timing portfolios.....	151

Appendix B: Indicators for each CSP dimension	152
Chapter 5.....	153
5 Pension Scheme Redesign and Wealth Redistribution Between the Members and Sponsor: The USS Rule Change in October 2011	154
5.1 USS	158
5.2 Methodology.....	159
5.3 Literature Review	162
5.4 Forecasting Asset Returns, Inflation and Salaries	163
5.5 Forecasting the Size of the Age Cohorts	167
5.6 Forecasting the Liabilities.....	168
5.7 Generating the Cohort Cash Flows	171
5.8 Redistribution.....	175
5.9 Results	181
5.10 Robustness Checks.....	186
5.11 Conclusions	188
Appendix 5.A - USS Rules Pre and Post-October 2011	191
Appendix 5.B.1 Final Salary (FS) Scheme	193
Appendix 5.B.2 Career Revalued Benefit (CRB) Scheme.....	195
Appendix 5.C Time Varying Price of Risk.....	196
Chapter 6.....	206
6 Conclusions and Future Research	207
6.1 Summary and Conclusions	207
6.2 Suggestions for Future Research	211
References.....	215

List of Figures

Figure 3.1.....	66
Figure 4.1.....	128
Figure 4.2.....	130
Figure 4.3.....	133
Figure 5.1.....	197
Figure 5.2.....	197
Figure 5.3.....	198
Figure 5.4.....	198
Figure 5.5.....	199
Figure 5.6.....	199
Figure 5.7.....	200
Figure 5.8.....	200
Figure 5.9.....	201
Figure 5.10.....	201
Figure 5.11.....	202
Figure 5.12.....	202
Figure 5.13.....	203

List of Tables

Table 2.1.....	37
Table 3.1.....	65
Table 3.2.....	68
Table 3.3.....	69
Table 3.4.....	70
Table 3.5.....	75
Table 3.6.....	77
Table 3.7.....	80
Table 3.8.....	86
Table 4.1.....	122
Table 4.2.....	122
Table 4.3.....	124
Table 4.4.....	132
Table 4.5.....	137
Table 4.6.....	139
Table 4.7.....	139
Table 4.8.....	139
Table 4.9.....	143
Table 5.1.....	204
Table 5.2.....	183

Chapter 1

1 Introduction and Overview

This thesis examines three different, but related problems in the broad area of portfolio management for long-term institutional investors, and focuses mainly on the case of pension funds. In what follows, we provide a comprehensive overview of the following chapters.

The main objective of Chapter 2 is to introduce the reader to the basic aspects of portfolio theory, asset liability management (ALM) modeling and pension scheme design – the key elements of Chapters 3, 4 and 5. In particular, the fundamental Markowitz (1952) portfolio theory (mean-variance portfolio optimization), the most important criticisms of the mean-variance portfolio optimization framework, the latest developments in portfolio techniques that deal with estimation errors in the input data as well as alternative metrics used for portfolio evaluation are discussed. Second, Chapter 2 describes the benefits of the asset-liability management (ALM) models for long-term institutional investors such as insurance companies, pension and endowment funds, instead of ignoring the corresponding liabilities and using asset-only portfolio strategies. In addition, Chapter 2 describes the close relation between Operations Research (OR) and ALM modelling, provides a review of the most important techniques used in the computation of optimal ALM strategies (e.g. stochastic programming, portfolio theory, stochastic simulation, dynamic programming and stochastic control) as well as the most popular methods for generating scenarios, and explains the issue of the computational intractability of scenario-based ALM methods in solving realistic asset-liability management

problems. Third, Chapter 2 provides a detailed description of the two dominant types of pension schemes (defined benefit and defined contribution) in the UK and US, and gives a comprehensive review of the relevant pension studies that compute the intergenerational redistributive effects between the stakeholders of a pension scheme, which occur after various pension rule changes.

In Chapter 3, a novel numerical portfolio optimization technique – robust (worst-case) optimization – is used to formulate and solve the asset-liability management (ALM) problem for a real-world pension scheme (USS). Robust optimization is particularly well suited to solving the ALM problem since it assumes that the uncertain input data are not known with certainty, but they lie within uncertainty sets. Hence, it adopts a maximin approach by solving the ‘worst-case’ optimization problem under the assumption that the stochastic input parameters used in the optimization problem take the worst-case values within the uncertainty structures defined by the modeler. As a result, robust optimization deals with estimation risk by ruling out possible optimal solutions that promise superior performance due to statistical misspecifications in the input data. Furthermore, the robust formulated ALM model presented in Chapter 3 incorporates additional important characteristics of the pension asset-liability management model such as upper and lower bounds for each asset class, prohibits borrowing and short sales and imposes a non-negative constraint on the expected value of the asset-liability portfolio return. Problems formulated with robust optimization techniques are easily solved and can handle problems with large data requirements in a more computationally tractable manner than scenario based approaches (e.g. stochastic programming and stochastic

simulation), and often requires the estimation of fewer stochastic parameters and reduces estimation risk.

This is the first application of a computationally tractable (easily solved) asset-liability management model based on robust optimization techniques to a real-world pension scheme (the Universities Superannuation Scheme, USS), and the first pension asset-liability management model that maximizes the Sharpe ratio. Also, the pension liabilities in Chapter 3 are split into three categories – active members, deferred members and pensioners, and the optimal asset allocation is transformed into the overall contribution rate. In addition, the proposed pension asset-liability management framework is benchmarked against various important benchmarks such as the actual USS performance computed by using the actual USS asset allocation decisions as well as the Sharpe and Tint, Bayes-Stein, and Black-Litterman models. The empirical results reveal that robust (worst-case) optimization has a clearly superior out-of-sample performance than the four benchmarks across 20 performance metrics that measure many different important portfolio characteristics such as risk, risk-adjusted performance, second-order stochastic dominance, diversification, stability, contribution rate, funding ratio and cumulative wealth amongst others. Finally, the conclusions remain unchanged by various robustness checks such as by testing different estimation and investment periods, relaxing the constraints on asset weights, and by using an alternative set of asset classes and different uncertainty sets with a smaller size.

Chapter 4 investigates whether the selection of the portfolio optimization strategy matters in the SRI industry and provides some answers to the question of which

portfolio approaches tend to create SRI portfolios with better out-of-sample performance, given certain socially responsible investment criteria. This issue is of great importance for institutional investors since it is well known that long-term institutional investors such as pension funds, life insurance firms and endowment funds are committed to corporate social performance (CSP) and socially responsible investment (SRI). Although the size of the SRI literature is very broad, the number of studies that have investigated different ways of optimal SRI portfolio construction is very limited. Most importantly, these studies are very limited as the portfolio methods used are very heavily based on the Markowitz (1952) portfolio theory (mean-variance portfolio framework), and hence ignore the negative effects of estimation risk and parameter uncertainty in the corresponding input data. Such studies simply use the mean-variance portfolio framework by adding additional constraints with SRI preferences or incorporate these preferences by altering the objective function of the optimization process. Although the large majority of modern portfolio optimization techniques are heavily based on Markowitz (1952) portfolio theory, their optimal portfolio solutions are highly sensitive to perturbations in the input parameters, see for instance Green and Hllofield (1992), and often lead to portfolios that are poorly diversified, unstable and with a weak out-of-sample performance.

To construct socially responsible investment (SRI) portfolios with only US companies in Chapter 4, this study uses corporate social performance (CSP) metrics based on the MSCI ESG STATS (MSCI KLD) database. It is well known that this database is the most popular in the relevant research, and according to Sharfman (1996) it is

described as a reliable and consistent database that contains about 3,000 US firms over a time horizon of over 20 years. This study attempts to contribute to the existing literature by employing three different optimal portfolio diversification methods (Markowitz, norm-constrained and Black-Litterman portfolios) and three more simplistic asset allocation techniques (equally weighted, risk-parity and reward-to-risk timing portfolios) to the same SRI-screened universe. Out-of-sample performance comparisons are made between these six different portfolio construction methods, and 14 different performance measures are used to capture several important characteristics such as risk and risk-adjusted performance, diversification and stability, amongst others. The empirical results show that more 'formal' portfolio optimization methods (Markowitz, Black-Litterman and norm-constrained portfolios) tend to construct less risky Socially Responsible investing (SRI) portfolios with superior risk-returns trade-offs and a significantly smaller number of 'active' assets than more simplistic asset allocation techniques (1/N, risk parity and reward-to-risk). The Black-Litterman portfolio approach often comes first, while naïve diversification (1/N) usually has the worst performance on these criteria. Finally, the main conclusions are robust to a variety of additional tests that include stricter screening criteria for the construction of socially responsible investment portfolios, the use of estimation windows with a different length than the base case as well as different ways of evaluating the out-of-sample portfolio performance.

Chapter 5 deals with the short, medium and long term performance of a real world pension scheme (the Universities Superannuation Scheme, USS) before and after the rule changes that took place in October 2011, as well as the wealth redistribution

between various age cohorts of the future and active members, pensioners, and the sponsor that occur as a result of the pension scheme rule changes. Specifically, in October 2011 USS closed the final salary (FS) scheme to the new members, where the sponsor bears all the risks such as investment, longevity, interest rate, inflation, salary growth and regulatory risk, and forced new members to join the newly established career average revalued earnings (CARE) scheme, while USS also introduced a 'cap and share' rule for setting contribution rates.

The study presented in Chapter 5 also includes many important aspects that have not previously been incorporated in pension studies such as lump sum payments, deferred members (members that have left the scheme and are not currently paying contributions), spouses' pensions, both final salary (FS) and career average revalued earnings (CRB) sections and a dynamic retirement age. Furthermore, the framework that simulates the pension scheme in Chapter 5 is modeled for a period longer than the working life. It also employs three different asset allocation strategies; the fixed-mix, risk-shifting and risk management approach, with the last two strategies responding to the funding ratio each time the portfolio is rebalanced. A stable over time vector auto-regressive (VAR) model with 13 variables is used to generate future asset returns, inflation rates and the factors of the Nelson-Siegel yield curve, while the population of active and deferred members of the pension scheme each year follows a stochastic process. Although this study mainly focuses on a particular pension fund (USS), the general methodology presented in Chapter 5 could also be applied by other schemes in countries such as the UK and USA that have changed their rules.

The empirical results presented in Chapter 5 reveal that the post-October 2011 USS scheme (the pension scheme after the rule changes in October 2011) is sustainable in the long run with some problems in the mid-term, in contrast to the pre-October 2011 scheme that is non-viable in the long run. Also, the fixed-mix and risk-shifting asset allocation strategies are more favorable than the risk-management approach for both the pre and post 2011 schemes. The quantification of the redistributive effects due to the pension rule changes in October 2011 shows that future members lose about the 65% of their pension wealth (or an 11% reduction in their overall compensation) with an increase in the risk of their pension wealth by about a third, in contrast to the older age-cohorts where the corresponding losses are insignificant and the risk of their pension wealth is almost the same after the rule changes. The sponsor's pension costs decrease by about 26%. Finally, the main conclusions remain unchanged by trying various robustness checks such as the replacement of the stochastic discount factors (SDFs) with the riskless discount rates for the computation of the NPVs and the use of different upper bounds on the total contribution rate.

Finally, Chapter 6 summarizes the main ideas, scientific contributions and the corresponding outputs of each chapter separately, and provides some possible directions that could be investigated and explored for future research.

Chapter 2

2 Literature Review

2.1 Introduction

The main purpose of this chapter is to familiarize the reader with the literature in portfolio theory, asset liability management (ALM) modelling and pension schemes design. In the following section (2.2), the fundamental Markowitz portfolio theory is discussed, while section 2.3 describes the criticisms of modern portfolio theory, which have been discussed in the literature the recent years, such as its high sensitivity to estimation risk and parameter uncertainty. Section 2.4 provides a comprehensive review of portfolio techniques dealing with estimation risk, while section 2.5 discusses alternative performance measures used for portfolio evaluations, such as measures based on lower partial moments, drawdown and value-at-risk. Section 2.6 gives the motivation behind the use of ALM techniques by long-term institutional investors (e.g. pension funds) instead of just using asset-only approaches. Section 2.7 explains how Operations Research (OR) is involved in the process of deriving optimal ALM strategies, describes the main methods used in ALM modelling and explains the disadvantages in applying scenario-based techniques to solve realistic ALM problems. Section 2.8 provides a review of techniques for generating scenarios. In addition, section 2.9 describes the two main types of pension schemes (defined benefit and defined contribution) that are dominant in the UK and US and provides a review of studies that deal with the intergenerational redistributive effects after pension scheme rule changes. Finally, a conclusion is provided in section 2.10.

2.2 Markowitz Portfolio Theory

Portfolio selection is the problem of capital allocation over a set of available assets by maximizing the portfolio return and minimizing the corresponding risk. Although the benefits that occur by diversifying a portfolio have been reported since the beginning of financial markets, i.e. reducing portfolio risk, Markowitz (1952, 1959) is the first that formulated a mathematical framework for optimal portfolio selection, the so called mean-variance portfolio framework. Markowitz (1952) uses the expected value and the variance of the random portfolio return to compute the total return and the associated risk respectively, showing that the mathematical problem can be represented as a convex quadratic program (the well known efficient frontier) by assuming either an upper cap on the variance, or a lower bound on the portfolio return. The Markowitz portfolio allocation model had an apparent effect on the asset pricing and financial economic modeling. For instance, the Capital Asset Pricing Model, Sharpe (1964), Lintner (1965) and Mossin (1966), was a direct and logical result of the Markowitz portfolio theory. Sharpe and Markowitz won the Nobel Prize in Economic Sciences in 1990 for their scientific contribution to asset allocation and asset pricing.

2.3 Modern Portfolio Theory Criticisms

In practice, the application of Markowitz (1952) portfolio theory requires the estimation of the mean and covariance matrix of the asset returns, and it has been widely reported that Markowitz theory is very sensitive to estimation errors in the input estimates (i.e. mean and covariance matrix), see for instance Green and

Hollifield (1992), Goldfarb and Iyengar (2003), Ceria and Stubbs (2006), DeMiguel et al. (2009a), Glasserman and Xu (2013) amongst others. Merton (1980) also points out that an accurate estimation of the means is a much harder task than estimating precisely the covariance matrix, while Kallberg and Ziemba (1984) mention that statistical errors in means are approximately 10 times as significant as errors in covariances. In other words, it means that if the sample means and covariances are not accurate, the optimal asset weights computed via the Markowitz mean-variance portfolio optimization model contains significant errors, since the optimal solutions are highly sensitive to disturbances in the input data of the problem. As a direct consequence, the Markowitz model leads to 'investment-irrelevant' portfolios according to Michaud (1999), which are unstable, poorly diversified and are characterized by a weak out-of-sample performance.

This situation has been very well described and investigated in the financial portfolio literature. For example, Michaud (1999) points out that, although Markowitz portfolio theory is a useful theoretical tool for portfolio optimization that can be applied easily in practice, many practitioners and academics have kept this theoretical framework at a distance. This is because it works like an 'error-prone procedure' and leads to unreliable portfolios with large errors due to its extreme sensitivity to perturbations in the input data of the mathematical optimization process. Goldfarb and Iyengar (2003), Ceria and Stubbs (2006), Glasserman and Xu (2013), Xing et al. (2014) and others also highlight that this phenomenon is a direct consequence of the fact that the solutions computed by the mean-variance portfolio method are very susceptible to disturbances in the parameters of the optimization

process, since these estimates are subject to large statistical errors. In addition, Kallberg and Ziemba (1984) investigate the case of misspecification in means, variances, covariances and investor's utility functions in normally distributed asset allocation problems, while Best and Grauer (1991) provide some empirical evidence on the susceptibility of optimal mean-variance asset allocations to disturbances in the means. Also, Broadie (1993) investigates the effect of the statistical errors in the input estimates on the construction of the efficient frontier, Chopra (1993) examines the relation between portfolio diversification and estimation risk, Chopra and Ziemba (1993) examine the equivalent loss by using the estimated instead of the true parameters in the mean-variance portfolio process, while a more inclusive assessment on the effect of statistical errors on portfolio selection can be found in Ziemba and Mulvey (1998).

2.4 Portfolio Approaches Dealing with Estimation Risk

Since portfolio theory is very sensitive to estimation risk that can often result in 'error-maximized' portfolios according to Michaud (1999), there are six main approaches that deal with estimation errors in the input parameters in an attempt to constructing superior portfolios with better characteristics according to DeMiguel et al. (2009b). These approaches are described below. The first approach involves alternative estimates of the means and covariance matrix, the second sets constraints on portfolio weights, the third uses simulations to generate alternative input data (portfolio re-sampling), the fourth computes optimal combinations of portfolios, the fifth uses moment restriction techniques to eliminate estimation risk,

and finally the sixth approach involves worst-case portfolio optimization (robust optimization).

2.4.1 Bayes' Estimators

The first approach tries to improve portfolio performance in the presence of estimation risk by altering the estimation of the means, variances and covariances (e.g. via alternative estimates). For instance, Black and Litterman (1992) as well as Drobetz (2001), Bessler et al. (forthcoming) and other studies combine neutral returns and subjective returns (views) by allowing investors to provide estimates for some asset returns or staying neutral on some others. The reliabilities of return estimates can be quantified and incorporated in the mean-variance portfolio optimization process. Furthermore, Jobson et al. (1979) try to enhance the performance of Markowitz portfolios by using James-Stein type estimators in the portfolio optimization problem, while Jorion (1986) and Frost and Savarino (1986) propose empirical Bayes estimators for the means, variances and covariances and try to eliminate extreme asset allocations (corner solutions) by reducing estimation risk. More recent studies such as Ledoit and Wolf (2003, 2004) propose a shrinkage approach to the covariance matrix estimator by using factor and Bayesian models, and Kourtis et al. (2012) attempt to improve portfolio performance by shrinking directly the inverse covariance matrix using two non-parametric methods.

2.4.2 Constraints on Portfolio Weights

The second approach attempts to eliminate estimation errors by setting constraints on the portfolio weights. Frost and Savarino (1988) provide evidence that estimation risk is significantly reduced by constraining asset weights in mean-variance portfolio

selection strategies. Board and Sutcliffe (1994) compare the Bayes-Stein estimation model with seven alternative estimation methods, finding that there is little to select amongst them when short selling is prohibited. Also, some recent studies use more sophisticated technicalities by imposing constraints on portfolio norms, such as DeMiguel et al. (2009a). For instance, Brodie et al. (2009) apply l_1 norm (the taxicab norm – defined as the sum of the absolute values of portfolio weights) constraints on asset weights within the portfolio optimization process to encourage the construction of sparse portfolios, e.g. portfolios with only a few active assets (assets with nonzero weights). In addition, Tola et al. (2008) and Xing et al. (2014) impose constraints on a combination of norms on portfolio weights l_1 and l_∞ (maximum norm), with the latter to be defined as the maximum absolute value of the portfolio weights, in an attempt to construct sparse-style portfolios, e.g. sparse portfolios as explained above by eliminating at the same time the possibility of large weights to be allocated in just a few assets. Finally, It has been reported in the literature that sparse-style portfolios are often very well diversified and usually have a better out-of-sample performance in terms of risk and risk adjusted return in comparison to naïve forms of portfolio optimization, see Xing et al. (2014) and others.

2.4.3 'Resampled Efficiency'

The third approach computes optimal portfolios by generating many data sets with simulation techniques (e.g. Monte Carlo), with the overall solution given by averaging these optimal portfolios. In particular, first Michaud (1999) proposes the so-called 'Resampled Efficiency' (RE) technique. 'Resampled Efficiency' uses Monte Carlo simulation methods to satisfactorily replicate parameter uncertainty in an

attempt to compute optimal mean-variance portfolios with better out-of-sample characteristics. Scherer (2002) provides a comprehensive review of the concept of portfolio resampling proposed by Michaud (1999), and indicates some possible weak points of the 'Resampled Efficiency' portfolio technique. In addition, Becker et al. (2015) carry out a complete simulation study with both constrained and unconstrained portfolio optimization processes for a variety of estimators. Although their empirical findings indicate that Markowitz (1952) overall performs better than Michaud (1999), they also provide strong evidence that the Markowitz mean-variance portfolio framework is more sensitive than Michaud (1999) to changes on asset weights' constraints and to different estimators used for the mean and variance of asset returns.

2.4.4 Optimal Mixture Portfolios

The forth approach computes portfolios that are actually combinations of other portfolios, such as the minimum variance, the mean-variance and the 1/N portfolios. Hence, the intuition behind these 'mixture' portfolios is to shrink directly the portfolio weights. For instance, Kan and Zhou (2007) are motivated by the idea that estimation risk may not be efficiently diversified by holding only a combination of the tangency portfolio (the portfolio with the highest Sharpe ratio) and the risk-free asset, and propose the 'three-fund' portfolio rule in the class of mixture portfolios that combine the mean-variance and minimum-variance portfolios, in which the functionality of the 'third' fund is to eliminate estimation risk. Furthermore, DeMiguel et al. (2009b) differentiate their position from Kan and Zhou (2007), and propose a mixture of equally weighted and minimum-variance portfolios. Their main

intuition is to put more emphasis on the estimation of covariances instead of the means given that it is well accepted that the estimation of expected returns is a much more difficult task than covariances.

2.4.5 Portfolios with Moments Restrictions

'The fifth approach attempts to decrease the negative effects of estimation risk by constructing portfolios with moment restrictions. DeMiguel et al. (2009b) describe 3 different portfolio strategies that set restrictions on the estimation of the statistical moments of the asset returns, and these are the well-known minimum variance portfolio, the value-weighted portfolio implied by the market model and portfolios constructed by asset-pricing models with unobservable factors. For the latter strategy, MacKinlay and Pastor (2000) show that the covariance matrix of the residuals error terms of the factor model contains any resulting mispricing due to unobservable factors.'

2.4.6 Robust (Worst-Case) Optimization

Robust optimization is a relatively new numerical method that has grown thanks to the rapid improvement of the computing technology in the last few years. Although robust optimization overlaps with stochastic and dynamic programming, it can be assigned to its own category and consists the forth approach of dealing with estimation risk in the broad area of portfolio optimization. Robust optimization adopts a maximin approach and formulates 'worst-case' optimization problems, the so called 'robust counterparts'. In particular, it assumes that the uncertain/stochastic input parameters of the portfolio optimization process are not known with certainty, but lie within uncertainty sets. Hence, it tries to eliminate the possibility of selecting

portfolios that promise good performance due to estimation errors by computing optimal portfolio solutions under the assumption that the uncertain input data of the optimization problem take the worst-case values within these uncertainty structures (worst-case scenario). At this point, we also have to point out that the size and shape (e.g. interval, ellipsoidal, polygonal) of the uncertainty sets play a major role in this numerical method since they alter the level of conservativeness of the asset allocation, characterize the risk preferences of investors and most importantly result in computationally tractable (e.g. easily solved) mathematical programming problems. For further technical details of the mathematics of robust optimization, we refer to Fabozzi et al. (2007) amongst others.

Robust optimization has attracted significant interest in recent years by institutional investors and academics since they consider it is a strong and very efficient method for computing optimal asset allocations subject to estimation risk in the input data, see for instance Ben-Tal and Nemirovski (1998). So far, robust optimization has been applied only in equity portfolio management by taking into account uncertainties due to estimation errors in the input parameters as far as mean-variance portfolio strategies concerned. Gabrel et al. (2014) describe the recent advances in worst-case optimization in all areas of science (e.g. Engineering, Medicine, Finance etc.) since 2007.

Ben-Tal and Nemirovski (1999) consider linear programs with stochastic data and by introducing ellipsoidal uncertainty sets, they show that the equivalent robust programming problem (robust counterpart) is easily solved in polynomial time via conic quadratic programming, while Ben-Tal and Nemirovski (1998) study the

computational efficiency of robust formulated convex optimization problems. Costa and Paiva (2002) derive computationally tractable robust optimal asset allocation problems for tracking errors by assuming polytopic uncertainty for the means and covariances, while Goldfarb and Iyengar (2003) derive easily solved robust mean-variance portfolio problems under a variety of interval and ellipsoidal uncertainty structures for the uncertain market parameters using factors models. Ceria and Stubbs (2006) propose a zero-net alpha-adjustment robust portfolio optimization technique under ellipsoidal uncertainty structures to construct robust mean-variance portfolios that are less conservative by assuming that the estimation errors on the input parameters do not necessarily have negative effects on the portfolio performance. In addition, Ben-Tal, Margalit and Nemirovski (2000) as well as Bertsimas and Pachamanova (2008) construct easily solved robust multi-period asset allocation problems using ellipsoidal uncertainty sets. El Ghaoui et al. (2002) develop tractable worst-case value-at-risk robust programming portfolio models via conic programming, while Quaranta and Zaffaroni (2008) use robust portfolio techniques to construct tractable robust models that minimize the conditional value at risk of a stock portfolio. Also, Huang et al. (2010) differentiate from the existing studies, most of which consider uncertainty just for the means and the covariance matrix, and provide a worst-case conditional value-at-risk portfolio framework by considering the worst-case scenario of the underlying distribution of portfolio returns. Furthermore, Glasserman and Xu (2013) construct a multistage robust portfolio control framework with transaction costs that takes into account both statistical

estimation and model errors by using both interval and ellipsoidal uncertainty structures.

2.5 Performance Measures and Portfolio Evaluation

It has been widely reported that institutional and individual investors depend on risk-adjusted performance measures to select amongst available assets and evaluate their portfolios. Without any doubt, the most famous risk-adjusted performance measure is the Sharpe ratio because it is easy to compute and has been widely investigated in the literature, see for instance Lo (2002) and Ledoit and Wolf (2008) and others. In particular, the Sharpe ratio is defined as mean excess portfolio return over a risk-free rate, divided by the standard deviation of the portfolio returns according to Sharpe (1966). Given the fact that the Sharpe ratio takes into account just the first two statistical moments: means, variances and covariances, it is only sufficient if portfolio returns are normally distributed, but portfolio and fund returns often exhibit fat tails and follow asymmetric distributions. Another important disadvantage of the Sharpe ratio is that it considers both positive and negative deviations from the expected return to compute risk.

For instance, Cumming et al. (2014) investigate expanded portfolios (portfolios with conventional and alternative asset classes), and highlight that such portfolios often generate fat tails. They use 6 performance measures to evaluate different portfolio optimization strategies on such portfolios: mean return, Sharpe ratio, Sortino ratio, Sterling ratio, value-at-risk and conditional value-at-risk. As a logical and direct consequence, the use of the Sharpe ratio or other measures that contain the

classical risk (standard deviation) may lead to irrelevant conclusions since it overestimates the portfolio performance by underestimating the corresponding risk. We follow Eling and Schuhmacher (2007) and split the performance measures into three categories: measures based on the lower partial moments, drawdown and value-at-risk. Auer (2015) also provides a comprehensive review on alternative reward to risk ratios.

Lower partial moments are particularly popular in the asset management industry because they only take into account negative deviations of returns, e.g. returns that are below a lower acceptable bound (which could be the average portfolio return, a risk-free rate or zero). According to Sortino and Van der Meer (1991), lower partial moments are often considered as a more reliable risk measure than the standard deviation, since the latter considers both positive and negative deviations from expected return, and may result in incorrect conclusions when distributions that are not symmetric around the mean (skewed distributions). The order $n=0,1,2,..$ in lower partial moments defines the weight of deviations from the lower acceptable return (bound). The shortfall probability, expected shortfall and semi-variance describe the lower partial moment of order 0, 1 and 2 respectively, see also Eling and Schuhmacher (2007).

2.5.1 Performance Measures based on Lower Partial Moments

Omega, the Sortino ratio and Kappa are performance measures that are based on lower partial moments of order 1, 2 and 3 respectively, see for instance Eling and Schuhmacher (2007). Omega is computed as the mean excess portfolio return (over the minimal acceptable bound), divided by the lower partial moment of order 1,

while the Sortino ratio and Kappa are defined as the mean excess portfolio return (over the lower acceptable portfolio return), divided by the lower partial moment of order 2 and 3 respectively. Furthermore, the higher partial moments constitute an alternative way of measuring excess returns since they measure positive deviations, the opposite of the lower partial moments. For instance, the upside potential ratio, see Eling and Schuhmacher (2007), is defined as the higher partial moment of order 1, divided by the lower partial moment of order 2.

2.5.2 Performance Measures based on Drawdown

Drawdown measures are popular in the asset management industry and are often used by commodity and hedge fund traders, see Eling and Schuhmacher (2007). Drawdown measures are also popular with institutional investors such as pension funds, see for instance Berkelaar and Kouwenberg (2010), who develop a drawdown approach within stochastic programming for pension asset liability management. The drawdown rate measures drops from the highest point in cumulative portfolio returns over a certain time horizon and is a measure that does not depend on assumptions of distributions. The Calmar ratio, the Sterling ratio and the Burke ratio are alternative risk-adjusted performance measures based on drawdown. In particular, the Calmar ratio is defined as the mean excess portfolio return (over a risk-free rate, for instance), divided by the maximum drawdown over a certain period. In a similar way, the Sterling ratio and the Burke ratio are computed by the mean excess portfolio return, divided by the expected value and the Euclidean norm of drawdowns over a specific time period respectively.

2.5.3 Performance Measures based on Value-at-Risk

Excess return divided by the value-at-risk, the conditional Sharpe ratio and the modified Sharpe ratio proposed by Dowd (2000), Agarwal (2004) and Gregoriou (2003) respectively, are performance measures based on the typical value-at-risk calculation, the conditional and the modified value-at-risk. They have the important advantage of being distribution free measures. The standard value-at-risk gives the possible portfolio loss with a given confidence level over a certain time period, while the conditional value-at-risk is defined as the expected value of portfolios returns that do not exceed the possible losses indicated by the standard value-at-risk over a certain period and confidence level. The modified value-at-risk is given by a more complex mathematical formula that takes into account higher moments (skewness and kurtosis) except for the means and variances, see for instance Eling and Schuhmacher (2007). Hence, the 'excess return on value-at-risk' is given by the mean excess portfolio return divided by the value-at-risk, the conditional Sharpe ratio is computed as the mean excess portfolio return divided the conditional value-at-risk and the modified Sharpe ratio is defined as the mean excess portfolio return divided by the modified value-at-risk.

2.6 Asset Liability Management (ALM)

In contrast to the asset-only portfolio approach where investment managers and decision makers do not take into account liabilities in the optimization process and determine their asset allocation by considering just the assets, asset liability management (ALM) attempts to provide optimal asset allocation strategies taking into account future commitments and goals that consist the so-called liabilities.

Hence any investment decision is taken in terms of both the assets and liabilities. More authoritatively, the Society of Actuaries (SOA (2013)) defines the ALM process as 'the ongoing process of formulating, implementing, monitoring and revising strategies related to assets and liabilities to achieve an organization's financial objectives'. Hence asset liability management can be seen as an asset allocation tool that could provide efficient and prudent management to institutional investors (e.g. pension funds, life insurance funds, banks, sovereign funds, endowments, etc.) and allows them to make investment decisions which meet their future obligations and the corresponding risks associated with them (e.g. interest rate, inflation, market, longevity, mortality, market, liquidity, credit risk and others) and remain trustworthy and solvent.

The application of the asset liability management models attracted significant attention after the de-regulation of interest rates in 1979 according to Rachev and Tokat (2000) (see also the case of the US Savings and Loans and the corresponding problems that loan holders faced during the period 1980 to 1982 due to the increase in interest rates). Hence, the risk of liabilities increased significantly and made their forecasting a much more difficult task, which was a direct and logical consequence of the de-regulation in interest rates. Hence, the use of more sophisticated and efficient portfolio management tools that could effectively help market agents optimally allocate their assets, while matching their future obligations (liabilities), became a necessity. Although the initial scope of asset liability management modeling was to hedge the interest rate risk of the liabilities, they now include more types of risk such as inflation, longevity, mortality, market, credit, liquidity,

operational, currency and demographic risk, amongst others. Although the main objective of the ALM framework is to provide a strategic risk management tool for institutional investors to help them meet their future obligations via an appropriate asset allocation, the process of deriving optimal ALM strategies is often too challenging to be applied in practice and for this reason different operational research techniques have been used to efficiently construct and solve ALM problems.

The literature on asset liability management is fruitful and covers a broad spectrum of institutional investors as mentioned above. For instance, Merton (1969) develops a dynamic programming approach to asset liability management in order to provide an efficient tool for personal wealth management. Kusy and Ziemba (1986) apply a stochastic linear programming methodology to asset liability management for banks in the presence of cash flow and investment uncertainty in order to provide more effective management given the complex operations involved in a banking system. ALM frameworks have also been widely used in recent years to help pension fund trustees make better decisions, given that the associated pension liabilities include different types of risk (e.g. interest rate, inflation, longevity and demographic risk amongst others), see for instance, Sharpe and Tint (1990), Ezra (1991) and Board and Sutcliffe (2007), Ang et al. (2013) and others. Furthermore, the ALM framework has been applied to other market agents such as university endowments, see Merton (1993), sovereign wealth funds with known government liabilities, see Scherer (2011), and other types of funds that manage different type of assets classes such as real estate funds, as in Chun et al. (2000).

2.7 Operational Research and ALM Techniques

Operational research (OR) (or management science/decision science) is a sub-field of applied mathematics that applies a variety of advanced numerical and statistical methods to make better and more efficient decisions in a number of real world problems, e.g. asset management, logistics, supply chain management, air-traffic management, waste management, health care management, energy management, water supply management and others. ALM policies fall into this category due to their complex numerical structure. Problems modeled with operational research techniques typically use mathematical programming, such as linear, quadratic, non-linear, stochastic, dynamic, goal, robust (worst-case) and cone programming, by considering an objective function that has to be optimized (e.g. maximized or minimized) in terms of some decision variables that are defined by the decision makers, and subject to a certain set of constraints due to the corresponding regulations and rules associated with each problem. Except for ALM modeling, the applications of OR in finance include the valuation of financial instruments such as options and asset-backed securities, market imperfection modeling and optimal asset-only portfolio selection, amongst others. For a more comprehensive and detailed review with practical examples from the finance literature, we refer to Board et al. (2003).

There are various OR techniques that have been employed to compute optimal asset liability management strategies, which are applied primarily to pension funds and secondly to life insurance funds, banks, individual wealth management, endowments, sovereign funds and others. In particular, they are divided into four

categories: stochastic programming, e.g. Kusy and Ziemba (1986), Carino et al. (1994) and Geyer and Ziemba (2008); portfolio theory, e.g. Sharpe and Tint (1990), Ezra (1991), Board and Sutcliffe (2007), Ang et al. (2013) and Chun et al. (2000); stochastic simulation, e.g. Boender (1997), Van Rooij et al. (2004), Mulvey et al. (2000) and Mulvey et al. (2005) and dynamic programming and stochastic control, e.g. Rudolf and Ziemba (2004), Dondi et al. (2008) and Giamouridis et al. (2014). Stochastic programming provides more flexibility in the assumptions required than do other models, but it requires a huge computational effort to solve. Dynamic programming and stochastic control involve the solutions of hard non-linear optimization programs. ALM models based on stochastic simulation usually do not contain a separate section that will generate optimal portfolios strategies. Finally, the asset liability management techniques that are based on mean-variance portfolio theory (e.g. Sharpe and Tint (1990) model) do not have significant data requirements, are more easily understood by practitioners, and are easily formulated, applied and solved, but they face different types of problems such as their high exposure to estimation risk and parameter uncertainty.

The procedure for deriving optimal ALM strategies is often computationally challenging, and the majority of the methods described above require too much computational effort to be widely used and applied in practical real world problems. For instance, although stochastic programming as well as scenario based approaches (dynamic and stochastic control) are popular techniques for ALM problems, Fabozzi et al. (2007) highlight that “unfortunately the dimension of realistic stochastic programming models is usually very large, and optimization is challenging, even with

today's advanced technology". For example, if we consider a real problem with 10 assets and liabilities, 10 evaluation periods and assume 10 possible independent asset and liability returns (outcomes) in each period, then the number of possible scenarios would be $10(10^{10})=100$ billion. If we further suppose that the computer can run 100 scenarios per second, it would take 3.17 years to solve the problem. It is obvious that such large problems cannot be handled by institutional investors that want to have a frequent review of their asset allocation. As a result, they may consider unrealistic problems with a small number of assets and evaluation periods to assess their strategies, and this can lead to seriously wrong conclusions. In addition, if the stochastic programming problem has 100 billion decision variables, the formulated problem would be extremely large, making it well beyond the RAM storage space of most computers. Gulpinar and Pachamanova (2013) use only 2 assets, 4 time periods and write that 'The only reason for selecting a small number of assets for investment and a small number of time periods is the stochastic programming formulation. It takes a very long time (hours) to obtain an optimal solution to stochastic programming formulations with even a small number of scenarios.' In what follows, we review a variety of ALM studies that fall into each of the four main categories (stochastic programming, dynamic programming and stochastic control, portfolio theory and simulation) as listed above.

2.7.1 Stochastic Programming ALM Models

Kusy and Ziemba (1986) propose a multi-stage stochastic linear programming framework for banks in order to provide effective asset liability management in the presence of uncertain cash flows and asset returns. The proposed model

incorporates some important legal, financial and bank related characteristics together with their uncertainties, and the empirical results reveal it to be superior to naïve benchmarks (deterministic linear programming). In addition, Carino et al. (1994) develop a multi-period stochastic linear programming asset liability management framework for the Yasuda Fire and Marine Insurance Co., Ltd. In particular, the model contains complicated Japanese insurance rules and practices and determines the optimal asset allocation in a multi-stage environment, while its objective is to generate an adequate portfolio return to repay interest on savings accounts related to insurance contracts, as well as maximizing the terminal wealth over the policy horizon. We refer to Carino and Ziemba (1998) for further technical details of the formulation of the Russell-Yasuda Kasai Model. Furthermore, Geyer and Ziemba (2008) develop the financial planning model of an Austrian pension fund (Siemens Austria) using a multi-stage linear stochastic programming model, which provides flexibility for the number of time periods and their length. They consider state dependent correlations across the different assets classes in alternative market conditions that give the opportunity to the asset liability model to foresee and respond to harsh as well as normal conditions. Their objective is to maximize the expected wealth at the horizon date minus penalty costs for benchmarks goals determined in each decision period.

2.7.2 Single Period ALM Models (Portfolio Theory)

It is well known that Sharpe and Tint (1990) formulated a portfolio optimization framework that takes into account liabilities as well as assets, the so-called surplus optimization framework, and is mainly based on the Markowitz (1952) mean-

variance portfolio framework. Their objective is to maximize utility, defined as the expected value of excess asset returns over the liabilities, minus the corresponding variance. The model also provides some flexibility by incorporating the ‘importance’ of liabilities in the surplus optimization framework (0 and 1 for assets only and full surplus optimization respectively). In addition, Ezra (1991) uses the surplus framework of Sharpe and Tint (1990) and highlights the need for pension managers to focus on surplus optimizations instead of assets only approaches due to the high uncertainty of the liability side. Board and Sutcliffe (2007) propose an ALM model that disaggregates the pension liabilities into three categories (pensioners, active and deferred members) and apply a generalization of the mean-variance portfolio framework which incorporates these liabilities in the computation of the optimal asset allocation. They use the mathematical formula of Haberman (1992) to estimate the mean and variance of the contribution rate for each point of the efficient frontier of the asset-liability portfolio. This provides a complete asset-liability management tool that can be applied very easily in practice by pension fund trustees. In addition, Ang et al. (2013) extend the Sharp and Tint (1990) model to incorporate a penalty for failing to meet the pension liabilities (downside risk) and they use options to value the deficit between asset and liabilities. Finally, Chun et al. (2000) apply the surplus optimization model to effectively manage real estate investment portfolios in an asset-liability framework.

2.7.3 Stochastic Simulation ALM Models

Boender (1997) proposes a hybrid simulation – optimization stochastic asset-liability model as a decision tool for more sustainable pension fund management. This

framework uses scenarios, instead of hypothetical probability distributions to model the uncertain future risk drivers, and non-linear optimization to compute the optimal asset allocation and contribution rate for each point of the efficient frontier. In addition, Van Rooij et al. (2004) develop a pension asset-liability management framework that is based on both defined benefit and career average revalued earnings pension systems. This model provides flexibility for various parameters such as the indexation, portfolio selection and retirement age amongst others, and replicates the uncertainty in a number of stochastic parameters via historical and stochastic simulations. Mulvey et al. (2000) propose an integrated stochastic asset-liability system that contains a scenario generator (stochastic simulator) together with a non-linear optimization programming model to help pension funds to better understand their risks and eliminate the possibility of becoming insolvent. Finally, Mulvey et al. (2005) propose a stochastic simulation approach to assess and prevent pension schemes from becoming insolvent.

2.7.4 Dynamic Programming and Stochastic Control ALM Models

Rudolf and Ziemba (2004) use dynamic programming and control and propose an inter-temporal portfolio selection framework that could be applied by both life insurers and pension funds. The model assumes that both asset and liability returns follow an Ito process that depends on specific state variables and its objective is to maximize the inter-temporal utility, which is defined as value of assets minus liabilities. Dondi et al. (2008) develop a dynamic asset-liability management model for Swiss pension funds. They use a dynamic factor structure that takes into account extreme events (fat tails) to model asset returns, while the pension liabilities are

based on the current and future cash flows. Their objective is to minimize any possible deficit subject to a certain goal for the future surplus.

2.8 Techniques for Generating Scenarios

There are different techniques one can use to generate future scenarios for risky variables. For example, Kouwenberg and Zenios (2006) split them into three categories; bootstrapping, time series analysis and statistical techniques. In a more comprehensive analysis, Mitra (2006) lists four different techniques that can be employed to generate future scenarios for representing parameter uncertainty in both asset and liability returns, and they are the sampling, simulation (very similar to sampling), statistical approach as well as alternative techniques (e.g. machine learning techniques).

2.8.1 Sampling Approach

The sampling (bootstrapping) technique is the most easily applicable method for generating scenarios. In the sampling technique, which is sometimes also called 'simulation technique', the modeller simply assumes a probability density function based on a particular distribution (e.g. normal, F-distribution) and generates future scenarios for asset returns or other uncertain parameters taking into account their historical correlations. For instance, Geyer and Ziemba (2008) employ different distributions for each of the asset classes used in their model. Specifically, they take into account fat tails for stocks, and assume t-distributions with five degrees of freedom for stock returns, and normal distributions for bond returns as returns on fixed income products do not often exhibit fat tails. The main advantage of

bootstrapping (or simulation) in the scenario generation process is its easy applicability to situations and problems where other numerical methods are computationally intractable and lead to infeasible solutions. Although this is often considered as the most popular and easy sampling approach, we refer to Mitra (2006) for extensions, such as the conditional sampling technique. Finally, it is obvious that the approximation of the uncertain/stochastic parameters with certain probability distributions could be very dangerous, leading to irrelevant and error-maximized conclusions.

2.8.2 Statistical Approach

Statistical techniques apply core econometric models, e.g. time series analysis, using a set of historical data as well as stochastic analysis in the scenario generation process. For example, vector autoregressive (VAR) models are amongst the most popular econometric time series tools currently used by the academic community and practitioners in institutional funds to generate future scenarios for the corresponding uncertain variables. For instance, Boender (1997), Ferstl and Weissensteiner (2011), Gulpinar and Pachamanova (2013), Hoevenaars and Ponds (2008) and Chen et al. (2014) use a vector autoregressive model of order 1, VAR(1), to generate future scenarios. In addition, stochastic diffusion processes, such as Wiener processes, as well as the well known principal component analysis (PCA) have also been used to generate scenarios.

2.8.3 Alternative Techniques

Although the sampling, simulation and statistical approaches are more popular techniques and have been more widely used in the literature, there are some less

common methods that have also been used in relevant studies to generate scenarios such as the Artificial Neural Network method, and the clustering and scenario reduction approach. The Artificial Neural Network (ANN) framework is a processing technique that significantly differs from the other methods due to its structure. In particular, ANNs consist of a large number of elements, which are connected with each other, (neurons). They receive input signals and generate output signals. This complex network structure can help decision makers to solve complicated problems which cannot be solved with classical numerical methods. The clustering approach is often based on an econometric (time-series) process that is used to create a large number of scenarios (e.g. paths). Then, a specific number of scenarios are 'clustered' by another technique, see for instance Rustem et al. (2004). Finally, the scenario reduction technique is another alternative approach applied to the scenario generation process. The extremely large size of some scenario trees makes the process computationally intractable, and some decision makers may desire to decrease the problem complexity by applying a variety of methods (e.g. sampling techniques) directly to the scenario trees, see Mitra (2006) for further details.

2.9 Pension Schemes Design and Intergenerational Transfers

There are two main types of pension schemes – defined benefit (DB) schemes and defined contribution (DC) schemes, as well as other more complex schemes that are the minority at the moment (e.g. hybrid schemes). According to ONS (2013), there were 8.1 million DB and DC active members (the members that currently contribute to a pension scheme) in the UK in 2011, with the 65% of the members in the public sector and the rest of them (35%) in the private sector. Defined benefit schemes

were the prevalent type with 7.2 million active members in 2011 in contrast to just 900 million active members in DC schemes. In recent years, the majority of DB schemes in the UK have gradually stopped accepting new members, and they automatically join the newly established DC schemes. As a direct and logical consequence of this, the number of active members in DC schemes exceeded that of DB schemes in 2014 for the first time.

2.9.1 Defined Benefit (DB) Pension Schemes

Defined Benefit pension schemes are managed by a body of trustees and pensions at retirement in the UK and US are often based on the employee's final salary and the accrued years (years of employment). For instance, if the final annual salary is £50,000, the accrual rate is 1/70 (1.43% - usually expressed as percentages) and the employer and member have paid contributions for 35 years (35 years of service), then the annual pension at retirement is equal to $(35/70) \times 50,000 = £25,000$, an amount that is promised by the sponsor that bears all the corresponding risks associated with DB schemes. It is a common policy for both the employers and employees to pay contributions into the fund. For example, the sponsors of the Universities Superannuation Scheme (USS), one of the largest funds in the UK and worldwide with over 300,000 members and 374 individual sponsors, pay 16% of salary, while the employers pay 7.5% making a total contribution of $16\% + 7.5\% = 23.5\%$ of the employees' current salary.

DB pension funds are very large institutional investors and their assets under management in the OECD countries in 2012 were roughly \$32 trillion, and the corresponding pension obligations (pension liabilities) were much higher (OECD,

2013). Pension schemes are long term institutional investors with long time horizons since the new active members will start receiving a pension many decades later, and the trustees should make long-term investment decisions in order to be able to meet their promises (e.g. paying future pensions). For this reason, the pension funds' trustees who run large pension funds often use asset liability management models (e.g. stochastic programming ALM models) to compute the optimal asset allocation, with the selection of the individual assets determined by the fund managers. For instance, USS carries out triennial valuations to estimate the surplus or deficit. Inevitably, these computations involve many forecasts of uncertain future events (e.g. number of future members, cash flows, portfolio returns and others), which are subject to significant estimation risk, and hence different methods may generate different valuations of the same fund. Once actuaries complete the valuations, then the trustees decide whether they should increase the employer's contribution rate to offset any deficit given the 'balance of cost' design of DB schemes where the sponsor bears all the risks.

The majority of UK defined benefit schemes are 'balance of cost', which means that only the sponsor is responsible for meeting the future pension obligations of the scheme, and as a direct consequence the sponsor alone bears all the risks associated with the pension promises. In other words, the members are obliged to pay a fixed contribution rate over time, which is independent of the financial state of the scheme, and the remainder of the cost associated with the future pension payments is covered by the employer. This can be unbearable if the scheme has a large deficit and may lead the sponsor to become insolvent. In particular, the sponsor of a DB

scheme is exposed to a variety of risks such as investment risk, interest rate risk, inflation risk, salary growth risk, longevity risk and regulatory risk. The investment risk is related to the uncertain asset returns, the interest rate risk is related to the fluctuations in government bond yields that are a key element to the valuation of pension liabilities, the inflation risk is associated with the valuation of index linked liabilities, and the salary growth risk is related to the uncertain final salary which is connected to inflation fluctuations. Finally, the longevity risk is related to the risky life expectancy of active members and pensioners and the regulatory risk to the changing pension rules over the time.

2.9.2 Defined Contribution (DC) Schemes

The main idea behind the defined contribution schemes (or 401(k) investments in the United States) is that contribution rates for both the sponsors and members are constant proportions of salaries. The members of a DC scheme have their own pension pot that is paid out when they retire. Since April 2015 the pensioners in the UK, who have accumulated a pot of money through a DC scheme, are no longer obliged to buy an annuity with this money, but any money not used to purchase an annuity is normally taxed according to each person's income tax rate. Such pension schemes can be considered as alternative 'savings accounts for retirement', since they do not necessarily give pensions but provide flexibility to the members to receive the money accumulated in the pot as a taxable lump sum.

We assume that a member of a DC scheme pays contributions of 7%, and its sponsor pays contributions of 10% per year (overall contribution rate of 17% per annum). The member's annual earnings as well as the corresponding annual returns on the

pension pot can be seen in table 2.1 below. We also assume that the member pays contributions for five years and that any contributions are paid at the end of the year. As we can observe in the following table, the pension pot after five years is worth £40,912. Then, each member of a DC scheme needs to make some important decisions such as whether to purchase an annuity or to receive a taxable lump sum that consists of all the money in the pot, and which annuity to buy and when.

	Earnings	Contribution	Open Fund	Return	Pension Pot
1	£40,000	£6,800	0	-	£6,800
2	£42,000	£7,140	£6,800	3%	£14,144
3	£44,000	£7,480	£14,144	6%	£22,473
4	£46,000	£7,820	£22,473	4%	£31,192
5	£48,000	£8,160	£31,192	5%	£40,912

Table 2.1: Example of member's annual earnings and pension pot's annual returns

In contrast to DB schemes where the sponsor bears all the risks associated with the pension liabilities and is responsible for meeting the future pension obligations, (e.g. the risks related to investments, interest rates, inflation, salary growth and longevity), these risks are borne by the members in defined contribution schemes until they retire. If a member of a DC scheme decides to buy an annuity, then these risks except for the salary growth risk are transferred to the insurance company that pays the pensions to the member (annuity provider). Although the contribution rates are constant over time (unchanged), the only risk the sponsor has to bear is the salary risk since any fluctuations on the salary level leads to changes in the contributions the sponsor pays into the pension pot.

Defined contribution schemes in the UK are managed in two different ways. The first involves a trustee board that manages the pension scheme in a similar way as in defined benefit schemes. The second way involves an insurance company (annuity provider), who is responsible for providing a defined contribution pension to the employees (contract-based DC pension schemes). The latter type of DC pension scheme is basically a group of separate individual pensions (group personal pensions) and it is easier, more flexible and less costly to install and manage for the employer than a trust-based pension scheme. Also, the control of the sponsor of the scheme is largely eliminated since the pensions are provided by the annuity providers.

2.9.3 Intergenerational Effects

The sponsors of UK defined benefit (DB) pension schemes are required to meet the pension promises (e.g. pension payments) according to certain rules and independently of the financial state of the scheme. As a direct and logical consequence, the sponsor (e.g. employer) bears all the risks such as the interest rate, inflation, longevity, regulation, investment and salary growth risk. But these risks can be shared between the sponsor and the members of the scheme such as by increasing the members' contribution rates and the retirement age or decreasing the accrual rate in case of an extensive deficit that cannot be handled by the sponsor alone. Because reduction of the members' accrued benefits is not permitted by the law in the UK and US, possible rule modifications will only affect future accruals. In practice, it means that the members of a scheme with less years of service will be more affected by such changes because these members will accrue benefits under

the new rules for more years, while the older members (e.g. close to retirement) are significantly less affected since their accrued benefits until the date of changes are fully protected by the law.

The members of a pension scheme pay contributions in anticipation of receiving a pension in the future and hence they have certain expectations of the net present value (NPV) of the cash flows with the pension scheme before a rule is modified. As a result, the redistributive effect of a rule change is defined as the difference between the net present values (NPVs) of the cash flows before and after the rule change for each cohort. For instance, pension wealth redistribution between the members (active, deferred and future members) and sponsor happens with an increase in the retirement age or with a drop in the accrual rate. Furthermore, the cash flow changes between the various members (active, deferred and future members), pensioners and sponsor is a zero sum game.

Significant consideration is only given to the absolute numbers of the new members' contribution rates and the retirement ages and others when rule changes are announced, but an explicit consideration of the redistributive effects on the wealth associated with these changes is often omitted. In addition, all the existing studies that investigate the redistribution of pension wealth have been conducted for hypothetical Dutch pension schemes, where the sponsor is not involved in the redistribution. In particular, the existing studies consider that the only obligation for the sponsor of the scheme is to pay a constant contribution rate, and hence the redistribution takes place between different age cohorts (intergenerational redistribution). But the 'balance of cost' scheme is the dominant scheme in the UK

and US and the sponsor bears the default risk in this case. In addition, the existing literature does not incorporate many important aspects associated with pension scheme design such as lump sum payments, a time-varying retirement age, upper bounds on the salary growth, deferred members (members that have left the scheme and do not pay contributions), spouses' pensions and asset classes such as alternative investments (hedge funds, commodities, real estates and others).

Chapman et al. (2011) carried out the first study that investigates the effects of pension scheme rule changes. In particular, they assume that potential changes in pension scheme rules will affect six stakeholders: the members, the government, external advisors, externals and the sponsor's share and debt holders and use numerical methods (Monte Carlo simulations) to simulate the cash flows between the different stakeholders for a policy horizon of just ten years. Then, they calculate the net present values (NPVs) of the cash flows via risk neutral valuation (SDFs) of a hypothetical scheme for the base case and different pension rules. The changes in the NPVs after the rule changes give the redistributive effects.

Ponds (2003) considers a hypothetical Dutch scheme where the sponsor pays a fixed contribution rate (the sponsor bears no risk) and hence the members of the scheme bear all the risks. The author uses the same methodology as in Chapman et al. (2001) with simulations, risk neutral valuations and stochastic discount factors (SDFs) to compute the redistributive effects between different age cohorts of active, future members and pensioners by applying different scheme rules. Hoevenaars and Ponds (2008) and Draper et al. (2014) have also investigated the intergenerational redistributive effects that result from rule changes in hypothetical Dutch pension

funds. Beetsma et al. (2014) investigate intergenerational redistributions between various active members' cohorts and the sponsor for a hypothetical state pension fund in the US, while Hoevenaars et al. (2009) quantify redistribution between the members and the sponsor, by ignoring the different age cohorts of the members of the scheme, for a hypothetical Dutch pension scheme.

2.10 Conclusions

To sum up, the main aim of this chapter was to provide a comprehensive literature review to the reader in three areas – portfolio theory, asset-liability modelling and pension schemes, and hence to prepare her/him for the main chapters of this thesis. In particular, all these three areas are involved in Chapter 3, where a novel numerical technique – robust (worst-case) optimization - is used to formulate an ALM model and is applied to a real world pension scheme (Universities Superannuation Scheme - USS) for first time, while it is also benchmarked against the Sharpe and Tint, Bayes-Stein and Black-Litterman models. Chapter 4 is mainly related to portfolio theory, since three different optimal portfolio diversification techniques (Markowitz, norm-constrained and Black-Litterman portfolio) and three different naïve diversification approaches (equally-weighted, risk-parity and reward-to-risk timing portfolios) are applied to the same Socially Responsible Investment screened investment universe and their actual (out-of-sample) performance evaluated using various performance measures. Furthermore, Chapter 5 combines elements from both the ALM modelling and pension schemes design literature, and evaluates the short, medium and long term performance of the pre and post – 2011 USS pension scheme under three different asset allocation strategies (fixed-mix, risk-management and risk-shifting).

The redistributive effects of the rule changes in 2011 are also quantified in Chapter 5.

Chapter 3

3 Asset Liability Modelling and Pension Schemes: The Application of Robust Optimization to USS¹

3.1 Introduction

Pension schemes are among the largest institutional investors, and in 2012 the OECD countries had pension assets of \$32.1 trillion (with liabilities several times larger), accounting for 41% of the assets held by institutional investors (OECD, 2013). Pension schemes have very long time horizons, with new members likely to be drawing a pension many years later, and therefore need to make long term investment decisions to meet their liabilities. To help them do this many of the larger defined benefit (DB) pension schemes model their assets and liabilities using asset-liability management (ALM). These models determine the scheme's asset allocation, with stock selection left to the fund managers. While a widespread switch to defined contribution schemes is underway, DB schemes will remain very large investors for decades to come as they continue to serve their existing members and pensioners.

There are a variety of techniques for deriving optimal ALM strategies for pension funds, and they fall into four main categories: stochastic programming, e.g. Kusy and Ziemba (1986), Kouwenberg (2001), Kouwenberg and Zenios (2006) and Geyer and Ziemba (2008); dynamic programming, e.g. Rudolf and Ziemba (2004) and Gao (2008); portfolio theory, e.g. Sharpe and Tint (1990), and stochastic simulation, e.g.

¹ The content of this Chapter was presented at the 4th International Conference of the Financial Engineering and Banking Society (Surrey, UK), 8th Portuguese Finance Network International Conference (Vilamoura, Portugal), 12th Annual International Conference on Finance (Athens, Greece) and the Annual Workshop of the Dauphine-Amundi Chair in Asset Management (Paris, France). The corresponding academic paper based on this Chapter has been published by the European Journal of Finance - DOI: <http://dx.doi.org/10.1080/1351847X.2015.1071714>

Boender (1997). The process of computing optimal ALM strategies can be challenging, and most of the existing techniques are too demanding to be widely applied in practice.

While stochastic programming is a popular technique for solving ALM problems, the present study uses the new technique of robust optimization. Robust optimization has attracted considerable attention in recent years and is considered by many practitioners and academics to be a powerful and efficient technique for solving problems subject to uncertainty (Ben-Tal and Nemirovski, 1998). It has been applied to portfolio optimization and asset management, allowing for the uncertainty that occurs due to estimation errors in the input parameters, e.g. Ben-Tal and Nemirovski (1999), Rustem et al. (2000), Ceria and Stubbs (2006) and Bertsimas and Pachamanova (2008). Robust optimization recognises that the market parameters of an ALM model are stochastic, but lie within uncertainty sets (e.g. upper and lower limits).

Portfolio theory is highly susceptible to estimation errors that overstate returns and understate risk, and three main approaches to dealing with such errors in portfolio and ALM models have been used. The first approach involves changing the way the mean and covariance matrix of returns are estimated; for example, James-Stein shrinkage estimation (e.g. Jobson, Korkie and Ratti, 1979), Bayes estimation (e.g. Black and Litterman, 1992), and the overall mean of the estimated covariances (Elton and Gruber, 1973). A second approach is to constrain the asset proportions to rule out the extreme solutions generated by the presence of estimation errors (e.g. Board and Sutcliffe, 1994). A third approach is to use simulation to generate many

data sets, each of which is used to compute an optimal portfolio, with the average of these optimal portfolios giving the overall solution (Michaud, 1999). Robust optimization offers a fourth approach to estimation errors in which the objective function is altered to try to avoid selecting portfolios that promise good results, possibly due to estimation errors. It adopts a maximin objective function, where the realized outcome has a chosen probability of being at least as good as the optimal robust optimization solution, which should rule out solutions based on optimistic estimation errors.

The only previous application of robust optimization to pension schemes is Gulpinar and Pachamanova (2013). Their hypothetical example uses two asset classes – an equity market index and the risk-free rate, 80 observations, non-stochastic liabilities and an investment horizon that consists of four investment periods of three months each. They imposed a lower bound on the funding ratio (i.e. assets/liabilities) which was constrained to never be less than 90%. The objective was to maximise the expected difference between the terminal value of scheme assets and contributions to the scheme. There was an experimental simulation, rather than out-of-sample testing. The chosen confidence level, i.e. protection against estimation errors, was set to be the same for each uncertainty set, and varied between zero and three. They found a clear negative relationship between the chosen confidence level and the value of the objective function.

We extend the Goldfarb and Iyengar (2003) model of robust optimization from portfolio problems to ALM by incorporating risky liabilities with their own fixed ‘negative’ weights, disaggregate the liabilities into three categories, and include

upper and lower bounds on the proportions of assets invested in the main asset classes. The objective function we use is the Sharpe ratio, which gives the solution that maximises the excess return per unit of risk, and is the first study to use the Sharpe ratio as the objective for an ALM study of a pension scheme. We compare these results with those of four benchmarks. The first is the actual portfolios chosen by the Universities Superannuation Scheme (USS), and the second is a modified version of the Sharpe and Tint (1990) model which has been widely used in ALM models (Ang, Chen & Sundaresan, 2013; Board & Sutcliffe, 2007; Chun, Ciochetti & Shilling, 2000; Craft, 2001, 2005; De Groot & Swinkels, 2008; Ezra, 1991; Keel & Muller, 1995; Nijman & Swinkels, 2008). The third is Bayes-Stein estimation of the inputs to the modified Sharpe and Tint ALM model, and the last is Black-Litterman estimation of the portfolio inputs to the modified Sharpe and Tint ALM model. Neither Bayes-Stein, nor Black-Litterman have previously been used in an ALM study of a pension scheme.

Scenario-based programming (dynamic programming, stochastic programming, etc.) is not used as a benchmark because of the enormous computational burden this would entail. In our application to USS with 14 assets and liabilities, and assuming five independent outcomes each three year period, the total number of scenarios to be evaluated for the four out-of-sample periods would be $4(5^{14})$, or 24.4 billion. While we do not use stochastic programming, in an experiment involving a hypothetical pension scheme and 100 scenarios, Gulpinar and Pachamanova (2013) found that robust optimization outperformed stochastic programming.

We use a single period portfolio model, and there are two alternative theoretical

justifications for the use of such models when investors can rebalance their portfolios (Campbell and Viceira, 2002, pp. 33-35). Assuming all dividends are reinvested, the first justification is that the investor has a logarithmic utility function. The second justification is that asset returns are independently and identically distributed over time, and investor risk aversion is unaffected by changes in wealth. The immediate reinvestment of dividends is common practice, while there is evidence that the share of household liquid assets allocated to risky assets is unaffected by wealth changes, implying constant risk aversion (Brunnermeier and Nagel, 2008). At the 1% level of significance the only assets that exhibit first order autocorrelation for the 1993-2008 period, which are used below in estimating the parameters of our empirical application, are UK 10-year bonds and UK property. Property index returns are well known to exhibit autocorrelation due to the use of stale prices in their construction. Since the assumptions for a theoretical justification are more or less satisfied and the multi-period alternative is computationally impractical, a single period model appears to be a reasonable simplification and has been used by many researchers for solving pension ALM problems; including Ang, Chen & Sundaresan (2013); Board and Sutcliffe (2007); Chun, Ciochetti & Shilling (2000); Craft (2001, 2005); De Groot & Swinkels (2008); Ezra (1991); Keel & Muller (1995); Nijman & Swinkels (2008); Sharpe & Tint (1990). For a more comprehensive discussion on the appropriateness of using single period models in a portfolio optimization context, we also refer to Ang (2014).

Our resulting ALM model is applied to USS using monthly data for 18 years (1993 to 2011), i.e. 216 months. The choice of USS has the advantages that, as the sponsors

are tax exempt, there is no case for 100% investment in bonds to reap a tax arbitrage profit. In addition, because the sponsors' default risk is uncorrelated with that of USS, there is no need to include the sponsor's assets in the ALM model. The data is adjusted to allow for USS's foreign exchange hedging from April 2006 onwards as this is when USS changed its benchmark onto a sterling basis. Finally, we use an actuarial model to transform the robust optimization solutions (i.e. asset proportions) into the scheme's projected contribution rate.

Ultimately pension scheme trustees are concerned about the scheme's funding ratio and contribution rate. Since the asset allocation has an important influence on the contribution rate and funding ratio, trustees need to make the asset allocation decision in the light of its effect on these two variables. Trustees wish to reduce the cost of the scheme to the sponsor and members by minimising the contribution rate. Trustees are also concerned with the regulatory limits on the funding ratio. UK legislation places upper and lower limits on the funding ratio of pension schemes, and the likelihood of breaching these requirements must be considered when making the asset allocation decision. MacBeth, Emanuel and Heatter (1994) report that pension scheme trustees prefer to make judgements in terms of the scheme's funding ratio and contribution rate, rather than the scheme's asset-liability portfolio. We use the model developed by Board and Sutcliffe (2007), which is a generalization of Haberman (1992), to transform the asset allocations to projected contribution rates. This generalization allows the discount rate to differ from the investment return, which improves the economic realism of the actuarial model.

The remainder of this paper is organized as follows. In section 3.2 we provide a brief

overview of robust optimization and explain our robust optimization ALM model and the four benchmarks. Section 3.3 describes USS, and in section 3.4 we describe the data set, explain how returns on the three liabilities were estimated, and use a factor model to estimate the three uncertainty sets. In section 3.5 we compute optimal out-of-sample investment policies for USS using robust optimization and the four benchmarks, together with the implications of the asset allocations for the funding ratio and projected contribution rate. Section 3.6 presents robustness checks on our results, and section 3.7 summarizes our findings.

3.2 Robust Optimization ALM Model

Robust optimization is a powerful and efficient technique for solving optimization problems subject to parameter uncertainty that has a number of advantages over other analytical techniques. First, it solves the worst-case problem by finding the best outcome in the most unfavourable circumstances (the maximin), i.e. each stochastic parameter is assumed to take the most unfavourable value in its uncertainty set. Since DB schemes must meet their pensions promise, investment strategies based on ALMs using robust optimization are well suited to the pension context where prudence and safety are important. Second, previous techniques such as stochastic programming and dynamic stochastic control are computationally demanding, while robust optimization is much easier to solve. The computational complexity of robust optimization is the same as that of quadratic programming in terms of the number of assets and time periods; while the complexity of scenario based approaches, such as stochastic programming and dynamic stochastic control, is exponential in the number of assets and time periods. Therefore realistic robust

optimization problems can be solved in a few seconds of computer processing time. Finally, in comparison with other techniques, robust optimization is less sensitive to errors in the input parameters, i.e. estimation errors. This tends to eliminate extreme (e.g. corner) solutions, leading to investment in more stable and diversified portfolios, with the benefit of lower portfolio transactions costs.

Robust optimization requires the specification of the maximum deviation from each of the expected values of the stochastic input parameters that the decision maker is prepared to accept. This maximum deviation is set to reflect the level of confidence (denoted by ω , where $0 < \omega < 1$) the decision maker requires that the optimal value of the objective function will be achieved when the optimal solution is implemented.

In the model used below there are three uncertainty sets corresponding to the three stochastic parameters in the factor model (see equation (3.1) below):- mean returns, factor coefficients and disturbances. These uncertainty sets are for (a) mean returns, (b) the coefficients of the factor model used to estimate the factor loadings matrix, and (c) the variance-covariance matrix of the disturbances. The decision-maker specifies the level of confidence (denoted by ω), that the Sharpe ratio given by the optimal robust optimization solution will be achieved or surpassed in the out-of-sample period. Two important characteristics of each uncertainty set are its size and shape. The size of an uncertainty set is governed by the parameter ω (confidence level), allowing the provision of a probabilistic guarantee that the constraints with uncertain/stochastic parameters are not violated. The shape of the uncertainty sets (e.g. ellipsoidal, box, etc.) is chosen to reflect the decision-maker's understanding of the probability distributions of the stochastic parameters; and an appropriate

selection of the shape of the uncertainty sets is essential for the robust optimization problem to be tractable (Natarajan, 2009). We chose to use elliptical (ellipsoidal) uncertainty sets which are very widely used when the constraints involve standard deviations; and in most cases result in a tractable and easily solved problem (Bertsimas, Pachamanova and Sim, 2004).

The initial formulation of the robust optimization problem with its uncertain parameters is transformed into the robust counterpart (or robust analog problem). This transformed problem has certain parameters (the worst-case value within its uncertainty set for each stochastic parameter) with only linear and second-order cone constraints. This second order cone problem (SOCP) can be easily solved, see Ben-Tal and Nemirovski (1998, 1999, 2000).

We divide the liabilities into three components: (i) members who are currently contributing to the pension fund (active members), (ii) deferred pensioners who have left the scheme but not yet retired and so currently do not generate any cash flows for the fund, and (iii) pensioners who are currently receiving a pension and so generate cash outflows from the fund. The liability for active members varies principally with salary growth until they retire, and then with interest rates, longevity and inflation. For deferred pensioners, their liability varies with the chosen revaluation rate until they retire (USS uses inflation as the revaluation rate), and then with interest rates, longevity and inflation. The liability for pensions in payment varies with interest rates, longevity and inflation as from the current date. We extend the Goldfarb and Iyengar model to include stochastic liabilities as well as stochastic assets in the factor model. More details are described later in this section,

as well as in Appendices 3.B and 3.C.

In the long run pension schemes do not take short positions, and so we rule out short selling by imposing non-negativity constraints on the asset proportions. We also do not permit borrowing money because UK pension schemes are prohibited from long-term borrowing. They are also prohibited from using derivatives, except for hedging or facilitating portfolio management. In addition, pension funds usually set upper and lower limits on the proportion of assets invested in asset classes such as equities, fixed income, alternative assets, property and cash. Therefore our ALM model includes upper and lower bounds on broad asset classes to rule out solutions that would be unacceptable to the trustees. This also tends to reduce the effects of estimation errors in the ALM inputs.

ALM studies of pension schemes using programming models have employed a wide variety of objective functions. For example, maximise the terminal wealth or surplus with penalties for risk and breaching constraints; or minimise the present value of contributions to the scheme with penalties for risk and breaching constraints. While some of the penalties can be quantified in monetary terms, a penalty for risk requires the specification of a risk aversion co-efficient. Many previous ALMs have assumed the pension scheme has a particular utility function of wealth or pension surplus, e.g. constant relative risk aversion (CRRA), constant absolute risk aversion (CARA) or quadratic utility, together with a specified risk aversion parameter. However, pension schemes are non-corporeal entities with an infinite life, and specifying their preferences in terms of a calibrated utility function is problematic.

An alternative approach is to use the market price of risk, which can be calculated from market data. The Capital Asset Pricing Model implies that the market's trade-off between risk (standard deviation) and return is given by the slope of the Capital Market Line. This leads to the use of the Sharpe ratio to select a risky portfolio, and has been widely used in academic studies. Following Sharpe (1994), we define the Sharpe ratio as the return on a fund in excess of that on a benchmark portfolio, divided by the standard deviation of the excess returns. The risk free rate is usually chosen as the benchmark portfolio, but we use the liability portfolio. So in our case the fund is the pension scheme's asset portfolio, and the benchmark is its liability portfolio. Therefore we divide the mean return on the asset-liability portfolio by the standard deviation of the asset-liability portfolio. Individual risk-return preferences will differ from this ratio but, since large pension schemes have well diversified portfolios, it gives the average value of individual preferences revealed in the capital market, and so offers a reasonable objective for a very large multi-employer scheme like USS.

Our approach is significantly different from the surplus optimization framework proposed by Sharpe and Tint (1994). Especially, the goal of surplus optimization is to maximize the utility function defined as the expected return on assets in excess of liabilities minus the variance of the excess returns. As a result, surplus optimization's objective function is actually a quadratic utility that uses excess returns instead of asset only returns, a risk aversion parameter equal to 1, and can be considered as a special case of mean-variance optimization. But the use of calibrated utility functions for defining the risk preferences for pension schemes is problematic for all the

reasons mentioned previously.

By using robust optimization to solve the problem, depending on the size of the chosen confidence level (ω), the scheme's optimal portfolio is effectively more risk averse than simply maximizing the Sharpe ratio. Use of the Sharpe ratio ignores the scheme's current funding ratio, but if the risk attitudes of pension schemes are wealth independent, the funding ratio is irrelevant. The empirical evidence on the effect of the funding ratio on the asset allocation, allowing for real world influences such as default insurance, taxation and financial slack, is mixed (An, Huang and Zhang, 2013; Anantharaman and Lee, 2014; Atanasova and Gatev, 2013; Bodie, Light, Morck and Taggart, 1985, 1987; Comprix and Muller, 2006; Guan and Lui, 2014; Mohan and Zhang, 2014; Li, 2010; McCarthy and Miles, 2013; Munro and Barrie, 2003; Petersen, 1996; Rauh, 2009). This is consistent with the view that pension schemes do not alter their asset allocation in a predictable way with their funding ratio, except when very close to default.

Following Goldfarb and Iyengar (2003) and others, we assume the stochastic asset and liability returns are described by the following factor model²:-

$$\tilde{\mathbf{r}}_{A,L} = \tilde{\boldsymbol{\mu}}_{A,L} + \tilde{\mathbf{V}}^T \mathbf{f} + \tilde{\boldsymbol{\varepsilon}}_{A,L} \quad (3.1)$$

² We have used factors not asset classes because it is common practice in robust optimization, see for instance Goldfarb and Iyengar (2003), Ling and Xu (2012) and other studies. It reduces the number of parameters to be estimated by about 20% in comparison with the classical approaches (e.g. the Sharpe and Tint model) and by much more for Bayes-Stein and Black-Litterman; and there is some evidence that it helps in creating more stable portfolios. Finally, the use of factors plays a significant role in making the robust optimization problem computationally tractable (e.g. a second order cone problem - SOCP), see for instance Goldfarb and Iyengar (2003), Glasserman and Xu (2013), and Kim et al. (2014).

where $\tilde{\mathbf{r}}_{A,L}$ is a joint column vector with $n_A + n_L$ elements that contains the uncertain asset and liability returns; $\tilde{\boldsymbol{\mu}}_{A,L}$, with $n_A + n_L$ elements, is the joint column vector of the random asset and liability mean returns; the column vector \mathbf{f} with m (number of factors) elements contains the factor returns that drive the risky assets and liabilities; the matrix $\tilde{\mathbf{V}}$ with m rows and $n_A + n_L$ columns contains the corresponding uncertain factor coefficients; and $\tilde{\boldsymbol{\varepsilon}}_{A,L}$ with $n_A + n_L$ elements is the column vector of uncertain disturbances. The covariance matrix of the factor returns is denoted by \mathbf{F} (m rows and m columns), and the diagonal covariance matrix of the disturbances by $\tilde{\mathbf{D}}$ ($n_A + n_L$ elements on its diagonal).

The matrix of uncertain factor coefficients ($\tilde{\mathbf{V}}$) belongs to an elliptical uncertainty set denoted by S_ν . The elements of the column vector of random mean returns ($\tilde{\boldsymbol{\mu}}_{A,L}$) and the diagonal elements of the covariance matrix of the disturbances ($\tilde{\mathbf{D}}$) lie within certain intervals which are represented by the uncertainty structures S_{mean} and S_d respectively. The uncertainty set S_{mean} is completely parameterized by the market data, while S_{mean} and S_ν also depend on the parameter ω . The parameter ω specifies the level of confidence, and hence allows the provision of probabilistic warranties on the performance of the robust portfolios. More details about the exact form of the uncertainty structures and their parameterization can be found in Appendix 3.B. The use of a factor model means there is no need to estimate the covariance matrix of the asset-liability returns, just the covariance matrix of the factor returns. This reduces the dimensionality of the problem from $[n(n-1)/2+n]$,

where n is the total number of assets and liabilities, to $[m(m-1)/2+m]$, where m is the number of factors.

The robust optimization problem we solve is given by the following maximin problem. See also Goldfarb and Iyengar (2003), section 3.3.

$$\begin{aligned}
& \underset{\Phi_A}{\text{maximize}} && \left[\frac{\min_{\{\tilde{\mu}_{A,L} \in S_{mean}\}} [\tilde{\mu}_{A,L}^T \Phi_{A,L}]}{\sqrt{\max_{\{\tilde{V} \in S_v\}} [\Phi_{A,L}^T \tilde{V}^T \mathbf{F} \tilde{V} \Phi_{A,L}] + \max_{\{\tilde{D} \in S_d\}} [\Phi_{A,L}^T \tilde{D} \Phi_{A,L}]}} \right] \\
& \text{s.t.} && \mathbf{1}^T \Phi_A = 1 \\
& && \Phi_{A,i} \geq 0, \quad \forall i = 1, \dots, n_A \quad (3.2) \\
& && - \sum_{i \in \text{class}X} \Phi_{A,i} + \theta_{max}^{\text{class}X} \mathbf{1}^T \Phi_A \geq 0, \quad \forall \text{class}X \\
& && \sum_{i \in \text{class}X} \Phi_{A,i} - \theta_{min}^{\text{class}X} \mathbf{1}^T \Phi_A \geq 0, \quad \forall \text{class}X
\end{aligned}$$

where $\Phi_{A,L}$ denotes the joint column vector of asset proportions Φ_A (decision variables) and liability proportions Φ_L . Liability proportions are fixed, negative, and sum to -1. $\theta_{min}^{\text{class}X}$ and $\theta_{max}^{\text{class}X}$ represent the minimum and maximum values for each broad asset class, classX (e.g. classX = equities, bonds, property, etc.). The objective in equation (3.2) is to maximise the Sharpe ratio under the worst circumstances (i.e. maximin)³. The worst case mean return in the numerator is divided by the square root of the worst case variance (two terms) in the denominator. The worst case mean return, as well as the worst-case variance, depend on the parameter ω (more details about their parameterization from market data and the confidence level can be found in Appendices 3.B and 3.C). As ω is increased the size of the uncertainty

³ For mathematical reasons the expected Sharpe ratio must be constrained to be strictly positive, and so the lower bound on expected returns of the asset-liability portfolio is set to 0.1%, rather than zero. This rules out asset allocations that are expected to worsen the scheme's funding position.

sets increases, which worsens the worst-case circumstances. The maximin optimization problem described in (3.2) can be converted to a second order cone program (SOCP) and hence is a tractable and easily solved mathematical optimization problem. Further details of this transformation appear in Appendix 3.C.

The first benchmark we use is the actual portfolios chosen by USS. The second benchmark is a modified version of the Sharpe and Tint (1990) model, where the objective function has been changed to maximise the expected return on the asset portfolio in excess of the liability portfolio, divided by the standard deviation of these excess returns (the Sharpe ratio). This model has the same additional constraints on the asset weights as the robust optimisation model, while the covariance matrix is estimated directly from the asset and liability returns. Hence, the modified Sharpe and Tint benchmark is :-

$$\begin{aligned}
& \underset{\Phi_A}{\text{maximize}} && \frac{\Phi_{A,L}^T \boldsymbol{\mu}_{A,L}}{\sqrt{\Phi_{A,L}^T \boldsymbol{\Sigma}_{A,L} \Phi_{A,L}}} \\
& \text{s.t.} && \mathbf{1}^T \Phi_A = 1 \\
& && \Phi_{A,i} \geq 0, && \forall i = 1, \dots, n_A \\
& && - \sum_{i \in \text{class}X} \Phi_{A,i} + \theta_{\max}^{\text{class}X} \mathbf{1}^T \Phi_A \geq 0, && \forall \text{class}X \\
& && \sum_{i \in \text{class}X} \Phi_{A,i} - \theta_{\min}^{\text{class}X} \mathbf{1}^T \Phi_A \geq 0, && \forall \text{class}X
\end{aligned} \tag{3.3}$$

where $\boldsymbol{\Sigma}_{A,L}$ denotes the sample covariance matrix of the asset and liability returns.

The third benchmark is Bayes-Stein (Jorion, 1986) where estimates of the inputs to a portfolio problem are based on the idea that estimated returns a long way from the norm have a higher chance of containing estimation errors than estimated returns close to the norm. To deal with estimation errors, the Bayes-Stein estimates of the

input parameters are the weighted sum of the historic return for each asset and a global estimate of returns (the norm). The global estimate of returns (the norm) is the return on the minimum variance portfolio when short sales are permitted, denoted $\mu_{A,L,min}$. The factor governing the extent to which the historic returns are shrunk towards the global norm is denoted by $0 \leq g \leq 1$. The Bayes-Stein estimator of the column vector of the sample mean asset and liability returns ($\boldsymbol{\mu}_{BS}$) is given by :-

$$\boldsymbol{\mu}_{BS} = (1-g)\boldsymbol{\mu}_{A,L} + g\mu_{A,L,min}\mathbf{1} \quad (3.4)$$

where the shrinkage coefficient is:-

$$g = \frac{(n_A + n_L) + 2}{((n_A + n_L) + 2) + p(\boldsymbol{\mu}_{A,L} - \mu_{A,L,min}\mathbf{1})^T \boldsymbol{\Sigma}^{-1} (\boldsymbol{\mu}_{A,L} - \mu_{A,L,min}\mathbf{1})} \quad (3.5)$$

The number of assets and liabilities and the sample size are denoted by $n_A + n_L$ and p respectively, and $\mathbf{1}$ is a column vector of ones. As in Jorion (1986) and Bessler,

Opfer and Wolff (forthcoming) we define $\boldsymbol{\Sigma}$ as $\boldsymbol{\Sigma} = \left(\frac{p-1}{p-(n_A + n_L) - 2} \right) \boldsymbol{\Sigma}_{A,L}$. The

Bayes-Stein estimator of the covariance matrix of the asset and liability returns ($\boldsymbol{\Sigma}_{BS}$) is given by:-

$$\boldsymbol{\Sigma}_{BS} = \frac{(p-1)(p+\varphi+1)}{(p+\varphi)[p-(n_A + n_L) - 2]} \boldsymbol{\Sigma}_{A,L} + \frac{\varphi(p-1)}{p(p+\varphi+1)[p-(n_A + n_L) - 2]} \frac{\mathbf{1}\mathbf{1}^T}{\mathbf{1}^T \boldsymbol{\Sigma}_{A,L}^{-1} \mathbf{1}} \quad (3.6)$$

where the scalar φ represents the precision of the prior distribution of returns, and is expressed as:-

$$\varphi = \frac{(n_A + n_L) + 2}{(\boldsymbol{\mu}_{A,L} - \mu_{A,L,min} \mathbf{1})^T \boldsymbol{\Sigma}^{-1} (\boldsymbol{\mu}_{A,L} - \mu_{A,L,min} \mathbf{1})} \quad (3.7)$$

We use the Bayes-Stein estimates of the returns vector and the covariance matrix in the modified Sharpe and Tint model.

Black-Litterman is the final benchmark and is another way of dealing with estimation errors which combines the subjective views of the investor concerning expected returns and risks with those of a reference portfolio, which is usually the market equilibrium asset proportions and covariance matrix, (Black and Litterman, 1992; Idzorek, 2005; and Bessler, Opfer and Wolff; forthcoming). The resulting (posterior) estimates of expected returns and the covariance matrix are then used in a portfolio model. We use the modified Sharpe and Tint model in equation (3.3), which maximises the Sharpe ratio, subject to various constraints.

The implied excess-returns used in the original Black and Litterman (1992) model are based on reverse-optimization, assuming that the reference portfolio is the result of a mean-variance optimization. In a similar way, Ardia and Boudt (2015) define implied returns as *'the returns for which a supposedly mean-variance efficient portfolio is effectively efficient given a covariance matrix'*. Given that we consider both risky assets and stochastic liabilities that behave as risky assets with their own negative weights, we define the column vector of implied excess returns for the reference portfolio as:-

$$\boldsymbol{\Pi} = \mathbf{R} \boldsymbol{\Sigma}_{A,L} \boldsymbol{\Phi}_{A,L}^R \quad (3.8)$$

where R is the investor's risk aversion parameter, which disappears in the Sharpe and Tint formulation, $\Phi_{A,L}^R$ is a column vector of the asset and (negative) liability weights of the reference portfolio estimated using the proportions for USS at the start of each out-of-sample period⁴. The posterior column vector of asset and liability returns (μ_{BL}) is:-

$$\mu_{BL} = \left[(c\Sigma_{A,L})^{-1} + P^T\Omega^{-1}P \right]^{-1} \left[(c\Sigma_{A,L})^{-1}\Pi + P^T\Omega^{-1}QV \right] \quad (3.9)$$

where c is the overall level of confidence in the vector of implied excess returns and is set to 0.1625⁵, P is a matrix of zeros and ones defining the assets and liabilities involved in each view, QV is a column vector of the views of returns specified by the investor, and Ω is a diagonal matrix of the reliability of each view, estimated following Meucci (2010), as:-

$$\Omega = \frac{1}{\delta} P\Sigma_{A,L}P^T \quad (3.10)$$

where δ is the overall level of confidence in the investor's views, which we set to unity⁶. In selecting investor views we follow Bessler, Opfer and Wolff (forthcoming) and use the mean return for each asset and liability during the estimation period. This is to avoid supplying Black-Litterman with more information than any other

⁴ Bessler, Opfer and Wolff (forthcoming) show that Black-Litterman results are robust to the choice of reference portfolio.

⁵ This is the mean of the range of values used by previous studies. Bessler, Opfer and Wolff (forthcoming) show that Black-Litterman results are robust to the choice of c over the 0.025 to 1.00 range.

⁶ We experimented with different values of δ and found it had little effect on the Black-Litterman performance, and so we followed Meucci (2010) and set δ equal to one.

technique. Following Satchell and Scowcroft (2000) and Bessler, Opfer and Wolff (forthcoming) the posterior covariance matrix (Σ_{BL}) is estimated by:-

$$\Sigma_{BL} = \Sigma_{A,L} + \left[(c\Sigma_{A,L})^{-1} + P^T \Omega^{-1} P \right]^{-1} \quad (3.11)$$

3.3 Application to USS

The robust optimization model will be used to derive optimal investment policies for USS. USS was created in 1974 as the main pension scheme for academic and senior administrative staff in UK universities and other higher education and research institutions (Logan, 1985). In 2014 USS was the second largest pension scheme in the UK, and the 36th largest in the world with 316,440 active members, deferred pensioners and pensioners. It is a multi-employer scheme with 379 separate sponsors (or institutions), and assets valued at £42 billion in 2014. There are two important advantages in using USS as the real world application. First, there is no need to include the assets and liabilities of the sponsors in the model; and second, because the USS sponsors are tax exempt, the optimal asset allocation is not the corner solution of 100% bonds which maximises the tax relief.

It is generally accepted that a pension scheme and its sponsor should be treated as a single economic entity, and this has a number of important implications for ALMs. The discount rate used in valuing pension liabilities must be increased above the riskless rate to reflect the risk of default by the sponsor. This sponsor default risk can be reduced by incorporating the assets and liabilities of the sponsor in the ALM model, allowing for correlations between the assets and liabilities of the sponsor and

the pension scheme. For example, if the sponsor makes cars, investment by the pension scheme in the shares of car producers increases the risk of default because, when the scheme has a deficit due to poor investment returns, the sponsor is also likely to be experiencing adverse business conditions. The value of the pension liabilities is an input to the ALM model, and this valuation depends on the discount rate used to value the liabilities, which in turn depends on the risk of default by the sponsor. The risk of default depends on the asset allocation of the pension scheme, which is the output from the ALM model. This leads to a circularity that Inkmann and Blake (2012) solve using simulation.

However, USS is a multi-employer scheme where the 379 institutions (sponsors) are funded largely by the UK government and student fees. Therefore the default risk of the sponsor is effectively uncorrelated with the assets and liabilities of the scheme. In addition, USS is a last-man-standing multi-employer scheme, and default would require the bankruptcy of every institution, i.e. the collapse of the UK university and research community. Therefore, for USS the default risk of the sponsor is both very low and independent of that of the scheme, and so need not be considered when setting the discount rate. In consequence the assets and liabilities of the 379 sponsors will not be incorporated in the ALM model of USS.

Taxation of the sponsor creates an arbitrage argument for a pension fund to invest 100% in bonds (as did Boots in 2000). There are two different arguments for 100% bond investment by a pension fund:- (a) Black (1980) (see also Surz, 1981; Black and Dewhurst, 1981, Frank, 2002), and (b) Tepper (1981) (see also Bader, 2003, Frank, 2002). However, UK universities (the sponsors of USS) are not liable to pay tax, and

so the tax arguments of Black and Tepper leading to 100% investment in bonds do not apply. The tax exemption of UK pension schemes also means there is no need to adjust their returns for taxation.

3.4 Data and Analysis

a. Assets. The main asset classes used by the trustees of UK pension funds, including USS, over the past two decades are UK, European and US equities, US and UK bonds, UK property and cash. In recent years interest in alternative assets has increased, and so we have also included this asset class (represented by hedge funds and commodities). We used 11 assets, and these are set out in the upper section of Table 3.1. As a robustness check we replaced three assets with S&P GSCI Light Energy, UK private equity and UK infrastructure (see section 3.6). Monthly data on these assets was collected from DataStream and the Bank of England for the period April 1993 to March 2011 (216 observations) corresponding to seven triennial actuarial valuations of USS, and we used monthly returns. Although all its liabilities are denominated in sterling, USS has substantial investments in non-sterling assets (about £15 billion in 2013). Until 2006 USS did not hedge any of this foreign exchange risk, but thereafter USS hedged all its foreign exchange risk. Therefore asset returns are adjusted onto a sterling basis for April 2006 onwards. Finally, we ignore the administrative expenses and transaction costs of the scheme.

Its maximin objective means that robust optimization will tend to perform better than the other techniques when the market falls; while the USS benchmark with its high equity investment will tend to perform better in a rising market. Figure 3.1

shows that the data period (1993-2011) covers a wide range of economic conditions, with three strong upward trends in the UK equity market, and two substantial falls. Property and hedge funds show a rise until mid-2007 followed by a sharp fall, and then a rise, while 20 year bond prices rise in the late 1990s and are then fairly steady. This suggests that the results below are not due to testing the ALM models on a falling market, which would favour robust optimization and penalize USS. Indeed, over the entire data period, an investment in the FTSE index basket rose by more than 200%, and hedge funds and property rose by an even larger percentage.

Type	Source
Asset Classes	
UK Equities	FTSE All Share Total Return
EU Equities	MSCI Europe excl. UK Total Return
US Equities	S&P500 Total Return
10 year UK Bonds	10-year UK Gov. Yields
20 year UK Bonds	20-year UK Gov. Yields
10 year US Bonds	10-year US Gov. Yields
20 year US Bonds	20-year US Gov. Yields
Hedge Funds	HFRI Hedge Fund Index
Commodities	S&P GSCI Total Return
UK Property	IPD Index Total Return
Cash	UK 3 Month Treasury Bills
Factors	
Global Equities	MSCI World Total Return
20 year UK Bonds	20-year Gov. Yields
UK Expected Inflation	UK 10-year Implied Inflation
UK Short Term Interest Rate	6-month UK Interbank Rate

Table 3.1: Dataset for Asset Classes and Factors. Monthly data on the 11 assets and four factors, together with their sources, for April 1993 to March 2011




Figure 3.1: Cumulative Wealth from an Investment in Each of the Four Main Asset Classes 1993-2011

b. Liabilities. The liabilities were split into three groups - active members, deferred pensioners and pensioners. Changes in their value are driven by changes in four main factors - long-term interest rates, expected salary growth, expected inflation and longevity expectations. The actuarial equations in Board and Sutcliffe (2007) were used to compute the monthly returns for each of the three types of liability (see Appendix 3.D). This was done using monthly 20-year UK government bond yields⁷, and the monthly index of UK 10-year implied inflation. The 20-year government bond yield was used as the discount rate because, while no cash flow forecasts are available, USS is an immature scheme and data on the age distribution of active members, deferreds and pensioners suggests the duration of USS liabilities

⁷ This abstracts from the effects on returns of the low liquidity of pension liabilities and the inflation risk inherent in government bond yields, as these effects tend to cancel out.

is over 20 years (USS, 2013). The USS actuary estimates expected salary growth as expected inflation plus one percent, and so monthly changes in expected salary growth were computed in this way. Monthly data on changes in longevity expectations is not available, and so these expectations were held constant throughout each three year period at the value used in the preceding actuarial valuation (see the last two rows of Table 3.2).

The computation of the liabilities also requires a number of parameters - expected age at retirement, life expectancy at retirement, and the average age of actives and deferreds. Although the USS normal retirement age is 65 years, expected retirement ages are earlier. Row 1 of Table 3.2 shows the expected retirement ages for actives and deferreds used for each triennial USS actuarial valuation, and row 2 has the expected longevity of USS members at the age of 65 (USS Actuarial Valuations). Since the average age of USS members throughout the period was 46 years (HEFCE, 2010), and the average age of USS pensioners was 70 years (USS, 2013), the number of years for which each group was expected to receive a pension are also shown in rows 3 and 4 of Table 3.2. The USS accrual rate in the final salary section is $1/80^{\text{th}}$ per year. In addition there is a lump sum payment, and using the USS commutation factor of 16:1, this increases the accrual rate to $1/67.37$.

Expected Values in Years	1993/6	1996/9	1999/02	2002/5	2005/8	2008/11
Retirement Age	60	60	60	60	60	62
Longevity at 65	20	20	21	21	21	24
Pension Period - Actives & Deferreds	25	25	26	26	26	27
Pension Period – Pensioners	15	15	16	16	16	19

Table 3.2: Demographic data for actives, deferreds and pensioners: expected retirement age for active and deferred members, expected longevity at age 65, expected number of years for which current active and deferred members and pensioners will receive a pension

c. Constraints. The upper and lower bounds on the asset proportions of the five main asset classes were set so as to rule out extreme and unacceptable asset proportions. This was done with reference to the benchmarks and the associated permitted active positions specified by USS over the data period. The bounds used were: (35% ≤ Equities ≤ 85%); (5% ≤ Fixed Income ≤ 30%); (0% ≤ Alternative Assets ≤ 30%); (2% ≤ Property ≤ 15%) and (0% ≤ Cash ≤ 5%). In addition, the expected return on the asset-liability portfolio was required to be non-negative, and short sales and borrowing were excluded.

Actuarial valuations of USS are carried out every three years, with the oldest available actuarial valuation on 31st March 1993, and the most recent valuation on 31st March 2011. The data is divided into six non-overlapping periods to coincide with these seven triennial actuarial valuations, as shown in Table 3.3.

Periods (t)	Start	End	Length (p)
Period 1	1993 M4	1996 M3	36
Period 2	1996 M4	1999 M3	36
Period 3	1999 M4	2002 M3	36
Period 4	2002 M4	2005 M3	36
Period 5	2005 M4	2008 M3	36
Period 6	2008 M4	2011 M3	36

Table 3.3: Non-overlapping three year data periods: the data is divided into six 3-year periods corresponding to USS actuarial valuations

To estimate and test the asset allocations we used four out-of-sample non-overlapping windows. USS is a very long term investor and sets its asset allocation every three years after conducting an actuarial valuation and commissioning an asset-liability study. Therefore we have used a three year out-of-sample period, together with a six year estimation period. In the literature, the length of the rolling estimation windows varies, but five to ten years is generally considered to be appropriate. For instance, Bessler, Opfer and Wolff (forthcoming) and Bessler and Wolff (2015) use rolling estimation windows of 1 year, 2 years, 3 years, 4 years and 5 years. Xing et al. (2014) use rolling estimation periods of 5 years, 10 years and 15 years, while DeMiguel et al. (2009a) use 10 years. Hence, the choice of a six year rolling estimation window lies within the range used by previous studies and seems to be a sensible choice. The robustness checks in section 3.6 include a one year out-of-sample period with a three year estimation period and show that our conclusions remain unchanged. Also, the findings that shorter or longer estimation windows do not change the main conclusions are in accordance to Bessler, Opfer and Wolff

(forthcoming) and Bessler and Wolff (2015). The data for the initial six years (periods one and two) was used to compute the optimal robust optimization asset allocation for the subsequent three years (period three). The estimation period was then rolled forward by 36 months, so that data for periods two and three was now used to compute the optimal asset allocation, which was tested on data for period four, and so on, giving four out-of-sample test periods of 36 months each, providing 144 out-of-sample months.

Each of the three liabilities was treated as a separate risky ‘asset class’ with ‘negative’ and fixed weights for each of the six three year periods. We calculated the proportions for each type of pension liability from the triennial actuarial valuations (see Table 3.4). For the six year estimation periods we used the average of the liability weights for the two 3-year periods concerned. In computing returns on the asset and liability portfolios, the assets were weighted by the funding ratio at the start of the relevant three year period.

Type	Period 1	Period 2	Period 3	Period 4	Period 5	Period 6
Active	59.50%	57.42%	55.23%	57.36%	52.51%	52.20%
Pensioners	36.05%	36.25%	37.89%	35.01%	39.56%	39.90%
Deferred	4.45%	6.33%	6.88%	7.63%	7.93%	7.90%

Table 3.4: Proportions of total pension liabilities

d. Uncertainty Sets. For each estimation period, we calculated the parameters involved in the three uncertainty sets, and hence in the final mathematical optimization problem, using a factor model (equation 3.1). The returns uncertainty set requires the estimation of 14 mean returns, and for each estimation period the

means of the 14 asset and liability returns were used. For each of the six year estimation periods natural log returns on the 11 assets and three liabilities were separately regressed on the natural log returns of the four factors listed in the lower section of Table 3.1, together with a constant term. These 14 regressions per estimation period generated 56 estimated coefficients (the factor loadings matrix) and 14 constant terms. In total, the factor loadings uncertainty set requires the estimation of 56 coefficients and 10 covariances. The disturbances uncertainty set has 14 parameters, and these were estimated using the residuals from the regressions. The modified Sharpe and Tint model requires 105 elements of the covariance matrix to be estimated. So overall this benchmark requires the estimation of $(105+14) = 119$ stochastic parameters per estimation period; while robust optimisation needs $(56+10+14+14) = 94$ stochastic parameters; a reduction of over 20%.

Bayes-Stein requires the estimation of 105 elements of the asset-liability covariance matrix, 14 elements of the mean asset-liability portfolio returns, as well as the estimation of the parameters g and $\mu_{A,L,min}$ (see section 3.2 for more details). Hence, the Bayes-Stein approach requires in total the estimation of $(105+14+2) = 121$ parameters per estimation period; 29% more than robust optimization. Black-Litterman requires the estimation of 105 elements of the asset-liability covariance matrix, the weights of the reference asset-liability portfolio (14), the vector of the investor's views (14), the overall level of confidence in the views (1), the factor that measures the reliability of the implied return estimates (1) and the reliability of each

view (14); giving a total of 149 parameters to be estimated, or 59% more than robust optimization.

Given the very large number of parameters to be estimated for each optimization model in each estimation period (over 2,000 parameters in total), It is not possible to include a table with the input parameters used in the optimization processes. Instead, Figure 3.1 gives a very clear picture of the variability of 4 main asset classes used in our analysis.

The factor model requires us to choose the number and identity of the factors, and we experimented with both the identity and number of factors before settling on those listed in Table 3.1. It is helpful if the ratio of the number of factors to the number of assets and liabilities is small. In previous robust optimization studies this ratio is 0.080 and 0.233 (Goldfarb & Iyengar, 2003); 0.125 and 0.238 (Ling & Xu, 2012) and 0.200 (Glasserman & Xu, 2013). With four factors and 14 assets and liabilities we have a ratio of 0.286, which is higher than previous studies. Five factors would increase the total number of parameters to be estimated by 19 and increase this ratio to 0.357, which would be appreciably higher than any previous study. Therefore we settled on a parsimonious model of four factors.

The universe of assets used by Goldfarb and Iyengar (2003) consists of 43 stocks ranked at the top of each of 10 industry categories by Dow Jones. Their robust optimization model uses 10 factors that consist of 5 major indices (US equity indices and US bonds) as well as 5 eigenvectors that have the 5 largest eigenvalues of the covariance matrix of the asset returns and hence describe very well their asset

universe. In a similar way, Ling and Xu (2012) use 21 stocks from Shanghai Stock 50 index and 5 major Shanghai equity indices as market factors in their robust formulated portfolio optimization model. Furthermore, Glasserman and Xu (2013) use 3 factors in their model, which are moving averages of prices changes over periods different length, for their investment universe that consist of 15 commodity futures.

Our investment universe differs significantly from that of previous studies on robust portfolio optimization since we include both conventional and alternative asset classes as well as pension liabilities. Hence, the selection of the factors should be done carefully in order to the factor model used in our robust optimization model to be as much accurate as possible for a reasonable number of factors that should be significantly smaller than the number of assets.

We chose 20 year UK government bond prices as one of the factors because the discount rate is a key determinant of the value of the three liabilities. It also helps to explain the prices of the four long term government bonds. We included the UK 10 year implied inflation rate as a factor because it is another important determinant of the three liabilities, and also helps to explain asset returns. The MSCI World Total Return index was added to explain returns on the three equity indices and, to a lesser extent, returns on the three alternative assets. Finally we used the 6 month UK interbank rate as the fourth factor to explain cash and other asset returns. Although alternative or additional factors could also have been used to improve the explanatory power, such as by adding liquidity to increase the explanatory power of the factor model on UK properties, it is also important the number of market factors

used to be significantly smaller from the number of assets and liabilities for reasons regarding the dimensionality of the mathematical programming model and its computationally tractability as discussed previously. Hence, the factors should be selected carefully to both describe well all the dependent variables and keep the dimensionality of the robust formulation in reasonable levels. Indeed, the R^2 values in Table 3.5 show that these four factors do a good job in explaining returns for all of the 14 assets and liabilities. In section 3.6 we present robustness checks where we use three different factors.

The adjusted R^2 values and significance of these 14 regressions for the entire data set appear in Table 3.5 (the results for each of the four 72 month estimation periods were broadly similar). This shows that for all the assets and liabilities, an F -test on the significance of the equation was significant at the 0.1% level, and the adjusted R^2 s were generally high. The 100% adjusted R^2 for 20 year UK bonds is because 20 year UK bonds was one of the four factors included in the factor model.

The equations in Appendix 3.B were then used to compute the three uncertainty sets. Robust optimization requires a value of ω to be chosen. Goldfarb and Iyengar (2003) point out that the correct selection of the parameter ω remains a problem since the level of uncertainty (noise) in the input parameters is not known *a priori*, Gulpinar and Pachamanova (2013) mention that the choice of ω should be done carefully in order to enhance the performance of the robust portfolio optimization model, while Scherer (2007) also points out that the calibration of ω is a difficult task. However, previous authors have used a value of $\omega = 0.99$ (e.g. Goldfarb and Iyengar, 2003; Delage and Mannor, 2010; Kim, Kim, Ahn and Fabozzi, 2013; and Ling

and Xu, 2012) and we also set $\omega = 0.99$. Goldfarb and Iyengar (2003) also state that setting $\omega \approx 1$ is a sensible choice. Furthermore, we experimented with other values of ω , such as 0.90 and 0.95, and obtained broadly similar results (see section 3.6). The value of ω was not set equal to unity because the required confidence level would become infinite.

	1993-2011	
	Adj.R ² %	p-value
UK Equities	33.34	0.000
EU Equities	91.04	0.000
US Equities	94.88	0.000
10 year UK Bonds	82.64	0.000
20 year UK Bonds	100.00	0.000
10 year US Bonds	43.90	0.000
20 year US Bonds	41.06	0.000
Hedge Funds	79.76	0.000
Commodities	40.90	0.000
UK Property	24.63	0.000
Cash	42.97	0.000
Actives	93.33	0.000
Deferreds	93.45	0.000
Pensioners	93.34	0.000

Table 3.5: Adjusted R^2 and Significance Levels of the 14 Regression Equations. Monthly returns on each of the assets and liabilities for 1993 to 2011 were regressed on monthly returns of four factors. These are the MSCI World total return index, the 20 year UK bonds, the implied UK 10 year inflation rate and the UK 6 month interbank rate.

3.5 Results

The asset allocations for robust optimization and the four benchmarks for the four

out-of-sample periods appear in Table 3.6⁸. This shows that, while the robust optimization solutions are subject to upper and lower bounds on the asset proportions, these constraints are never binding. However varying these bounds changes the optimal solutions because the robust counterpart is a nonlinear convex optimisation problem, and its optimal solution need not be at a corner point. Therefore the solutions were affected by the upper and lower bounds. Sharpe and Tint, Bayes-Stein and Black-Litterman are constrained in every period by the upper bound of 15% on property. In periods 4 and 5 Sharpe and Tint and Bayes-Stein are constrained by the lower bound on equities of 35%, while in period 5 Sharpe and Tint, Bayes-Stein and Black-Litterman are constrained by the upper bound of 30% on alternatives. In contrast to the vast majority of literature where portfolios are rebalanced every month, see for instance DeMiguel et al. (2009b), Bessler and Wolff (2015), and hence the portfolio rebalancing is smoother over the policy horizon, the 6 years (72 months) estimation periods used in our study are rolled forward by 3 years (36 months) for each new rebalancing, and hence two consecutive estimation windows are just overlapped for the half of their length. As a result, it is not very surprising that the asset allocation for the non-RO techniques is volatile over the investment horizon. In section 3.6 we also examine the effects of relaxing these bounds.

⁸ The robust optimization ALM model was solved using SeDuMi 1.03 within MATLAB (Sturm, 1999), and this took 0.67 seconds for each out-of-sample period on a laptop computer with a 2.0 GHz processor, 4 GB of RAM and running Windows 7. The modified Sharpe and Tint model was solved using the *fmincon* function in MATLAB for constrained nonlinear optimization problems (interior point algorithm) and took less than a second, as did the Bayes-Stein and Black-Litterman models.

	Period 3 - 1999-2002					Period 4 - 2002-2005					Period 5 - 2005-2008					Period 6 - 2008-2011				
	RO	S&T	USS	BS	BL	RO	S&T	USS	BS	BL	RO	S&T	USS	BS	BL	RO	S&T	USS	BS	BL
UK Equities	17.33	0.00	57.38	0.00	24.91	18.13	31.44	38.81	30.43	26.51	15.46	0.00	38.36	0.00	0.00	18.68	0.00	23.06	0.00	0.00
EU Equities	15.44	10.22	11.46	10.62	12.12	13.27	0.00	21.50	0.00	0.00	13.65	35.00	22.52	35.00	38.55	11.62	50.00	18.32	50.00	50.00
US Equities	15.63	57.91	11.46	60.69	38.37	15.52	3.56	21.50	4.57	8.49	18.00	0.00	22.52	0.00	2.89	16.52	0.00	18.32	0.00	0.00
Total Equities	48.39	68.13	80.30	71.32	75.41	46.92	35.00	81.81	35.00	35.00	47.10	35.00	83.40	35.00	41.44	46.82	50.00	59.70	50.00	50.00
10-year UK Bonds	6.14	0.00	2.15	0.00	0.00	5.90	0.00	2.72	0.00	0.00	6.12	0.00	0.22	0.00	0.00	6.46	0.00	1.90	0.00	0.00
20-year UK Bonds	5.73	13.87	2.15	13.68	9.59	5.33	24.45	2.71	24.29	20.00	6.83	0.00	0.22	0.00	0.00	6.48	0.00	1.90	0.00	0.00
10-Year US Bonds	5.29	0.00	2.35	0.00	0.00	5.99	0.00	1.33	0.00	0.00	5.23	11.18	4.78	10.05	13.56	5.32	0.00	4.25	0.00	0.00
20-year US Bonds	4.52	0.00	2.35	0.00	0.00	5.49	0.00	1.33	0.00	0.00	4.42	8.82	4.78	9.95	0.00	4.61	5.00	4.25	5.00	5.00
Total Bonds	21.68	13.87	9.00	13.68	9.59	22.71	24.45	8.09	24.29	20.00	22.59	20.00	10.00	20.00	13.56	22.87	5.00	12.30	5.00	5.00
Commodities	7.64	0.00	0.00	0.00	0.00	7.26	0.00	0.00	0.00	0.00	6.12	30.00	0.00	30.00	17.96	7.90	30.00	8.00	30.00	22.00
Hedge Funds	8.87	3.00	0.00	0.00	0.00	9.03	25.55	0.00	25.71	30.00	10.15	0.00	0.00	0.00	12.04	8.80	0.00	8.00	0.00	8.00
Total Alternatives	16.52	3.00	0.00	0.00	0.00	16.29	25.55	0.00	25.71	30.00	16.26	30.00	0.00	30.00	30.00	16.70	30.00	16.00	30.00	30.00
UK Property	9.51	15.00	8.40	15.00	15.00	10.01	15.00	8.30	15.00	15.00	9.74	15.00	3.10	15.00	15.00	9.54	15.00	7.00	15.00	15.00
Cash	3.90	0.00	2.30	0.00	0.00	4.08	0.00	1.80	0.00	0.00	4.30	0.00	3.50	0.00	0.00	4.08	0.00	5.00	0.00	0.00

Table 3.6: Asset Proportions for Robust Optimization (RO), Sharpe and Tint (S&T), USS, Bayes-Stein (BS) and Black-Litterman (BL). Six year estimation period and 144 months out-of-sample. Optimal asset allocations for each of the four out-of-sample 36 month periods using a 72 month estimation period. The original bounds, assets, factors and constraints were used, and $\omega = 0.99$.

Table 3.6 shows that robust optimization leads to remarkably stable asset proportions across the 12 out-of-sample years, with between 47% and 48% in equities, 22% to 23% in bonds, 4% in cash, 10% in property, and between 16% and 17% in alternatives. The asset allocations for the four benchmarks are much more variable. The modified Sharpe and Tint asset allocations vary for equities between 35% and 68%, bonds between 5% and 24%, and alternatives between 3% and 30%. Property is always constrained at the upper bound of 15%, and cash is always constrained at the lower bound of zero. The Bayes-Stein and Black-Litterman allocations also show considerable variability. For example the Bayes-Stein and Black-Litterman allocation to alternatives varies from zero to 30%, while bond allocations vary from 5% to over 20%. The USS allocations are also variable. For the first three periods USS had over 80% of the assets invested in equities, with no investment in alternative assets until the last period, when it jumped to 16% and the equity proportion fell to 60%.

Table 3.7 compares the results for the five methods over the 144 out-of-sample months. The 144 months comprised the four out-of-sample 36 month periods, with different solutions applying for each 36 month period. Where relevant, the monthly figures were adjusted to give annualized figures. In Table 3.7 the score for the technique with the best performance on each measure is in bold. Robust optimization is the best technique on all the performance measures.

The Sharpe ratio is the mean excess return on the asset-liability portfolio divided by the standard deviation of returns of the asset-liability portfolio. Since the Sharpe ratio uses the standard deviation to measure risk, to provide an alternative

perspective we also used a wide range of additional performance measures which do not rely on the standard deviation. The annualized conditional Sharpe ratio (Eling and Schuhmacher, 2007) is computed at the 99% level, and the annualized mean excess return is the mean of the asset minus liability returns. The Sortino ratio is the mean return on the asset-liability portfolio, divided by the standard deviation of returns on the asset-liability portfolio computed using only negative returns, with the minimum acceptable return set to zero (Prigent, 2007). The Dowd ratio is the mean return on the asset-liability portfolio, divided by the value at risk for a chosen confidence level (deflated by the initial value of the asset-liability portfolio), (Prigent, 2007). We used the 99% confidence level for the value at risk when computing the Dowd ratio, and have also included the value at risk and the conditional value at risk, both at the 99% level, in Table 3.7 as measures of tail risk. Another distribution-free performance measure - second order stochastic dominance - is also included in Table 3.7.

Performance Measures		Robust Optimization	Sharpe & Tint	USS	Bayes -Stein	Black-Litterman
1	Annualized Mean Sharpe Ratio	0.0793	-0.0221	0.0282	-0.0275	-0.0028
2	Annualized Conditional Sharpe Ratio (99%)	0.0052	-0.0013	0.0020	-0.0017	-0.0002
3	Annualized Mean Excess Return	1.353%	-0.4476%	0.5368%	-0.5580%	-0.0556%
4	Annualized Mean SD (Negative Returns)	11.92%	14.70%	13.49%	14.74%	14.35%
5	Annualized Sortino Ratio	0.1135	-0.0304	0.0398	-0.0379	-0.0039
6	Dowd Ratio	0.0057	-0.0014	0.0020	-0.0018	-0.0002
7	2 nd Order Stochastic Dominance	First	Fourth	Second	Fifth	Third
8	Value at Risk (99%)	0.1986	0.2638	0.2189	0.2638	0.2586
9	Conditional VaR (99%)	0.2158	0.2798	0.2293	0.2798	0.2697
10	Omega Ratio	1.0682	0.9813	1.0237	0.9768	0.9976
11	Maximum Drawdown	0.4551	0.6611	0.5325	0.6612	0.6452
12	Sterling Ratio	0.0095	-0.0021	0.0027	-0.0025	-0.0003
13	Mean Diversification	0.1135	0.3141	0.2544	0.3208	0.2631
14	Entropy-based Diversification	9.786	2.823	5.624	3.7236	4.4075
15	Mean Stability	0.0016	0.3203	0.0347	0.3276	0.1961
16	Mean Funding Ratio	93.01%	87.89%	89.98%	87.69%	89.10%
17	SD (Funding Ratio)	9.700%	11.95%	11.61%	12.04%	12.33%
18	Mean Projected Contribution Rate	17.68%	18.11%	17.93%	18.28%	18.11%
19	SD (Projected Contribution Rate)	0.8345%	1.026%	1.592%	1.350%	1.387%
20	Cumulative Asset Return	32.21%	-11.97%	16.56%	-13.32%	-3.466%

Table 3.7: Out-of-Sample Performance Measures – Six Year Estimation Period. The results for each of the four out-of-sample 36 month periods were adjusted, where relevant, on to an annualised basis. The original bounds, assets, factors and constraints were used, and $\omega = 0.99$.

The Omega ratio is the ratio of the average gain to the average loss, and is an additional distribution-free measure of performance (Bessler, Opfer and Wolff, forthcoming). Gains are the positive excess returns of assets over liabilities, and

losses are the negative excess returns of assets over liabilities. The drawdown rate, which is also distribution-free, measures declines from peaks in cumulative wealth over a specific time horizon (t). The drawdown rate is defined as:-

$$DD_t = \frac{\max_{0 \leq \tau \leq t} \{W_\tau\} - W_t}{\max_{0 \leq \tau \leq t} \{W_\tau\}} \quad (3.12)$$

where W_t is the cumulative wealth (assets only) at time t . Maximum drawdown is the largest drawdown rate. If a strategy has a lower mean drawdown rate over a specific time horizon in comparison with others, it tends to have lower volatility and value-at-risk. We also included some further performance measures based on the drawdown rate. The Sterling ratio is the mean asset return divided by the average drawdown rate (Eling and Schuhmacher, 2007).

Diversification of the asset-only portfolios was measured as the average across the four periods of the sum of the squared portfolio proportions for each period (Blume and Friend, 1974). For zero diversification the score is one, while for full diversification it is $1/n_A$ (or 0.091 when $n_A = 11$). Following Bera and Park (2008) we also used entropy to measure asset diversification. We modified Shannon's entropy by taking its exponent giving the measure Z in equation (3.13), so that for zero diversification $Z = 1$, and for full diversification $Z = n_A$, (which is 11).

$$Z = \exp\left(-\sum_{i=1}^{n_A} \Phi_{A,i} \ln(\Phi_{A,i})\right) \quad (3.13)$$

The stability of portfolio proportions from one triennial period to the next was measured as the average value across the three changes in asset allocation of the

sum of squares of the differences between the portfolio proportion for each asset in adjacent time periods (Goldfarb and Iyengar, 2003). This measure can be viewed as a proxy for transactions costs under the assumption that the cost functions are linear and similar across assets. Robust optimization adopts a maximin objective, and with an ω value of 0.99, it selects portfolios that are very likely to deliver at least their expected Sharpe ratio. Therefore it is pessimistic/conservative ($\omega \approx 1$) enough, tending to select very cautious portfolios. During bull markets it still selects portfolios that will deliver at least the expected Sharpe ratio, even if there is a market downturn. Thus robust optimization asset allocations tend to be more stable than those of other asset selection techniques. Our stability results are in accordance with the literature. For instance, Gulpinar and Pachamanova (2013) and Goldfarb and Iyengar (2003) report that the size of changes in asset proportions is significantly smaller for the robust portfolio formulation than for the classical approaches. The use of a four factor model, rather than 14 assets and liabilities, may also play a role in ensuring the stability of the robust optimization portfolios. For instance, MacKinlay and Pastor (2000) use factor models in an attempt to construct estimators of means and covariances of asset returns that are more stable and reliable than estimators obtained using traditional methods (e.g. sample-based moments), and hence to produce portfolios that eliminate the negative effects of estimation risk. Hence, we conclude that both the maximin objective with an $\omega \approx 1$ as well as the use of a factor model drive the stable asset allocation of the robust optimization strategy.

The monthly out-of-sample asset and liability returns, in conjunction with the values of USS assets and liabilities at the previous actuarial valuation, were used to compute the funding ratio each month. These monthly ratios were averaged to give the mean funding ratio. The mean projected contribution rate was computed using the actuarial formulae in Board and Sutcliffe (2007) with a spread period (M) of 15 years (see Appendix 3.A). The number of years accrued by the average member (P) was 18 years, while administrative expenses were set to zero. The term $N_A S$ cancels out with terms in AL_A . The values of the discount rate (d') and salary increase (e) were the average values over the preceding two triennial periods. Finally the cumulative returns over the 144 out-of-sample months in Table 3.7 are for just the asset portfolio.

Although many studies in the literature, see for instance DeMiguel et al. (2009b), Bessler and Wolff (2015), go one step further by exploring whether the differences in the performance metrics produced by different portfolio models are statistically significant, this is not the case for our study. Given the triennial valuation of USS and the very long-term horizon of pension funds in general, our analysis consists of only 4 out-of-sample (investment) periods, and hence a further analysis for statistical comparisons of the performance metrics between the different optimization models used is not possible to be conducted.

3.6 Robustness Checks

In the literature, it has been widely reported that portfolio optimization models are often sensitivity to a number of dimensions such as the length of estimation and investment periods, the rebalancing frequency, the constraints on portfolio weights,

and the universe of assets used, see for instance Becker et al. (2015), Bessler, Opfer and Wolff (forthcoming) and Bessler et al. (2015), amongst others, as well as the quality of factors used in robust formulated portfolios and the size of the uncertainty sets (uncertainty structures), see for instance Goldfarb and Iyengar (2003), Gulpinar and Pachamanova (2013) and Kim et al. (2014).

We varied the base case above along six dimensions. For each robustness check except the last we changed one aspect of the base case, while keeping the others at their values in the base case. We:- (i) reduced the estimation period from six to three years; (ii) used two alternative sets of factors in the robust optimization - in the first set UK expected inflation was replaced by RPI, and in the second set UK 6 month rates were replaced by UK 3 month rates, UK 20 year bonds by UK 10 year bonds, and UK expected inflation by RPI; (iii) the S&P GSCI total return index was replaced by the S&P GSCI Light Energy total return index, and UK private equity and UK infrastructure were included as alternative assets, replacing 20 year UK and US bonds; (iv) ω was reduced from 0.99 to 0.90; (v) the upper and lower bounds on the asset allocations were relaxed by 5%, becoming ($30\% \leq \text{equities} \leq 90\%$); ($0\% \leq \text{fixed income} \leq 35\%$); ($0\% \leq \text{alternative assets} \leq 35\%$); ($0\% \leq \text{property} \leq 20\%$); and ($0\% \leq \text{cash} \leq 10\%$), and (vi) the estimation period was reduced to three years and the out-of-sample period was reduced to one year.

Bessler, Opfer and Wolff (forthcoming) suggest that the reliability of the views incorporated in the Black-Litterman model is time-varying. For each of our out-of-sample periods in Table 3.7 we estimated the reliability of the views for the subsequent out-of-sample period using the entire estimation period of 72 months.

As a further robustness check, we compared the base case with five versions of the Black-Litterman model where we used the 12, 18, 24, 30 and 36 months immediately prior to the start of each out-of-sample period to estimate the reliability of the views, measured as the variance of the historic forecast errors. For all five of these shorter estimation periods, robust optimization remains superior on every performance measure.

The results for these six alternative formulations of the problem are summarised in Table 3.8. For the base case and the six alternative cases, the best technique for each performance measure is indicated. Where robust optimization is not the best technique, Table 3.8 also shows the ranking of robust optimization.

In the literature, there is some empirical evidence that alternative reward-to-risk ratios yield similar rankings across 2 specific asset classes (hedge funds and commodities) according to Eling and Schuhmacher (2007) and Auer (2015). However, in an environment where pension funds invest in a variety of asset classes (both conventional and alternative investments) with different statistical properties and take into account stochastic pension liabilities into their asset allocation process by using ALM models, we consider that a variety of discrete performance measures should be used and evaluated as in our case and hence all the performance metrics presented in sections 3.5 and 3.6 are important for the pension fund trustees given the existing complex investment situation.

Performance Measures		Base-Case	1	2	3	4	5	6
1	Annualized Mean Sharpe Ratio	RO	RO	RO	RO	RO	RO	RO
2	Annualized Conditional Sharpe Ratio	RO	USS-2	RO	RO	RO	RO	RO
3	Annualized Mean Excess Return	RO	BS-5	RO	BL-2	RO	RO	RO
4	Annualized Mean SD (Negative Returns)	RO	RO	RO	RO	RO	RO	RO
5	Annualized Sortino Ratio	RO	RO	RO	RO	RO	RO	RO
6	Dowd Ratio	RO	RO	RO	RO	RO	RO	RO
7	2 nd Order Stochastic Dominance	RO	RO	RO	RO	RO	RO	RO
8	VaR(99%)	RO	RO	RO	RO	RO	RO	RO
9	Conditional VaR	RO	RO	RO	RO	RO	RO	RO
10	Omega Ratio	RO	RO	RO	RO	RO	RO	RO
11	Maximum Drawdown	RO	RO	RO	RO	RO	RO	RO
12	Sterling Ratio	RO	RO	RO	RO	RO	RO	RO
13	Mean Diversification	RO	RO	RO	RO	RO	RO	RO
14	Entropy-based Diversification	RO	RO	RO	RO	RO	RO	RO
15	Mean Stability	RO	USS-2	RO	RO	RO	RO	RO
16	Mean Funding Ratio	RO	RO	RO	S&T-4	RO	RO	RO
17	SD (Funding Ratio)	RO	RO	RO	RO	RO	RO	RO
18	Mean Projected Contribution Rate	RO	BS-5	RO	S&T-4	RO	RO	ST-4
19	SD (Projected Contribution Rate)	RO	RO	RO	RO	RO	RO	RO
20	Cumulative Asset Return	RO	RO	RO	RO	RO	RO	RO

Table 3.8: Robustness Checks for the Base Case. Base case - 6 year estimation period, 3 year out-of-sample test period, $\omega = 0.99$, standard constraints, standard factors, standard assets. Column 1 - 3 year estimation period, 3 year out-of-sample test period, $\omega = 0.99$, standard constraints, standard factors, standard assets. Column 2 - 6 year estimation period, 3 year out-of-sample test period, $\omega = 0.99$, standard constraints, two sets of alternative factors versus the standard factors. The performance of robust optimization was compared using the three different sets of factors. Column 3 - 6 year estimation period, 3 year out-of-sample test period, $\omega = 0.99$, standard constraints, standard factors, alternative assets. USS was not used in the comparison as they did not invest in alternative assets until the last of the four periods. Column 4 - 6 year estimation period, 3 year out-of-sample test period, $\omega = 0.90$, standard constraints, standard factors, standard assets. Column 5 - 6 year estimation period, 3 year out-of-sample test period, $\omega = 0.99$, relax each of the constraints by 5%, standard factors, standard assets. Column 6 - 3 year estimation period, 1 year out-of-sample test period, $\omega = 0.99$, standard constraints, standard factors, standard assets. The asset proportions and performance measures for each robustness check are available on request.

Table 3.8 demonstrates that across a wide range of performance measures and robustness checks robust optimization is clearly the best technique for portfolio formation⁹. This is true for every performance measure and every robustness check.

Indeed, none of the four benchmarks is the best on more than three occasions, while

⁹ In every case robust optimization is also the superior technique when the Calmar ratio, Burke ratio and average drawdown performance measures are used (Eling and Schuhmacher, 2007).

robust optimization is superior on over 130 occasions.

3.7 Conclusions

In this paper we extended the robust mean-variance portfolio framework proposed by Goldfarb and Iyengar (2003) by incorporating risky pension liabilities as separate stochastic assets with their own fixed 'negative' weights. This framework uses a factor loadings matrix, rather than a covariance matrix, which reduces the number of parameters to be estimated by over 20%. Robust optimization ensures that the probability of the actual outcome being worse than the optimal robust solution is equal to $1-\omega$, where $0 < \omega < 1$. With its maximin objective function and confidence level (ω), robust optimization tends to rule out solutions based on favourable errors in estimating the stochastic input parameters, so tackling estimation risk. We modelled extra features of the pension ALM problem by adding additional linear constraints which set upper and lower bounds for each asset sub-class, ruled out short selling and borrowing, and required the expected return on the asset-liability portfolio to be non-negative. Our final formulation is computationally tractable and easily solved, which is not the case for ALM models using stochastic programming or dynamic stochastic control.

Our study is unusual because we disaggregate the pension liabilities into three components (active members, deferred members and pensioners). We also use the Sharpe ratio as the objective of a pension ALM model. Furthermore, we use an actuarial model to transform the optimal asset allocation to the scheme's projected contribution rate. Finally, we derive optimal investment policies using the robust optimization ALM framework for a real-world pension scheme - USS. The choice of

this scheme has the advantages that, as the sponsors are tax exempt, the tax arbitrage investment of 100% in bonds is irrelevant; and that, since default risk of the sponsors is both very low and independent that of the scheme, there is no need to include the assets and liabilities of the sponsors in the ALM. The performance of the robust optimization framework was benchmarked against the modified Sharpe and Tint, Bayes-Stein, Black-Litterman models, and USS's actual investments. This analysis allowed for USS switching to fully hedging its foreign exchange risk from April 2006.

The 20 performance measures and six robustness checks, indicate that robust optimization is clearly superior. Hence, our conclusions are not consistent with the various criticisms against robust optimization such as that this numerical technique is equivalent to a Bayesian shrinkage estimator and, therefore, offers no additional marginal value, see Scherer (2007) for instance. Despite the big changes in market conditions during the data period (see Figure 3.1), the robust optimization asset allocation is remarkably stable (47% equities, 22% bonds, 10% property, 17% alternatives and 4% cash), and is close to being a fixed-mix strategy. These asset proportions and performance measure rankings are robust to variations in the estimation period, the out-of-sample period, the factors used in forecasting asset returns, the assets included in the ALM, the asset bounds, and the value of ω . We conclude that robust optimization is a promising technique for solving pension scheme ALMs.

Appendix 3.A – Projected Contribution Rate

The projected contribution rate (CR) is given by:-

$$CR = SCR + kAL_A(1-FR)/(N_A S) \quad (A.3.1)$$

Where

$$k = \frac{1}{\sum_{z=0}^{M-1} (1+d)^{-z}} \quad (A.3.2)$$

$$1+d = \frac{1+d'}{1+e} \quad (A.3.3)$$

$$SCR = \frac{AL_A}{PY \times N_A \times S \times a_{\overline{1}|}} + AE \quad (A.3.4)$$

d' is the discount rate for liabilities

$a_{\overline{1}|}$ is an annuity to give the present value of earnings by the member over the next year

AE is the administrative expenses of the scheme, expressed as a proportion of the current salaries of the active members.

M is the life in years of a compound interest rate annuity - the spread period

FR denotes the funding ratio

AL_A is the actuarial liability for the active members of the scheme (see equation D.3,1),

PY is the average member's past years of service as at the valuation date,

S is the average member's annual salary at the valuation date,

e is the forecast nominal rate of salary growth per annum between the valuation date and retirement,

N_A is the current number of active members of the scheme.

Appendix 3.B - Uncertainty Sets

Given the following factor model for asset and liability returns over a single period:-

$$\tilde{\mathbf{r}}_{A,L} = \tilde{\boldsymbol{\mu}}_{A,L} + \tilde{\mathbf{V}}^T \mathbf{f} + \tilde{\boldsymbol{\varepsilon}}_{A,L} \quad (B.3.1)$$

the uncertainty sets for the column vector of the random asset and liability mean returns ($\tilde{\boldsymbol{\mu}}_{A,L}$), the matrix of uncertain factor coefficients ($\tilde{\mathbf{V}}$) and the diagonal covariance matrix of the uncertain disturbances ($\tilde{\mathbf{D}}$) are described as follows:-

$$S_{mean} = \left\{ \tilde{\boldsymbol{\mu}}_{A,L} : \tilde{\boldsymbol{\mu}}_{A,L} = \boldsymbol{\mu}_{A,L,0} + \tilde{\boldsymbol{\xi}}, \left| \tilde{\xi}_i \right| \leq \gamma_i(\omega), i=1, \dots, n_A + n_L \right\} \quad (\text{B.3.2})$$

$$S_v = \left\{ \tilde{\mathbf{V}} : \tilde{\mathbf{V}} = \mathbf{V}_0 + \tilde{\mathbf{W}}, \left\| \tilde{\mathbf{W}}_i \right\|_g \leq \rho_i(\omega), i=1, \dots, n_A + n_L \right\} \quad (\text{B.3.3})$$

$$S_d = \left\{ \tilde{\mathbf{D}} : \tilde{\mathbf{D}} = \text{diag}(\tilde{\mathbf{d}}), 0 \leq \tilde{d}_i \leq d_{upper,i}, i=1, \dots, n_A + n_L \right\} \quad (\text{B.3.4})$$

where $\tilde{\mathbf{W}}_i$ is the i^{th} column of $\tilde{\mathbf{W}}$ and $\left\| \mathbf{w} \right\|_g = \sqrt{\mathbf{w}^T \mathbf{G} \mathbf{w}}$ defines the elliptic norm of a column vector \mathbf{w} with respect to a symmetric and positive definite matrix \mathbf{G} . Each component of $\tilde{\boldsymbol{\mu}}_{A,L}$ and $\tilde{\mathbf{d}}$ is assumed to lie within a certain interval, while $\tilde{\mathbf{V}}$ belongs to an elliptical uncertainty set. The least square estimates ($\boldsymbol{\mu}_{A,L,0}, \mathbf{V}_0$) of ($\boldsymbol{\mu}_{A,L}, \mathbf{V}$) are computed with multivariate linear regression.

The parameterization of the parameters of the uncertainty sets is:-

$$\rho_i(\omega) = \sqrt{(m+1)c_{m+1}(\omega)s_i^2}, \quad i=1, \dots, n_A + n_L \quad (\text{B.3.5})$$

$$\gamma_i(\omega) = \sqrt{(m+1)(\mathbf{A}^T \mathbf{A})_{11}^{-1} c_{m+1}(\omega)s_i^2}, \quad i=1, \dots, n_A + n_L \quad (\text{B.3.6})$$

$$\mathbf{G} = \mathbf{B}\mathbf{B}^T - \frac{1}{p}(\mathbf{B}\mathbf{1})(\mathbf{B}\mathbf{1})^T \quad (\text{B.3.7})$$

$$d_{upper,i} = s_i^2 = \frac{\sum_{k=1}^p (r_{A,L,i}^k - \mu_{A,L,0,i} - V_{0,1,i} f_1^k - \dots - V_{0,m,i} f_m^k)^2}{p-m-1}, \quad i=1, \dots, n_A + n_L \quad (\text{B.3.8})$$

where m is the number of factors, $c_{m+1}(\omega)$ is the ω -critical value of the F -distribution with $m+1$ degrees of freedom in the numerator and $p-m-1$ degrees of freedom in the denominator. p is the number of sub-periods within the estimation period. In addition, $\mathbf{B} = [\mathbf{f}^1, \mathbf{f}^2, \dots, \mathbf{f}^p]$ (m rows and p columns), $\mathbf{A} = [\mathbf{1} \mathbf{B}^T]$ (p rows and $m+1$ columns) and $\mathbf{1}$ denotes the column vector of ones.

Finally, s_i^2 is an unbiased estimate of the variance of the residual return of the i^{th} asset/liability (see also section 5, equation 51 (page 16) in Goldfarb and Iyengar (2003)). It is obvious that $r_{A,L,i}^k$ refers to the i^{th} asset/liability in the k^{th} sub-period, $\mu_{A,L,0,i}$ to the least-squares estimate for the i^{th} asset/liability, $V_{0_m,i}$ to the least-squares estimate for the m^{th} factor and the i^{th} asset/liability. More technical details for the statistical justification of the parameters involved in the uncertainty sets can be found in Goldfarb and Iyengar (2003) (section 5, page 15 to 17).

Appendix 3.C – Conversion to a Second Order Cone Problem

In this Appendix, we derive the expressions for the worst-case mean return and worst-case variance and follow the same form as in Goldfarb and Iyengar (2003) (see also model (3.2) in section 3.2):-

$$\min_{\{\tilde{\mu}_{A,L} \in S_{mean}\}} \left[\tilde{\mu}_{A,L}^T \Phi_{A,L} \right] \quad (\text{worst case mean return})$$

$$\max_{\{\tilde{\mathbf{v}} \in S_v\}} \left[\Phi_{A,L}^T \tilde{\mathbf{V}}^T \mathbf{F} \tilde{\mathbf{V}} \Phi_{A,L} \right] \quad (\text{worst case variance 1})$$

$$\max_{\{\tilde{\mathbf{D}} \in S_d\}} \left[\Phi_{A,L}^T \tilde{\mathbf{D}} \Phi_{A,L} \right] \quad (\text{worst case variance 2})$$

and show how the maximin problem described in section 3.2 (model 3.2) can be formulated as a computationally tractable and easily solved second order cone program (SOCP).

Worst-Case Mean Return: The mean return is expressed in terms of the column vector of asset proportions (Φ_A) as follows ($\mathbf{1}^T \Phi_A = 1$, Φ_L is fixed):-

$$\tilde{\mu}_{A,L}^T \Phi_{A,L} = \tilde{\mu}_A^T \Phi_A + \tilde{\mu}_L^T \Phi_L = \left[\tilde{\mu}_A^T + \tilde{\mu}_L^T \Phi_L \mathbf{1}^T \right] \Phi_A \quad (\text{C.3.1})$$

since $\tilde{\mu}_{A,L}$ is a joint column vector of $\tilde{\mu}_A$ and $\tilde{\mu}_L$ as well as $\mu_{A,L,0}$ is a joint column vector of $\mu_{A,0}$ and $\mu_{L,0}$ (see least square estimates in Appendix 3.B). Since the uncertainty in $\tilde{\mu}_{A,L}$ is specified by (B.3.2) and $\Phi_A \geq \mathbf{0}$ and $\Phi_L < \mathbf{0}$, it can be easily

shown using some algebra that the worst-case mean return is given by the following expression:-

$$\min_{\{\tilde{\boldsymbol{\mu}}_{A,L} \in \mathcal{S}_{mean}\}} \left[\tilde{\boldsymbol{\mu}}_{A,L}^T \boldsymbol{\Phi}_{A,L} \right] = \left[\boldsymbol{\mu}_{A,0}^T - \boldsymbol{\gamma}_A^T(\omega) + (\boldsymbol{\mu}_{L,0}^T + \boldsymbol{\gamma}_L^T(\omega)) \boldsymbol{\Phi}_L \mathbf{1}^T \right] \boldsymbol{\Phi}_A, \quad (\text{C.3.2})$$

where $\gamma_{A,i}(\omega) = \gamma_i(\omega)$, for $i = 1, \dots, n_A$ and $\gamma_{L,i-n_A}(\omega) = \gamma_i(\omega)$, for $i = n_A + 1, \dots, n_A + n_L$ (see also Appendix 3.B for $\gamma(\omega)$).

Worst-Case Variance 1: Using some matrix algebra and the fact that $\mathbf{1}^T \boldsymbol{\Phi}_A = 1$, the term $\boldsymbol{\Phi}_{A,L}^T \tilde{\mathbf{V}}^T \mathbf{F} \tilde{\mathbf{V}} \boldsymbol{\Phi}_{A,L}$ can be rewritten in terms of the vector of assets ($\boldsymbol{\Phi}_A$) as follows:-

$$\boldsymbol{\Phi}_{A,L}^T \tilde{\mathbf{V}}^T \mathbf{F} \tilde{\mathbf{V}} \boldsymbol{\Phi}_{A,L} = \boldsymbol{\Phi}_A^T (\tilde{\mathbf{V}}^{sub} + \tilde{\mathbf{V}}')^T \mathbf{F} (\tilde{\mathbf{V}}^{sub} + \tilde{\mathbf{V}}') \boldsymbol{\Phi}_A = \left\| (\tilde{\mathbf{V}}^{sub} + \tilde{\mathbf{V}}') \boldsymbol{\Phi}_A \right\|_f^2 \quad (\text{C.3.3})$$

where

$$\tilde{\mathbf{V}}_{(i,j)}^{sub} = \tilde{\mathbf{V}}_{(i,j)}, \text{ for } i = 1, \dots, m, \quad j = 1, \dots, n_A,$$

$$\tilde{\mathbf{V}}'_{(i,j)} = \sum_{k=1}^{n_L} \boldsymbol{\Phi}_{L,k} \tilde{\mathbf{V}}_{(i,n_A+k)}, \text{ for } i = 1, \dots, m, \quad j = 1, \dots, n_A.$$

Also,

$$\mathbf{V}_{0,(i,j)}^{sub} = \mathbf{V}_{0,(i,j)}, \text{ for } i = 1, \dots, m, \quad j = 1, \dots, n_A,$$

$$\mathbf{V}'_{0,(i,j)} = \sum_{k=1}^{n_L} \boldsymbol{\Phi}_{L,k} \mathbf{V}_{0,(i,n_A+k)}, \text{ for } i = 1, \dots, m, \quad j = 1, \dots, n_A,$$

$$\tilde{\mathbf{W}}_{(i,j)}^{sub} = \tilde{\mathbf{W}}_{(i,j)}, \text{ for } i = 1, \dots, m, \quad j = 1, \dots, n_A,$$

$$\tilde{\mathbf{W}}'_{(i,j)} = \sum_{k=1}^{n_L} \boldsymbol{\Phi}_{L,k} \tilde{\mathbf{W}}_{(i,n_A+k)}, \text{ for } i = 1, \dots, m, \quad j = 1, \dots, n_A,$$

(see also the uncertainty structure for $\tilde{\mathbf{V}}$ in (B.3.3)).

Since the uncertainty in $\tilde{\mathbf{V}}$ is specified by (B.3.3), it can be shown following the same process as in Goldfarb and Iyengar (2003) (see page 7 to 8 and equations 16 to 20 in Goldfarb and Iyengar (2003)) that the worst-case variance 1 is less than ν , if and only if:-

$$\max_{\{\mathbf{y}; \|\mathbf{y}\|_g \leq r\}} \|\mathbf{y}_0 + \mathbf{y}\|_f^2 \leq \nu \quad (\text{C.3.4})$$

with $\mathbf{y}_0 = (\mathbf{V}_0^{\text{sub}} + \mathbf{V}_0') \Phi_A$, $\mathbf{y} = (\tilde{\mathbf{W}}^{\text{sub}} + \tilde{\mathbf{W}}') \Phi_A$, $r = \sum_{i=1}^{n_A} \Phi_{A,i} \left[\rho_i(\omega) - \sum_{j=1}^{n_L} \Phi_{L,j} \rho_{n_A+j}(\omega) \right]$.

Constraint (C.3.4) can be reformulated as a set of linear equalities, linear inequalities and second order cone constraints (SOCC) in exactly the same manner as in LEMMA 1 in Goldfarb and Iyengar (2003) (page 8, part i). One can also see the corresponding DEFINITION 1 that is motivated by LEMMA 1 on page 10 at the bottom in Goldfarb and Iyengar (2003). We represent this collection of linear equalities, linear inequalities and second order cone constraints as follows:-

$$\left(\sum_{i=1}^{n_A} \Phi_{A,i} \left[\rho_i(\omega) - \sum_{j=1}^{n_L} \Phi_{L,j} \rho_{n_A+j}(\omega) \right]; \nu; \Phi_A \right) \in \text{Def}(\mathbf{V}_0^{\text{sub}} + \mathbf{V}_0', \mathbf{F}, \mathbf{G}), \quad (\text{C.3.5})$$

defining $\text{Def}(\mathbf{V}_0^{\text{sub}} + \mathbf{V}_0', \mathbf{F}, \mathbf{G})$ to be the set $\left(\sum_{i=1}^{n_A} \Phi_{A,i} \left[\rho_i(\omega) - \sum_{j=1}^{n_L} \Phi_{L,j} \rho_{n_A+j}(\omega) \right]; \nu; \Phi_A \right)$

that satisfy the following:-

There exist $\sigma, \tau \geq 0$ and a column vector \mathbf{t} with m positive elements that satisfy:-

$$\tau + \mathbf{1}^T \mathbf{t} \leq \nu$$

$$\sigma \leq \frac{1}{\lambda_{\max}(\mathbf{H})}$$

$$\left\| \begin{bmatrix} 2r \\ \sigma - \tau \end{bmatrix} \right\| \leq \sigma + \tau$$

$$\left\| \begin{bmatrix} 2\mathbf{w}_i \\ 1 - \sigma\lambda_i - \mathbf{t}_i \end{bmatrix} \right\| \leq 1 - \sigma\lambda_i + \mathbf{t}_i, \quad i=1, \dots, m,$$

where $r = \sum_{i=1}^{n_A} \Phi_{A,i} \left[\rho_i(\omega) - \sum_{j=1}^{n_L} \Phi_{L,j} \rho_{n_A+j}(\omega) \right]$, $\mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^T$ is the spectral decomposition of

$$\mathbf{H} = \mathbf{G}^{-\frac{1}{2}} \mathbf{F} \mathbf{G}^{-\frac{1}{2}}, \quad \mathbf{\Lambda} = \text{diag}(\lambda_i) \text{ and } \mathbf{w} = \mathbf{Q}^T \mathbf{H}^{\frac{1}{2}} \mathbf{G}^{\frac{1}{2}} (\mathbf{V}_0^{\text{sub}} + \mathbf{V}_0') \Phi_A.$$

Worst-Case Variance 2: The term $\Phi_{A,L}^T \tilde{\mathbf{D}} \Phi_{A,L}$ is expressed in terms of the vector of asset weights (Φ_A), using some matrix algebra and the fact that $\mathbf{1}^T \Phi_A = 1$, as follows:-

$$\Phi_{A,L}^T \tilde{\mathbf{D}} \Phi_{A,L} = \Phi_A^T (\tilde{\mathbf{D}}^{\text{sub}} + \tilde{\mathbf{D}}') \Phi_A \quad (\text{C.3.6})$$

where $\tilde{\mathbf{D}}^{\text{sub}} = \text{diag}(\tilde{\mathbf{d}}^{\text{sub}})$, $\tilde{d}_i^{\text{sub}} = \tilde{d}_i$, for $i = 1, \dots, n_A$,

$$\tilde{D}'_{(i,j)} = \sum_{k=1}^{n_L} \tilde{d}_{n_A+k} \Phi_{L,k}^2, \text{ for } i = 1, \dots, n_A, \quad j = 1, \dots, n_A.$$

Since the uncertainty in $\tilde{\mathbf{D}}$ is specified by (B.3.4), it can be easily shown using some matrix algebra that:-

$$\begin{aligned} \max_{\{\tilde{\mathbf{d}} \in \mathcal{S}_d\}} \left[\Phi_{A,L}^T \tilde{\mathbf{D}} \Phi_{A,L} \right] &= \max_{\{\tilde{\mathbf{d}} \in \mathcal{S}_d\}} \left[\Phi_A^T (\tilde{\mathbf{D}}^{\text{sub}} + \tilde{\mathbf{D}}') \Phi_A \right] \\ &= \Phi_A^T (\mathbf{D}_{\text{upper}}^{\text{sub}} + \mathbf{D}'_{\text{upper}}) \Phi_A \\ &= \left\| \left(\mathbf{D}_{\text{upper}}^{\text{sub}} + \mathbf{D}'_{\text{upper}} \right)^{\frac{1}{2}} \Phi_A \right\|^2 \end{aligned} \quad (\text{C.3.7})$$

since $\Phi_A^T \tilde{\mathbf{D}}^{\text{sub}} \Phi_A \leq \Phi_A^T \mathbf{D}_{\text{upper}}^{\text{sub}} \Phi_A$ and $\Phi_A^T \tilde{\mathbf{D}}' \Phi_A \leq \Phi_A^T \mathbf{D}'_{\text{upper}} \Phi_A$ ($\tilde{d}_i \leq d_{\text{upper},i}$),

where $\mathbf{D}_{\text{upper}}^{\text{sub}} = \text{diag}(\mathbf{d}_{\text{upper}}^{\text{sub}})$, $d_{\text{upper},i}^{\text{sub}} = d_{\text{upper},i}$, for $i = 1, \dots, n_A$,

$$\mathbf{D}'_{\text{upper},(i,j)} = \sum_{k=1}^{n_L} d_{\text{upper},n_A+k} \Phi_{L,k}^2, \text{ for } i = 1, \dots, n_A, \quad j = 1, \dots, n_A.$$

The constraint $\max_{\{\tilde{\mathbf{d}} \in \mathcal{S}_d\}} \left[\Phi_{A,L}^T \tilde{\mathbf{D}} \Phi_{A,L} \right] = \Phi_A^T (\mathbf{D}_{\text{upper}}^{\text{sub}} + \mathbf{D}'_{\text{upper}}) \Phi_A \leq \delta$ can be rewritten as a

second order cone constraint (SOCC) as follows:-

$$\left\| \begin{bmatrix} 2 \left(\mathbf{D}_{\text{upper}}^{\text{sub}} + \mathbf{D}'_{\text{upper}} \right)^{\frac{1}{2}} \Phi_A \\ 1 - \delta \end{bmatrix} \right\| \leq 1 + \delta \quad (\text{C.3.8})$$

since $\mathbf{z}^T \mathbf{z} \leq xy \Leftrightarrow 4\mathbf{z}^T \mathbf{z} + (x-y)^2 \leq (x+y)^2 \Leftrightarrow \left\| \begin{bmatrix} 2\mathbf{z} \\ x-y \end{bmatrix} \right\| \leq x+y, x, y \geq 0$ which is straightforward from the definition of the Euclidian norm and from the fact that the matrix $\mathbf{D}_{upper}^{sub} + \mathbf{D}'_{upper}$ is symmetric.

Final Formulation: The linear constraints with certain parameters for non-negativity and asset lower/upper bounds as in model (3.2) in section 3.2 are as follows:-

$$\Phi_{A,i} \geq 0, \quad \forall i = 1, \dots, n_A \quad (\text{C.3.9})$$

$$-\sum_{i \in classX} \Phi_{A,i} + \theta_{max}^{classX} \mathbf{1}^T \Phi_A \geq 0, \quad \forall classX \quad (\text{C.3.10})$$

$$\sum_{i \in classX} \Phi_{A,i} - \theta_{min}^{classX} \mathbf{1}^T \Phi_A \geq 0, \quad \forall classX \quad (\text{C.3.11})$$

Bringing together (C.3.2), (C.3.5), (C.3.8), (C.3.9), (C.3.10), (C.3.11), model (3.2) in section 3.2 is given by the following easily solved and computationally tractable SOCP:-

$$\begin{aligned} & \text{minimize} \quad v + \delta \\ & \text{s.t.} \quad \left\| \begin{bmatrix} 2(\mathbf{D}_{upper}^{sub} + \mathbf{D}'_{upper})^{\frac{1}{2}} \Phi_A \\ 1 - \delta \end{bmatrix} \right\| \leq 1 + \delta \\ & \quad \left(\sum_{i=1}^{n_A} \Phi_{A,i} \left[\rho_i(\omega) - \sum_{j=1}^{n_L} \Phi_{L,j} \rho_{n_A+j}(\omega) \right]; v; \Phi_A \right) \in \text{Def}(\mathbf{V}_0^{sub} + \mathbf{V}_0', \mathbf{F}, \mathbf{G}) \\ & \quad \left[\boldsymbol{\mu}_{0,A}^T - \boldsymbol{\gamma}_A^T(\omega) + (\boldsymbol{\mu}_{0,L}^T + \boldsymbol{\gamma}_L^T(\omega)) \Phi_L \mathbf{1}^T \right] \Phi_A \geq 0.001 \\ & \quad \Phi_A \geq \mathbf{0}, \\ & \quad -(\mathbf{A}_{classX}^T + \theta_{max}^{classX} \mathbf{1}^T) \Phi_A \geq 0, \quad \forall classX \\ & \quad (\mathbf{A}_{classX}^T - \theta_{min}^{classX} \mathbf{1}^T) \Phi_A \geq 0, \quad \forall classX \end{aligned} \quad (\text{C.3.12})$$

Note that the last 2 sets of linear constraints in (C.3.12) are the same as (C.3.10) and (C.3.11).

\mathbf{A}_{classX} is a column vector of ones and zeros indicating which assets participate in the broad asset classes $\left(\sum_{i \in classX} \Phi_{A,i} = \mathbf{A}_{classX}^T \Phi_A \right)$. Given the expressions (C.3.1), (C.3.3) and (C.3.6), the nominator as well as the denominator in model (3.2) (in section 3.2)

are homogeneous in terms of the column vector of asset weights (Φ_A) and thus (Φ_A) is no longer normalized in (C.3.12). This is a key issue for the SOCP in (C.3.12) since the fact that (Φ_A) is no longer normalized (the normalization condition $\mathbf{1}^T \Phi_A = 1$ is dropped) increases the computational tractability of the problem. Homogenization is maintained even if we add extra linear affine constraints in terms of Φ_A .

Appendix 3.D - Actuarial Liability Models

The actuarial liability for active members is:-

$$AL_A = N_A \times \left(\frac{PY \times S}{A} \right) \times \left\{ \frac{(1+e)}{(1+h)} \right\}^{RA-G} \times \left\{ \left[1 - \left(\frac{1+h}{1+pl} \right)^{-W} \right] / \left(\frac{1+h}{1+pl} - 1 \right) \right\} \quad (D.3.1)$$

where

A is the accrual rate,

h is the nominal discount rate between now and retirement,

RA is the average member's forecast retirement age,

G is the average age of the member at the valuation date,

W is the life expectancy of members at retirement,

pl is the rate of growth of the price level,

AL_A , P , e , N_A and S are defined in Appendix A

The final term in equation (D.3.1) is the capital sum required at time RA to purchase an index-linked annuity of £1 per year.

A simple model for the computation of the actuarial liability for pensioners is:-

$$AL_P = N_p \times PEN \times \left\{ \left[1 - \left(\frac{1+h}{1+pl} \right)^{-q} \right] / \left(\frac{1+h}{1+pl} - 1 \right) \right\} \quad (D.3.2)$$

where

AL_p is the actuarial liability for pensioners,

N_p is the current number of pensioners,

PEN is the average current pension;

The final term is the capital sum required now to purchase an index-linked annuity of £1 per year for the life expectancy, q , of pensioners. Adjustments to this simple model are required for dependents' pensions, death lump sum, etc.

A similar expression for the liability of deferred pensioners is:-

$$AL_D = N_D \times \left(\frac{P_D \times S_D}{A} \right) \times \left\{ \frac{(1+pl)}{(1+h)} \right\}^{RA-G} \times \left\{ \left[1 - \left(\frac{1+h}{1+pl} \right)^{-W} \right] / \left(\frac{1+h}{1+pl} - 1 \right) \right\} \quad (D.3.3)$$

where

AL_D is the actuarial liability for the deferred pensioners of the scheme,

N_D is the current number of deferred pensioners of the scheme,

S_D is the average deferred pensioners' leaving salary, compounded forwards to the valuation date at the inflation rate (p), and

P_D is the average deferred pensioner's past years of service as at the valuation date.

The total actuarial liability (AL_T) is:-

$$AL_T = AL_A + AL_p + AL_D \quad (D.3.4)$$

which is the sum of the actuarial liabilities for every active member, pensioner and deferred pensioner.

When computing monthly changes in the actuarial liability, the terms A , N_A , N_D , N_p , S , S_D and PEN are constant within each triennial period, while the initial values of AL_A , AL_D and AL_p come from the actuarial valuations. Therefore the values of A , N_A , N_D , N_p , S , S_D and PEN are not required.

Chapter 4

4 Socially Responsible Investment Portfolios: Does the Optimization Process Matter?¹⁰

4.1 Introduction

Corporate Social Responsibility (CSR) and Corporate Social Performance (CSP)¹¹ have become crucially important concepts in the modern business world. Broadly defined as “a management concept whereby companies integrate social and environmental concerns in their business operations and interactions with their stakeholders”¹², it has gained traction over the past 20 years. A growing number of stakeholders have increased societal demands that corporations perform well financially, while operating in a responsible and ethical manner.

This trend is noticeable in the latest surveys. Grant Thornton’s International Business Report¹³ in 2014 surveyed 2,500 firms in 34 countries and showed that more and more businesses are adopting socially and environmentally sustainable practices and initiatives. These range from charitable donations and active participation in local community causes to improving energy efficiency and applying more effective waste management. The majority of these firms cite client/consumer demand as one of the dominant driving forces behind their decision to move to more sustainable business formats. Similarly, the Nielsen Global Survey on Corporate Social Responsibility

¹⁰ The content of this Chapter was presented at the 5th International Conference of the Financial Engineering and Banking Society (Nantes, France), 4th European Business Research Conference (London, UK) and the ICMA Centre Internal Research Seminar (Reading, UK).

¹¹ The two terms have been used interchangeably in relevant empirical research. In this paper, we use CSP.

¹² United Nations Industrial Development Organization, retrieved October 2014 from <http://www.unido.org/en/what-we-do/trade/csr/what-is-csr.html>

¹³ For additional information, the interested reader is directed at <http://www.grant-thornton.co.uk/en/Media-Centre/News/2014/Global-survey-finds-good-CSR-makes-good-business-sense-British-businesses-reacting-to-stakeholders-demands/>, retrieved October 2014.

(2013) used a sample of 29,000 participants from 58 countries and found that at least half of global consumers are willing to “walk the talk” and pay a premium for goods and services produced by socially responsible firms.

In line with these developments, demand for CSP in financial markets, also known as Socially Responsible Investing (SRI)¹⁴, has also been growing rapidly. According to the Global Sustainable Investment Review 2012, which is a product of the collaboration of a variety of organizations and sustainable investment forums across the world, approximately US\$13.6 trillion of assets under professional management incorporate environmental, social or governance considerations into the investment selection process. This represents more than 20% of the total assets under professional management in the areas covered in the report, and includes positive and negative screening, shareholder activism strategies, norm-based screening, best-in-class approaches and other forms of SRI. While the criteria for an investment to be deemed socially responsible are not strict, it is undeniable that SRI is nowadays a large and expanding segment of the financial markets.

As a result, a significant amount of scholarly research has been dedicated to the investigation of the nature of the relationship between CSP and firm financial performance. Meta-studies focusing on this area (Margolis et al., 2009; Orlitzky et al., 2003) demonstrate both its depth and breadth. Using data from hundreds of relevant papers going as far back as 1972, these studies provide evidence of an overall positive link between the two concepts. At the portfolio level of analysis, comparing SRI funds and indices with “conventional” funds and indices with

¹⁴Also referred to as Environmental, Social, and Governance Investing, Sustainable Investing and Impact Investing, though there are some conceptual differences between these terms.

otherwise similar characteristics commonly points to statistically indistinguishable performance (Renneboog et al., 2008; Schroder, 2007; Statman, 2000; Statman, 2006), although there are indications of SRI outperformance in certain contexts (Derwall and Koedijk, 2009; Kempf and Osthoff, 2007).

Despite the size of this literature, a very small number of studies have investigated optimal ways to construct SRI portfolios, either in the sense of the screening criteria used to narrow the investment universe, or the optimization process employed to determine the asset proportions. Barnett and Salomon (2006) is one of the few papers that focuses on the effects of screening intensity in SRI funds, and provides evidence of a U-shaped relationship between the number of social/environmental screens used and fund performance. Similarly, there are only a handful of papers (Ballesterio et al. 2012; Drut, 2012; Utz et al. 2014) which explore the portfolio optimization frameworks used in SRI. Although such studies contribute significantly to this underdeveloped part of the literature, they are limited in that they do not go far beyond the Markowitz (1952) mean-variance optimisation framework. They simply extend it by adding SRI preferences as an additional constraint, or incorporate them in the objective function. Although Markowitz optimization is the basis for the vast majority of modern portfolio optimization methods, it suffers from significant estimation risk (Green and Hlofield, 1992; DeMiguel et al., 2009a), and this leads to solutions that are very sensitive to the inputs, and the generation of unstable and poorly diversified portfolios.

This omission of estimation risk is unfortunate as, compared to conventional portfolios, SRI portfolios are characterised by a greater level of uncertainty in their

inputs. This is due to the inherent complexity in measuring CSP, and the largely discretionary nature of CSP reporting. So, within the SRI framework it is important to consider alternative optimization techniques, and to investigate the extent to which they lead to the construction of substantially different portfolios in terms of risk, risk-return trade-off, diversification and the stability of the constituent assets. Our study contributes to the literature by applying six different optimization methods¹⁵ to the same SRI-screened investment universe, and comparing their out-of-sample performance as captured by 14 different metrics indicative of various important portfolio characteristics. In this way our study is the first to answer the question of whether the portfolio optimization process matters in SRI, and to further contextualise this answer by indicating which methods tend to lead to better results. The potential practical usefulness of this study is also significant. If different optimization techniques lead to different SRI portfolio performance, this would indicate that, apart from the social and environmental screening criteria, investors and fund managers also need to carefully consider the choice of asset allocation method. Financially savvy investment techniques and moral objectives need not be mutually exclusive. In fact, recognition of which optimization methods yield better results within SRI may enhance the growth of the SRI sector, leading to a larger share within the financial markets, and a lower cost of capital for the CSP champions. This, in turn, will strengthen the pressure from the financial markets for the adoption of sustainable practices by companies.

¹⁵ Although all of these techniques are frequently referenced as optimization methods, the broader term “portfolio construction” is more accurate in some cases. We follow the norm and hereafter refer to the entire set of alternative methods as optimization methods.

The remainder of the paper is structured as follows: Section 4.2 reviews previous studies of the application of optimization techniques¹⁶ to forming SRI portfolios, and discusses the alternative portfolio construction methods which we compare and contrast. Section 4.3 contains details of the CSP database we use and the portfolio evaluation methods we employ. Section 4.4 presents our empirical results, and Section 4.5 concludes.

4.2 Related literature and motivation of the study

The vast majority of scholarly research dedicated to SRI portfolios focuses on identifying the ways in which they are different from (or similar to) conventional investments in terms of their constituents, the performance they achieve and the risks they bear. A surprisingly small number of academic papers have investigated ways in which the portfolio construction process, be it through the use of alternative security selection criteria or different optimization techniques, can lead to the generation of better performing, more efficient and stable SRI portfolios.

Barnett and Salomon (2006) shed some light on the optimal number and type of screening criteria used by SRI funds. Their findings depict a non-linear link between screening intensity and fund/portfolio performance. SRI portfolios where just a few or many social screens are employed outperform portfolios with an intermediate number of such screens. In addition, the authors also investigate the financial contribution of particular types of screening, and they find that community relations screening increases financial performance, whereas environmental and labour

¹⁶We make use of the terms “technique”, “approach”, “model”, “process” and “method” interchangeably in this regard.

relations screens tend to decrease financial performance. Capelle-Blancard and Monjon (2014) on the other hand, investigate French SRI funds and find that sectoral screens (i.e. avoiding investing in the so-called 'sin' stocks) decrease financial performance, while other types of CSP screen do not have a noticeable financial impact on fund performance.

In an effort to investigate the common claim that SRI funds are in reality nothing more than conventional funds in disguise, Kempf and Osthoff (2008) compare the sustainability characteristics of the portfolio holdings of SRI funds with those of conventional funds. Their investigation focuses on US equity funds and demonstrates that the social and environmental ratings of their constituent stocks are indeed higher than those of otherwise similar conventional funds. Thus any outperformance of these funds can be attributed to the higher CSP levels of the securities they include.

Complimentary to this line of academic research is the small, and fairly new, literature dedicated to the use of alternative optimisation approaches to construct well-diversified and efficient SRI portfolios. Hallerbach et al. (2004) were the first to point out that the SRI literature lacked suggestions for combining the social characteristics of risky assets and the standard financial information in the portfolio optimization process. They presented an interactive multiple goal programming approach for managing an investment portfolio where the decision criteria include social effects.

An alternative approach was suggested by Drut (2012), who investigated whether adding restrictions regarding CSP when deriving optimal investment strategies leads

to portfolios that underperform otherwise similar conventional investments. He uses the classical mean-variance model of Markowitz (1952), and imposed an extra constraint for the CSP rating. Drut concludes that the effects of adding a CSP constraint depend “on the link between the returns and the responsible ratings and on the strength of the constraint” (p.28). Hence, including additional CSP considerations may not necessarily lead to suboptimal portfolio performance.

Ballesterio et al. (2012) used goal programming within the framework of classical Markowitz mean-variance optimization to allow investors to take account of ethical issues, in addition to the standard financial information. They considered both “green” and “conventional” assets, and used a two-dimensional objective function (financial and environmental). Their numerical analysis revealed that substantial green investment is generally outperformed by modest green investment, a rare result within the core empirical literature, and hence discourage investors from investing a large part their portfolio in green assets.

The most recent relevant work in the area comes from Utz et al. (2014) who extended the Markowitz model by adding a social responsibility objective, in addition to the portfolio return and variance, causing the traditional efficient frontier to become a three dimensional surface. When applying their framework to both conventional and SRI mutual funds they did not find any evidence that social responsibility, used as a third criterion and measured by CSP scores, plays an important role in the financial outcome of asset allocation. The authors did, however, find a modestly lower volatility associated with socially responsible compared to conventional funds.

In short, the studies of the optimisation techniques for SRI portfolios tend to focus, not on the effectiveness of the techniques themselves in creating well-performing, stable and diversified portfolios, but rather on providing generic frameworks that integrate financial with social and environmental considerations. They investigate whether there is a financial cost to including these additional CSP considerations, and whether SRI portfolios tend to outperform or underperform otherwise similar conventional portfolios. Contrary to the above, our study explores whether different methodologies which are applied in the generic professional investing arena lead to the construction of SRI portfolios with superior characteristics.

The current study attempts to answer questions of the following type. Are some approaches to portfolio selection superior in creating the less volatile SRI portfolios sought by particularly risk averse investors such as pension schemes and insurance funds? Which asset allocation approaches lead to SRI portfolios which remain reasonably stable in terms of their constituent assets, thereby minimizing transaction costs? Does optimizing a different measure of risk and returns change the results? Or is performance of the various optimization methods broadly similar when forming SRI portfolios?

A common denominator of previous studies is the use of the Markowitz framework (or extensions of it) in the formation of SRI portfolios, and this has several important drawbacks. The application of Markowitz mean-variance optimisation requires the estimation of the means, variances and covariances of the asset returns for the investment universe under consideration. In practice this means that, if the sample means and covariances are subject to estimation error, optimal portfolios

constructed via Markowitz optimization can be unstable, and characterised by poor diversification and out-of-sample performance. This phenomenon has been well-substantiated in the portfolio selection literature. For instance, Michaud (1999) states that, although Markowitz theory provides a convenient framework for portfolio optimization, in practice it is an “error-prone process” that often leads to the construction of portfolios with problematic properties. Broadie (1993) has also studied the effects of estimation risk on the construction of the Markowitz efficient frontier, while a more comprehensive review of the influence of estimation errors on portfolio selection can be found in Ziemba and Mulvey (1998). This is why it is important to study portfolios constructed using approaches that allow for estimation risk, and to compare their characteristics and performance.

The above argument applies to both conventional and socially responsible investing, but a strong case can be made that estimation errors in the input parameters are a more important issue when constructing SRI portfolios.

There is a plethora of studies showing that CSP influences both asset returns (Brammer et al, 2006; Galema et al., 2008; Edmans, 2011; Hillman and Keim, 2001; Von Arx and Ziegler, 2014), and financial risk (Bouslah et al., 2013; Lee and Faff, 2009; Oikonomou et al., 2012). Both qualitative literature reviews (Margolis and Walsh 2003) and statistical meta-analyses (Margolis et al., 2009; Orlitzky et al., 2003) broadly substantiate this conclusion. Hence, CSP contributes to the estimation risk of the input parameters used in constructing portfolios. However CSP scores are subject to considerable estimation error, and this is for four reasons.

First, CSP is a concept which has proved very hard to define. Many definitions have been vague or too inclusive. In the words of Votaw (1973) 'the term is a brilliant one; it means something, but not always the same thing, to everybody'. The work of Carroll (1991) has been influential in defining CSP, and makes reference to a variety of tiers or levels of firm responsibilities (economic, legal, ethical and philanthropic) that taken together constitute CSP. The European Commission on the other hand simply refers to CSP as a concept whereby "companies are taking responsibility for their impact on society"¹⁷.

Second, CSP is characterised by a large amount of variability and heterogeneity in its various dimensions making its accurate measurement a problematic task (Abbott and Monsen, 1979; Griffin and Mahon, 1997). CSP may be related to, inter alia, issues involving a firm's treatment of the natural environment, employee welfare, philanthropic activity, engagement with local societies and interaction with controversial industries.

Third, subjective judgements are involved, not only in assessing a company's performance in all of the above, but in measuring the relative importance of each CSP dimension for a firm belonging to a particular industry and operating within a specific socio-cultural environment. For example, it could be judged that oil and energy companies should put more emphasis on the environmental aspects of their CSP due to their significant footprint, whereas firms in the financial services sector should be more concerned about product quality and ethical business practices. Hence, the quantification of CSP is a complex task which requires the collection and

¹⁷ http://europa.eu/rapid/press-release_MEMO-11-730_en.htm

assessment of information both internal and external to the firm by sophisticated, independent assessors such as MSCI, Sustainalytics, Oekom and other agencies producing social ratings for companies.

Finally, CSP disclosures remain a discretionary part of corporate reporting in most countries (Orlitzky, 2013). Due to this, voluntary CSP reports are not subject to the same government oversight and regulatory scrutiny which applies to compulsory company reporting. Hence, erroneous or misleading CSP reporting may not lead to legal and financial sanctions, making such disclosures more susceptible to unintentional errors and deliberate manipulation by opportunistic firm managers (Edwards, 2008). This further complicates the issue of the accurate measurement of CSP.

Overall, whether it is due to the inherent definitional complexity and heterogeneity of CSP, or the subjectivity and misinformation surrounding CSP issues and a lack of regulatory scrutiny, there is an additional degree of ambiguity when considering CSP as criterion in portfolio creation. Orlitzky (2013) even goes as far as suggesting that CSP may increase the overall level of “noise trading”¹⁸, with the noise associated with CSP scores leading to additional noise in the pricing of assets in SRI portfolios. High (or low) reported CSP scores are likely to be subject to greater estimation error than more average scores. If CSP scores are priced positively by financial markets, such over (under) estimation of the CSP scores biases the company’s expected returns upwards (downwards), and may also bias its estimated variance downwards (upwards). Therefore companies with high (or low) CSP scores have higher

¹⁸Noise trades are based on false signals, not on the underlying economic fundamentals.

estimation risk in their returns and risk. Because SRI portfolios are characterised by a greater degree of estimation errors in the input parameters, i.e. return and risk, an optimisation method which is less sensitive to these values should be employed. However, the SRI literature is lacking in providing meaningful suggestions, and this is the gap our study attempts to fill.

There are various alternative portfolio optimisation frameworks which could be used for the construction of SRI portfolios with desirable properties. We apply six of these frameworks, all of which are widely known and commonly considered by the professional portfolio management community. Each framework has a different underlying rationale, and may lead to the construction of portfolios with different characteristics. Though dozens of different optimization techniques are available, we believe the six we use are an appropriate representation of the broad alternative rationales behind asset allocation mechanisms. In addition, previous research has been conducted on each of them which allow us to compare and contrast the findings of this study (within the SRI framework) with those of general portfolio optimization studies. The first three of the approaches below (Markowitz, robust estimation, and Black-Litterman) are “classical”, quantitatively sophisticated models, whereas the other three (naïve diversification, risk parity, and reward-to-risk) are more recent approaches with a less solid mathematical basis, and draw largely on basic investing intuition. Below, we provide an explanation of the rationale behind these approaches and a broad outline of their implementation. The technical details of each framework are in Appendix 4.A.

i) Markowitz portfolio optimization

In spite of the various problems we have outlined, the Markowitz (1952) portfolio optimization technique is the forefather of the vast majority of modern portfolio construction methods, and usually serves as the basis for the comparison of the performance of different models. Markowitz was the first to formally recognize the importance of diversification, and to create a method whose principal premise is that only the first two moments (mean and variance) of the return distribution are important to investors. Hence the ultimate goal is to create a portfolio by optimizing the risk-return trade-off. As self-evident as this may seem for today's investment professionals, Markowitz's work provided the foundation on which the mathematical modelling of portfolio construction was established. Furthermore, there have recently been calls for a return to Markowitz's model of portfolio construction (Kaplan, 2014), with explicit risk and expected return assumptions, instead of the implicit assumptions made by many of the alternative methods.

ii) Robust estimation

A sophisticated set of portfolio construction practices, which has been used when considering "conventional" (i.e. non-SRI) assets, involves imposing norm constraints on the portfolio weights to obtain the desired characteristics (see, for instance, Ledoit and Wolf, 2003 and 2004; Fan et al., 2008). We elect to use a technique that falls within this category, and adopt a robust portfolio technique, inspired by Xing et al. (2014) among others, to construct superior portfolios in the presence of estimation risk which, as we noted above, is higher when creating SRI portfolios. This approach encourages the creation of sparse portfolios with relatively few active

positions and significantly reduced associated transaction costs. It is also particularly well suited to the preferences of the SRI investing community as it tends to generate cautious, low risk portfolios. It has been documented that long-term institutional investment is greater in companies with high CSP scores (Johnson and Greening, 1999; Cox et al., 2004). This demand arises principally from pension funds and life assurance companies, who are characterized by high levels of risk aversion, and who consider worst-case scenarios to ensure their investment decisions are guided by prudence and safety. Insurance companies in many countries must comply with prudential regulations, such as Solvency II for countries in the European Union, while defined benefit pension schemes must satisfy their regulators that they will meet their pensions promise. So both these large groups of institutional investor have a low tolerance for risk.

iii) Black-Litterman

The Black-Litterman (1992) asset allocation model is another approach commonly employed by a variety of financial institutions. It is particularly popular among active money managers “who believe they hold information superior to that of other market participants, but wish to update their beliefs using market prices” (Gofman and Manela, 2012). The main advantage of this model is that it allows the investor to combine the market equilibrium with the views of the investor. In the words of He and Litterman (1999), the intuition underlying this approach can be summed up as: “the user inputs any number of views, which are statements about the expected returns of arbitrary portfolios, and the model combines the views with equilibrium, producing both the set of expected returns of assets as well as the optimal portfolio

weights". In this way, optimal portfolios start from a set of "neutral" weights which are then tilted in the direction of investor views.

iv) Naïve diversification

The naive diversification approach is based on the very simple rule whereby $1/N$ of the investor's wealth is allocated to each of the N assets available in the investment universe being considered. In other words, it leads to the construction of an equally weighted portfolio of the available set, or screened subset, of assets. Unlike the mean–variance framework of portfolio optimization, it does not attempt to assign asset weights to optimize the risk-return trade-off. Instead, the most appealing feature of the naive approach lies in its simplicity, as it does not require the estimation of expected returns, covariances, or higher moments of asset returns. In addition, the previous literature provides evidence that the naive diversification ($1/N$) approach is not inferior to sample-based mean-variance models (Bloomfield, Leftwich, and Long, 1977), or even to most of the extensions of the Markowitz optimization framework (DeMiguel et al., 2009b). Therefore, it is considered a reasonable approach to portfolio formation.

v) Risk-parity portfolios

In recent years the risk-parity portfolio approach has attracted significant interest from academics and practitioners, and is widely applied by long term institutional investors such as pension funds, and insurance companies, as well as mutual funds (Anderson, Bianchi and Goldberg, 2012). In its simplest form, it leads to a portfolio of risky assets where the weights are anti-proportional to each asset's variance of returns (i.e. total risk). The emphasis the approach places on risk increased its

popularity in the post-crisis period, as models related to the Markowitz framework were accused of not providing effective risk controls when they were most needed. In addition, the risk-parity approach benefits from the fact that assets with high volatility usually earn a lower premium per volatility unit than those with lower volatility (Baker et al., 2011; and Frazzini and Pederson, 2014).

vi) Reward-to-risk timing portfolios

The reward-to-risk timing portfolio strategy has been proposed by Kirby and Ostdiek (2012). Its development was motivated by the finding that naive diversification portfolios tend to outperform mean-variance optimization approaches. Kirby and Ostdiek (2012) argue that the main reason behind this is the greater instability of the portfolios created by Markowitz-style methods. Hence, they created an alternative method which keeps the essential rationale of the importance of the risk-return trade-off intact, but leads to more stable portfolios with lower transaction costs. The reward-to-risk timing strategy allocates asset weights in proportion to the contribution of each asset's mean-variance ratio to the mean-variance ratio of the entire universe of assets.

Further techniques for deriving optimal portfolio strategies which might have been considered include: stochastic programming, e.g. Geyer and Ziemba (2008); dynamic programming, e.g. Rudolf and Ziemba (2004); and stochastic simulation, e.g. Boender (1997). However, they are computationally challenging, making them inappropriate for use in practice for the sizeable portfolios we consider. For instance, Platanakis and Sutcliffe (forthcoming) mention that the number of scenarios required by stochastic programming exceeds 24 billion for a portfolio with just 14

assets, four non-overlapping investment periods and five independent outcomes for each uncertain parameter per estimation period. With 100 assets this figure rises to 3.1554×10^{70} . As a result these techniques are not used in our study due to the computational load they would entail. Goal programming (Hallerbach et al., 2004; Ballesterro et al. 2012) includes SRI preferences in the objective function so that the investor optimizes some combination of both financial and social performance. Our study investigates the impact of financial optimization on an investment universe screened according to CSP criteria, and so goal programming lies beyond the scope of this study.

To summarise, we consider a variety of widely applied modern portfolio construction approaches with different points of emphases and supporting rationales, and conduct a horse race between them using a socially responsibly screened universe of stocks. The next section discusses the portfolio evaluation measures we use, and then describes the CSP data which allows us to identify sustainable/responsible equity investments.

4.3 Model and dataset

4.3.1 Portfolio evaluation metrics

We compare the impact of the different portfolio construction techniques on socially responsible investments along the following dimensions: risk, risk-adjusted returns, level of diversification and stability of asset weights. We use different metrics to capture alternative aspects of the first two of these dimensions and to ensure the convergent validity of these comparisons. The performance evaluation metrics we use will now be explained.

i) Risk

We use the annualized mean standard deviation of portfolio returns as it is the most common measure of total risk. However, although the standard deviation is an appropriate measure of risk for normal (or at least symmetric around the mean) distributions of returns, it may lead to erroneous conclusions in skewed distributions. This is because it treats deviations above and below the mean in the same way, although only the latter should be a source of concern for investors. Hence, we use the annualized mean standard deviation only for negative returns, which is the semi-standard deviation that Markowitz (1991) identified as a “more plausible measure of risk”.

We also use the Value at Risk (VaR) measure which is commonly used for financial risk management purposes. VaR captures the maximum monetary (or percentage) loss for a given investment horizon and a specified probability level, indicating the loss for outcomes in the extreme left tail of the distribution, i.e. the worst outcomes. We use a 99% probability level (i.e. focusing on the worst 1% scenarios) and an investment horizon equal to our out-of-sample period (2001 until 2011). Along similar lines we use the 99th percentile conditional value-at-risk, which is defined as the expected value of the portfolio's returns that do not exceed the possible losses, as indicated by the standard VaR.

Finally, drawdown measures are popular in the asset management industry, and are often used by commodity and hedge fund traders (Eling and Schuhmacher, 2007), as well as by institutional investors such as pension funds (Berkelaar and Kouwenberg, 2010) to assess the magnitude of large potential drops in portfolio returns. The

maximum drawdown rate measures the drop from the highest point in cumulative portfolio returns over a certain time horizon (we use the entire out-of-sample period of twelve years), and is a measure that does not depend on distributional assumptions.

ii) Risk-adjusted performance

Optimization methods maximise the portfolio's risk-adjusted performance. However, since there is no consensus on the most appropriate way to measure returns or risk or how to combine the two in order to measure their trade-off, many different metrics have been proposed and used. The simplest ratio to calculate is the ratio of mean portfolio returns divided by their standard deviation - effectively a Sharpe ratio with a zero risk-free rate (Sharpe, 1994). A more advanced metric, which is an extension of the Sharpe ratio, has been proposed by Dowd (2000). This measure is calculated by dividing the mean return by the VaR of the portfolio, and Dowd (2000) provides several numeric examples which demonstrate its superiority over the Sharpe ratio. Another version of the Sharpe ratio is the Sortino ratio which uses only downside risk (as captured by the semi-standard deviation) instead of total risk (Rollinger and Hoffman, 2014). This modification avoids the paradoxical investment choices brought about by non-normality of the distribution of asset returns. We also calculate the Omega ratio (Shadwick and Keating (2002) which is defined as the probability weighted ratio of gains versus losses for some threshold return target (we use zero, as is common practice). One of the main benefits of this metric over the alternatives is that, by construction, it considers all the moments of the empirical distribution of returns.

In a final set of portfolio performance metrics we divide portfolio returns by the average drawdown to capture significant price falls from previous peaks. A few similar, but distinct, measures have been used for this purpose. The standard metric is the Sterling ratio, which measures the average return divided by the average drawdown for an investment period, Bacon (2008). We also make use of the Calmar ratio, which is the average annual return divided by the maximum drawdown for the entire out of sample period. Young (1991) concludes that the Calmar ratio is superior because it changes gradually, leading to a smoothing of the portfolio's risk-adjusted performance, especially when compared to the Sterling and Sharpe ratios. As a final variation, we use the Burke ratio by taking the difference between the portfolio return and the risk free rate, and dividing it by the square root of the sum of the square of the drawdowns (Burke, 1994). Although these three measures are positively correlated, they are distinct, and can lead to moderately different empirical evaluations of portfolios produced via different approaches.

iii) Diversification and stability

SRI requires additional screening of the universe of investable assets using non-financial criteria (positive, negative, and best-in-class screening are indicative approaches), and the ambiguity in companies' CSP scores makes it more likely that this process will lead to greater estimation risk in their inputs to portfolio models. Therefore it may be harder to create SRI portfolios which effectively reduce idiosyncratic risks through diversification than it is for non-SRI portfolios. So examining the way in which the portfolio optimization process influences this characteristic is important for our analysis. We measure the diversification of the

portfolios by summing up the squared portfolio weights for each constituent and each estimation period, following Blume and Friend (1975).

In addition, a portfolio construction approach which results in substantial rebalancing each period leads to significant transaction costs that reduce returns. So the stability of the resulting portfolio also needs to be examined. Following Goldfarb and Iyengar (2003), the portfolio stability between two successive investment periods is measured by summing the squares of the differences between each asset's portfolio weights in adjacent investment periods.

4.3.2 Dataset

To create SRI portfolios we use CSP metrics constructed using the MSCI ESG STATS database¹⁹. In the relevant research this dataset is the most frequently used, and has been characterised as *“the best-researched and most comprehensive”* (Wood and Jones, 1995) in this field, as well as *“the de facto research standard at the moment”* for measuring CSP (Waddock, 2003, p. 369). It is a multi-dimensional CSP database rich in both the cross section of firms analysed (currently about 3,000 US firms) and the timespan covered (23 years), and has been shown to be characterised by reliability, consistency and construct validity (Sharfman, 1996).

The MSCI ESG STATS data contains annual assessments of the societal and environmental policies and practices of US corporations since 1991. Firms from every sector and industry are assessed on a plethora of indicators relevant to distinct aspects of CSP, which are referred to as *“qualitative issue areas”*. These are: community relations, diversity in the workplace, treatment of employees,

¹⁹Known as KLD STATS before the acquisition of KLD (as part of RiskMetrics) by MSCI in 2010.

environmental issues, product (or services) level of safety and quality, corporate governance framework, and respect for human rights. The relevant assessment is done separately on positive aspects (“strengths”) and controversial aspects (“concerns”) for each qualitative issue area. Sources both internal to the companies (e.g. proxy statements, quarterly reports and other firm documentation) and external to them (e.g. articles in the business and financial press, periodicals, and general media) are used to conduct the assessments of their social performance. In 1991 the dataset covered 650 firms, including all the firms listed in the S&P 500 Composite Index and the Domini 400 Social Index (now the MSCI KLD 400 Social Index). In 2001 this number grew as the relevant universe incorporated the largest 1,000 US companies in terms of market value. Expansion continued in 2003 with the inclusion of the 3,000 largest US firms. Since 2003 the number of firms in the dataset has remained stable at approximately 3,000, and this dataset is available to us until 2011.

We follow the relevant empirical work which uses the MSCI ESG STATS database (Hillman and Keim, 2001; Oikonomou et al., 2012) and focus solely on those qualitative business issues that can be directly connected with primary stakeholder groups. This is based on the stakeholder theory framework developed by Clarkson (1995) which broadly posits that strong collaborative links with those stakeholder groups that are essential to the firm’s viability and operational well-being (i.e. the primary stakeholders) are the only ones that will produce tangible financial benefits to the firm. Hence, the CSP measures used to create SRI portfolios are based on those qualitative issue areas considered important for effective stakeholder

management with local communities, employees (including diversity issues), customers and environmental groups/activists (Hillman and Keim, 2001). An outline of the five indicators used in the assessment of each CSP issue area we are interested in can be found in Appendix 4.B.

For the core part of our analysis we construct aggregate measures of CSP for each firm-year observation in the MSCI ESG STATS universe between 1991 and 2011. For each of the five issue areas of interest we sum all the indications for social strengths and deduct the sum of the respective indications for social concerns for a given firm in a given year. Then we calculate the arithmetic average of all five of these scores in order to create a single, multidimensional CSP rating indicative of the firm's overall social and environmental profile²⁰. Our approach follows previous scholarly work in the area of CSP and finance (Jo and Harjoto, 2012 and Deng, Kang and Low, 2013 being two notable examples). Finally, based on these aggregated CSP scores, we estimate the ranking of each firm across the entire universe covered by MSCI (formerly, KLD) in a given year, and average this relative ranking across the years when the firm is included in the database. We exclude firms for which we cannot construct aggregate scores for at least 10 years out of the 22 in our sample, which helps to ensure the robustness and consistency of the CSP standing of each company. This process results in the estimation of average, aggregate, CSP rankings for 1,362 US firms. We identified the 100 firms with the highest CSP scores as the

²⁰Creating such a multidimensional CSP measure raises questions about the appropriate way to weight each dimension (i.e. the relative importance of each dimension). The common practice in the literature is to use equal weighting (Deng et al., 2013; Oikonomou et al, 2012) which is what we do. In addition, as a robustness check in subsection 4.4.3 we look at robust SRI portfolios based on individual CSP dimensions to investigate whether our results can be replicated using each of the five individual CSP measures.

sub-set of CSP screened firms. This ensures that we have a large enough number of stocks to benefit from the risk reducing effects of diversification when we form portfolios which consist entirely of the top CSP performers. We match this dataset with total returns (i.e. returns that include dividends) for these firms from Thomson Reuters DataStream.

4.4 Results

4.4.1 Main results

Due to the smaller coverage of firms by KLD during its earlier stages, as well as missing observations for quite a few firms over that period, it is not feasible to include years prior to 1993 in the data. Furthermore, KLD data is available to us up to 2011 (inclusive). Tables 4.1 and 4.2 depict the details of the estimation and investment periods (in months) we use to evaluate the SRI portfolios. Each three year out-of-sample period is preceded by its six year estimation period.

Periods(t)	Start	End	Length
Estimation Period 1	1994M1	1999M12	72
Estimation Period 2	1997M1	2002M12	72
Estimation Period 3	2000M1	2005M12	72
Estimation Period 4	2003M1	2008M12	72

Table 4.1: Six-Year Estimation Periods

Periods(t)	Start	End	Length
Investment Period 1	2000M1	2002M12	36
Investment Period 2	2003M1	2005M12	36
Investment Period 3	2006M1	2008M12	36
Investment Period 4	2009M1	2011M12	36

Table 4.2: Non-Overlapping Three-Year Investment Periods

In the literature, the length of the estimation period varies, but five to ten years is generally considered to be appropriate. For instance, Xing et al. (2014) use rolling windows of five years (60 months), ten years (120 months) and 15 years (180 months) to evaluate out-of-sample performance; DeMiguel et al. (2009a) use ten years (120 months); DeMiguel et al. (2009b) use ten years (120 months), 30 years (360 months) and 500 years (6,000 months) to evaluate the out-of-sample performance of simulated data, while Platanakis and Sutcliffe (forthcoming) use a six year (72 months) rolling window. The choice of a six year estimation window which starts in 1994M1 and ends in 1999M12 for the first estimation period lies within the range used by previous studies. We use out-of-sample investment periods of three years. We believe this to be a reasonable investment span within the area of SRI for two reasons. First, we know that a significant part of the demand for SRI comes from long term institutional investors (pension funds and insurance funds, as noted by Cox et al., 2004). These investors have long investment horizons, and generally apply buy-and-hold strategies for long periods (Ryan and Schneider, 2002). Second, the extant literature argues that CSP leads to the creation of comparative advantages that become economically valuable in the long run (Cox et al, 2004; Hillman and Keim, 2001; Waddock and Graves, 1997); while in the short run it may not yield any tangible financial benefits to the firm and investor.

We use data for the first estimation period to compute the optimal portfolio for each method for the following three years (first investment period). Then we roll the data forward by 36 months, so that the second estimation period is now used to compute the optimal portfolio for the second investment period, and so on; providing a total

of four out-of-sample test periods of three years each, or 144 out-of-sample months (12 years) in total. Notice that we have restricted our analysis to include only positive CSP weights, as our investment universe is solely comprised of stocks with a strong CSP track record. We have also ruled out negative asset weights (or short sales), as institutional investors (who are responsible for the majority of demand for SRI) do not engage in short selling.

Performance Measures	Markowitz	Robust	Black-Litterman	1/N	Risk Parity	Reward-to-Risk
Risk Measures						
Mean standard deviation	0.1352	0.1317	0.1748	0.2119	0.1657	0.1547
Mean downside standard deviation	0.0910	0.0894	0.1016	0.1483	0.1193	0.1159
VaR(99%)	0.1216	0.1216	0.1323	0.2305	0.1844	0.1800
Conditional VaR(99%)	0.1434	0.1420	0.1341	0.2432	0.2071	0.2107
Maximum Drawdown Rate	0.3078	0.3056	0.2112	0.6382	0.5237	0.5439
Risk-Return Trade-Off						
Mean Risk-Adjusted Returns	0.4002	0.4183	0.5009	0.0667	0.1438	0.0736
Dowd Ratio	0.0371	0.0378	0.0551	0.0051	0.0108	0.0053
Sortino Ratio	0.5943	0.6162	0.8614	0.0953	0.1998	0.0983
Omega Ratio	1.3717	1.3926	1.5144	1.0557	1.1253	1.0614
Sterling Ratio	0.0937	0.0974	0.2030	0.0102	0.0211	0.0083
Calmar Ratio	0.0146	0.0150	0.0345	0.0018	0.0038	0.0017
Burke Ratio	0.0017	0.0018	0.0032	0.0003	0.0005	0.0002
Diversification and Stability						
Mean Diversification	0.0940	0.0813	0.1282	0.0100	0.0152	0.0248
Mean Stability	0.0910	0.0690	0.2083	0.0000	0.0019	0.0155

Table 4.3: Comparison of the performance of six different portfolio construction approaches across fourteen different metrics. Initial investment universe of 100 consistently high performing CSP firms. Estimation period of six years and out-of-sample period of twelve years. VaR stands for Value at Risk, and $1/N$ is the naive diversification approach

Table 4.3 contains the core of our empirical results, and compares the performance of the six portfolio construction methods we employ (Markowitz, robust estimation, Black-Litterman, naive diversification ($1/N$), risk parity, and reward-to-risk) on the

universe of the best 100 CSP performers. We restrict the universe to 100 firms as a reasonable compromise between having firms which do not really constitute the “cream of the crop” in terms of CSP, and having insufficient firms to effectively study the difference in the impact of the optimization methods on the performance of the SRI portfolios. The performance of these portfolios, formed in six different ways, is compared using 14 criteria which examine risk (five measures), risk-adjusted returns (seven measures), diversification, and portfolio stability. The comparisons are made over the 144 out-of-sample months (12 years) which include four investment periods (4×36 months), with different optimal portfolios applying for each three year period (36 months). The results are adjusted to present annualized figures (where applicable), as is the norm in the asset management industry.

Focusing on risk, the robust estimation approach produces the least risky SRI portfolios in terms of total risk (mean standard deviation), total downside risk (mean downside standard deviation) and VaR, while it comes second to the Black-Litterman approach in terms of conditional VaR and maximum drawdown. The Markowitz model also performs well, finishing second or third in almost all of the risk metrics, and ties first on VaR. On the other hand, the naïve diversification ($1/N$) approach consistently produces the riskiest portfolios across all the measures, with the risk parity and reward-to-risk approaches also producing high risk portfolios. The differences between the scores of the most and least risky portfolios are substantial. In terms of total risk, the robust estimation approach leads to an SRI portfolio with an average annualised standard deviation of returns of 13.17%, whereas the equivalent number for the naive diversification approach is 21.19%, i.e. over 60%

higher; while the VaR score for the $1/N$ portfolios is 90% higher than for robust estimation. The maximum drawdown for the “risky” naive diversification SRI portfolios over 100% larger than for the “safe” Black-Litterman SRI portfolio. These observations are particularly important for the risk-averse, long-term institutional investors who form a significant portion of the demand for SRI.

Focusing on the risk-return trade-off, analysis of the extensive array of metrics we have used produces a very clear picture. In terms of portfolio risk-adjusted returns (Dowd ratio, Sortino ratio, Omega ratio, Sterling ratio, Calmar ratio and Burke ratio), the Black-Litterman model leads to the best out-of-sample performance, with robust estimation ranking second. At the other end of the spectrum, the naive diversification and reward-to-risk methods produce the worst risk-return ratios. Once more, the differences in the extremes are quantitatively large. For example, the value of the Dowd ratio for the portfolio produced using the Black-Litterman method is 0.0551, while for the “naive” portfolio it is just 0.0051 (i.e. less than one tenth of the value of the former). Similarly, looking at the Sterling ratio, the Black-Litterman approach again produces the best result with a value of 0.2030, which is more than 20 times larger than the corresponding result of 0.0083 for the reward-to-risk method. Comparisons across the other risk-return metrics corroborate this conclusion.

The picture changes when looking at the diversification and stability of portfolio constituents. By construction, the naive approach leads to an equal weighting of all the assets in the investment universe (a 1% investment in all 100 stocks in our case), and this remains stable in every period. Hence it leads to the optimal diversification

and stability scores for the metrics we utilize. What is interesting is that the second best approach with regard to these aspects is risk parity, whereas Black-Litterman (which led all the other models in terms of riskiness and risk-return tradeoff) performs the worst, and the Markowitz model is the second worst. So, although the more quantitative portfolio optimization techniques (Black-Litterman, robust estimation and Markowitz) lead to less risky portfolios which provide higher returns per unit of risk taken, they are also associated with less diversification and require more significant rebalancing of their constituent assets. The exact opposite is true for the more simplistic portfolio construction techniques which are based on fundamental investment intuition (naive diversification, risk parity and reward-to-risk).

We now examine some key characteristics the 144 month time series of returns for the different approaches. We focus first on risk, and look at drawdown rates. As can be seen in Figure 4.1, for the majority of the 12 year evaluation period, all the approaches lead to portfolios with reasonably similar drawdown rates. However, from the start of the global financial crisis (late 2007), the drawdown rates of the different models diverge significantly. The Black-Litterman approach is consistently associated with the lowest drawdown, followed by the Markowitz and robust estimation approaches (with almost identical drawdown), whereas naive diversification, risk parity and reward-to-risk have much higher drawdown rates during this period.

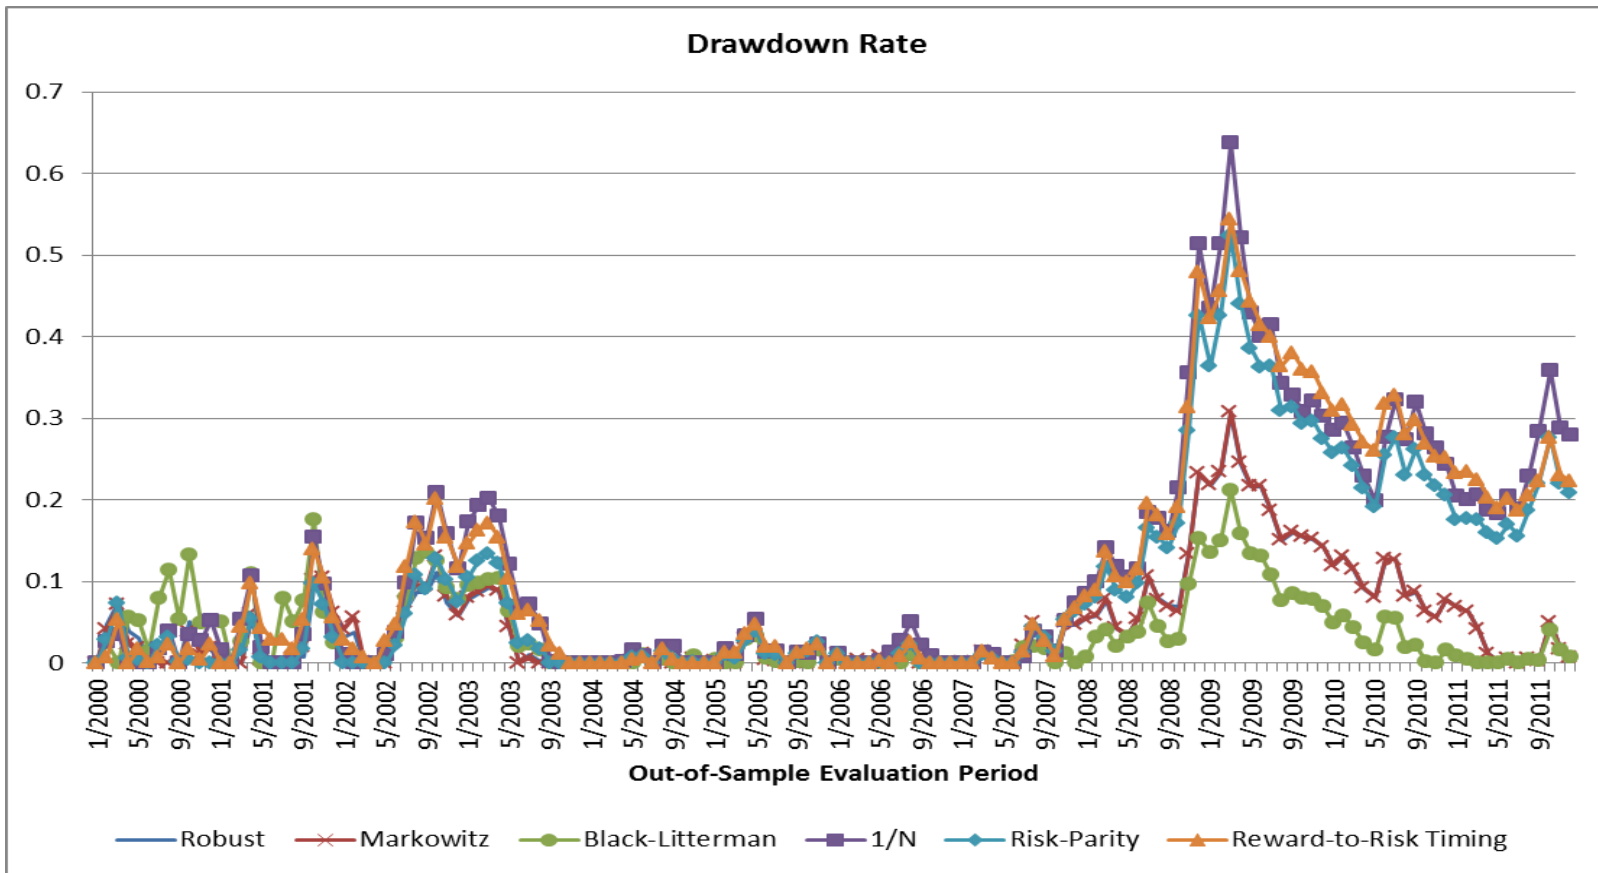


Figure 4.1: Comparison of the drawdown rate of SRI portfolios constructed using the six different portfolio construction approaches over the 12 years of the out-of-sample period.

Figure 4.2 shows the cumulative wealth associated with the different strategies. The Black-Litterman approach dominates all the other strategies in terms of cumulative wealth throughout the entire 12 years (2000-2011). After the first two years the Black-Litterman portfolio clearly moves ahead of its rivals, and over time gains a significant advantage which it maintains irrespectively of the overall direction of the market. Once more, the naive diversification, risk parity and reward-to-risk approaches perform worst, while the robust estimation and Markowitz models are somewhere in between the best and worst performing strategies (and trend so closely together that are nearly indistinguishable as in Figure 4.1). Given that all these portfolios are “long only” (i.e. no short-selling of assets, or negative weights), it is to be expected that the cumulative wealth for all strategies falls in 2008 and 2009 when the financial markets were collapsing, before it starts climbing again.




Figure 4.2: Comparison of the cumulative wealth of SRI portfolios constructed using the six different portfolio construction approaches over the 12 years of the out-of-sample period.

Finally, in Figure 4.3 we provide a comparison of the distribution of asset weights for the SRI portfolios constructed using five different portfolio approaches (we do not include the assets of the $1/N$ approach as they are all assigned a 1% weight). Only weights of 1% or more are presented in Figure 4.3, as this is a rule-of-thumb cutoff point for professional asset managers. The weights are average values across the four investment periods. The Markowitz, robust estimation and Black-Litterman models all lead to portfolios with exactly 26 assets. However, the identity of these assets is not the same across these three approaches, and the size distribution of asset weights is also different. The distribution of asset weights is very similar for the Markowitz and robust estimation approaches, with comparable maxima (10.88% and 11.13% respectively), and six assets with weights of 4% or more in each portfolio. On the other hand, the Black-Litterman portfolio has a lower maximum weight (8.79%), and eight assets with weights of approximately 4% or more. The risk parity and reward-to-risk techniques lead to portfolios with more assets and lower average weights. The risk parity portfolio comprises 42 assets with a weight of 1% or more, and a maximum weight of just 3.56%; while the reward-to-risk portfolio contains 39 assets with a maximum weight of 3.95%. So, although the three less formal optimization models create portfolios which are more stable and require less rebalancing across investment periods, they also contain a greater number of assets compared to the more formal optimization methods. Hence, no clear conclusion can be drawn about the overall impact that transaction costs would have from this analysis.

To investigate this further, we computed the value of shares traded at the start of each of the four out-of-sample periods for each of the six methods. Assuming transactions costs to be 1% of the value of shares traded, we computed the total transactions costs for each method across the 12 years, expressed as a proportion of the initial investment. These percentages appear in table 4.4, along with the corresponding cumulative percentage increase in wealth for each method from Figure 4.3. This shows that the three simple methods have lower transactions costs than the three optimisation methods. But it also shows that the cumulative increases in wealth for the three optimisation methods are very much larger than for the simple methods, so that the net increases in wealth are much larger for the three optimisation methods. This suggests that overall, the three quantitative optimisation methods are preferable to the three simple methods after allowing for transactions costs.

	Total Transactions Costs as a % of Initial Wealth	Cumulative % Increase in Initial Wealth	Differences
Markowitz	5.51%	64.92%	59.41%
Robust	5.22%	66.13%	60.91%
Black-Litterman	8.82%	105.03%	96.21%
1/N	2.40%	16.95%	14.55%
Risk Parity	2.35%	28.61%	26.26%
Reward-to-Risk	3.10%	13.66%	10.56%

Table 4.4: Comparison of the cumulative transactions costs and cumulative increases in wealth. Initial investment universe of 100 consistently high performing CSP firms. Estimation period of six years and out-of-sample period of twelve years. *1/N* is the naive diversification approach.

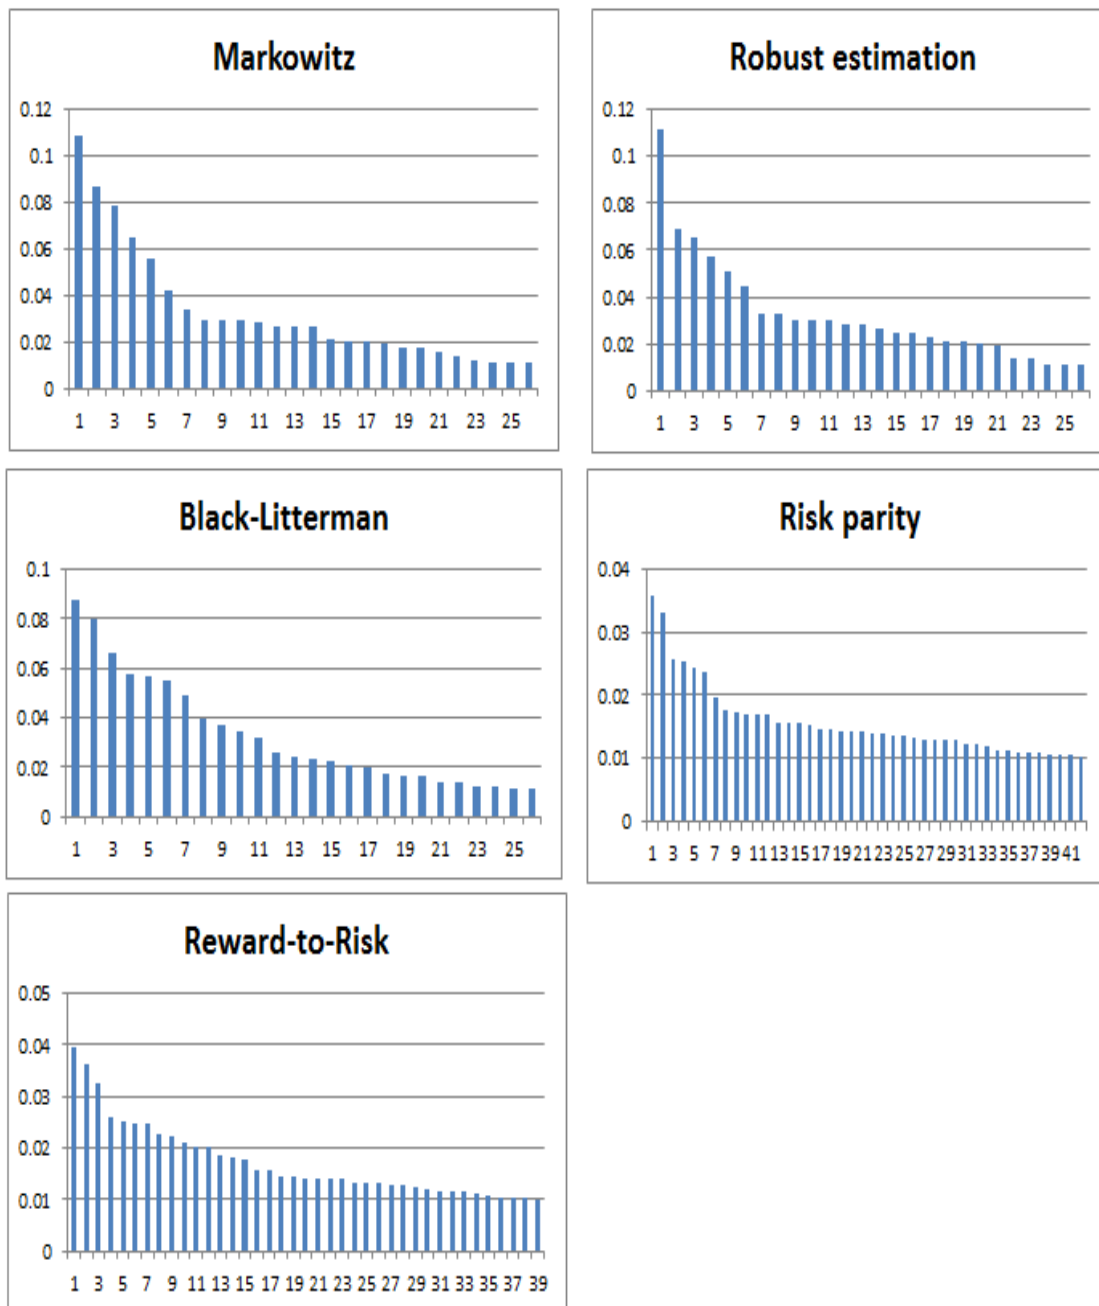


Figure 4.3: Comparison of the asset weight distributions of SRI portfolios constructed using five different portfolio construction approaches. Only weights of assets which are allocated 1% or more are presented.

Before continuing with the robustness tests and additional analyses, we will compare our key findings with the main conclusions from the general literature on asset allocation. In a nutshell, while contradictory results exist regarding the relative effectiveness of different optimization techniques, there is considerable evidence which supports simple portfolio selection methods such as $1/N$. When comparing the

performance of a range of different methods, including $1/N$, Markowitz, risk parity and minimum variance; Board and Sutcliffe (1994), Zhu (2015) and Jacobs et. al. (2014) found mixed results with no clear winner. However, there is more positive evidence.

Bloomfield, Leftwich and Long (1977) found that naive portfolio allocation methods are superior to more sophisticated methods, and that $1/N$ performs well; while Jorion (1991) demonstrated that, for NYSE stocks, $1/N$ is superior to the sophisticated techniques of Markowitz, Bayes-Stein and minimum variance. More recently, Jagannathan and Ma (2003) show that $1/N$ is superior to Markowitz for US stocks. DeMiguel et al. (2009b) compare 14 different methods using US equity data sets, and show that none of them consistently outperforms the $1/N$ approach in terms of risk-adjusted returns. Tu and Zhou (2011) reach similar conclusions. Brown et al. (2013) show that this outperformance is compensation for the increased tail risk (i.e. extreme loss risk) that the naïve diversification portfolio bears. Kirby and Ostdiek (2012) also point out that the stability of naïve diversification is one of the main causes behind its strong performance, and that the reward-to-risk approach can yield stronger results, even in the presence of high transaction costs. Finally, Chaveset. al. (2011) find that the simple methods of $1/N$ and risk parity are superior to Markowitz and minimum variance, while Ang (2014) shows that $1/N$ is preferable to minimum variance and Markowitz. Therefore, the literature for general portfolios tends to support the use of simple, rather than sophisticated portfolio selection techniques.

Our results show that, within the SRI framework, the Black-Litterman approach produces portfolios with the strongest out-of-sample risk-adjusted returns. The robust estimation approach generally produces good results, and the reward-to-risk approach beats the naïve diversification method, as Kirby and Ostdiek (2012) have shown. Our results conflict with those of the general asset allocation literature surveyed above. For an investment universe screened for CSP, the simple methods ($1/N$, risk parity and reward-to-risk) consistently yield the poorest results in terms of risk, maximum possible losses and risk-adjusted returns; while the sophisticated methods (Black-Litterman, robust estimation and Markowitz) yield the best results. This is an important and interesting finding for the SRI community.

One way that could possibly explain the fact that our empirical results contradict a significant part of the existing literature, which supports that naïve diversification methods such as $1/N$ is doing very well, is by employing state-of-the-art multi-factor risk models to assess risk profiles of the portfolio's constituent assets such as the Fama-French three-factor model (Fama and French (1993)) or the Carhart four-factor model (Carhart (1997)). Factor-based risk assessment for portfolios is outside the scope of this research, however we recognise this is a developing area, and future research could investigate the question why the portfolio optimization/construction methods used in our study generate SRI portfolios with significantly different characteristics in comparison to a significant part of the existing literature that supports that more simplified portfolio construction techniques are doing very well in constructing portfolios with unscreened stocks in terms of CSP performance.

4.4.2 Robustness tests

To test the robustness of our results, we narrow our investment universe to the top 80 firms (a significant shrinkage of 20% in the number of assets) in terms of aggregate CSP score. According to traditional finance theory, further restricting the investment universe should lead to inferior portfolio performance. On the other hand, given the strong empirical link between higher CSP and lower financial risk (Orlitzky and Benjamin, 2001; Godfrey et al., 2009; Oikonomou et al., 2012), applying more intense CSP screening criteria may improve the performance of SRI portfolios. Hence, we decrease the number of equities to 80 and use the same estimation and investment periods as in our previous analysis to compare the performance of SRI portfolios according to their construction method. The results are summarized in Table 4.5.

Performance Measures	Markowitz	Robust	Black-Litterman	1/N	Risk Parity	Reward-to-Risk
Risk Measures						
Mean standard deviation	0.1281	0.1242	0.1514	0.2135	0.1670	0.1622
Mean downside standard deviation	0.0853	0.0839	0.0933	0.1488	0.1208	0.1229
VaR(99%)	0.1180	0.1168	0.1320	0.2438	0.2021	0.2094
Conditional VaR(99%)	0.1389	0.1404	0.1447	0.2487	0.2170	0.2319
Maximum Drawdown Rate	0.2879	0.2954	0.2530	0.6493	0.5430	0.6168
Risk-Return Trade-Off						
Mean Risk-Adjusted Returns	0.4631	0.4580	0.4584	0.0890	0.1582	0.0598
Dowd Ratio	0.0419	0.0406	0.0438	0.0065	0.0109	0.0039
Sortino Ratio	0.6952	0.6775	0.7439	0.1277	0.2186	0.0789
Omega Ratio	1.4365	1.4229	1.3999	1.0764	1.1395	1.0505
Sterling Ratio	0.1064	0.1041	0.1162	0.0143	0.0231	0.0063
Calmar Ratio	0.0172	0.0160	0.0229	0.0024	0.0041	0.0013
Burke Ratio	0.0019	0.0019	0.0022	0.0004	0.0006	0.0002
Diversification and Stability						
Mean Diversification	0.1012	0.0649	0.2312	0.0125	0.0187	0.0301
Mean Stability	0.0706	0.0413	0.4892	0.0000	0.0023	0.0190

Table 4.5: Comparison of the performance of six different portfolio construction approaches across fourteen different metrics. Initial investment universe of 80 consistently high performing CSP firms. Estimation period of six years and out-of-sample period of twelve years. VaR stands for Value at Risk and 1/N is the naive diversification approach.

The reduction in the number of assets included in the SRI portfolios does not change the core of our previous conclusions. The robust estimation approach still produces the least risky portfolios (having the lowest average standard deviation, downside standard deviation and VaR, and the second lowest conditional VaR), with Markowitz usually creating the second best portfolios in this regard, with the Black-Litterman approach following next. At the other end of the spectrum the naïve diversification approach leads to the riskiest portfolios, while the risk parity and reward-to-risk approaches also have high risk. Compared to the core results, the situation is a bit different for the risk-return metrics, although the traditional, more quantitative optimization methods still outperform the mathematically less formal alternatives.

The Black-Litterman technique still ranks first in this dimension according to every metric (except for the Omega ratio). The Markowitz model usually finishes second best, and tends to outperform the robust portfolio. The naïve diversification and reward-to-risk portfolios still have the lowest risk-adjusted returns on every relevant metric. As previously, the rank order is reversed when looking at the diversification and stability measures, with the $1/N$ approach producing the best results, followed by risk parity. The Black-Litterman model finishes last, with Markowitz as second worst. Overall, even when significantly reducing the investment universe, the rank order of the different approaches remains largely unchanged. The sophisticated approaches have lower risk and a superior risk-return trade-off than the unsophisticated approaches, but the simpler techniques are more diversified and stable.

As a second robustness test we keep the number of assets at 100, but change the length of the estimation periods to nine years (108 months) instead of six years (72 months). Tables 4.6 and 4.7 provide the relevant details of the new estimation and investment periods. We now have only three estimation periods and three investment periods.

Periods(t)	Start	End	Length
Estimation Period 1	1994M1	2002M12	108
Estimation Period 2	1997M1	2005M12	108
Estimation Period 3	2000M1	2008M12	108

Table 4.6: Nine-Year Estimation Periods

Periods(t)	Start	End	Length
Investment Period 1	2003M1	2005M12	36
Investment Period 2	2006M1	2008M12	36
Investment Period 3	2009M1	2011M12	36

Table 4.7: Non-Overlapping Three-Year Investment Periods

Performance Measures	Markowitz	Robust	Black-Litterman	1/N	Risk Parity	Reward-to-Risk
Risk Measures						
Mean standard deviation	0.1249	0.1250	0.1247	0.2237	0.1846	0.1619
Mean downside standard deviation	0.0822	0.0826	0.0820	0.1597	0.1345	0.1244
VaR(99%)	0.1354	0.1363	0.1346	0.2559	0.2255	0.2318
Conditional VaR(99%)	0.1354	0.1363	0.1346	0.2559	0.2255	0.2318
Maximum Drawdown Rate	0.2543	0.2578	0.2519	0.6531	0.5689	0.5483
Risk-Return Trade-Off						
Mean Risk-Adjusted Returns	0.4799	0.4795	0.4840	0.0659	0.1101	0.1272
Dowd Ratio	0.0369	0.0366	0.0374	0.0048	0.0075	0.0074
Sortino Ratio	0.7288	0.7253	0.7365	0.0923	0.1511	0.1655
Omega Ratio	1.4492	1.4484	1.4536	1.0565	1.0971	1.1121
Sterling Ratio	0.1223	0.1204	0.1258	0.0093	0.0144	0.0140
Calmar Ratio	0.0196	0.0194	0.0200	0.0019	0.0030	0.0031
Burke Ratio	0.0024	0.0024	0.0024	0.0003	0.0005	0.0005
Diversification and Stability						
Mean Diversification	0.0869	0.0798	0.0871	0.0100	0.0150	0.0233
Mean Stability	0.0605	0.0542	0.0600	0.0000	0.0009	0.0107

Table 4.8: Comparison of the performance of six different portfolio construction approaches across fourteen different metrics. Initial investment universe of 100 consistently high performing CSP firms. Estimation period of nine years.

The results are summarized in Table 4.8. All our previous conclusions remain valid, and in some cases are even stronger than those drawn from the original results. The Black-Litterman model dominates all the alternative SRI portfolios according to every metric of risk and risk-adjusted performance. The Markowitz and robust approaches are second and third best respectively according to the same criteria. The naïve diversification technique produces the riskiest portfolios with the worst risk-return ratios, while the risk-parity and reward-to-risk portfolios do not fare much better. Once more, the $1/N$ approach leads to the most stable and well-diversified portfolios, followed by the risk-parity portfolios; whereas the Markowitz and Black-Litterman portfolios perform worst on both these dimensions.

4.4.3 Additional analyses

It has been documented that different measures of CSP based on different aspects or dimensions of corporate sustainability relate to distinct stakeholder groups (Griffin and Mahon, 1997; Mattingly and Berman, 2006) and may have different impacts on financial performance. This is especially relevant when looking at samples of firms from different industries, where the social and environmental issues and key performance indicators can be significantly different. So far in our analysis we avoided this issue by using an aggregate, multidimensional measure of CSP to construct SRI portfolios. In this subsection, we create five different SRI investment data sets, each based on one of the CSP qualitative issue areas from which the aggregate CSP measure was constructed; i.e. relationships with local communities, diversity in the workplace, employee relations, environmental considerations, and product safety and quality.

To construct these SRI portfolios we follow the principles outlined in subsection 4.4.1. Thus, we use the top 100 firms for each of the qualitative issue areas, and the estimation and investment periods described in Tables 4.1 and 4.2. The performance metrics and the optimization approaches employed also remain the same. The results appear in Table 4.9 which contains five different panels, each of which focuses on one of the five CSP dimensions.

Community relations	Markowitz	Robust	Black-Litterman	1/N	Risk Parity	Reward-to-Risk
Mean standard deviation	0.1433	0.1458	0.1649	0.2134	0.1784	0.1605
Mean downside standard deviation	0.1042	0.1063	0.1166	0.1508	0.1332	0.1234
VaR(99%)	0.1939	0.1962	0.2031	0.2387	0.2032	0.1869
Conditional VaR(99%)	0.1994	0.2008	0.2132	0.2513	0.2298	0.2168
Maximum Drawdown Rate	0.3806	0.4000	0.3742	0.6010	0.5781	0.5811
Mean Risk-Adjusted Returns	0.3793	0.3365	0.3945	0.1542	0.1417	0.0995
Dowd Ratio	0.0234	0.0208	0.0267	0.0115	0.0104	0.0071
Sortino Ratio	0.5219	0.4615	0.5579	0.2182	0.1899	0.1294
Omega Ratio	1.3680	1.3203	1.3840	1.1368	1.1259	1.0860
Sterling Ratio	0.0766	0.0617	0.0815	0.0277	0.0200	0.0116
Calmar Ratio	0.0119	0.0102	0.0145	0.0046	0.0036	0.0023
Burke Ratio	0.0016	0.0013	0.0018	0.0007	0.0005	0.0003
Mean Diversification	0.0805	0.0724	0.1421	0.0100	0.0147	0.0226
Mean Stability	0.0741	0.0731	0.2710	0.0000	0.0026	0.0154
Diversity	Markowitz	Robust	Black-Litterman	1/N	Risk Parity	Reward-to-Risk
Mean standard deviation	0.1429	0.1437	0.1628	0.2097	0.1648	0.1650
Mean downside standard deviation	0.1023	0.1027	0.1073	0.1482	0.1220	0.1309
VaR(99%)	0.1326	0.1428	0.1487	0.2314	0.2116	0.2293
Conditional VaR(99%)	0.1491	0.1542	0.1593	0.2453	0.2122	0.2299
Maximum Drawdown Rate	0.4899	0.4966	0.3954	0.7746	0.6642	0.8273
Mean Risk-Adjusted Returns	0.1519	0.1501	0.2319	-0.0074	0.0257	-0.1238
Dowd Ratio	0.0136	0.0126	0.0211	-0.0006	0.0017	-0.0074
Sortino Ratio	0.2123	0.2100	0.3517	-0.0104	0.0347	-0.1560
Omega Ratio	1.1272	1.1257	1.2043	0.9939	1.0217	0.9003
Sterling Ratio	0.0169	0.0166	0.0363	-0.0009	0.0027	-0.0086
Calmar Ratio	0.0037	0.0036	0.0080	-0.0002	0.0005	-0.0021
Burke Ratio	0.0005	0.0005	0.0009	0.0000	0.0001	-0.0003
Mean Diversification	0.0900	0.0610	0.1071	0.0100	0.0145	0.0244

Mean Stability	0.0562	0.0478	0.0999	0.0000	0.0016	0.0149
Employee relations	Markowitz	Robust	Black-Litterman	1/N	Risk Parity	Reward-to-Risk
Mean standard deviation	0.1537	0.1479	0.1475	0.2031	0.1726	0.1722
Mean downside standard deviation	0.1125	0.1088	0.1059	0.1454	0.1276	0.1329
VaR(99%)	0.1371	0.1376	0.1414	0.2141	0.1763	0.2005
Conditional VaR(99%)	0.1952	0.1933	0.1877	0.2434	0.2215	0.2475
Maximum Drawdown Rate	0.3675	0.3719	0.3453	0.5345	0.4713	0.5287
Mean Risk-Adjusted Returns	0.3445	0.3499	0.3819	0.1297	0.2078	0.1396
Dowd Ratio	0.0322	0.0313	0.0332	0.0103	0.0170	0.0100
Sortino Ratio	0.4710	0.4756	0.5319	0.1812	0.2813	0.1809
Omega Ratio	1.3214	1.3296	1.3568	1.1113	1.1851	1.1223
Sterling Ratio	0.0576	0.0577	0.0683	0.0205	0.0311	0.0166
Calmar Ratio	0.0120	0.0116	0.0136	0.0041	0.0063	0.0038
Burke Ratio	0.0013	0.0013	0.0015	0.0006	0.0008	0.0005
Mean Diversification	0.0825	0.0665	0.0803	0.0100	0.0168	0.0267
Mean Stability	0.0897	0.0543	0.0816	0.0000	0.0030	0.0166
Environment	Markowitz	Robust	Black-Litterman	1/N	Risk Parity	Reward-to-Risk
Mean standard deviation	0.1381	0.1376	0.1392	0.2190	0.1699	0.1585
Mean downside standard deviation	0.0997	0.1012	0.0973	0.1511	0.1199	0.1149
VaR(99%)	0.1395	0.1388	0.1444	0.2298	0.1939	0.1792
Conditional VaR(99%)	0.1504	0.1496	0.1470	0.2306	0.1973	0.1890
Maximum Drawdown Rate	0.4469	0.4621	0.3792	0.6466	0.5077	0.5016
Mean Risk-Adjusted Returns	0.1752	0.1426	0.2242	0.0592	0.1772	0.1419
Dowd Ratio	0.0145	0.0118	0.0180	0.0047	0.0129	0.0105
Sortino Ratio	0.2429	0.1939	0.3205	0.0858	0.2513	0.1958
Omega Ratio	1.1454	1.1169	1.1875	1.0482	1.1531	1.1203
Sterling Ratio	0.0208	0.0163	0.0307	0.0085	0.0281	0.0168
Calmar Ratio	0.0045	0.0035	0.0069	0.0017	0.0049	0.0037
Burke Ratio	0.0005	0.0004	0.0007	0.0003	0.0007	0.0005
Mean Diversification	0.0856	0.0790	0.0937	0.0100	0.0149	0.0247
Mean Stability	0.0729	0.0752	0.1153	0.0000	0.0017	0.0150
Product safety and quality	Markowitz	Robust	Black-Litterman	1/N	Risk Parity	Reward-to-Risk
Mean standard deviation	0.1610	0.1609	0.1938	0.2449	0.2014	0.1931
Mean downside standard deviation	0.1218	0.1242	0.1440	0.1739	0.1490	0.1452
VaR(99%)	0.1642	0.1675	0.2170	0.2630	0.2418	0.2346
Conditional VaR(99%)	0.2103	0.2217	0.2406	0.2836	0.2741	0.2719
Maximum Drawdown Rate	0.5993	0.6195	0.8073	0.8290	0.6744	0.6969
Mean Risk-Adjusted Returns	0.0686	0.0662	-0.1149	-0.0224	0.1141	0.0782
Dowd Ratio	0.0056	0.0053	-0.0086	-0.0017	0.0079	0.0054

Sortino Ratio	0.0907	0.0857	-0.1547	-0.0316	0.1542	0.1039
Omega Ratio	1.0574	1.0559	0.9109	0.9819	1.1002	1.0662
Sterling Ratio	0.0067	0.0062	-0.0097	-0.0028	0.0152	0.0087
Calmar Ratio	0.0015	0.0014	-0.0023	-0.0006	0.0028	0.0018
Burke Ratio	0.0002	0.0002	-0.0004	-0.0001	0.0005	0.0003
Mean Diversification	0.0955	0.0845	0.0983	0.0100	0.0155	0.0254
Mean Stability	0.0761	0.0629	0.1346	0.0000	0.0030	0.0154

Table 4.9: Comparison of the performance of six different portfolio construction approaches across fourteen different metrics for five different dimensions of corporate social and environmental performance. Initial investment universe of 100 consistently high performing CSP firms. Estimation period of six years and out-of-sample period of twelve years. VaR stands for Value at Risk and $1/N$ is the naive diversification approach.

Table 4.9 reveals that, although there is variability, the conclusions drawn from the aggregate measure of CSP are verified for the majority of the individual CSP dimensions. More specifically, the Black-Litterman approach consistently produces the highest risk-return trade-offs (and the minimum drawdown rate) for the community relations, diversity and employee relations aspects of CSP. Markowitz also does very well in these CSP dimensions, creating portfolios with the lowest volatility, downside risk and VaR, while also usually finishing second in terms of risk-adjusted returns. The robust approach ranks second or third in terms of both riskiness and risk-return trade-offs, whereas the naïve diversification and reward-to-risk approaches invariably lead to portfolios with the worst values on these measures. So once again the formal optimization models outperform the less strict portfolio construction approaches. However, in line with the core findings of this study, the $1/N$ approach still produces the most well-diversified and stable portfolios, while the Black-Litterman and Markowitz models finish last on these criteria.

The picture is qualitatively similar, although not identical, when focusing on corporate environmental performance. The key differences are that the robust

portfolio performs less well, finishing next to last on most of the risk-adjusted return ratios, while the risk-parity portfolio does better, being ranked second best on most risk-return metrics. Black-Litterman is still the model of choice according to most criteria, while $1/N$ is last.

Things are quite different when using product safety and quality as the CSP feature guiding portfolio construction. This is the only CSP dimension where the Black-Litterman approach leads to poorly performing portfolios with the worst risk-adjusted returns, diversification and stability characteristics. On the other hand, the risk parity and reward-to-risk techniques are the methods with the best and second best risk-return trade-offs respectively, something that has not been the case in any of our previous analysis.

The distinctiveness of CSP dimensions and the variability of the financial impacts of each has been well documented in the empirical CSP literature (Hillman and Keim, 2001; Mattingly and Berman, 2006; Oikonomou et al., 2012). Hence, our results are compatible with previous findings. Overall, Table 4.9 shows that the results from the solo use of the CSP measures produces only slightly different conclusions. This suggests that our results are not highly sensitive to the weighting scheme involved in computing the aggregate CSP. Finally, for the same reason as in Chapter 3, it was not possible to investigate whether the differences in performance metrics between the portfolio construction models are statistically significant.

4.5 Conclusions

We expand the SRI literature by moving beyond the question of whether portfolios comprising “sustainable equities” outperform conventional investments, and focus

on finding optimal ways to construct SRI portfolios. We have found that the optimization process for forming SRI portfolios matters. There are large, economically significant differences in the risk, risk-adjusted returns, diversification and intertemporal stability of the SRI portfolios, depending on which optimization technique is used. Formal optimization techniques (Markowitz, Black-Litterman and robust estimation) tend to produce less risky SRI portfolios with higher risk-adjusted returns and a smaller total number of constituent assets compared to less formal techniques (naïve diversification, risk parity and reward-to-risk). The Black-Litterman model is usually the best technique, while naïve diversification is usually the worst on these criteria. These conclusions are robust to different lengths of the estimation and investment periods, and to the use of more stringent CSP screening criteria. Our conclusions using CSP-screened assets are in contrast to studies of unscreened assets, which have found that naive diversification is one of the best techniques.

When implementing the various portfolio construction approaches for single CSP dimensions, we have found that our key findings for aggregate CSP scores also hold for the community, diversity and employee relations dimensions of CSP. But they are less applicable for environmental performance and product safety and quality. This is in line with previous work in the wider literature on the financial effects of CSP, and demonstrates the contextualization required for an analysis to be complete.

Overall, our study shows that just applying stringent SRI criteria to restrict the investment universe to the best socially and environmentally performing companies is insufficient. The optimization process is also very important. It further demonstrates to fund managers, institutional and retail investors (especially those

who are more risk-averse and have longer-term investment horizons) that the more quantitative approaches to portfolio construction typically lead to better results. The appropriate selection of an optimization technique is an issue which needs to be taken into serious consideration for anyone placing their funds in SRI. Choosing the correct optimization method for the creation of SRI portfolios will lead to stronger financial performance, which in turn will generate greater demand for this kind of investment. Through this mechanism, the cost of equity for any corporation that applies sustainable/ responsible/ethical practices will be reduced, incentivizing them to engage in such behavior, while penalizing companies involved in various social or environmental controversies by increasing their cost of capital. In short, the selection of the most suitable optimization method for SRI portfolios will have an effect on the bottom line of companies and, through this, on the promotion of societal well-being and environmental conservation. Hence, the results of this study are of interest and importance to a variety of constituents including investors, fund managers, corporate executives, social and environmental activists and overall concerned citizens.

Although our study is innovative within the SRI field, it is limited by considering only one asset class (equities) and the geographic coverage of the markets considered (US). Future studies can extend our analysis in either of these directions. Furthermore, the selection of the CSP criteria and dataset used is always an important issue within the literature (Griffin and Mahon, 1997). Using different social and environmental sources of data and alternative CSP metrics would provide a useful test for the reliability of different optimization methods for SRI.

Appendix 4.A: Mathematical definition of implemented optimization approaches

4.A.1 Markowitz portfolio optimization

The optimization framework proposed by Markowitz (1952) assumes that the expected value ($\boldsymbol{\mu}$) and the covariance ($\boldsymbol{\Sigma}$) of asset returns are known with certainty. Specifically, if $\boldsymbol{\Phi}$ denotes the column vector of portfolio weights (decision variables) defined as $\boldsymbol{\Phi} = [\Phi_1, \Phi_2, \dots, \Phi_N]^T$ with N assets in the portfolio, a sample variance-covariance matrix of asset returns ($\boldsymbol{\Sigma}$) and a column vector of mean asset returns ($\boldsymbol{\mu} = [\mu_1, \mu_2, \dots, \mu_N]^T$), then the minimum variance portfolio selection problem is expressed as follows:

$$\begin{aligned} \min_{\boldsymbol{\Phi}} \quad & \boldsymbol{\Phi}^T \boldsymbol{\Sigma} \boldsymbol{\Phi} \\ \text{s.t.} \quad & \boldsymbol{\mu}^T \boldsymbol{\Phi} \geq \alpha \\ & \mathbf{1}^T \boldsymbol{\Phi} = 1 \\ & \Phi_i \geq 0, \quad \forall i=1, \dots, N \end{aligned} \tag{A.4.1}$$

where the objective is the selection of a portfolio $\boldsymbol{\Phi}$ that minimizes the risk (variance) among all feasible portfolios. The constraint $\mathbf{1}^T \boldsymbol{\Phi} = 1$ requires that the portfolio weights sum to one. The constraint $\boldsymbol{\mu}^T \boldsymbol{\Phi} \geq \alpha$ sets a lower bound on portfolio mean return²¹. We also rule out short selling by imposing non-negativity constraints ($\Phi_i \geq 0$) on the asset weights.

4.A.2 Robust estimation approach

To deal with the effects of parameter uncertainty we also apply a robust estimation strategy which is inspired by previous studies of robust asset allocation with norm constraints on the portfolio weights, such as DeMiguel et al. (2009a). Specifically, we follow Xing et al. (2014) and impose a constraint of an l_1 norm, $\|\boldsymbol{\Phi}\|_1$ (taxicab or

²¹In our analysis, we set the lower bound of the mean portfolio return (parameter α) to 1% on an annual basis. We assume this is the minimum expected return an investor would be willing to accept in order to invest in a risky portfolio. The selection of the exact value of α is not crucial in this framework and does not influence the conclusions drawn from our results.

Manhattan norm) and an l_∞ norm, $\|\Phi\|_\infty$ (maximum norm) on the portfolio weights. The taxicab norm (l_1) is the sum of the absolute values of a vector, and setting an upper bound on l_1 encourages sparse solutions, i.e. portfolios with active positions in only a few assets (sparse portfolios), see for instance Brodie et al. (2008).

Having active positions in only a few assets leads to the significant practical benefit of lower transaction costs. However, after applying the l_1 norm constraint, some of the positions may be very large. The additional use of the l_∞ norm addresses this issue, see for instance Brondell and Reich (2008). The maximum norm (l_∞) of a vector is the largest absolute value of the elements in the vector, and an upper bound on l_∞ prevents large positions in any asset. Therefore a combination of the l_1 and l_∞ upper bounds tends to produce sparse portfolios without very large individual weights.

By additionally employing the three optimization constraints of problem (A.4.1), the optimization problem can be written as follows:

$$\begin{aligned}
 \min_{\Phi} \quad & \Phi^T \Sigma \Phi \\
 \text{s.t.} \quad & \|\Phi\|_1 + \|\Phi\|_\infty \leq c \\
 & \mu^T \Phi \geq \alpha \\
 & \mathbf{1}^T \Phi = 1 \\
 & \Phi_i \geq 0, \quad \forall i=1, \dots, N
 \end{aligned} \tag{A.4.2}$$

where $\|\Phi\|_1 = \sum_{i=1}^N |\Phi_i|$ denotes the l_1 norm, $\|\Phi\|_\infty = \max_{1 \leq i \leq N} \{|\Phi_i|\}$ represents the l_∞ norm ($|\Phi_i|$ denotes the absolute value of Φ_i), while $\mathbf{1}$ is a column vector of ones. Furthermore, $c \geq 1 + \frac{1}{N}$ denotes the upper bound²² of the constraint that involves the l_1 and l_∞ norms. We run the robust estimation strategy using time varying

²² The lowest feasible value of c occurs when it is equal to $1 + 1/N$ and all the asset weights are equal to $1/N$. Simulations of robust portfolio models usually start with a value of c just above 1. Setting c some way above 1 permits the optimization process to make a trade-off between preventing short sales and allowing large asset weights.

values for c . For the first period we set $c = 1.1$, but for subsequent periods we set c equal to the value in the range 1.1 to 9.0 that gave the best out-of-sample performance during the previous period, Fan et al. (2008), Fan et al. (2012) and Xing et al. (2014). The conclusions are unaltered when different values of c are used in the first period.

4.A.3 Black-Litterman approach

The Black-Litterman portfolio framework combines the subjective views of the investor, in terms of returns and risk, with those of a benchmark portfolio (e.g. the equilibrium market portfolio), and is an alternative way of dealing with estimation risk in the input parameters (Black and Litterman, 1992). The posterior estimates of expected returns and covariances are then used in the portfolio optimization process. We use the portfolio optimization model described in problem (A.4.1), which minimizes the portfolio risk (variance), subject to three linear constraints on asset weights.

The column vector of implied excess returns ($\mathbf{\Pi}$) for the benchmark portfolio is expressed as follows:

$$\mathbf{\Pi} = \lambda \mathbf{\Sigma} \mathbf{\Phi}^{benchmark} \quad (\text{A.4.3})$$

where λ denotes the risk aversion coefficient²³ and $\mathbf{\Phi}^{benchmark}$ is a column vector of the asset weights of the benchmark portfolio²⁴. The column vector of the posterior asset returns ($\mathbf{\mu}_{BL}$) is given by:

$$\mathbf{\mu}_{BL} = \left[(\tau \mathbf{\Sigma})^{-1} + \mathbf{P}^T \mathbf{\Omega}^{-1} \mathbf{P} \right]^{-1} \left[(\tau \mathbf{\Sigma})^{-1} \mathbf{\Pi} + \mathbf{P}^T \mathbf{\Omega}^{-1} \mathbf{Q} \right] \quad (\text{A.4.4})$$

where τ represents the overall level of confidence in the column vector $\mathbf{\Pi}$. We set this parameter to 0.1625, which is the mean of the values used in the literature (see

²³ The investor's risk aversion parameter disappears in the optimization process, since we use just the portfolio variance in the objective function.

²⁴ We have used the equally-weighted portfolio ($1/N$) as the benchmark portfolio. Bessler, Opfer and Wolff (forthcoming) show that the effect of the choice of benchmark portfolio on the Black-Litterman results is minimal.

for instance Bessler, Opfer and Wolff, forthcoming; and Platanakis and Sutcliffe, forthcoming).

Bessler, Opfer and Wolff (forthcoming) experimented with different values of the parameter τ , and showed that the Black-Litterman results are robust to the choice of τ in the range between 0.025 to 1.00. In addition, \mathbf{P} denotes a binary matrix defining the assets involved in each view, \mathbf{Q} is a column vector that contains the views (subjective returns), and $\mathbf{\Omega}$ is a diagonal matrix that quantifies the reliability of each view. The latter is estimated following Meucci (2010), as follows:

$$\mathbf{\Omega} = \frac{1}{\delta} \mathbf{P} \mathbf{\Sigma} \mathbf{P}^T \quad (\text{A.4.5})$$

where δ represents the overall level of confidence in the investor's views. We follow Meucci (2010), setting δ to one. We tried different values of the parameter δ and found that the impact on the Black-Litterman results is negligible. We follow Bessler, Opfer and Wolff (forthcoming) and use the sample means as subjective return estimates. Following Satchell and Scowcroft (2000), Bessler, Opfer and Wolff (forthcoming) and other studies, we estimate the posterior covariance matrix ($\mathbf{\Sigma}_{\text{BL}}$) as follows:

$$\mathbf{\Sigma}_{\text{BL}} = \mathbf{\Sigma} + \left[(\tau \mathbf{\Sigma})^{-1} + \mathbf{P}^T \mathbf{\Omega}^{-1} \mathbf{P} \right]^{-1} \quad (\text{A.4.6})$$

Finally, Bessler, Opfer and Wolff (forthcoming) suggest that the reliability of the views incorporated in the Black-Litterman model is time-varying. For each of our out-of-sample periods, we estimate the reliability of the views for the subsequent out-of-sample period using the entire estimation period of 72 months.

4.A.4 Risk-parity portfolio construction

The risk-parity portfolio approach is based on the idea that portfolio components (i.e. assets) contribute to the same extent to portfolio risk. In its simplified version, the risk-parity approach ignores correlations between asset returns, and the asset weights are anti-proportional to their sample variance. Hence, the portfolio weights are computed as follows:-

$$\Phi_i = \frac{1/\sigma_i^2}{\sum_{i=1}^N (1/\sigma_i^2)}, \quad \forall i=1, \dots, N \quad (\text{A.4.7})$$

4.A.5 Reward-to-risk timing portfolios

The reward-to-risk timing portfolio strategy has been proposed by Kirby and Ostdiek (2012) and is based on the reward-to-risk ratio, which is defined as the mean return divided by the variance of each asset. Specifically, the reward-to-risk timing strategy takes into account both risk and return, and allocates more weight to assets with higher risk-adjusted returns. The portfolio weights are given by:

$$\Phi_i = \frac{\mu_i^+ / \sigma_i^2}{\sum_{i=1}^N (\mu_i^+ / \sigma_i^2)}, \quad \forall i=1, \dots, N, \quad (\text{A.4.8})$$

where $\mu_i^+ = \max(\mu_i, 0)$ to prevent short selling. In the very rare case when all asset returns are negative, an equally-weighted portfolio ($1/N$) is considered.

Appendix B: Indicators for each CSP dimension

MSCI KLD Qualitative Issue Areas	Strengths	Concerns
Community	<ul style="list-style-type: none"> - Charitable Giving - Innovative Giving - Non-US Charitable Giving - Support for Housing - Support for Education - Indigenous Peoples Relations - Volunteer Programs - Other Strength 	<ul style="list-style-type: none"> - Investment Controversies - Negative Economic Impact - Indigenous Peoples - Tax Disputes - Other Concern
Diversity	<ul style="list-style-type: none"> - CEO's identity - Promotion - Board of Directors - Work/Life Benefits - Women & Minority Contracting - Employment of the Disabled - Gay & Lesbian Policies - Other Strength 	<ul style="list-style-type: none"> - Controversies - Non-Representation - Other Concern
Employee	<ul style="list-style-type: none"> - Union Relations - No-Layoff Policy - Cash Profit Sharing - Employee Involvement - Retirement Benefits Strength - Health and Safety Strength - Other Strength 	<ul style="list-style-type: none"> - Union Relations - Health and Safety Concern - Workforce Reductions - Retirement Benefits - Other Concern
Environment	<ul style="list-style-type: none"> - Beneficial Products and Services - Pollution Prevention - Recycling - Clean Energy - Communications - Property, Plant, and Equipment - Management Systems - Other Strength 	<ul style="list-style-type: none"> - Hazardous waste - Regulatory Problems - Ozone Depleting - Substantial Emissions - Agricultural Chemicals - Climate Change - Other Concern
Product Safety & Quality	<ul style="list-style-type: none"> - Quality - R&D/Innovation - Benefits to Economically - Other Strength 	<ul style="list-style-type: none"> - Product Safety - Marketing/Contracting Concern - Antitrust - Other Concern

Chapter 5

5 Pension Scheme Redesign and Wealth Redistribution Between the Members and Sponsor: The USS Rule Change in October 2011²⁵

On retirement the sponsor of a UK defined benefit (DB) pension scheme promises to pay a pension according to the rules of the scheme, regardless of the scheme's financial state. This appears to place all the risks (investment, interest rates, inflation, salaries, longevity, regulation, etc.) on the sponsor, who is usually the employer. But the sponsor can share these risks with active and future members of the scheme by altering the rules applying to future accruals. For example, a large deficit may lead to rule changes such as an increase in the members' contribution rate, the introduction of limited price indexation, a later retirement age, or a reduction in the accrual rate. Because UK law does not allow accrued benefits to be reduced, rule changes only apply to future accruals. This means that the youngest scheme members are the hardest hit by such action as they will be accruing benefits under the new rules for many years, while those near retirement are largely unaffected since their substantial accrued benefits are legally protected.

Before a rule change the various scheme participants have both accrued benefits and expectations of the net present value (NPV) of their future interactions with the scheme, i.e. contributions to be made and pensions to be received²⁶. After a rule change these expectations are altered, and the difference between NPVs of the cash flows before and after the rule change for each age cohort quantifies the

²⁵ The content of this Chapter was presented at the ICMA Centre Internal Research Seminar (Reading, UK) and has been accepted for revise and resubmit by the Insurance: Mathematics and Economics journal.

²⁶ The resulting changes in cash flows between the members and sponsor are zero sum.

redistributive effect of the rule change. For example, an increase in the member contribution rate redistributes pension wealth from active and future members to the sponsor. Therefore a rule change leads to the redistribution of pension wealth and risk between the main groups of participant - the sponsor, active members, deferred members²⁷, pensioners and future members.

When rule changes are proposed, attention usually focuses on the details of these changes such as contribution rates, accrual rates and retirement ages, but with no detailed valuation of the size of the wealth transfer. Almost no explicit consideration is given to the effects of a rule change on the wealth of the different age cohorts, or to the riskiness of this wealth, and these can be substantial. Therefore an important objective of this paper is to stimulate a greater awareness of the redistributive effects on wealth and risk of pension scheme redesign, particularly the generational effects. While this paper deals with a particular pension scheme and rule change, the methodology can be applied to investigate the redistributive effects of rule changes by other DB schemes where the sponsor remains responsible for meeting the pension promise, as in countries such as the UK and USA. It can also be used to investigate the long run viability of such DB pension schemes.

Previous investigations of the redistribution of pension wealth by rule changes have been of hypothetical schemes. This is the first paper to quantify the redistributive effects of a major package of rule changes by a large real-world DB pension scheme - the UK Universities Superannuation Scheme (USS). Almost all previous studies have been of hypothetical Dutch schemes where the sponsor has no obligation beyond

²⁷Members who are no longer active contributors, but who have not yet retired.

paying a fixed contribution rate. Therefore the sponsor is not involved, and all the redistribution is between different generations of member, i.e. inter-generational redistribution. In 2011 USS was a 'balance of cost' scheme where, unlike Dutch schemes, the sponsor bears the default risk, and so any redistribution of wealth and risk is primarily between the sponsor and members.

To quantify redistribution stemming from the October 2011 rule change, a benchmark must be specified. One possible benchmark is to compute the 'true' funding position of USS in October 2011, and then to distribute any deficit among the sponsor and the cohorts of members and pensioners. However, there would be a considerable degree of uncertainty and subjectivity attached to such a benchmark. In October 2011 USS had a well-defined set of rules, the main features of which had remained unchanged since 1975, when USS began. Therefore a reasonable expectation for members in October 2011 was that the pension promises enshrined in the USS rules would be honoured, and so the benchmark we use is the pre-October 2011 scheme.

This paper incorporates many aspects of the problem not included in previous studies - lump sum payments on retirement, deferred pensioners, limited price indexation, spouses' pensions, increases in the retirement date, both final salary and career revalued benefits (CRB) sections, and consumer price indexation (CPI) of the accrued benefits of the CRB section active members and the accrued benefits of deferred pensioners, as well as pensions in payment. In addition, we compute final salaries using the retail price index (RPI), see Appendix 5.A. This is also the first study of redistribution by a scheme moving to 'cap and share' contribution rates. We

model the pension scheme for longer than a working lifetime to avoid the problem of back-loading, where contributions made when young represent worse value than those made when old²⁸. If the effects of a rule change are quantified for a period shorter than a working lifetime, the presence of back-loading is likely to show that the young receive a less favourable outcome than the old. We also employ a dynamic asset allocation strategy by allowing the asset allocation to respond to the current funding ratio (assets/liabilities), rather than use a fix-mix investment strategy as have most previous studies. With 13 factors the vector auto-regression (VAR) model we use to forecast asset returns and inflation includes many more assets than previous studies, and is only the second study to include the three factors of the yield curve (level, slope and curvature) in the VAR model, rather than selected interest rates. Finally, we model the numbers of new active and deferred scheme members each year as stochastic processes.

Section 5.1 describes USS, and section 5.2 outlines our methodology. Section 5.3 has a literature review, followed in section 5.4 by details of the data and methodology used to forecast the yield curve, asset returns, inflation and academic salaries each period until the horizon date. Section 5.5 contains the procedure for forecasting the size of each age cohort, and section 5.6 explains how the liabilities (i.e. the accrued benefits) of each age cohort are estimated at the end of each period. Section 5.7 then brings together all these forecasts to calculate the triennial values of the USS funding ratio, revisions to the member and sponsor contribution rates, and adjustments to the asset allocation. In section 5.8 these are used to generate the

²⁸ Back-loading occurs when the scheme uses age-independent contribution and accrual rates (as does USS) and the rate of return on the scheme's assets exceeds the rate of salary growth.

cash flows to and from the various participants each time period until the horizon date. The NPVs of these cash flows are valued using stochastic discount factors (SDF) to give the redistribution of wealth generated by the October 2011 rule changes. The results appear in section 5.9, with robustness checks in section 5.10, where the use of riskless discount rates also permits estimates of the changes in risk²⁹. Finally, section 5.11 has the conclusions.

5.1 USS

In 2014 USS was the second largest pension scheme in the UK, and the 36th largest in the world with 316,440 active members, deferred pensioners and pensioners. It is a multi-employer scheme with 374 separate sponsors (or institutions), and assets valued at £42 billion in 2014. Until the rule change implemented in October 2011, USS was an open final salary scheme. In October 2011 USS was split into two sections - a final salary section that was closed to new members in October 2011, and a CRB section, which operates on a career average revalued earnings (CARE) basis, and started operation in October 2011. The rule changes in October 2011 were a matter of heated public controversy between the institutional sponsors of USS, represented by the Employers Pension Forum; and the members and pensioners of USS, represented by the University and College Union (UCU), leading to lengthy industrial action by members of the UCU³⁰.

USS is a very large and complicated scheme with a 295 page rule book, and so any model of USS is bound to be a gross simplification. This study captures the financially

²⁹It is not possible to use SDFs to measure changes in risk.

³⁰No explicit concerns were expressed for the distributional implications of the rule change.

important features of USS, including all the rules that changed. The other important changes implemented in October 2011, besides new members joining the CRB section, were (a) an increase in the contribution rate for the final salary section, (b) the introduction of a 'cap and share' rule for deficits and surpluses, (c) linking the normal retirement age to the state pension age, and (d) limiting the indexation of pensions and deferred pensions. The rules pre and post-October 2011 are set out in Appendix 5.A. This appendix also details some of the other USS rules incorporated in our model, including lump sum payments, spouses' pensions, deferred pensioners, and the computation of final salary. Unchanged rules tend to be less important because they have a similar effect on pension wealth before and after the rule change, and so tend not to create redistribution.

5.2 Methodology

We modelled the effects of the six rule changes as a single package, rather than examining the effects of each rule change separately. This is because we are primarily interested in the effects of the package, the rule changes interact and so the effects of the October 2011 rule changes are only available by treating them as a package, and because repeating the analysis another six times would be a considerable undertaking. The analysis of the redistributive effects of the USS rules change in October 2011 is divided into two main steps. The first step is to model the evolution of USS over the horizon period, permitting forecasts of the cash flows between each age cohort, the sponsors and USS under two alternative sets of rules - those pre and post-October 2011. This will be done using three year time periods, as this is the frequency of USS actuarial valuations and contribution rate reviews. In the

second step, the NPV of the forecast cash flows for each age cohort and the sponsors is computed for both the pre and post-October 2011 rules. This allows the calculation of the NPV of the change in expected pension wealth for each cohort caused by the October 2011 rule changes, which is the standard way of measuring pension redistribution, Bonenkamp (2009).

We concentrate on expected pension wealth, although the October 2011 rule change may well have other effects on members and the sponsor. Pension contributions are an important component of university expenditure. Since the government no longer raises university funding to compensate for increases in the cost of USS, any additional sponsor contributions must be funded by the universities themselves. Apart from raising additional revenue, universities might make cost savings by cutting expenditure on capital projects, reducing salaries or increasing workloads, with an adverse effect on active members. It is also possible that after October 2011 employers used their reduction in pension contributions, relative to the benchmark, to raise salaries to compensate for the drop in expected pension wealth of members. Many empirical studies have tried to quantify the compensating wage differential, i.e. the size of the trade-off between pension benefits and wages, and recent examples of this literature include Disney, Emmerson and Tetlow (2009), Gerakos (2010) and Haynes and Sessions (2013). Attempts to quantify the wage-pension trade-off have encountered substantial econometric and data problems (Allen and Clark, 1987), but subject to these reservations, the empirical evidence suggests the trade-off is well below one-for-one. Consistent with this evidence UK academic salaries have showed no obvious response to the USS rule change of

October 2011³¹. Such consequential effects on the membership such as higher or lower salaries, worse conditions of service etc. are outside the scope of this research.

The members, future members, pensioners and deferred pensioners of the two sections of USS (final salary and CRB) are disaggregated into age-based cohorts, where the age range covered by each cohort is five years. We use five years because this is the period used by USS when they supplied us with some cohort data, and provides computational tractability. For the post-October 2011 final salary section there are eight cohorts each of active members and deferred pensioners aged between 25 and 65 years, and six cohorts of pensioners aged between 65 and 95. In addition, the post-October 2011 CRB section has three cohorts of actives and deferreds aged between 25 and 35 years. This is because new active members (future cohorts) enter directly into the four youngest cohorts each year. The post-October 2011 CRB section also has 11 cohorts of both future actives and deferreds aged from minus 30 to 25 years of age in 2011. This makes a total of 50 cohorts for the post-October 2011 scheme. The continuation of an unchanged pre-October 2011 final salary scheme has a total of 44 cohorts, all of which appear as part of the 50 post-October 2011 scheme cohorts. The other participants in the scheme are the sponsors of USS, i.e. the 374 UK universities and related institutions, who are treated as a single group.

Forecasting asset returns, yield curves, inflation, longevity and salaries for the horizon period (54 years) is a daunting task; as is projecting the membership in each USS age cohort during this period. Therefore the resulting cash flow forecasts are

³¹Between October 2011 and October 2015 academic salary scales rose by only about 1% per year.

inevitably subject to a considerable degree of estimation risk. Because of the heroic forecasts required, a range of financial and demographic forecasts are employed to generate a distribution of outcomes, and the sensitivity of the conclusions to some of the important assumptions is investigated as a robustness check. In contrast to Chapter 3 where out-of-sample tests are conducted to compare different portfolio construction methods in an environment where there is no knowledge of the future as it often happens in practice, the scope of the study presented in this chapter is different and the model requires heroic and even 'dangerous' forecasts of a number of parameters in order the analysis to be conducted.

5.3 Literature Review

In 2001 Chapman, Gordon and Speed suggested taking a much wider view of the effects of changes in pension scheme rules than had previously been the case. They identified six stakeholders who are affected by a change in pension scheme rules - the sponsor's shareholders, the sponsor's debt holders, the employees, externals (the sponsor's suppliers and customers), consultants and advisors, and the government. For a hypothetical UK scheme, they simulated the cash flows between these six stakeholders over a ten year period, and then used SDFs to compute the NPV of the cash flows for each stakeholder. This was done for both a base case and various alternative pension rules, and the average NPV for each set of rules for each group of stakeholders computed. The changes in these averages gave the redistributive effects of the rule change on the wealth of each stakeholder group.

Ponds (2003) proposed using the approach of Chapman, Gordon and Speed (2001) to quantify the intergenerational redistributive effects of different pension scheme

rules. He considered a hypothetical Dutch scheme where the sponsor bears no risk, and analysed redistribution between active members, pensioners and future members; and between age cohorts of these three groups. Using the same methodology, Hoevenaars and Ponds (2007, 2008), Lekniute (2011) and Draper, Van Ewijk, Lever and Mehlkopf (2014) have also illustrated intergenerational redistribution among age cohorts arising from changes in scheme rules for hypothetical Dutch pension schemes; while Hoevenaars, Kocken and Ponds (2009) and Hoevenaars (2011) have investigated redistribution between the sponsor and members (but not between age cohorts of members) for hypothetical Dutch DB schemes. Finally, for a hypothetical US state pension scheme, Lekniute, Beetsma and Ponds (2014) and Beetsma, Lekniute and Ponds (2014) simulated intergenerational redistribution between cohorts of active members and the sponsor (the state's tax payers) due to rule changes.

5.4 Forecasting Asset Returns, Inflation and Salaries

In this section we use the Nelson and Siegel (1987) model to estimate the parameters of the yield curve for the data period (1993-2010)³². We then estimate a VAR(1) model to enable us to forecast asset returns, inflation, and the yield curve for the horizon period. Finally we generate forecasts of the salaries of the various age cohorts of USS active members until the horizon date. In making these forecasts we only use data that would have been available to USS at the time of the rule change.

³²The data we use to estimate the VAR(1) model starts in 1993. This is because earlier data is not available for some of the maturities involved in the estimation of the three Nelson-Siegel yield curve factors: θ_1 , θ_2 and θ_3 . The length of the estimation period we use for the VAR(1) model is in line with that used by Ferstl and Weissensteiner (2011) and Gulpinar and Pachamanova (2013).

A. *Yield Curves.* Diebold and Li (2006) have developed a variation of the Nelson-Siegel model for forecasting yield curves which allows the entire yield curve to be represented by only three parameters:-

$$y_n(t) = \beta_{1,t} + \beta_{2,t} \left(\frac{1 - e^{-\lambda n}}{\lambda n} \right) + \beta_{3,t} \left(\frac{1 - e^{-\lambda n}}{\lambda n} - e^{-\lambda n} \right) \quad (5.1)$$

where $y_n(t)$ denotes the spot rate (zero coupon) at time t for a maturity of n periods, and β_1 , β_2 and β_3 are the level, slope and curvature respectively of the yield curve. Following Diebold and Li, we set the annual decay rate λ in equation (5.1) to 0.1827. We use end-of-quarter yields from 1993 to 2010 for zero coupon UK government bonds with maturities of 3 months, 1 year, 2 years, 3 years, 5 years, 10 years and 20 years. Following Diebold and Li (2006), who estimate the Nelson-Siegel yield curve using data for a 15 year period, we apply linear interpolation to compute the nearby maturities and estimate a time series for each of the three parameters in equation (5.1). These three time series of the parameters of the yield curve are then included in a VAR(1) model to generate forecasts of the yield curve in future years, (Ferstl and Weissensteiner, 2011).

B. *VAR(1) Model.* We use a VAR(1) model to generate the future scenarios, as have Hoevenaars and Ponds (2008), Hoevenaars, Kocken and Ponds (2009), Hoevenaars, Molenaar and Ponds (2010), Hoevenaars (2011), Lekniute (2011), Lekniute, Beetsma, and Ponds (2014). The financial data included in our VAR(1) model consists of quarterly excess returns from 1993 to 2010 for UK equities (FTSE All Share Total Return index), European equities (MSCI Europe excluding the UK Total Return index), US equities (S&P500 Composite Total Return index), hedge funds (HFRI Hedge Fund

index), commodities (S&P GSCI Total Return index), UK property (UK IPD Index Total Return index), together with quarterly values for UK dividends (FTSE All Share Dividend Yield), US dividends (S&P500 Composite DS Dividend Yield), and UK inflation rates (UK RPI and UK CPI). In addition, we include the three estimated parameters of the Nelson-Siegel yield curve factors, β_1 , β_2 and β_3 , in the VAR(1) model in equation (5.2):-

$$\mathbf{x}_{t+1} = \mathbf{c} + \mathbf{B}\mathbf{x}_t + \boldsymbol{\zeta}_{t+1} \text{ where } \boldsymbol{\zeta}_{t+1} \sim N(\mathbf{0}, \boldsymbol{\Sigma}) \quad (5.2)$$

where \mathbf{x}_t is a column vector of economic factors at time t , \mathbf{c} is a column vector of constants, \mathbf{B} is a square matrix of coefficients, $\boldsymbol{\zeta}_{t+1}$ is a column vector of disturbances at time $t+1$, and $\boldsymbol{\Sigma}$ is the variance-covariance matrix of the column vector of disturbances. The estimated VAR(1) model with 13 variables appears in Table 5.1, with the estimated covariance matrix of the disturbances in Appendix 5.D³³. The VAR(1) model shows strong evidence of predictability of stock returns and hedge funds. These findings are in accordance to Berkelaar and Kouwenberg (2010) and Hoevenaars et al. (2008), among others. The largest eigenvalue of the estimated coefficient matrix (\mathbf{B}) is 0.9428, and since this is less than one the system is stable and shocks to the system dampen over time. We also tried including jumps in the VAR(1) model, but the results were inferior.

Following Hoevenaars and Ponds (2008), Hoevenaars, Kocken and Ponds (2009),

³³VAR(1) is preferable to VAR(2) according to the Schwarz criterion for model selection. *F*-tests for the regressions of the 13 variables show that one lag is preferable to two lags for ten of the variables. Finally, the VAR(1) model has a lower maximum eigenvalue than the VAR(2) model. For these reasons we prefer the VAR(1) model to the VAR(2) model. All of the 13 variables are stationary. None of the sets of regression residuals for the 13 explanatory variables displayed serial correlation. Ten of the sets of residuals are normally distributed, with plots of the remaining three sets looking approximately normal.

Hoevenaars, Molenaar and Ponds (2010) and Hoevenaars (2011), we forward iterate this model for the out-of-sample period to produce 5,000 sets of forecasts (i.e. scenarios) of asset returns, inflation rates, and the yield curve until the horizon date³⁴. To produce a scenario we generate an \mathbf{x} vector at time $t+1$ (\mathbf{x}_{t+1}^*) for each out-of-sample period, where the superscript * indicates a value estimated using equation (5.3):-

$$\mathbf{x}_{t+1}^* = \mathbf{c} + \mathbf{B}\mathbf{x}_t^* + \boldsymbol{\zeta}_{t+1}^* \quad (5.3)$$

The value of the $\boldsymbol{\zeta}_{t+1}^*$ vector each period is generated by Monte Carlo simulation using the estimated multivariate normal distribution of the disturbances in equation (5.2). We computed the maximum and minimum yields for each maturity (3 months to 20 years) for the Nelson-Siegel yield curves estimated in section 5.4A. The maximum yield for all maturities is almost flat at 8.5%, while the minimum yield rises in a more or less linear manner from 0.4% for 3 months to 3.8% for 20 years. When generating future scenarios we impose these upper and lower bounds on the forecast yield curves.

C. Salaries. Salaries rise for two reasons - general increases in the salary scale (S_s), and incremental pay rises (S_i) as active members age or are promoted. The rise in salaries for an age cohort of active members is the product of these two sources of wage rises, i.e. $(1+S_i)(1+S_s)$. The USS scheme actuary estimates general salary increases as the forecast rate of RPI inflation plus one percent (USS, 2014). Table E in Higher Education Statistics Agency (2014) provides academic salary levels for each age cohort in 2012-13. We use this data to compute the relationship between age

³⁴In a similar context Chen, Pelsser and Ponds (2014) also used 5,000 scenarios.

and salary for UK academics, with the slope of this curve giving the rate of incremental pay rises for different age cohorts. Together with the RPI inflation forecasts from the VAR(1) model, this allows S_s and S_i to be forecast until the horizon. USS (2014) gives the total pensionable salaries of active members in 2011 as £5.845 billion, and the number of active members in 2011 as 139,931. This implies an average salary of £41,771 for USS active members in 2011, and we use this to calibrate the cohort salary data from HESA (2014) to match the USS active membership average.

5.5 Forecasting the Size of the Age Cohorts

The forecasts of the numbers of actives, deferreds and pensioners in each age cohort for the pre and post-October 2011 schemes use membership data from the USS annual reports, as well as data for 2014 on the size of each age cohort in the final salary and CRB sections supplied directly to us by USS. Changes in the total number of active members of USS for the years 1997 to 2014 are regressed on a time dummy. By modelling the total number of active members we allow for early leavers and late joiners. When computing the pension wealth effects of the rule changes on each age cohort of active members, we assume they expect to stay in USS until retirement. The estimated slope of this regression is zero, with a highly significant constant term of 5,050, indicating that USS total active membership is increasing by 5,050 per year. We use the residuals from this regression to estimate the standard deviation of annual changes in the total active membership of the pre-October 2011 scheme at 2,366. When forecasting the annual increase in the total active membership we choose a value at random from a normal distribution and add it to

the forecast increase to allow for year to year fluctuations. We assume that each year equal proportions of the new active members enter the four youngest age cohorts. This assumption is chosen to ensure that the average distribution of active members by age cohort across all the years until the horizon date approximates the age distribution of active members in 2014 supplied to us by USS.

We follow a similar procedure to forecast the annual changes in the number of deferred pensioners for the pre and post-October 2011 schemes, except that new deferreds enter with a five year lag³⁵. In a regression of changes in the total number of deferred pensioners for 1997 to 2014 on a time dummy the estimated slope coefficient is zero, and the highly significant constant term is 4,292, with a standard deviation of 1,359. The resulting average age distribution of deferred members across all the years until the horizon date approximates the age distribution of deferred members in 2014 supplied to us by USS.

5.6 Forecasting the Liabilities

We use an actuarial model based on Board and Sutcliffe (2007) to forecast the values of the scheme's liabilities for each age cohort until the horizon date, and this requires the specification of a number of parameters, see Appendix 5.B. We are modelling the performance and decisions of USS over the horizon period, where the value of the liabilities is a key input to computing the funding ratio and revising the contribution rates and asset allocation. Therefore, to model the decisions of USS, we need to use the same inputs as USS when valuing the liabilities.

³⁵So we implicitly assume that, on average, deferreds have previously been actives for five years.

We base the life expectancy for each cohort of pensioners on the National Life Tables for the UK, 2011-13 (ONS, 2014), with the number of members of each pensioner cohort reducing each year in accordance with this life table. We generated a blended mortality table using the weights of 55.5% male and 44.5% female, taken from the HESA (2014) data on the gender of academic staff. In 2011 USS members had a blended life expectancy of 25 years at age 65, which incorporates the actuary's estimate of future improvements in USS longevity, USS (2012). As USS members have a greater life expectancy than the general population, the national life tables are uprated by six years, so that at age 65 USS pensioners are expected to live until they are 90 years of age. Following Carnes and Olshansky (2007), we do not allow for any further increases in longevity over the horizon period. For simplicity, those pensioners who reach the age of 90 years are assumed to die at the age of 95, which matches their expected longevity of five years at the age of 90.

Based on USS (2014) we assume that two thirds of pensioners have an eligible beneficiary (usually a spouse) at the time of their death, and that surviving beneficiaries live for another three years. During this time eligible beneficiaries receive a pension equal to half that of the deceased pensioner. So the extra cost of a spouse's pension is approximately equal to $0.667 \times 0.5 \times 3 = 1$ year of the deceased member's pension. To account for this additional liability, we increase longevity by one year.

The computation of the liabilities also requires estimates of the number of accrued years for each age cohort. The age of each cohort is taken as the mid-point of the cohort's age range, and their accrued years are computed from their current age,

assuming they joined USS at an average age of 32.5 years. We then adjust these numbers in 2011 for each cohort to match the average accrued years across all the cohorts of active members of 10.4 years given in USS (2014). On a technical provisions basis, the USS liability for deferred pensions in the 2011 actuarial valuation was £2.792 billion (USS, 2012). We used equation B.5.1 in Appendix 5.B to compute the implied number of accrued years for deferred pensioners at four years, and adjusted the accrued years for deferred members to match the estimated average number of accrued years.

Economic theory indicates that the cash flows in each future year should be discounted using the rate of return on a portfolio that replicates the risk and return of this cash flow, leading to the use of a different discount rate for each year. However, to model USS decisions on contribution rates and asset allocation we need to use the same discount rate as USS, which is the average of the current yield curve for the next 20 years plus 1.7%, and so we do likewise, (USS, 2014). However, when discounting the cash flows to compute the wealth changes for each cohort and the sponsor in section 5.8, we will use SDFs. SDFs have been widely used to compute the Net Present Values (NPVs) of pension scheme cash flows, see for instance section 5.8. Especially, using SDFs to discount the cash flows to compute the wealth changes between the different stakeholders of USS, see section 5.8 for more details, is a separate process from the computation of the pension liabilities. Actuarial liabilities are used to model USS decisions (asset allocation and contribution rates), and the average of the current yield curve with maturity up to 20 years plus 1.7% is used as the nominal discount rate (the same as USS) in the actuarial formulas provided in

Appendix 5.B. Finally the total liabilities computed using our model for 2011 were calibrated to equal their value on a USS technical provisions basis of £35.3437 billion in 2011, with the liabilities for each of the constituent age cohorts correspondingly adjusted, (USS, 2012).

The number of people in each cohort was estimated in section 5.5, and the current salary for the members of each cohort is their initial salary increased by the forecasts of salary growth from section 5.4. Our estimates of salary increases follow the USS methodology. Section 5.4 supplies the forecasts of CPI and RPI, and the estimates of the number of members and their salaries in each cohort are then calibrated to match the 2011 aggregate numbers published by USS (2014).

In contrast to the Dutch research, because the sponsor remains liable for the pension promise, the liabilities for the various cohorts of USS take no account of the overall scheme surplus or deficit. However, in section 5.8 when computing the future cash flows for pensioners and sponsor, a share of the scheme surplus or deficit at the horizon date is allocated to the sponsor.

5.7 Generating the Cohort Cash Flows

The cash flows for each cohort per time period are the pensions and lump sums pensioners receive, less the contributions active members make to the scheme, while the sponsor just pays contributions per time period to the scheme. Contributions to the scheme for each cohort are the number of people in the cohort, times the average cohort salary, times the sum of the current contribution rates for active members and the sponsor. The size and average salary of each cohort were

computed in section 5.4. The contribution rate for members of the final salary section increased from 6.35% to 7.5% in October 2011, while the sponsor's contribution rate remained at 16%. The members' contribution rate for the CRB section is 6.5%, and that for the CRB sponsor is 16%. The total pension payment to each cohort is the number of people in the cohort times their initial pension, adjusted for subsequent limited price indexation, computed using the rules in Appendix 5.A. The lump sum calculation for each cohort also follows Appendix 5.A.

The cash flow calculations use the contribution rates and asset allocation for that period, both of which can change over time in response to the scheme's funding ratio. The liabilities were estimated in section 5.6. The total value of the scheme's assets at the end of each time period is the value of the investments at the start of the period, plus asset returns, the contributions received from the active members and sponsor during the period, less the lump sums and pensions paid out. Asset returns are computed using the forecasts of asset returns in section 5.4. The USS contribution rates and asset allocation are adjusted each period in response to the current value of the funding ratio, whose initial value in 2011 was 92% on a technical provisions basis (USS, 2012).

7A. Adjusting the Contribution Rates. Given the volatility of the funding ratio and the costs of change, we only adjust the contribution rates for the final salary and CRB sections when the funding ratio is below 90% or above 120%. They are adjusted so as to extinguish any surplus or deficit over a 15 year spread period, leading to a funding ratio of unity. For the post-October 2011 scheme, the difference between the final salary and CRB contribution rates for active members remains fixed at 1%.

At the request of USS, Ernst and Young assessed the sponsor's covenant and concluded that the maximum contribution rate the majority of universities can pay is 25% (USS, 2014). Given the 'cap and share' rule, this implies a member contribution rate of 10.27% for members of the final salary section, making a total contribution rate of 35.27%. To prevent the contribution rate reaching unrealistically high levels, we impose an upper bound of 35% on the total final salary contribution rate (34% for the CRB section) for both the pre and post-October 2011 schemes. We also investigate a maximum contribution rate of 29%, or 20.275% for the sponsor and 8.725% for members of the final salary section, (28% for the CRB section) as a robustness check in section 5.10.

7B. Adjusting the Asset Allocation. As well as changing the contribution rates in response to the funding ratio, the asset allocation may also be altered. There are two rival theories of how a scheme's funding ratio and the probability of default affect its asset allocation. The risk management view is that as the probability of default rises, e.g. the funding ratio falls, schemes shift out of high risk assets into low risk assets; while the risk shifting view is the opposite³⁶. The risk management approach is motivated by the view that riskier cash flows can lead financially distressed firms faster or closer to bankruptcy (Rauh (2009)). Hence, the risk management view suggests that pension fund managers can decrease the probability of default by investing into low risk assets instead of riskier assets and hence ensuring sufficient funds to avoid financial distress. In contrast to the risk management view, the risk

³⁶ Addoum, van Binsbergen and Brandt (2010), Amir and Benartzi (1999), An, Huang and Zhang (2013), Anantharaman and Lee (2014), Atanasova and Gatev (2013), Bodie, Light, Morck and Taggart (1985, 1987), Comprix and Muller (2006), Coronado, Liang and Orszag (2006), Friedman (1983), Guan and Lui (2014), LCP (2014), Li (2010), McCarthy and Miles (2013), Mohan and Zhang (2014), Munro and Barrie (2003), Petersen (1996) and Rauh (2009).

shifting view is based on the theory of asset substitution (Jensen and Meckling (1976)), and suggests that pension fund managers can increase the value of pension funds' portfolio and improve the funding ratio by increasing the volatility (risk) of the fund's assets when there is a high probability of default (e.g. low funding ratio), since the newer and riskier investment potentially improves the portfolio return. The empirical evidence on these rival views is mixed, and so alternative sets of results are generated for each of these views. Based on their past volatility, we divide the assets into three groups: high risk (UK, European and US equities and commodities), medium risk (real estate and hedge funds) and low risk (10 year bonds and cash). Adopting a very simple rule, the difference between the current funding ratio and the initial funding ratio in 2011 is used to adjust the money allocated to high and low risk assets, subject to the constraints of not allocating more than the available funds, or negative funds, to the high and low risk assets.

For example, suppose the initial asset allocation in 2011 is 65% to high risk assets, 15% to medium risk assets and 20% to low risk assets; the initial funding ratio in 2011 is 0.80 and that by 2014 this has risen to 0.95. Using the risk management approach the asset allocation to high risk assets rises by $(0.95-0.80) = 0.15$ to 80%, the allocation to low risk assets drops by a corresponding amount to 5%, and the allocation to medium risk assets is unchanged at 15%. The allocations to the individual asset classes within each risk category are rebalanced in proportion to the change in the total funds available for that risk category. A similar rule applies for risk shifting, except the directions of change are reversed. Therefore we investigate the effects of three alternative asset allocation strategies - the actual USS allocation

in 2011 in conjunction with a fix-mix strategy of rebalancing the asset weights back to the initial allocation every three years, risk management and risk shifting. The initial asset allocation in 2011 is UK equities 23.06%, EU equities 18.32%, US equities 18.32%, cash 5%, 10-year UK government bonds 12.3%, UK property 7%, hedge funds 8%, and commodities 8%.

5.8 Redistribution

The objective is to estimate the magnitude of the pension wealth transfer from each cohort of scheme members to the sponsor. The NPV of each series of cash flows is computed using risk-neutral valuation. Risk neutrality is based on the assumption that each share price is exactly equal to the discounted expectation of the share price under this measure, while it has the important advantage that it does not require any knowledge of the risk preferences of the stakeholders of a pension scheme (e.g. members, pensioners and sponsor). For instance, if offered either 100 or a 50% chance each of 200 and 0, a risk neutral person would have no preference, in contrast to a risk averse or a risk seeking person. As a result, SDFs reflect the fact that the present values of cash flows are computed by discounting the future cash flows by the corresponding stochastic factor and then taking the expectation. For this reason SDFs have previously been used to compute the NPVs of pension scheme cash flows by Chapman, Gordon and Speed (2001), Hoevenaars and Ponds (2008), Hoevenaars, Kocken and Ponds (2009), Hoevenaars, Molenaar and Ponds (2010), Hoevenaars (2011), and Draper, Van Ewijk, Lever and Mehlkopf (2014). Lekniute (2011) and Lekniute, Beetsma and Ponds (2014) used risk neutral probabilities to value pension cash flows, which is logically equivalent to using SDFs. We follow these

previous authors in using risk neutral valuation, and use SDFs (or pricing kernels) to compute the NPVs of the cash flows for the member and pensioner cohorts and the sponsor.

In using SDFs to value the pension liabilities we are treating all the liabilities as potentially risky. Although the Pension Protection Fund (PPF) insures roughly 90% of UK DB pension liabilities in the event of default, this does not make USS liabilities riskless. Default by USS requires all UK universities to be in default, and in such circumstances it is likely that the PPF, which has no explicit government guarantee, will also be in default (Blake, Cotter and Dowd, 2007). While pensions in payment have priority in the event of default, we have used the same risky discount rate for all liabilities. Splitting the liabilities into two tranches does not alter the total default risk, although it would lead to slightly higher discount rates for actives and deferreds, and slightly lower discount rates for pensioners. However the effects of such an adjustment on the conclusions would be minimal.

The use of a unique set of positive SDFs to discount stochastic cash flows relies on the assumptions of complete and arbitrage-free markets in which the law of one price applies. Where markets exist, competition tends to ensure an absence of arbitrage opportunities and the validity of the law of one price. But if markets are incomplete, many alternative sets of positive SDFs exist. The valuation of DB pension liabilities faces the problem of a missing market for trading or hedging future salaries, and an imperfect market for hedging longevity. This leads to the contradiction that the use of risk neutral valuation to value pension scheme liabilities relies on complete markets, but if markets were complete pension schemes would

lose their reasons for existence, e.g. risk sharing, economies of scale, low transactions costs, etc. (McCarthy, 2005).

Our model of USS assumes no cohort longevity risk, leaving only diversifiable longevity risk. Since USS has a very large number of members, diversifiable longevity risk is roughly zero, Aro (2014), Donnelly (2014). Future salary increases for USS members are split into two components: a general uplift in the salary scale which is assumed by USS to be RPI inflation plus 1%, and a promotional salary increase which is specific to each age cohort. The general RPI-linked uplift in salaries can be hedged using RPI-linked bonds or RPI-linked swaps, leaving just the promotional increases. Our model assumes that, apart from the annual RPI+1% uplift to all scale points, the salary scale remains constant over time, making the average promotional increase for each age cohort highly predictable. So, while no instrument exists for trading or hedging promotional increases, they are probably low risk, and can more or less be replicated using the riskless asset, making the market approximately complete.

Hoevenaars and Ponds (2008), Hoevenaars, Molenaar and Ponds (2010), Hoevenaars (2011) and Draper, Van Ewijk, Lever and Mehlkopf (2014) simply assumed zero real wage growth to circumvent the incomplete markets problem. This is the same as our model, if promotional salary increases are excluded. De Jong (2008) discusses four methods to value salary-indexed stochastic future cash flows in the presence of incomplete markets, and advocates utility-based valuation assuming that individuals have a specified utility function. We prefer not to assume utility functions for actives, deferreds, pensioners and the sponsor, but to rely on the observed market prices used in the SDF computation, recognising that the assumptions required for the use

of SDFs are not fully met. Pukthuanthong and Roll (2015) have recently found that “the SDF theory’s main prediction, that the same SDF prices all assets during the same time period, cannot be rejected with our tests, data, or time periods. ... These results are consistent with complete markets and an absence of arbitrage.” This empirical finding supports our view that, while the assumptions may not be fully met, SDFs are still a useful way of valuing a sequence of risky cash flows. As a robustness check, in section 5.10 we also compute the NPVs using the forecast riskless rates, relying on the assumption that USS is backed by the UK university system, and so has minimal default risk. Using the riskless interest rate as the discount rate for valuing the liabilities of DB schemes is advocated by Broeders, Chen and Rijsbergen (2013), and has the advantage of allowing us to estimate the risks attached to expected changes in pension wealth.

If markets are complete, the law of one price applies, and current asset prices are arbitrage free; then a unique set of positive state prices exists such that each asset’s current price is the sum of the cash flows from the asset in each future state multiplied by the corresponding state price. This is the fundamental theorem of asset pricing (see, for instance, Cochrane, 2001), and no knowledge of individual preferences is required to compute the state prices. Following Ang, Bekaert and Wei (2008) and Cochrane and Piazzesi (2005); as well as Nijman and Koijen (2006), Hoevenaars and Ponds (2008), Hoevenaars, Kocken and Ponds (2009), Hoevenaars, Molenaar and Ponds (2010) and Hoevenaars (2011) from the pensions literature, we define SDFs (m_{t+1}) as:-

$$-\log(m_{t+1}) = y_t^{3-month} + \frac{1}{2} \boldsymbol{\varphi}_t^T \boldsymbol{\Sigma} \boldsymbol{\varphi}_t + \boldsymbol{\varphi}_t^T \boldsymbol{\zeta}_{t+1} \quad (5.4)$$

where $\zeta_{t+1} \sim N(\mathbf{0}, \Sigma)$ denotes a column vector of disturbances from the VAR(1) model, $y_t^{3\text{-month}}$ is the 3 month UK interest rate at time t estimated using the Nelson-Siegel yield curve (see section 5.4), and φ_t is a column vector of the time-varying prices of risk which is defined as in Cochrane and Piazzesi (2005), see appendix 5.C. The column vector φ_t (prices of risks) has 13 rows as the number of the economic variables used in the VAR(1) model. They are scenario-dependent and hence time-varying. SDFs are also scenario dependent, and for a given scenario the SDF for a cash flow in year k (denoted m_{t+k}^*) is the product of the SDFs for each of the first k years, i.e. $m_{t+k}^* = m_{t+1} \times m_{t+2} \times \dots \times m_{t+k}$.

For every scenario, i.e. each sequence of future returns, salaries, inflation rates, cohort size, and contribution rates until the horizon date, we generate annual cash flows. We then discount these back to the present using the set of SDFs specific to that cash flow sequence to get an NPV for each cohort. Finally, for each cohort we compute the average NPV across all cash flow sequences to place a value on this risky asset.

At the horizon date some age cohorts will still have future cash payments to make or receive, i.e. they are pensioners, actives or deferreds. These terminal obligations, which have not yet become cash flows, must be valued. Rather than forecast the cash flows until all the new joiners in the horizon year have died, i.e. until $(2065+70) = 2135$, we forecast the cash flows for another 25 years until 2090. This allows us to compute the subsequent cash flows for all those who are pensioners in the horizon year, but not for those who are actives or deferreds. To avoid the problem of back-loading we only compute NPVs for cohorts whose members are pensioners or

deceased at the horizon date, and not for cohorts with active or deferred members.

Given the 'cap and share' rule, the terminal surplus or deficit at the horizon date for the post-October 2011 scheme is shared between the active members and the sponsor. The scheme liabilities in 2065 are estimated as the present value in 2065 of the cash flows between 2065 and 2090 for actives, deferreds and pensioners in 2065, plus the USS valuation using Appendix 5.B of the scheme's liabilities to these actives and deferreds in 2090³⁷. Using these liabilities, together with the total value of USS assets in 2065, the sponsor is allocated 65% of this horizon year surplus or deficit. There is no need to allocate the remaining 35% between the active members at the horizon date as the NPVs of their cash flows are not being computed.

For a given set of rules, assumptions and forecasts, the average NPV represents the expected increase or decrease in pension wealth in October 2011 for each age cohort from the continued operation of USS according to a specified set of rules. We compute the wealth effects of the October 2011 rule change by examining differences in the NPVs for the pre and post-October 2011 schemes. This assumes that each of the alternative scheme designs remains unchanged for the horizon period. In the present case, if the pre-October 2011 scheme is not reformed there is an increased risk of financial distress for the sponsor, which may then impact on active and future members in the form of lower salaries, fewer jobs and subsequent scheme redesign. In common with previous studies of the redistributive effects of

³⁷The use of the USS valuation tends to understate the liabilities as it uses a high discount rate. It also ignores subsequent investment, salary, inflation, and contribution rate risk. However, since these cash flows occur at least 79 years in the future, their discounted value will be relatively small.

pension rule changes, we have not valued this risk, and assumed that members expect the chosen scheme to be unchanged.

5.9 Results³⁸

The time series of the forecast mean funding ratios, contribution rates and asset allocations are plotted in figures 5.1 to 5.6. They are interdependent, as the asset allocation and contribution rates are adjusted in response to the current funding ratio. The contribution rates then affect the funding ratio in subsequent periods. Cohorts with a mean age of 42 years and older in 2011 contain only members of the final salary section, while cohorts aged 22 years and younger in 2011 are all members of the CRB scheme. The cohorts initially aged 27, 32 and 37 years contain members of both the final salary and the CRB sections.

For the post-October 2011 scheme figure 5.1 shows that, while the mean funding ratio at first declines to below 80%, it steadily recovers. The improvement in the funding ratio from the mid-2020s onwards is due to the steady shift in the active membership from the final salary section to the cheaper CRB section. By 2053 all the active members are in the CRB section and the funding ratio has stabilized. At this time the long run funding ratio for fix-mix and risk shifting is over 115%, and for risk management it is over 100%. Figure 5.2 indicates that in the long run the pre-October 2011 scheme has an inadequate funding ratio for risk shifting and fix-mix of

³⁸We modelled the problem using Visual Basic for Applications (VBA). On average it takes 50 seconds to run each scenario on a desktop computer with a 3.2 GHz processor, 12 GB of RAM and running in Windows 7 Professional. Therefore it took 17 days of CPU time to run 5,000 scenarios with three different asset-allocation strategies and two robustness checks.

about 80%, and for risk management it is a disastrous 65%.

The results in figures 5.3 and 5.4 for the mean final salary contribution rates lead to broadly similar conclusions to those from figures 5.1 and 5.2. For the post-October 2011 scheme, after a rise to 29%, the total contribution rates for fix-mix and risk shifting steadily decline to a long run rate of below 23%, which is less than the 23.5% rate in 2011. The risk management contribution rate does not drop below 29% until roughly 2030, and then declines to just below 25%. For the pre-October 2011 scheme, risk shifting and fix-mix generate a long run contribution rate of around 27%, and for risk management it is about 30%. These results suggest that the pre-October 2011 scheme was not viable in the long run, irrespective of the asset allocation strategy adopted. Table 5.2 has the means and standard deviations of the funding ratios and contribution rates from 2011 to 2065. Given the considerable allocation to risky assets for the risk shifting and fix-mix strategies, it is not surprising that these two strategies lead both the schemes to higher funding ratios and lower contribution rates. On the other hand, we recognize that the improved performance may be accompanied with more volatile contribution rates and funding ratios across the 5,000 scenarios.

	Fix-Mix		Risk Mgt.		Risk Shifting	
	Post	Pre	Post	Pre	Post	Pre
Mean Funding Ratio	98.52%	80.94%	89.17%	70.96%	100.97%	84.23%
SD of the Funding Ratio	14.66%	5.10%	11.82%	6.80%	15.95%	5.86%
Mean final salary contribution rate	25.21%	27.09%	26.65%	28.76%	24.80%	26.49%
SD of the FS contribution rate	2.28%	1.38%	1.94%	1.67%	2.48%	1.44%
Mean allocation to risky assets	67.70%	67.70%	47.88%	38.38%	62.19%	68.36%
SD of allocation to risky assets	0.00%	0.00%	8.04%	10.28%	6.37%	1.93%
Mean allocation to low risk assets	17.30%	17.30%	37.12%	46.62%	22.81%	16.64%
SD of allocation to low risk assets	0.00%	0.00%	8.04%	10.28%	6.37%	1.93%
Mean asset return	3.87%	3.87%	2.43%	2.09%	4.23%	4.39%
SD of asset return	1.76%	1.76%	1.67%	1.68%	1.82%	1.80%

Table 5.2: Summary Statistics for the Three Asset Allocation Strategies

Figures 5.5 and 5.6 show the very different mean asset allocations for the risk shifting and risk management asset allocation strategies. For risk shifting in figure 5.5, both of the pre and post-October 2011 schemes retain their high allocation to risky assets and low allocation to low risk assets, with the difference being more extreme for the pre-October 2011 scheme³⁹. For risk management in figure 5.6 the

³⁹ Table 5.2 shows that for risk shifting the mean post-October 2011 allocation to risky assets (62.19%) is lower than for the fix-mix strategy. This is because, when the funding ratio is below its initial value of 92%, the risk shifting model seeks to increase the allocation to risky assets. However the risky asset proportion cannot rise above 85%, i.e. not more than 17.3% above its opening level of 92%. So if the funding ratio drops below 74.7% there can be no further increase in the allocation to risky assets, and this lowers the average risky asset proportion so that it is lower than for the fix-mix strategy.

pre-October 2011 scheme moves to a long run allocation of over 50% in low risk assets, and about 35% in risky assets; while the post-October 2011 scheme allocation moves to over 50% in risky assets, and about 35% in low risk assets.

Figures 5.1 to 5.6 present information on the generation of the cash flows until the horizon date, and we can now compute the NPVs of these cash flows for the age cohorts and sponsor. Figure 5.7 shows the mean percentage drop in the NPV for each of the age cohorts due to the rule change. There are separate lines for the CRB and final salary sections, which overlap for the 27 to 37 years cohorts. This shows that young cohorts lose a much higher percentage of their pension wealth than do older cohorts, and that CRB members lose much more than do final salary members when they are the same age. Future members (cohorts 17 and 22) lose about 65% of their pension wealth in 2011, which is the present value of their future cash flows with USS (i.e. member contributions, pension and lump sum), while the youngest final salary members lose about a quarter of their NPV, and pensioners lose virtually nothing. This difference is because in 2011 older members had accrued substantial benefits based on contributions made in previous years. These benefits are legally protected from the rule change, and so their NPVs do not drop, while future members are subject to a much larger NPV drop. In addition, because the NPV computation excludes contributions made before 2011 while the accrued benefits these contributions created are carried forwards, the NPVs for older members are larger than those for younger members. So a given absolute NPV reduction represents a smaller proportionate drop for older members.

We converted the proportionate losses per age cohort in figure 5.7 into approximate

losses per head in monetary terms for both actives ($AL = LY \times AY$) and deferreds ($DL = LY \times DY$), using equation (5.5) to estimate the mean loss of NPV per accrued year (LY):-

$$LY = TL / \{TN[AP(AY) + DP(DY)]\} \quad (5.5)$$

where TL is the mean total loss of NPV for the age cohort, TN is the total number of members of the age cohort, AP is the proportion of the age cohort who are actives, DP is the proportion of the age cohort who are deferreds, AY is the mean number of accrued years at retirement for each active member, and DY is the mean number of accrued years at retirement for each deferred pensioner. Figure 5.8 shows the results for the actives and these have a broadly similar shape to figure 5.7, although figure 5.7 has percentages, and figure 5.8 has losses in £ per head. Figure 5.8 shows that future active members of the CRB scheme lose about £100,000 per head, while the youngest members of the final salary scheme lose £40,000 per head. The corresponding diagram for the deferreds has an identical shape to figure 5.8, but with losses that are $(27.5/32.5) = 84.6\%$ lower.

Figure 5.9 shows the present values of the total loss in £bn. to each cohort of the two sections. The total loss in October 2011 to all the cohorts of the final salary section is £3.87 billion, or 12% of the total loss to all cohorts of both sections; and so 88% of the total loss fell on the CRB section. The total loss for the age cohorts we have analysed (17 to 92) is £18.0 billion, and the total loss for age cohorts -28 to 12 until 2065 is approximately £14.5 billion.

Figure 5.10 plots the mean percentage drop in the pension received at age 65 for each age cohort. Future scheme members experience a drop of about 38% in the

pension they will receive at age 65. Those in the age 42 cohort and older are all in the final salary scheme and will make higher contributions post-October 2011, make contributions for an extra one or two years, and draw their pension for one or two years less due to an increase in their retirement age. They will also benefit from an extra one or two years of additional accrual and salary increases, which will increase their pension at retirement. These factors lead to the total loss of £3.87 bn. for active members of the final salary scheme shown Figure 5.9, but none of these factors alter their pension at age 65. As figure 5.10 shows, the pension at age 65 for members of the final salary scheme is unaffected by the rule changes.

In addition to estimating the monetary loss to the members of each age cohort, we compute the corresponding total monetary gain to the sponsor resulting from the rule change. Since our horizon date is 2065, we quantify the gain for the 2011 to 2065 period. This is the present value of the reduction in the sponsor's contributions to USS until 2065 plus the present value of the sponsor's 65% share of the surplus or deficit in that year. The resulting change in NPV for the sponsor is a gain of £32.5 billion, equivalent to 26% of their pension cost for the pre-October 2011 scheme.

5.10 Robustness Checks

We have previously studied three alternative asset allocation strategies, and we now investigate two additional changes to the base model. So far the sponsor contribution rate for the final salary section has been capped at 25%, but some universities would be unhappy with contributing this much, and so we investigated capping the sponsor's contribution rate at just over 20%, to give a cap on the total final salary contribution rate of 29% (or 28% for the CRB section). For the post-

October 2011 scheme and the fix-mix strategy, the funding ratio drops to about 75% in the medium term, and the final salary contribution rate rises to roughly 26%. In the long run the funding ratio rises to almost 110%, and the final salary contribution rate falls to just above 22%. These results suggest that, even if the total contribution rate is capped at 29%, the post-October 2011 scheme is viable, albeit with an uncomfortable period of low funding before the long run equilibrium is reached.

Another robustness check involves replacing the SDFs with the riskless discount rates from the VAR(1) model when computing the NPVs. We use riskless rates because USS is backed by the UK university system on a last-man-standing basis, making default very unlikely. Figure 5.11 shows the expected percentage drop in the NPVs; with broad agreement between the SDF and riskless rate estimates of the percentage drop in the NPVs of pension wealth. For example, as for SDFs, the drop is 65% for future members. Figure 5.11 also shows the 10% and 90% percentiles of the estimated percentage drop in the NPVs, and these range from about 55% to almost 90% for future cohorts, with progressively less variation for older cohorts. Figure 5.12 shows the mean loss in £ per head using the riskless rate. For future members this is about £90,000, compared with £100,000 for SDFs, and ranges from £40,000 to £155,000. As before, the loss per head declines rapidly with age, and is markedly lower for members of the final salary section.

Using the riskless rates, the sponsor's total gain for the 2011-2065 period is £30.0 billion, compared to £32.5 billion using SDFs, with a 10% percentile of £10.5 billion, and a 90% percentile of £55.7 billion. As for SDFs, the sponsor's total gain using riskless rates is 26% of their costs for the pre-October 2011 scheme. The total loss

for the cohorts aged 17 to 92 is £16.1 billion for riskless rates (compared with £18.0 billion when using SDFs), with 10% and 90% percentiles of £6.6 billion and £27.5 billion respectively; while the total loss for the -28 to 12 cohorts is £13.8 billion (compared with £14.5 billion when using SDFs). Therefore the use of riskless rates to discount the cash flows rather than SDFs, supports the general conclusions reached using SDFs.

Finally Figure 5.13 compares the riskiness (coefficient of variation) of the NPVs computed using riskless rates before and after the rule change. For the younger members of the final salary scheme in October 2011 there is a modest increase in the coefficient of variation of their NPVs with the move to the new scheme, but no change for the older members and pensioners. For future members in 2011 the riskiness of their pension wealth increases by about one third, relative to the risks they would have faced if they had joined the old final salary scheme. The October 2011 rule changes reduce the coefficient of variation for the sponsor by about 10%.

5.11 Conclusions

For members close to retirement the value of their pension wealth may be their largest single asset. So changes in a scheme's rules can have an important effect on a member's total wealth. When redesigning DB pension schemes, modelling the long run effects on the sponsor and members is generally neglected. What is needed is a dynamic long-term model that incorporates the interactions between the funding ratio, contribution rate, asset allocation and asset returns, as well as the differential effects on the various age cohorts. This research has built and estimated such a model for the USS rule change in October 2011. Although we modelled all the rules

that changed, as well as other important rules, the complexity of the problem necessitates ignoring inconsequential rules. It also requires making heroic forecasts of asset returns, salaries, numbers of members and inflation far into the future. The actual situation of USS in 2065 will inevitably be substantially different from the mean forecast of our model, but since pension schemes have very distant horizons, such long term forecasts are necessary when analyzing the effects of a rule change. Therefore the model can only give broad indications, rather than precise estimates. However, when comparing two alternative sets of rules using exactly the same model and forecasts, we have a level playing field.

This is the first such study for a real scheme (USS), and also the first where the sponsor bears all or part of the risk, e.g. 'balance of cost' or 'cap and share'. It is also the first to incorporate a range of real world pension scheme features - lump sums, deferred pensioners, limited price indexation, spouses' pensions, an increase in the retirement age, two sections (final salary and CARE), 'cap and share' contribution rates, and an uncertain number of new members each year. It also examines three different asset allocation strategies - fix-mix, risk shifting and risk management - over a 54 year out-of-sample period, which is long enough to avoid the back-loading problem. This has the advantage that none of the actives in 2011 are still active in 2065, allowing the scheme to reach a new equilibrium by the horizon date.

For both schemes the performance of the risk shifting and fix-mix asset allocation strategies is similar, mainly because fix-mix involves a substantial allocation to risky assets, and both strategies are clearly superior to risk management. The results indicate that in the long run the pre-October 2011 scheme was not viable. Using the

two best asset allocation strategies (risk shifting and fix-mix) the long run funding ratio would be about 80%, and the contribution rate for the final salary scheme around 27%. For the risk management strategy the long run outcomes are markedly worse - a long run funding ratio of 65% and a contribution rate of 30%. The post-October 2011 scheme appears reasonably viable in the long run for the two best asset allocation strategies, with a funding ratio above 115% and a final salary contribution rate of about 23%, which is slightly below the 2011 rate of 23.5%. However, before this long run state is reached, the post-October 2011 scheme experiences funding ratios of 80% and contribution rates of about 29%, which would be problematic.

So the decision to redesign USS in October 2011 was justified, creating a post-October 2011 scheme that appears to be sustainable in the long run, although with medium term difficulties that are gradually solved as the active membership switches from the final salary section to the cheaper CRB section. These results indicate that a further redesign of USS is needed in the medium term to cope with progressively higher contribution rates and lower funding ratios⁴⁰. However, in the long run, when all the active members are in the cheaper CRB section, USS will become a well-funded scheme with a total contribution rate just above the pre-October 2011 value of 22.35%. Subsequent rule changes to deal with the medium term problems will only increase the long run strength of USS. The robustness checks broadly support these conclusions.

⁴⁰ In April 2016 there was another major rule change - the final salary section was closed to future accruals and the CRB section offered to these members. Contribution rates to the CRB scheme increased to 8% for members and 18% for the sponsor, i.e. 26% in total. Pensionable salary for the CRB section was capped (initially £55,000), with earnings above this cap eligible for a new defined contribution section.

The rule change in October 2011 resulted in the transfer of about £32.5 billion of wealth from the members to the sponsor during the 2011 to 2065 period. This is equivalent to about £600 million per year, or over 60% of the sponsor's contribution in 2011 of £938.4 million. The reduction in the present value in 2011 of the sponsor's pension contributions over this period is 26% using either SDFs or riskless rates. The cost of this wealth transfer is very unevenly distributed across the various age cohorts and sections, with the burden rising from near zero for pensioners and those close to retirement in 2011, to about 65% of their pension wealth for future members. Since pensions are deferred pay, this represents a substantial pay cut. Before the October 2011 rule change the total annual contributions to the scheme were 22.35% of salaries, but the above analysis suggests that the long run annual cost of providing this scheme was closer to 27% of salaries, of which 6.35% was paid by the members from their salaries, leaving 20.65% to be paid by the sponsor of this 'balance of cost' scheme. We have estimated that future members have experienced a drop in their pension wealth of 65%, which is equivalent to a drop of approximately $0.65 \times 0.2065 / 1.2065 \approx 11\%$ in their total compensation, or $0.65 \times 0.2065 = 13\%$ in their salaries.

Appendix 5.A - USS Rules Pre and Post-October 2011

The final salary section was closed to new members in October 2011, but remains open for accruals by existing members, while the CRB section has been open to new members and accruals since October 2011.

Rules that Changed in October 2011

1. Contribution Rate - Final Salary. The contribution rate for active members of the final salary section increased from 6.35% to 7.5%.

2. *Contribution Rate – ‘Cap and Share’*. Before October 2011 USS was a ‘balance of cost’ scheme, where the sponsor is ultimately responsible for meeting the pensions promise. Post-October 2011 for both the final salary and CRB sections, any increase in the contribution rate is shared between the sponsor and active members in the proportions 65% to the sponsor and 35% to the members. If there is a large surplus, contribution rates are reduced in the same proportions.

3. *Inflation Indexation for Pensions in Payment*. Until April 2011 RPI was used to fully uprate USS pensions in payment and deferred pensions, but the government changed this to CPI in April 2011. So before October 2011 there was full indexation of pensions in payment using CPI. Post-October 2011 for both the final salary and CRB sections if inflation, as measured by the CPI, is less than 5% there is full indexation. For inflation between 5% and 15% indexation is 5% plus half of the excess over 5%. If inflation is more than 15% indexation is capped at 10%. In periods of negative inflation pensions are not reduced, but no increase is applied. Benefits accrued before October 2011 in the final salary section increase fully in line with official pensions, i.e. uncapped CPI.

4. *Up-rating of the Accrued Benefits of Deferred Pensioners*. Before October 2011 there was no cap on the up-rating of the accrued benefits of deferred pensioners. After the rule change in October 2011 the accrued benefits of deferred pensioners are uprated by CPI, capped at 2.5%.

5. *Normal Retirement Age (NRA)*. Before October 2011 the NRA was 65 years. In October 2011 this was changed so that the USS NRA for the final salary and CRB sections increases with the state retirement age. This will rise to 66 years in about 2020, 67 years in about 2028, and 68 in about 2046.

Some Other Rules Which Did Not Change in October 2011

6. *Accrual Factor*. For both the final salary and CRB sections the accrual rate is $1/80^{\text{th}}$ plus a lump sum of three times the pension. Using a commutation factor of 1:16 to convert the lump sum into a pension, the accrual factor including the lump sum is $1/67.37$ for both sections.

7. *Lump Sum.* On retirement, pensioners can choose to take up to 25% of the value of their pension as a tax free lump sum. Pensioners are assumed to follow the USS default and take three times the annual value of their pension as a lump sum. Their subsequent pension payments are then based on an accrual factor of $1/80^{\text{ths}}$, rather than $1/67.37^{\text{ths}}$.

8. *Revaluation Rate for the CRB Section.* The revaluation rate used to uprate the average salary for active members of the CRB section each year to allow for inflation is the same as the inflation rate used to up-rate pensions in payment.

9. *Pensionable Salary for the Final Salary Section.* The pensionable final salary is the greater of: (a) the member's highest salary for any period of 12 complete months ending on the last day of a month during the last three years before retirement, and (b) the highest yearly average of the total salary of the member for any three consecutive years ending at the end of any month within the last ten years before retirement. Both amounts are increased, except for the last year before retirement, in proportion to any increase in the RPI between that published at the last day of the relevant year and that published at retirement.

10. *Spouses Pensions.* When a pensioner dies their spouse, civil partner or dependant partner (regardless of sex) receives a pension for life. The spouse's or civil partner's pension is $1/160^{\text{th}}$ times pensionable salary at retirement times pensionable service at retirement, plus pension increases from retirement to death. Note that this calculation ignores the actual lump sum chosen by the pensioner, and assumes they took the standard amount of three times their initial pension.

Appendix 5.B.1 Final Salary (FS) Scheme

The actuarial liability for active and deferred members of the cohort x at time t is given by:-

$$L_{A/D,t}^{x,FS} = N_{A/D,t}^{x,FS} \times \left(\frac{P_{A/D,t}^{x,FS} \times S_{A/D,t}^{x,FS}}{A} \right) \times \left\{ \frac{(1 + e_{A/D,t}^{x,FS})}{(1 + h_{A/D,t}^{x,FS})} \right\}^{R_{A/D,t}^{x,FS} - G_{A/D,t}^{x,FS}} \times \left[1 - \left(\frac{1 + h_{A/D,t}^{x,FS}}{1 + p_{A/D,t}^{x,FS}} \right)^{-W_{A/D,t}^{x,FS}} \right] / \left(\frac{1 + h_{A/D,t}^{x,FS}}{1 + p_{A/D,t}^{x,FS}} - 1 \right) \quad (\text{B.5.1})$$

where A is the accrual rate (constant),

$h_{A/D,t}^{x,FS}$ is the annual nominal discount rate at time t ,

$R_{A/D,t}^{x,FS}$ is the forecast retirement age of the active/deferred members in cohort x at time t ,

$G_{A/D,t}^{x,FS}$ is the average age of the active/deferred members in cohort x at time t ,

$W_{A/D,t}^{x,FS}$ is the life expectancy at retirement of the active/deferred members in cohort x at time t ,

$p_{A/D,t}^{x,FS}$ is the annual rate of growth of the price level at time t ,

$P_{A/D,t}^{x,FS}$ is the past years of service of active/deferred members in cohort x at time t ,

$S_{A/D,t}^{x,FS}$ is the annual salary of the active/deferred members in cohort x at time t ,

$e_{A/D,t}^{x,FS}$ is the expected nominal rate of salary growth per annum between time t and retirement of the active/deferred members in cohort x ,

$N_{A/D,t}^{x,FS}$ is the number of the active/deferred members in cohort x at time t .

The actuarial liability for pensioners in cohort x at time t is given by:-

$$L_{P,t}^{x,FS} = N_{P,t}^{x,FS} \times PEN_{P,t}^{x,FS} \times \left\{ \left[1 - \left(\frac{1 + h_{P,t}^{x,FS}}{1 + p_{P,t}^{x,FS}} \right)^{q_{P,t}^{x,FS}} \right] / \left(\frac{1 + h_{P,t}^{x,FS}}{1 + p_{P,t}^{x,FS}} - 1 \right) \right\} \quad (B.5.2)$$

where $N_{P,t}^{x,FS}$ is the number of the pensioners in cohort x at time t ,

$PEN_{P,t}^{x,FS}$ is the annual pension of the pensioners in cohort x at time t ,

$p_{P,t}^{x,FS}$ is the annual rate of growth of the price level at time t ,

$h_{P,t}^{x,FS}$ is the annual nominal discount rate at time t ,

$q_{P,t}^{x,FS}$ is the life expectancy of the pensioners in cohort x at time t .

The total actuarial liability of the FS scheme is given by:-

$$L_t^{FS, Total} = \sum_{x \in Z} (L_{A,t}^{x,FS} + L_{D,t}^{x,FS} + L_{P,t}^{x,FS}) \quad (B.5.3)$$

Appendix 5.B.2 Career Revalued Benefit (CRB) Scheme

The actuarial liability for active and deferred members in cohort x at time t is given by:-

$$L_{A/D,t}^{x,CRB} = N_{A/D,t}^{x,CRB} \times \left(\frac{P_{A/D,t}^{x,CRB} \times S_{A/D,t}^{x,CRB}}{A} \right) \times \left\{ \frac{(1 + e_{A/D,t}^{x,CRB})}{(1 + h_{A/D,t}^{x,CRB})} \right\}^{R_{A/D,t}^{x,CRB} - G_{A/D,t}^{x,CRB}} \times \left\{ 1 - \left[\frac{(1 + h_{A/D,t}^{x,CRB})^{-W_{A/D,t}^{x,CRB}}}{(1 + p_{A/D,t}^{x,CRB})} \right] \right\} / \left\{ \frac{(1 + h_{A/D,t}^{x,CRB})}{(1 + p_{A/D,t}^{x,CRB})} - 1 \right\} \quad (B.5.4)$$

where A is the accrual rate (constant),

$h_{A/D,t}^{x,CRB}$ is the annual nominal discount rate at time t ,

$R_{A/D,t}^{x,CRB}$ is the forecast retirement age of the active/deferred members in cohort x at time t ,

$G_{A/D,t}^{x,CRB}$ is the average age of the active/deferred members in cohort x at time t ,

$W_{A/D,t}^{x,CRB}$ is the life expectancy at retirement of the active/deferred members in cohort x at time t ,

$p_{A/D,t}^{x,CRB}$ is the annual rate of growth of the price level at time t ,

$P_{A/D,t}^{x,CRB}$ is the past years of service of active/deferred members in cohort x at time t ,

$S_{A/D,t}^{x,CRB}$ is the average annual revalued earnings of the active/deferred members in cohort x at time t ,

$e_{A/D,t}^{x,CRB}$ is the expected nominal rate of salary growth per annum between time t and retirement of the active/deferred members in cohort x ,

$N_{A/D,t}^{x,CRB}$ is the number of the active/deferred members in cohort x at time t .

The actuarial liability for the pensioners in cohort x at time t is given by:-

$$L_{P,t}^{x,CRB} = N_{P,t}^{x,CRB} \times PEN_{P,t}^{x,CRB} \times \left\{ \left[1 - \left(\frac{1 + h_{P,t}^{x,CRB}}{1 + p_{P,t}^{x,CRB}} \right)^{q_{P,t}^{x,CRB}} \right] / \left(\frac{1 + h_{P,t}^{x,CRB}}{1 + p_{P,t}^{x,CRB}} - 1 \right) \right\} \quad (B.5.5)$$

where $N_{P,t}^{x,CRB}$ is the number of the pensioners in cohort x at time t ,

$PEN_{P,t}^{x,CRB}$ is the annual pension of the pensioners in cohort x at time t ,

$p_{P,t}^{x,CRB}$ is the annual rate of growth of the price level at time t ,

$h_{P,t}^{x,CRB}$ is the annual nominal discount rate at time t ,

$q_{P,t}^{x,CRB}$ is the life expectancy of the pensioners in cohort x at time t .

Finally, the total actuarial liability of the CRB scheme is given by:-

$$L_t^{CRB, Total} = \sum_{x=\mathcal{X}} (L_{A,t}^{x,CRB} + L_{D,t}^{x,CRB} + L_{P,t}^{x,CRB}) \quad (B.5.6)$$

Appendix 5.C Time Varying Price of Risk

Following Cochrane and Piazzesi (2005, equation 8) and other studies, we compute the column vector $\boldsymbol{\varphi}_t$ for equation (5.4) using the following expression:-

$$\boldsymbol{\varphi}_t = \boldsymbol{\Sigma}^{-1} \left[\mathbf{c} + \frac{1}{2} \text{diag}(\boldsymbol{\Sigma}) \right] + \boldsymbol{\Sigma}^{-1} \mathbf{B} \mathbf{x}_t \quad (C.5.1)$$

If investors are risk neutral, an absence of arbitrage opportunities requires the spot rate expected next period to equal the implied forward rate for next period. Following Hoevenaars (2011), Hoevenaars, Molenaar and Ponds (2010), and Hoevenaars and Ponds (2008), we set the implied forward interest rate next period ($t+1$) (i.e. the first two right hand terms in equation 5.4) equal to the spot interest rate for next period from the VAR(1) model. The parameters $\boldsymbol{\Sigma}$ and \mathbf{B} come from the VAR(1) model in equation (5.2), while the column vector \mathbf{x}_t contains the state variables of the VAR(1) model at time t .




Figure 5.1: Post-October 2011 Scheme Mean Funding Ratios for the Three Asset Allocation Strategies




Figure 5.2: Pre-October 2011 Scheme Mean Funding Ratios for the Three Asset Allocation Strategies




Figure 5.3: Post-October 2011 Mean Contribution Rates for the Three Asset Allocation Strategies




Figure 5.4: Pre-October 2011 Mean Contribution Rates for the Three Asset Allocation Strategies




Figure 5.5: Mean Risk Shifting Asset Allocation for the Pre and Post-October 2011 Schemes




Figure 5.6: Mean Risk Management Asset Allocation for the Pre and Post-October 2011 Schemes




Figure 5.7: Percentage Drop in the NPV for Each Age Cohort Due to the Rule Change Using SDF




Figure 5.8: £s Loss Per Head for Actives in Each Age Cohort Using SDF

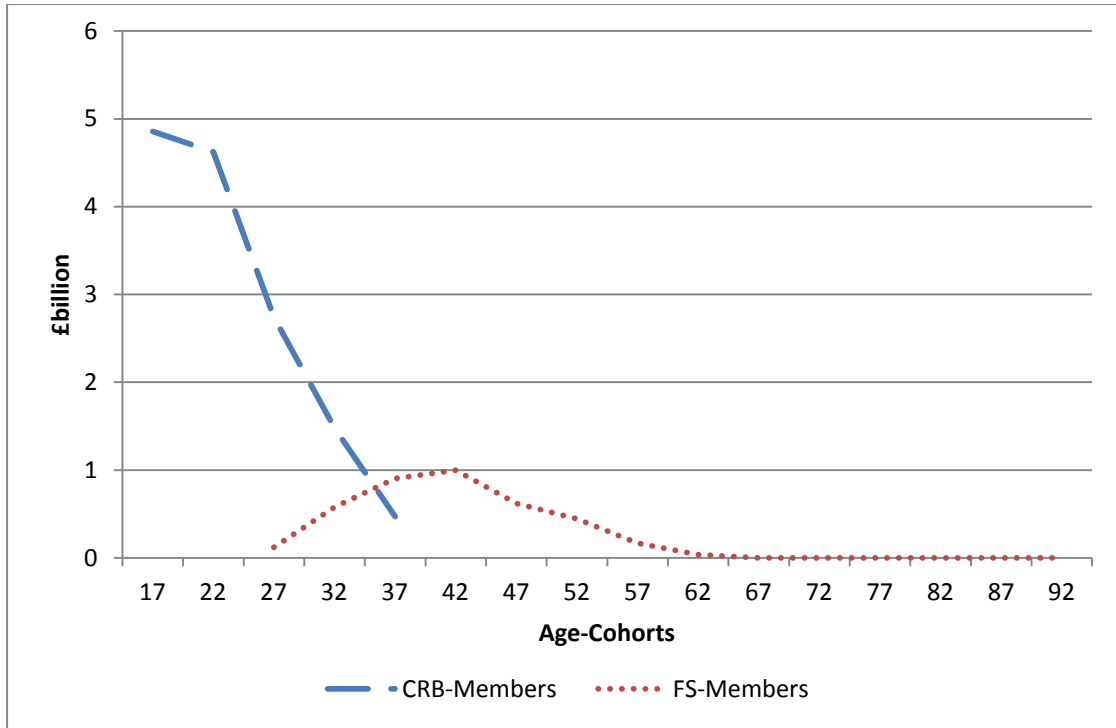


Figure 5.9: Losses Per Age Cohort in £bn. Using SDF




Figure 5.10: Mean Percentage Drop Per Head for Actives in Pension Received at Age 65.




Figure 5.11: Mean Percentage Drop in the NPV for Each Age Cohort Due to the Rule Change Using Riskless Rates




Figure 5.12: Mean £s Expected Loss Per Head for Actives in Each Age Cohort Using Riskless Rates

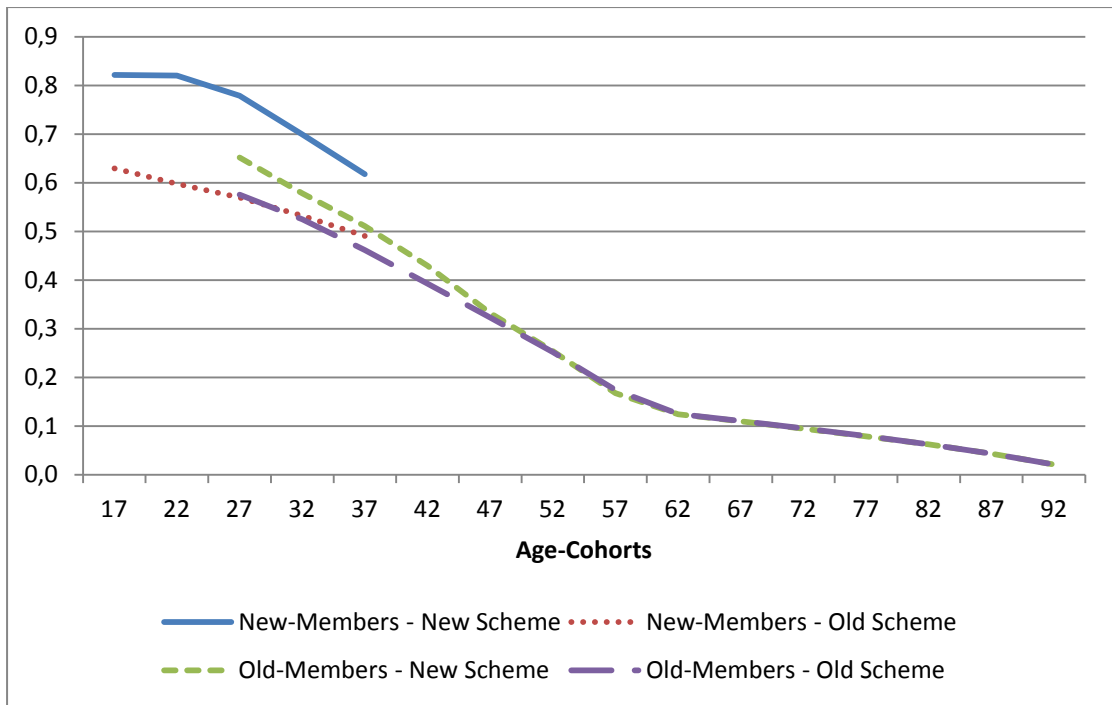


Figure 5.13: Coefficient of Variation of the NPVs of the Pre and Post-October 2011 Schemes Using Riskless Rates

Independent Variables	Dependent Variables												
	UK Equities	EU Equities	US Equities	β_1	β_2	β_3	RPI	CPI	US Div. Yield	UK Div. Yield	Real Estate	Hedge Funds	Commodities
UK Equities-1	-0.3612	-0.1091	-0.1422	-0.0018	0.0109	0.0214	0.0195	0.0139	0.0039	-0.0034	-0.0286	-0.0981	-0.3404
EU Equities-1	-0.0713	-0.2966	-0.0390	-0.0014	0.0046	-0.0118	0.0130	0.0100	0.0016	0.0053	0.0288	0.0035	-0.0884
US Equities-1	0.4112	0.2087	0.0848	-0.0009	-0.0028	-0.0206	-0.0119	-0.0054	-0.0029	0.0001	-0.0086	0.0867	0.3040
β_1 -1	-0.4952	-0.7877	0.0688	0.9448***	0.0285	0.2721	-0.0002	0.0120	-0.0059	0.0057	-0.3170	-0.4678	-0.9858
β_2 -1	1.4921	2.3739**	1.0690	0.0280	0.8903***	0.1488	-0.0556	-0.0452	-0.0212	-0.0685**	-0.3659	-0.2860	-1.8285
β_3 -1	0.1642	1.0112*	0.4500	0.0032	0.1225***	0.7355***	0.0095	-0.0481	-0.0148*	-0.0136	-0.1362	0.3279	0.4958
RPI-1	-2.4706	2.1284	0.2694	-0.1250	0.6941**	-1.0608	0.4349	0.1389	0.0375	-0.0509	-1.0724	-2.7645	0.6554
CPI-1	2.3009	-1.5969	1.6955	-0.0677	-0.3465	-0.1462	-0.8695**	-0.5552**	-0.0547	0.0524	0.6412	3.1756	-0.2533
US Div. Yield-1	-2.2987	-0.0907	-6.4177	0.5929**	-0.7719*	0.6549	0.7525*	0.8630**	1.1043***	0.4032**	-3.0353**	-3.2437	3.2957
UK Div. Yield-1	8.5757**	11.1708**	10.2069**	-0.2728	0.5048*	-0.7728	-0.4421	-0.5186**	-0.1840**	0.4519***	2.0274*	3.9016*	-6.2207
Real Estate-1	1.3405***	1.6201***	1.2460***	0.0227	-0.0048	0.13058	0.0093	-0.0228	-0.0191***	-0.0335***	0.7855***	0.5234***	-0.1845
Hedge Funds-1	0.0461	0.3166	-0.0197	-0.0008	-0.0126	0.0248	-0.0070	-0.0119	-0.0026	-0.0137	0.0670	0.0706	0.9494*
Commodities-1	-0.0569	-0.1527	-0.1455	0.0033	-0.0054	-0.0162	0.0028	0.0058	0.0027	0.0043	-0.0617**	-0.0368	-0.0026
Constant	-0.1815**	-0.2814***	-0.2174***	0.0023	-0.0080	0.0098	0.0089	0.0064	0.0037***	0.0083***	0.0033	-0.0467	0.1586
R ²	0.3274	0.3435	0.3005	0.9543	0.9268	0.6950	0.3788	0.3488	0.9246	0.8813	0.7435	0.3276	0.2002
Adj. R ²	0.1740	0.1938	0.1410	0.9438	0.9101	0.6255	0.2372	0.2003	0.9074	0.8543	0.6850	0.1743	0.0178
F statistic	2.1339	2.2941	1.8840	91.4837	55.4779	9.9930	2.6740	2.3484	53.7874	32.5666	12.7083	2.1364	1.0976

Table 5.1: VAR(1) Model Used to Generate the Forecasts

***, ** and * represent significance at the 1%, 5% and 10% levels respectively

	UK Equities	EU Equities	US Equities	β_1	β_2	β_3	RPI	CPI	US Div. Yield	UK Div. Yield	Real Estate	Hedge Funds	Commodities
UK Equities	0.004450	0.004534	0.004095	0.000034	-0.000011	0.000173	0.000046	0.000004	-0.000067	-0.000131	0.000318	0.001792	0.001023
EU Equities	0.004534	0.006844	0.004626	0.000039	0.000007	0.000181	0.000076	-0.000012	-0.000079	-0.000132	0.000427	0.002192	0.001609
US Equities	0.004095	0.004626	0.004966	0.000031	0.000005	0.000140	0.000090	0.000021	-0.000082	-0.000117	0.000155	0.001773	0.001213
β_1	0.000034	0.000039	0.000031	0.000010	-0.000013	0.000023	-0.000000	-0.000000	-0.000000	-0.000001	0.000006	0.000017	0.000052
β_2	-0.000011	0.000007	0.000005	-0.000013	0.000026	-0.000042	0.000006	0.000002	-0.000001	0.000002	-0.000007	-0.000001	0.000047
β_3	0.000173	0.000181	0.000140	0.000023	-0.000042	0.000227	-0.000004	-0.000001	-0.000002	-0.000005	0.000115	0.000121	0.000038
RPI	0.000046	0.000076	0.000090	-0.000000	0.000006	-0.000004	0.000029	0.000021	-0.000002	0.000002	0.000003	0.000051	0.000277
CPI	0.000004	-0.000012	0.000021	-0.000000	0.000002	-0.000001	0.000021	0.000020	-0.000001	0.000003	-0.000005	0.000014	0.000209
US Div. Yield	-0.000067	-0.000079	-0.000082	-0.000000	-0.000001	-0.000002	-0.000002	-0.000001	0.000002	0.000002	-0.000006	-0.000031	-0.000032
UK Div. Yield	-0.000131	-0.000132	-0.000117	-0.000001	0.000002	-0.000005	0.000002	0.000003	0.000002	0.000005	-0.000015	-0.000048	0.000038
Real Estate	0.000318	0.000427	0.000155	0.000006	-0.000007	0.000115	0.000003	-0.000005	-0.000006	-0.000015	0.000390	0.000225	0.000364
Hedge Funds	0.001792	0.002192	0.001773	0.000017	-0.000001	0.000121	0.000051	0.000014	-0.000031	-0.000048	0.000225	0.001424	0.001159
Commodities	0.001023	0.001609	0.001213	0.000052	0.000047	0.000038	0.000277	0.000209	-0.000032	0.000038	0.000364	0.001159	0.011842

Appendix 5.D: Covariance Matrix of the VAR(1) Model Residuals

Chapter 6

6 Conclusions and Future Research

6.1 Summary and Conclusions

The main purpose of the thesis is to examine three different but interconnected problems, which are presented in Chapters 3, 4 and 5, and lie in the broad area of portfolio and investment management for long-term institutional investors with the main emphasis on pension funds. Initially, Chapter 1 (Introduction) describes the framework and structure of the thesis and highlights the scientific contributions and the most important outcomes of each of the main chapters.

Chapter 2 is a literature review and its main goal is to help the reader get familiar with the fundamental aspects of portfolio theory, asset liability management (ALM) modelling and pension schemes design – the three most important elements of the Chapters that follow the second Chapter. Specifically, in Chapter 2 there is discussion of Markowitz (1952) portfolio theory and the most important criticisms of it (e.g. high sensitivity to the input parameters). Furthermore, Chapter 2 provides a comprehensive review of the latest portfolio techniques that try to address the estimation risk in the input parameters and construct portfolios with better out-of-sample performance. Chapter 2 also describes the close relation between Operations Research (OR) and ALM modelling, and explains the usefulness of asset-liability management models for long-term investors such as pension funds, insurance companies and endowments. It also describes the most popular methods used to solve the ALM problem such as stochastic programming, portfolio theory, stochastic simulation and dynamic programming, as well as the computational issues

(computational intractability) the majority of the existing techniques face. Chapter 2 provides a detailed description of the most important types of pension schemes in the UK and US such as defined benefit and defined contribution pension schemes. In addition it gives a comprehensive review of the methodology and the corresponding outcomes of the most recent pension studies that quantify the redistributive effects between the different stakeholders of a pension scheme due to various pension rule changes.

Chapter 3 deals with the application of a computationally tractable pension asset-liability management (ALM) model, which is based on robust optimization techniques, to a real world pension scheme (USS). Robust optimization is a numerical approach that takes into account the worst-case value of the stochastic input parameters within their uncertainty sets by adopting a maximin approach in the optimization process. Portfolio optimization models based on robust optimization methods are computationally tractable and can more easily solve large scale optimization problems in comparison to other popular asset-liability management model techniques that are based on the generation of future scenarios such as stochastic programming and stochastic simulation. Hence robust optimization is particularly well suited to solving the ALM problem. This study is also unusual since it considers three different types of pension liabilities for active members, deferred members and pensioners and uses the Sharpe ratio as the objective function in the optimization process. Using four different benchmarks; the Sharpe and Tint, Bayes-Stein and Black-Litterman models, as well as the actual USS investment decisions, the proposed robust formulated ALM model has a clearly better out-of-sample

performance over 20 portfolio metrics that measure various important portfolio characteristics (risk, risk-adjusted performance, stochastic dominance, diversification, stability, contribution rate, funding ratio and cumulative wealth, amongst others), and six robustness checks such as a different set of estimation periods, a relaxation of the constraints on the asset weights, a different set of asset classes and the use of uncertainty sets with smaller size, amongst others.

Chapter 4 considers the construction of socially responsible investment portfolios. A limited number of studies try to construct optimal SRI portfolios, and the majority of the existing studies are mainly based on the Markowitz (1952) mean-variance portfolio optimization process that ignores the serious negative effects of estimation risk in the input parameters. Chapter 4 explores whether the selection of the optimization method matters in the SRI industry and attempts to give some answers to the question of which portfolio construction techniques tend to create superior socially responsible investment portfolios. More precisely, three 'formal' portfolio optimization techniques (Markowitz mean-variance, norm-constrained and Black-Litterman based portfolios) and three more simplistic portfolio asset allocation methods (1/N, risk-parity and reward-to-risk timing portfolios) are applied to the same SRI-screened universe of US stocks using 14 performance metrics that measure a variety of different important portfolio characteristics such as risk, risk-adjusted performance, diversification and stability. The out-of-sample performance evaluations show that SRI portfolios based on more formal optimization techniques (Markowitz, norm-constrained and Black-Litterman portfolios) are less risky and have superior risk-return trade-offs to SRI portfolios constructed with the less quantitative

portfolio techniques (1/N, risk-parity and reward-to-risk timing portfolios). SRI portfolios based on the Black-Litterman portfolio technique usually have the best performance, in contrast the equally-weighted portfolio strategy (1/N) which usually has the worst performance on these criteria. Finally, the main conclusions drawn of Chapter 4 remain unchanged by various robustness checks such as the use of different estimation windows, the employment of more demanding criteria for the construction of SRI portfolios as well as alternative evaluation of portfolio performance.

The main objective of Chapter 5 is to investigate the long term performance of a real world pension scheme (USS) before and after the recent pension rule changes of October 2011 (pre-October and post-October 2011 scheme), and to quantify the redistributive effects between the different stakeholders of the pension scheme. In October 2011 USS closed the final salary (FS) scheme, where the sponsor bears all the risks such as investment, longevity, interest rate, inflation, salary growth and regulatory risk, and opened new a career average revalued earnings (CARE) scheme for the new members, and introduced a 'cap and share' rule for pension contributions. This study also incorporates many significant features not previously mentioned in the relevant pension literature such as lump sum payments, deferred members, spouses' pensions, a pension scheme that contains both a final salary and a CARE section, as well as an increasing retirement age over time, amongst others. Furthermore, the model applies three asset allocation strategies (fixed-mix, risk-shifting and risk management) and a VAR(1) model with 13 state variables is applied

to generate future asset returns, inflation rates and the factors of the Nelson-Siegel yield curve.

The corresponding empirical analysis shows that the pre-October 2011 scheme is not viable in the long run, in contrast to the post-October 2011 scheme that seems to be sustainable in the long term, although with some problems in the middle run, while the fixed-mix and risk-shifting approaches are preferable to risk management for both the schemes. The quantitative analysis of the redistributive effects shows that future members lose about 65% of their pension wealth, with an increase in the corresponding risk of a third, while the older members lose nothing with an insignificant increase in risk. The costs associated to the sponsor decrease by about a quarter. Finally, the core conclusions drawn in Chapter 5 remain the same when conducting different robustness checks such as the use of riskless discount factors instead of SDFs and the employment of a different upper bound on the contribution rate.

6.2 Suggestions for Future Research

Inevitably, research is always subject to further improvement and extensions. Hence, the second half of this chapter is devoted to briefly provide some research directions for potential future research.

Although stochastic programming is one of the most popular techniques for solving ALM problems, the robust asset-liability management model formulated in Chapter 3 is benchmarked against ALM models that are mainly based on portfolio theory techniques such as the Sharpe and Tint, Bayes-Stein and Black-Litterman portfolios,

as well as the actual USS performance. As discussed in Chapter 3, stochastic programming faces problems of computational intractability and cannot be used to solve the USS problem presented in Chapter 3. However, for someone who may be interested in a comparison between stochastic programming and robust optimization, future research could set up a small hypothetical example with a limited number of time periods for portfolio evaluation, and a pension fund that invests in a significantly smaller number of assets than those used in Chapter 3. However, it should be stressed again that stochastic programming cannot efficiently solve real-world ALM models and only comparison between the two methods can only be carried out in a non realistic environment.

Chapter 3 uses five asset classes; equities (UK, EU and US equities), bonds (UK and US government bonds), UK property, alternative investments (commodities and hedge funds) and cash. As an additional analysis, we also replaced the 20 year UK and US government bonds with UK private equity and UK infrastructure, and the S&P GSCI total return index with the S&P GSCI Light Energy return index as a robustness check. One could argue that the inclusion of futures and options positions could enhance the risk-return trade-off by hedging various positions and deal with risks on the liability side, such as sudden outflows. This study does not use futures and options because the Universities Superannuation Scheme is a long term institutional that does not consider short term profit, and hence futures and options are not used for speculative reasons. However, we recognize that the use of futures and options by pension schemes is a developing area, and future research could explore the

benefits from long term hedge positions in the Over-the-Counter (OTC) markets (e.g. interest rate swaps).

In Chapter 4, three different optimal diversification and three alternative naive diversification approaches are applied to the same SRI-screened universe subject to non-negativity constraints on the asset weights to rule out short selling. Although SRI portfolios mainly belong to the mutual fund category and hence non-negativity constraints on the portfolio weights should automatically be applied as in Chapter 4, Drut (2012), Ballestro et al. (2012) and Utz et al. (2014) allow short sales in the SRI context. Future research could investigate the effectiveness of short sales restrictions in the process of constructing optimal SRI portfolios, in line with Board and Sutcliffe (1994).

The empirical results in Chapter 4 also show that SRI portfolios based on the Black-Litterman approach have a better out-of-sample performance than the other portfolio techniques. However, there are more advanced versions of Black-Litterman than the one used in Chapter 4, and future research could be carried out to investigate the actual performance of more sophisticated versions of the Black-Litterman approach (e.g. Bessler et al. (forthcoming)) on the construction of SRI portfolios, see for instance Bessler, Opfer and Wolff (forthcoming).

Chapter 5 investigates the performance of USS before and after the recent pension rule changes of October 2011 and quantifies the corresponding redistributive effects between the various stakeholders of the scheme. Although the post-October 2011 scheme seems to be viable in the long run and has a superior performance to the pre-October 2011 scheme in terms of the funding ratio, it faces medium-term

problems due to the frozen defined benefit scheme that will continue to serve the existing members for many decades in the future. For this reason, pension trustees plan to implement more reforms to enhance the sustainability of USS, e.g. by transferring the members that are currently served by the defined benefit (DB) scheme to a career average revalued earnings (CARE) scheme and diversifying more their portfolio by adding more asset classes. As a result, future research could consist of the implementation of additional reforms for the sustainability of the scheme, as well as the investigation of how additional asset classes (e.g. alternatives) can add value in the portfolio and improve the mid and long term performance of the scheme.

References

- Abbott, W. F. and Monsen, R. J. (1979) On the Measurement of Corporate Social Responsibility: Self-Reported Disclosures as a Method of Measuring Corporate Social Involvement, *Academy of Management Journal*, vol. 22, pp. 501-515.
- Addoum, J.M., Van Binsbergen, J.H. and Brandt, M.W. (2010) Asset Allocation and Managerial Assumptions in Corporate Pension Plans, Working paper, Duke University.
- Agarwal, V. and Naik, N. Y. (2004) Risks and Portfolio Decisions Involving Hedge Funds, *Review of Financial Studies*, vol. 17, no. 1, pp. 63-98.
- Allen, S.G. and Clark, R.L. (1987) Pensions and Firm Performance. In *Human Resources and the Performance of the Firm*, edited by M.M. Kleiner, R.N. Block, M. Roomkin and S.W. Salsburg, Industrial Relations Research Association, pp. 195-242.
- Amir, E. and Benartzi, S. (1999) Accounting Recognition and the Determinants of Pension Asset Allocation, *Journal of Accounting, Auditing and Finance*, vol. 14, no. 3, Summer, pp. 321-343.
- An, H., Huang, Z. and Zhang, T. (2013) What Determines Corporate Pension Fund Risk-Taking Strategy?, *Journal of Banking and Finance*, vol. 37, no. 2, February, pp. 597-613.
- Anantharaman D. and Lee, Y.G. (2014) Managerial Risk Taking Incentives and Corporate Pension Policy, *Journal of Financial Economics*, vol. 111, no. 2, February, pp. 328-351.
- Anderson, R.M., Bianchi, S.W. and Goldberg, L.R. (2012) Will My Risk Parity Strategy Outperform?, *Financial Analysts Journal*, vol. 68, no. 6, pp. 75-93.
- Ang, A., Bekaert, G. and Wei, M. (2008) The Term Structure of Real Rates and Expected Inflation, *Journal of Finance*, vol. 63, no. 2, April, pp. 797-849.
- Ang, A., Chen, B. and Sundaresan, S. (2013) Liability-Driven Investment with Downside Risk, *Journal of Portfolio Management*, vol. 40, no. 1, pp. 71-87.
- Ang, A. (2014) *Asset Management: A Systematic Approach to Factor Investing*, Oxford University Press.
- Ardia, D. and Boudt, K. (2015) Implied Expected Returns and the Choice of a Mean-Variance Efficient Portfolio Proxy, *The Journal of Portfolio Management*, vol. 41, no. 4, pp. 68-81.
- Aro, H. (2014) Systematic and Nonsystematic Mortality Risk in Pension Portfolios, *North American Actuarial Journal*, vol. 18, no. 1, pp. 59-67.

Atanasova, C. and Gatev, E. (2013) Pension Plan Risk-Taking: Does it Matter if the Sponsor is Publicly-Traded? *Journal of Pension Economics and Finance*, vol. 12, no. 2, April, pp. 218-249.

Auer, B. R. (2015) Does the Choice of Performance Measure Influence the Evaluation of Commodity Investments?, *International Review of Financial Analysis*, vol. 38, pp. 142-150.

Bacon, C. *Practical Portfolio Performance Measurement and Attribution*, 2nd edition, Wiley 2008.

Bader, L.N. (2003) The Case Against Stock in Corporate Pension Funds, *Pension Section News*, no. 51, February, pp. 17-19.

Baker, M., Bradley B. and Wurgler, J. (2011) Benchmarks as Limits to Arbitrage: Understanding the Low Volatility Anomaly, *Financial Analysts Journal*, vol. 67, pp. 40-54.

Ballesteros, E., Bravo, M., Pérez-Gladish, B., Arenas-Parra, M. and Plà-Santamaria, D. (2012) Socially Responsible Investment: A Multicriteria Approach to Portfolio Selection Combining Ethical and Financial Objectives, *European Journal of Operational Research*, vol. 216, pp. 487–494.

Barnett, M. L. and Salomon, R. M. (2006) Beyond Dichotomy: The Curvilinear Relationship between Social Responsibility and Financial Performance, *Strategic Management Journal*, vol. 27, pp. 1101-1122.

Becker, F., Gürtler, M. and Hibbeln, M. (2015) Markowitz versus Michaud: Portfolio Optimization Strategies Reconsidered, *The European Journal of Finance*, vol. 21, no. 4, pp. 269-291.

Beetsma, R. M., Lekniute, Z. and Ponds, E. H. (2014) Reforming American Public-Sector Pension Plans: Truths and Consequences, *Rotman International Journal of Pension Management*, vol. 7, no. 2, pp. 66-74.

Ben-Tal, A., Margalit, T. and Nemirovski, A. (2000) Robust Modeling of Multi-Stage Portfolio Problems, *High performance optimization*, pp. 303-328.

Ben-Tal, A. and Nemirovski, A. (1998) Robust Convex Optimization, *Mathematics of Operations Research*, vol. 23, no. 4, pp. 769-805.

Ben-Tal, A. and Nemirovski, A. (1999) Robust Solutions of Uncertain Linear Programs, *Operations research letters*, vol. 25, no. 1, pp. 1-13.

Ben-Tal, A. and Nemirovski, A. (2000) Robust Solutions of Linear Programming Problems Contaminated with Uncertain Data, *Mathematical Programming*, vol. 88,

no. 3, pp. 411-424.

Bera, A.K. and Park, S.Y. (2008) Optimal Portfolio Diversification Using the Maximum Entropy Principle, *Econometric Reviews*, vol. 27, nos. 4-6, pp. 484-512.

Berkelaar, A. and Kouwenberg, R. (2010) A Liability-Relative Drawdown Approach to Pension Asset Liability Management, *Journal of Asset Management*, vol. 11, no. 2, pp. 194-217.

Bertsimas, D., Pachamanova, D. and Sim, M. (2004) Robust Linear Optimization under General Norms, *Operations Research Letters*, vol. 32, no. 6, pp. 510-516.

Bertsimas, D., and Pachamanova, D. (2008) Robust Multiperiod Portfolio Management in the Presence of Transaction Costs, *Computers & Operations Research*, vol. 35, no. 1, pp. 3-17.

Bessler, W. and Wolff, D. (2015) Do Commodities Add Value in Multi-Asset-Portfolios? An Out-of-Sample Analysis for Different Investment Strategies, *Journal of Banking and Finance*, vol. 60, pp. 1-20.

Bessler, W., Opfer, H. and Wolff, D. (forthcoming) Multi-Asset Portfolio Optimization and Out-of-Sample Performance: An Evaluation of Black-Litterman, Mean Variance and Naïve Diversification Approaches. *The European Journal of Finance*.

Best, M. J. and Grauer, R. R. (1991) On the Sensitivity of Mean-Variance-Efficient Portfolios to Changes in Asset Means: Some Analytical and Computational Results, *Review of Financial Studies*, vol. 4, no. 2, pp. 315-342.

Black, F. (1980) The Tax Consequences of Long Run Pension Policy, *Financial Analysts Journal*, vol. 36, no. 4, July-August, pp. 21-28.

Black, F. and Dewhurst, M.P. (1981) A New Investment Strategy for Pension Funds, *Journal of Portfolio Management*, vol. 7, no. 4, Summer, pp. 26-34.

Black, F. and Litterman, R. (1992) Global Portfolio Optimization, *Financial Analysts Journal*, vol. 48, no. 5, pp. 28-43.

Blake, D., Cotter, J. and Dowd, K. (2007) Financial Risks and the Pension Protection Fund: Can it Survive Them?, *Pensions: An International Journal*, vol. 12, no. 3, June, pp. 109-130.

Bloomfield, T., Leftwich, R. and Long, J.B. (1977) Portfolio Strategies and Performance, *Journal of Financial Economics*, vol. 5, no. 2, November, pp. 201-218.

Blume, M.E. and Friend, I. (1975) The Asset Structure of Individual Portfolios and Some Implications for Utility Functions, *Journal of Finance*, vol. 30, no. 2, May, pp. 585-603.

Board, J. L. and Sutcliffe, C. M. (1994) Estimation Methods in Portfolio Selection and the Effectiveness of Short Sales Restrictions: UK Evidence, *Management science*, vol. 40, no. 4, pp. 516-534.

Board, J.L.G. and Sutcliffe, C.M.S. (2007) Joined-Up Pensions Policy in the UK: An Asset Liability Model for Simultaneously Determining the Asset Allocation and Contribution Rate, *Handbook of Asset and Liability Management*, vol. 2, edited by S.A. Zenios and W.T. Ziemba, North Holland Handbooks in Finance, Elsevier Science B.V., pp. 1029-1067.

Board, J., Sutcliffe, C. and Ziemba, W. T. (2003) Applying Operations Research Techniques to Financial Markets, *Interfaces*, vol. 33, no. 2, pp. 12-24.

Bodie, Z., Light, J.O., Morck, R. and Taggart, R.A. (1985) Corporate Pension Policy: An Empirical Investigation, *Financial Analysts Journal*, vol. 41, no. 5, September-October, pp. 10-16.

Bodie, Z., Light, J.O., Morck, R. and Taggart, R.A. (1987) Funding and Asset Allocation in Corporate Pension Plans: An Empirical Investigation. In *Issues in Pension Economics*, edited by Z. Bodie, J.B. Shoven and D.A. Wise, University of Chicago Press, pp. 15-47.

Boender, G. C. (1997) A Hybrid Simulation/Optimisation Scenario Model for Asset/Liability Management, *European Journal of Operational Research*, vol. 99, no. 1, pp. 126-135.

Bondell, H.D. and Reich, B.J. (2008) Simultaneous Regression Shrinkage, Variable Selection, and Supervised Clustering of Predictors with OSCAR, *Biometrics*, vol. 64, pp. 115-123.

Bonenkamp, J. (2009) Measuring Lifetime Redistribution in Dutch Occupational Pensions, *De Economist*, vol. 157, no. 1, March, pp. 49-77.

Bouslah, K., Kryzanowski, L. and M'Zali, B. (2013) The Impact of the Dimensions of Social Performance on Firm Risk, *Journal of Banking and Finance*, vol. 37, pp. 1258-1273.

Brammer, S., Brooks, C. and Pavelin, S. (2006) Corporate Social Performance and Stock Returns: UK Evidence from Disaggregate Measures, *Financial Management*, vol. 35, pp. 97-116.

Broadie, M. (1993) Computing Efficient Frontiers Using Estimated Parameters, *Annals of Operations Research*, vol. 45, pp. 21-58.

Brodie, J., Daubechies, I., De Mol, C., Giannone, D. and Loris, I. (2009) Sparse and Stable Markowitz Portfolios, *Proceedings of the National Academy of Sciences*, vol. 106, no. 30, pp. 12267-12272.

Broeders, D., Chen, A. and Rijsbergen, D. (2013) Valuation of Liabilities in Hybrid Pension Plans, *Applied Financial Economics*, vol. 23, no. 15, pp. 1215-1229.

Brown, S. J., Hwang, I. and In, F. H. (2013) Why Optimal Diversification Cannot Outperform Naive Diversification: Evidence from Tail Risk Exposure. Available at SSRN: <http://ssrn.com/abstract=2242694>

Brunnermeier, M.K. and Nagel, S. (2008) Do Wealth Fluctuations Generate Time-Varying Risk Aversion? Micro-Evidence on Individuals' Asset Allocation, *American Economic Review*, vol. 98, no. 3, June, pp. 713-736.

Burke, G. (1994) A Sharper Sharpe Ratio, *Futures*, vol. 23, no. 3, p. 56.

Carhart, M. M. (1997) On Persistence in Mutual Fund Performance, *Journal of Finance*, vol. 52, no. 1, pp. 57-82.

Campbell, J.Y. and Viceira, L.M. (2002) *Strategic Asset Allocation: Portfolio Choice for Long Term Investors*, Oxford University Press.

Candelon, B., Hurlin, C. and Tokpavi, S. (2012) Sampling Error and Double Shrinkage Estimation of Minimum Variance Portfolios, *Journal of Empirical Finance*, vol. 19, September, pp. 511–527.

Capelle-Blancard, G. and Monjon, S. (2014) The Performance of Socially Responsible Funds: Does the Screening Process Matter?, *European Financial Management*, vol. 20, pp. 494–520.

Carino, D. R. and Ziemba, W. T. (1998) Formulation of the Russell-Yasuda Kasai Financial Planning Model, *Operations Research*, vol. 46, no. 4, pp. 433-449.

Carino, D. R., Kent, T., Myers, D. H., Stacy, C., Sylvanus, M., Turner, A. L. and Ziemba, W. T. (1994) The Russell-Yasuda Kasai Model: An Asset/Liability Model for a Japanese Insurance Company Using Multistage Stochastic Programming, *Interfaces*, vol. 24, no. 1, pp. 29-49.

Carnes, B.A. and Olshansky, S.J. (2007) A Realist View of Aging, Mortality and Future Longevity, *Population and Development Review*, vol. 33, no. 2, June, pp. 367-381.

Carroll, A. B. (1991) The Pyramid of Corporate Social Responsibility: Toward the Moral Management of Organizational Stakeholders, *Business Horizons*, vol. 34, pp. 39-48.

Ceria, S. and Stubbs, R. A. (2006) Incorporating Estimation Errors into Portfolio Selection: Robust Portfolio Construction, *Journal of Asset Management*, vol. 7, no. 2, pp. 109-127.

Chapman, R. J., Gordon, T. J. and Speed, C. A. (2001) Pensions, Funding and Risk, *British Actuarial Journal*, vol. 7, no. 4, pp. 605-662.

Chaves, D., Hsu, J., Li, F. and Shakernia, O. (2011) Risk Parity Portfolio vs. Other Asset Allocation Heuristic Portfolios, *Journal of Investing*, vol. 20, no. 1, Spring, pp. 108-118.

Chen, Z., Pelsser, A. and Ponds, E. (2014) Evaluating the UK and Dutch Defined-Benefit Pension Policies Using the Holistic Balance Sheet Framework, *Insurance: Mathematics and Economics*, vol. 58, pp. 89-102.

Chopra, V.K. and Ziemba, W.T. (1993) The Effects of Errors in Means, Variances and Covariances on Optimal Portfolio Choice, *Journal of Portfolio Management*, vol. 19, no. 2, pp. 6-11.

Chopra, V.K. (1993) Improving Optimization, *Journal of Investing*, vol. 2, no. 3, pp. 51-59.

Chun, G.H., Ciochetti, B.A. and Shilling, J.D. (2000) Pension-Plan Real Estate Investment in an Asset–Liability Framework, *Real Estate Economics*, vol. 28, no. 3, pp. 467-491.

Clarkson, M.B.E. (1995) A Stakeholder Framework for Analyzing and Evaluating Corporate Social Performance, *Academy of Management Review*, vol. 20, pp. 92-117.

Cochrane, J.H. (2001), *Asset Pricing*, Princeton University Press.

Cochrane, J.H. and Piazzesi, M. (2005) Bond Risk Premia, *American Economic Review*, vol. 95, no. 1, March, pp. 138-160.

Comprix, J. and Muller, K.A. (2006) Pension Funding Rules and Sponsor Firms' Risk Taking with Defined Benefit Pension Plan Assets, Working paper, Pennsylvania University.

Coronado, J., Liang, N. and Orszag, M. (2006) Moral Hazard from Government Pension Insurance: Evidence from US and UK Firm Finances, Working Paper, Federal Reserve Board, 39 pages.

Costa, O.L.V. and Paiva, A.C. (2002) Robust Portfolio Selection Using Linear-Matrix Inequalities, *Journal of Economic Dynamics and Control*, vol. 26, no. 6, pp. 889-909.

Cox, P., Brammer, S. and Millington, A. (2004) An Empirical Examination of Institutional Investor Preferences for Corporate Social Performance, *Journal of Business Ethics*, vol. 52, pp. 27-43.

Craft, T.M. (2001) The Role of Private and Public Real Estate in Pension Plan Portfolio Allocation Choices, *Journal of Real Estate Portfolio Management*, vol. 7, no. 1, January-March, pp. 17-23.

Craft, T.M. (2005) How Funding Ratios Affect Pension Plan Portfolio Allocations, *Journal of Real Estate Portfolio Management*, vol. 11, no. 1, January-April, pp. 29-36.

Cumming, D., Helge Haß, L. and Schweizer, D. (2014) Strategic Asset Allocation and the Role of Alternative Investments, *European Financial Management*, vol. 20, no. 3, pp. 521-547.

De Groot, W. and Swinkels, L. (2008) Incorporating Uncertainty About Alternative Assets in Strategic Pension Fund Asset Allocation, *Pensions: An International Journal*, vol. 13, April, nos 1-2, pp. 71-77.

De Jong, F. (2008) Valuation of Pension Liabilities in Incomplete Markets, *Journal of Pension Economics and Finance*, vol.7, no. 3, November, pp. 277–294.

Delage, E. and Mannor, S. (2010) Percentile Optimization for Markov Decision Processes with Parameter Uncertainty, *Operations Research*, vol. 58, no. 1, January-March, pp. 203-213.

DeMiguel, V., Garlappi, L., Nogales, F.J. and Uppal, R. (2009a) A Generalized Approach to Portfolio Optimization: Improving Performance by Constraining Portfolio Norms, *Management Science*, vol. 22, no. 5, pp. 1915-1953.

DeMiguel, V., Garlappi, L. and Uppal, R. (2009b) Optimal Versus Naive diversification: How Inefficient is the 1/N Portfolio Strategy?, *The Review of Financial Studies*, vol. 22, no. 5, pp. 1915–1953.

Deng, X., Kang, J.K. and Low, B.S. (2013) Corporate Social Responsibility and Stakeholder Value Maximization: Evidence from Mergers, *Journal of Financial Economics*, vol. 110, pp. 87-109.

Derwall, J. and Koedijk, K. (2009) Socially Responsible Fixed Income Funds, *Journal of Business Finance and Accounting*, vol. 36, pp. 210-229.

Diebold, F.X. and Li, C. (2006) Forecasting the Term of Government Bond Yields, *Journal of Econometrics*, vol. 130, no. 2, December, pp. 337-364.

Disney, R., Emmerson, C. and Tetlow, G. (2009) What is a Public Sector Pension Worth?, *Economic Journal*, vol. 119, no. 541, November, pp. 517-534.

Dondi, G., Herzog, F., Schumann, L.M. and Geering, H.P. (2006) Dynamic Asset and Liability Management for Swiss Pension Funds, *Handbook of Asset and Liability Management*, vol. 2, pp. 963-1028.

Donnelly, C. (2014) Quantifying Mortality Risk in Small Defined Benefit Pension Schemes, *Scandinavian Actuarial Journal*, vol. 2014, no. 1, pp. 41-57.

Dowd, K. (2000) Adjusting for Risk: An Improved Sharpe Ratio, *International Review of Economics and Finance*, vol. 9, no. 3, pp. 209-222.

Draper, N., Van Ewijk, C., Lever, M. and Mehlkopf, R. (2014) Stochastic Generational Accounting Applied to Reforms of Dutch Occupational Pensions, *De Economist*, vol. 162, no. 3, pp. 287-307.

Drobetz, W. (2001) How to Avoid the Pitfalls in Portfolio Optimization? Putting the Black-Litterman Approach at Work, *Financial Markets and Portfolio Management*, vol. 15, no. 1, pp. 59-75.

Drut, B. (2012) Social Responsibility and Mean-Variance Portfolio Selection, *Amundi Working Paper*.

Edmans, A. (2011) Does the Stock Market Fully Value Intangibles? Employee Satisfaction and Equity Prices, *Journal of Financial Economics*, vol. 101, pp. 621-640.

Edwards, M. (2008), *Small change: Why Business Won't Save the World*, San Francisco: Berrett-Koehler.

Eling, M. and Schuhmacher, F. (2007), Does the Choice of Performance Measure Influence the Evaluation of Hedge Funds?, *Journal of Banking and Finance*, vol. 31, no. 9, pp. 2632-2647.

Elton, E.J. and Gruber, M.J. (1973) Estimating the Dependence Structure of Share Prices - Implications for Portfolio Selection, *Journal of Finance*, vol. 28, no. 5, December, pp. 1203-1232.

Ezra, D.D. (1991) Asset Allocation by Surplus Optimization, *Financial Analysts Journal*, vol. 47, no. 1, pp. 51-57.

Fabozzi, F.J., Kolm, P.N., Pachamanova, D. and Focardi, S.M. (2007) *Robust portfolio optimization and management*, John Wiley & Sons.

Fama, E.F. and French, K.R. (1993) Common Risk Factors in the Returns on Stocks and Bonds, *Journal of Financial Economics*, vol. 33, no. 1, pp. 3-56.

Fan, J., Zhang, J. and Yu, K. (2008) Asset Allocation and Risk Assessment with Gross Exposure Constraints for Vast Portfolios, *The Annals of Statistics*, vol. 25, pp. 1425-1432.

- Fan, J., Zhang, J. and Yu, K. (2012) Vast Portfolio Selection with Gross-Exposure Constraints, *Journal of the American Statistical Association*, vol. 107, pp. 592-606.
- Ferstl, R. and Weissensteiner, A. (2011) Asset-Liability Management under Time-Varying Investment Opportunities, *Journal of Banking & Finance*, vol. 35, no. 1, pp. 182-192.
- Frank, M.M. (2002) The Impact of Taxes on Corporate Defined Plan Asset Allocation, *Journal of Accounting Research*, vol. 40, no. 4, September, pp. 1163-1190.
- Frazzini, A. and Pedersen, L.H. (2014) Betting Against Beta, *Journal of Financial Economics*, vol. 111, no. 1, pp. 1–25.
- Friedman, B.M. (1983) Pension Funding, Pension Asset Allocation and Corporate Finance: Evidence from Individual Company Data. In *Financial Aspects of the United States Pension System* edited by Z. Bodie and J.B. Shoven, University of Chicago Press, pp. 107-152.
- Frost, P.A., and Savarino, J.E. (1986) An Empirical Bayes Approach to Efficient Portfolio Selection, *Journal of Financial and Quantitative Analysis*, vol. 21, no. 3, pp. 293-305.
- Frost, P.A. and Savarino, J.E. (1988) For Better Performance: Constrain Portfolio Weights, *Journal of Portfolio Management*, vol. 15, no. 1, pp. 29-34.
- Gabrel, V., Murat, C. and Thiele, A. (2014) Recent Advances in Robust Optimization: An Overview, *European Journal of Operational Research*, vol. 235, no. 3, pp. 471-483.
- Galema, R., Plantinga, A. and Scholtens, B. (2008) The Stocks at Stake: Return and Risk in Socially Responsible Investment, *Journal of Banking and Finance*, vol. 32, pp. 2646-2654.
- Gao, J. (2008) Stochastic Optimal Control of DC Pension Funds, *Insurance: Mathematics and Economics*, vol. 42, no. 3, pp. 1159-1164.
- Gerakos, J. (2010) Chief Executive Officers and the Pay-Pension Tradeoff, *Journal of Pension Economics and Finance*, vol. 9, no. 2, April, pp. 303-319.
- Geyer, A. and Ziemba, W.T. (2008) The Innovest Austrian Pension Fund Financial Planning Model InnoALM, *Operations Research*, vol. 56, no. 4, pp. 797-810
- Ghaoui, L.E., Oks, M. and Oustry, F. (2003) Worst-Case Value-At-Risk and Robust Portfolio Optimization: A Conic Programming Approach, *Operations Research*, vol. 51, no. 4, pp. 543-556.

Giamouridis, D., Sakkas, A. and Tessaromatis, N. (2014) Dynamic Asset Allocation with Liabilities.

Available at SSRN: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2049950

Glasserman, P. and Xu, X. (2013) Robust Portfolio Control with Stochastic Factor Dynamics, *Operations Research*, vol. 61, no. 4, pp. 874-893.

Godfrey, P.C., Merrill, C.B. and Hansen, J.M. (2009) The Relationship between Corporate Social Responsibility and Shareholder Value: An Empirical Test of the Risk Management Hypothesis, *Strategic Management Journal*, vol. 30, pp. 425-445.

Gofman, M. and Manela, A. (2012) An Empirical Evaluation of the Black-Litterman Approach to Portfolio Choice. Available at SSRN: <http://ssrn.com/abstract=1782033>

Goldfarb, D. and Iyengar, G. (2003) Robust Portfolio Selection Problems, *Mathematics of Operations Research*, vol. 28, no. 1, pp. 1-38.

Green, R. and Hollifield, B. (1992) When will Mean-Variance Efficient Portfolios be well Diversified?, *Journal of Finance*, vol. 47, no. 5, pp. 1785-1809.

Gregoriou, G.N. and Gueyie, J.P. (2003) Risk-Adjusted Performance of Funds of Hedge Funds Using a Modified Sharpe Ratio, *The Journal of Wealth Management*, vol. 6, no. 3, pp. 77-83.

Griffin, J. J. and Mahon, J. F (1997) The Corporate Social Performance and Corporate Financial Performance Debate, *Business and Society*, vol. 36, pp. 5-31.

Grossman, S. and Zhou, Z. (1993) Optimal Investment Strategies for Controlling Drawdowns, *Mathematical Finance*, vol. 3, pp. 241–276.

Guan, Y. and Lui, D. (2014) Risk Shifting in Pension Investment, Working paper, Hong Kong Baptist University, 48 pages.

Gülpinar, N. and Pachamanova, D. (2013) A Robust Optimization Approach to Asset-Liability Management under Time-Varying Investment Opportunities, *Journal of Banking & Finance*, vol. 37, no. 6, pp. 2031-2041.

Gülpinar, N., Rustem, B. and Settergren, R. (2004) Simulation and Optimization Approaches to Scenario Tree Generation, *Journal of economic dynamics and control*, vol. 28, no. 7, pp. 1291-1315.

Haberman, S. (1992) Pension Funding with Time Delays: A Stochastic Approach, *Insurance: Mathematics and Economics*, vol. 11, no. 3, pp. 179-189.

Hallerbach, W., Ning, H., Soppe, A., and Spronk, J. (2004) A Framework for Managing a Portfolio of Socially Responsible Investments, *European Journal of Operational Research*, vol. 153, pp. 517–529.

Haynes, J.B. and Sessions, J.G. (2013) Work Now, Pay Later? An Empirical Analysis of the Pension-Pay Trade Off, *Economic Modelling*, vol. 30, January, pp. 835-843.

HEFCE (2010) *Staff Employed at HEFCE-Funded HEIs - Trends and Profiles 1995-96 to 2008-09*, HEFCE, February.

Higher Education Statistics Agency (2014) *Staff in Higher Education, 2012/13*, HESA.

Hillman, A.J. and Keim, G.D. (2001) Shareholder Value, Stakeholder Management, and Social Issues: What's the Bottom Line?, *Strategic Management Journal*, vol. 22, pp.125-139.

Hoevenaars, P.M.M. (2011) Pricing Embedded Options in Value-based Asset Liability Management. In *The Oxford Handbook of Quantitative Asset Management*, edited by B. Scherer and K. Winston, Oxford University Press, 2011, pp. 449-469.

Hoevenaars, P.M.M. (2009) Pricing Risk in Corporate Pension Plans: Understanding the Real Pension Deal, *Rotman International Journal of Pension Management*, vol. 2, no. 1, Spring, pp. 56-62.

Hoevenaars, P.M.M., Kocken, T. and Ponds, E. (2009) Pricing Risk in Corporate Pension Plans: Understanding the Real Pension Deal, *Rotman International Journal of Pension Management*, vol. 2, no. 1, Spring, pp.

Hoevenaars, P.M.M., Molenaar, R.D.J. and Ponds, E.H.M. (2010) Public Investment Funds and Value-based Generational Accounting. In *Central Bank Reserves and Sovereign Wealth Management*, edited by A. Berkelaar, J. Coche and K. Nyholm, Palgrave, pp. 328-348.

Hoevenaars, P.M.M. and Ponds, E.H.M. (2007) Inter-generational Value Transfers Within an Industry-wide Pension Fund: A Value-based ALM Analysis. In *Costs and Benefits of Collective Pension Systems*, edited by O.W. Steenbeck and S.G. Van der Lecq, Springer, pp. 95-117.

Hoevenaars, R.P., Molenaar, R.D., Schotman, P.C. and Steenkamp, T.B. (2008) Strategic Asset Allocation with Liabilities: Beyond Stocks and Bonds, *Journal of Economic Dynamics and Control*, vol. 32, no. 9, pp. 2939–2970.

Hoevenaars, R. P. and Ponds, E. H. (2008) Valuation of Intergenerational Transfers in Funded Collective Pension Schemes, *Insurance: Mathematics and Economics*, vol. 42, no. 2, pp. 578-593.

Huang, D., Zhu, S., Fabozzi, F.J. and Fukushima, M. (2010) Portfolio Selection under Distributional Uncertainty: A Relative Robust CVaR Approach, *European Journal of Operational Research*, vol. 203, no. 1, pp. 185-194.

Idzorek, T. M. (2005) *A Step-by-Step Guide Through the Black-Litterman Model, Incorporating User Specified Confidence Levels*, Chicago: Ibbotson Associates, 32 pages.

Inkmann, J. and Blake, D. (2012) *Managing Financially Distressed Pension Plans in the Interest of the Beneficiaries*, Working Paper, The Pensions Institute, November.

Jacobs, H., Muller, S. and Weber, M. (2014) How Should Individual Investors Diversify? An Empirical Evaluation of Alternative Asset Allocation Policies, *Journal of Financial Markets*, vol. 19, June, pp. 62-86.

Jagannathan, R. and Ma, T. (2003) Risk Reduction in Large Portfolios: Why Imposing the Wrong Constraints Helps, *Journal of Finance*, vol. 58, no. 4, August, pp. 1651-1683.

Jensen, M.C. and Meckling, W.H. (1976) Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure, *Journal of Financial Economics*, vol. 3, no. 4, pp. 305-360.

Jo, H., and M. Harjoto. (2012) The Causal Effect of Corporate Governance on Corporate Social Responsibility, *Journal of Business Ethics*, vol. 106, pp. 53-72.

Jobson, J.D., Korkie, B. and Ratti, V. (1979) Improved Estimation for Markowitz Portfolios Using James-Stein Type Estimators, *Proceedings of the American Statistical Association, Business and Economic Statistics Section*, pp. 279-284.

Johnson, R.D. and Greening, D.W. (1999) The Effects of Corporate Governance and Institutional Ownership Types on Corporate Social Performance, *Academy of Management Journal*, vol. 42, pp. 564-576.

Jorion, P. (1986) Bayes-Stein Estimation for Portfolio Analysis, *Journal of Financial and Quantitative Analysis*, vol. 21, no. 3, pp. 279-292.

Jorion, P. (1991) Bayesian and CAPM Estimators of the Means: Implications for Portfolio Selection, *Journal of Banking and Finance*, vol. 15, no. 3, June, pp. 717-727.

Kallberg, J.G. and Ziemba, W.T. (1984) Mis-Specifications in Portfolio Selection Problems, *Risk and capital*, pp. 74-87.

Kan, R. and Zhou, G. (2007) Optimal Portfolio Choice with Parameter Uncertainty, *Journal of Financial and Quantitative Analysis*, vol. 42, no. 3, pp. 621-656.

Kaplan, P.D. (2014) Back to Markowitz. Morningstar Indexes 2014/15 Yearbook. Available at: <http://corporate.morningstar.com/us/asp/detail.aspx?xmlfile=8383.xml>

Keel, A. and Muller, H. (1995) Efficient Portfolios in the Asset Liability Context, *Astin Bulletin*, vol. 25, no. 1, May, pp. 33-48.

Kempf, A. and Osthoff, P. (2007) The Effect of Socially Responsible Investing on Portfolio Performance, *European Financial Management*, vol. 13, pp. 908-922.

Kempf, A. and Osthoff, P. (2008), SRI Funds: Nomen est Omen, *Journal of Business Finance and Accounting*, vol. 35, pp. 1276-1294.

Kim, W.C., Kim, J.H., Ahn, S.H. and Fabozzi, F.J. (2013) What Do Robust Equity Portfolio Models Really Do?, *Annals of Operations Research*, vol. 205, no. 1, May, pp. 141-168.

Kim, W.C., Kim, J.H. and Fabozzi, F.J. (2014) Deciphering Robust Portfolios, *Journal of Banking and Finance*, vol. 45, pp. 1-8.

Kirby, C. and Ostdiek B. (2012) It's All in the Timing: Simple Active Portfolio Strategies that Outperform Naive Diversification, *Journal of Financial and Quantitative Analysis*, vol. 47, no. 2, pp. 437-467.

Kourtis, A., Dotsis, G. and Markellos, R. N. (2012) Parameter Uncertainty in Portfolio Selection: Shrinking the Inverse Covariance Matrix, *Journal of Banking and Finance*, vol. 36, no. 9, pp. 2522-2531.

Kouwenberg, R. (2001) Scenario Generation and Stochastic Programming Models for Asset Liability Management, *European Journal of Operations Research*, vol. 134, no. 2, pp. 279-292.

Kouwenberg, R. and Zenios, S.A. (2006) Stochastic Programming Models for Asset Liability Management, *Handbook of asset and liability management*, vol. 1, pp. 253-303.

Kusy, M.L. and Ziemba, W.T. (1986) A Bank Asset and Liability Management Model, *Operations Research*, vol. 34, no. 3, pp. 356-376.

LCP (2014) LCP Accounting for Pensions 2014, Lane, Clark and Peacock, August, 66 pages.

Ledit, O. and Wolf, M. (2003) Improved Estimation of the Covariance Matrix of Stock Returns With an Application to Portfolio Selection, *Journal of empirical finance*, vol. 10, no. 5, pp. 603-621.

Ledit, O. and Wolf, M. (2004) A Well-Conditioned Estimator For Large-Dimensional Covariance Matrices, *Journal of multivariate analysis*, vol. 88, no. 2, pp. 365-411.

Ledit, O. and Wolf, M. (2008) Robust Performance Hypothesis Testing with the Sharpe Ratio, *Journal of Empirical Finance*, vol. 15, no. 5, pp. 850-859.

Lee, D.D. and Faff, R.W. (2009) Corporate Sustainability Performance and

- Idiosyncratic Risk: A Global Perspective, *Financial Review*, vol. 44, pp. 213–237.
- Lekniute, Z. (2011) A Value-based Approach to Pension Plan Redesign in the Netherlands: Who Will Gain and Who Will Lose? Netspar MSc Thesis 2011-035, 62 pages.
- Lekniute, Z., Beetsma, R. and Ponds, E. (2014) A Value-based Approach to the Redesign of US State Pension Plans?, Working paper, May, 39 pages.
- Li, Y. (2010) Corporate Risk Management and Pension Asset Allocation. In *Pension Fund Risk Management: Financial and Actuarial Modelling*, edited by M. Micocci, G.N. Gregoriou and G. B. Masala, Chapman and Hall, pp. 365-387.
- Ling, A. and Xu, C. (2012) Robust Portfolio Selection Involving Options under a “marginal + joint” Ellipsoidal Uncertainty Set, *Journal of Computational and Applied Mathematics*, vol. 236, no. 14, August, pp. 3373–3393.
- Lintner, J. (1965) The valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets, *The review of economics and statistics*, vol. 47, no. 1, pp. 13-37.
- Lo, A. W. (2002) The Statistics of Sharpe Ratios, *Financial Analysts Journal*, vol. 58, no. 4, pp. 36-52.
- Logan, D. (1985) *The Birth of a Pension Scheme: A History of the Universities Superannuation Scheme*, Liverpool University Press.
- MacBeth, J.D., Emanuel, D.C. and Heatter, C.E. (1994) An Investment Strategy for Defined Benefit Plans, *Financial Analysts Journal*, vol. 50, no. 3, pp. 34-41.
- MacKinlay, A. C. and Pastor, L. (2000) Asset Pricing Models: Implications for Expected Returns and Portfolio Selection, *The Review of Financial Studies*, vol. 13, no. 4, pp. 883–916.
- Margolis, J.D. and Walsh, J. P. (2003) Misery Loves Companies: Rethinking Social Initiatives by Business, *Administrative Science Quarterly*, vol. 48, pp. 268-305.
- Margolis, J., Elfenbein, H., and Walsh, J. P. (2009) Does It Pay to Be Good and Does It Matter? A Meta-Analysis of the Relationship between Corporate Social and Financial Performance, Harvard Business School Working Paper No. 229.
- Markowitz, H. (1952) Portfolio Selection, *Journal of Finance*, vol. 7, no. 1, pp. 77-91.
- Markowitz, H. (1959) *Portfolio Selection*, Wiley, New York.
- Markowitz, H. (1991) *Portfolio Selection: Efficient Diversification of Investment (2e)*. Malden, MA: Blackwell Publishers Inc.

Mattingly, J.E. and Berman, S.L. (2006) Measurement of Corporate Social Action, *Business and Society*, vol. 45, pp. 20-46.

McCarthy, D. (2005) The Optimal Allocation of Pension Risks in Employment Contracts, Department of Work and Pensions, Research Report 272.

McCarthy, D. and Miles, D. (2013) Optimal Portfolio Allocation for Corporate Pension Funds, *European Financial Management*, vol. 19, no. 3, June, pp. 599-629.

Merton, R.C. (1969) Lifetime Portfolio Selection under Uncertainty: The Continuous Time Case, *The review of Economics and Statistics*, vol. 51, pp. 247-257.

Merton, R.C. (1980) On Estimating the Expected Return on the Market, *Journal of Financial Economics*, vol. 8, pp. 323-361.

Merton, R.C. (1993) Optimal Investment Strategies for University Endowment Funds, *Studies of supply and demand in higher education*, University of Chicago Press, pp. 211-242.

Meucci, A. (2010) *The Black-Litterman Approach: Original Model and Extensions*, 17 pages. Download from:-

http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1117574

Michaud, R. (1999) *Efficient Asset Management: A Practical Guide to Stock Portfolio Optimization and Asset Allocation*, Harvard Business School Press.

Mitra, S. (2006) A White Paper on Scenario Generation for Stochastic Programming, *Optirisk Systems: White Paper Series*.

Mohan, N. and Zhang, T. (2014) An Analysis of Risk-Taking Behaviour for Public Defined Benefit Pension Plans, *Journal of Banking and Finance*, vol. 40, March, pp. 403-419.

Mossin, J. (1966) Equilibrium in a Capital Asset Market, *Econometrica*, vol. 34, no. 4, pp. 768-783.

Mulvey, J.M., Gould, G. and Morgan, C. (2000) An Asset and Liability Management System for Towers Perrin-Tillinghast, *Interfaces*, vol. 30, no. 1, pp. 96-114.

Mulvey, J.M., Fabozzi, F.J., Pauling, W.R., Simsek, K.D. and Zhang, Z. (2005) Modernizing the Defined-Benefit Pension System, *The Journal of Portfolio Management*, vol. 31, no. 2, pp. 73-82.

Munro, J. and Barrie, A. (2003) New Solutions for Pension Fund Risks: A Review of Defined Benefit Investment Strategies, Barrie and Hibbert, Pension Fund Strategy Research Report No. 1.

- Natarajan, K., Pachamanova, D. and Sim, M. (2009) Constructing Risk Measures from Uncertainty Sets, *Operations Research*, vol. 57, no. 5, pp. 1129-1141.
- Nelson, C.R. and Siegel, A.F. (1987) Parsimonious Modelling of Yield Curve, *Journal of Business*, vol. 60, pp. 473–489.
- Nijman, T.E. and Koijen, R.S.J. (2006) Valuation and Risk Management of Inflation-Sensitive Pension Rights. In *Fair Value and Pension Fund Management*, edited by N. Kortleve, T. Nijman and E. Ponds, E. Elsevier, pp. 85-117.
- Nijman, T. and Swinkels, L. (2008) Strategic and Tactical Allocation to Commodities for Retirement Savings Schemes. In *The Handbook of Commodity Investing*, edited by F. Fabozzi, R. Fuss and D.G. Kaiser, Wiley, pp. 522-546.
- OECD (2013) *Pension Markets in Focus*, OECD.
- Office for National Statistics (2013) Occupational Pension Schemes Survey, United Kingdom 2012, ONS.
- Office for National Statistics (2013) National Life Tables, United Kingdom 2011-2013, ONS.
- Oikonomou, I., Brooks, C., and Pavelin, S. (2012) The Impact of Corporate Social Performance on Financial Risk and Utility: A Longitudinal Analysis, *Financial Management*, vol. 41, pp. 483–515.
- Orlitzky, M. and Benjamin, J. D. (2001) Corporate Social Performance and Firm Risk: A Meta-Analytic Review, *Business and Society*, vol. 40, pp. 369-396.
- Orlitzky, M., Schmidt, F. L. and Rynes, S.L. (2003) Corporate Social and Financial Performance: A Meta-Analysis, *Organization Studies*, vol. 24, pp. 403-441.
- Orlitzky, M. (2013) Corporate Social Responsibility, Noise and Stock Market Volatility, *The Academy of Management Perspectives*, vol. 27, pp. 238-254.
- Petersen, M.A. (1996) Allocating Assets and Discounting Cash Flows: Pension Plan Finance. In *Pensions, Savings and Capital Markets*, edited by P.A. Fernandez, J.A. Turner and R.P. Hinz, US Department of Labour; Pension and Welfare Administration, Office of Research and Economic Analysis, April, pp. 1-26.
- Platanakis, E and Sutcliffe, C. (forthcoming) Asset Liability Modelling and Pension Schemes: The Application of Robust Optimization to USS, *The European Journal of Finance*.
- Ponds, E. H. (2003) Pension Funds and Value-Based Generational Accounting, *Journal of Pension Economics and Finance*, vol. 2, no. 3, pp. 295-325.

- Prigent, J.L. (2007) *Portfolio Optimization and Performance Analysis*, Chapman and Hall.
- Pukthuanthong, K. and Roll, R. (2015) Agnostic Tests of Stochastic Discount Factor Theory, Social Science Working paper No. 1405, California Institute of Technology, April, 52 pages.
- Quaranta, A.G. and Zaffaroni, A. (2008) Robust Optimization of Conditional Value At Risk and Portfolio Selection, *Journal of Banking & Finance*, vol. 32, no. 10, pp. 2046-2056.
- Rachev, S.T. and Tokat, Y. (2000) *Asset and Liability Management: Recent Advances*, University of California, Santa Barbara, Department of Economics.
- Rauh, J.D. (2009) Risk Shifting Versus Risk Management: Investment Policy in Corporate Pension Plans, *Review of Financial Studies*, vol. 22, no. 7, July, pp. 2687-2733.
- Renneboog, L., Ter Horst, J. and Zhang, C. (2008) The Price of Ethics and Stakeholder Governance: The Performance of Socially Responsible Mutual Funds, *Journal of Corporate Finance*, vol. 14, pp. 302-322.
- Rollinger, T. and Hoffman, S (2014) Sortino: A “Sharper” Ratio, Red Rock Capital, http://www.redrockcapital.com/Sortino_A_Sharper_Ratio_Red_Rock_Capital.pdf
- Rudolf, M. and Ziemba, W.T. (2004) Intertemporal Surplus Management, *Journal of Economic Dynamics and Control*, vol. 28, no. 5, pp. 975-990.
- Rustem, B., Becker, R.G. and Marty, W. (2000) Robust min–max Portfolio Strategies for Rival Forecast and Risk Scenarios, *Journal of Economic Dynamics and Control*, vol. 24, no. 11, pp. 1591-1621.
- Ryan, L.V. and Schneider, M. (2002) The Antecedents of Institutional Investor Activism, *Academy of Management Review*, vol. 27, no. 4, pp. 554-573.
- Satchell, S. and Scowcroft, A. (2000) A Demystification of the Black-Litterman Model: Managing Quantitative and Traditional Portfolio Construction, *Journal of Asset Management*, vol. 1, no. 2, September, pp. 138-150.
- Scherer, B. (2002) Portfolio Resampling: Review and Critique, *Financial Analysts Journal*, vol. 58, no. 6, pp. 98-109.
- Scherer, B. (2007) Can Robust Portfolio Optimization Help to Build Better Portfolios?, *Journal of Asset Management*, vol. 7, no. 6, pp. 374-387.

Scherer, B. (2011) An Integrated Approach to Sovereign Wealth Risk Management, *Edhec-Risk Institute Publication*.

Schroder, M. (2007) Is There a Difference? The Performance Characteristics of SRI Equity Indices, *Journal of Business Finance and Accounting*, vol. 34, pp. 331–348.

Shadwick, W. and Keating, C. (2002) A Universal Performance Measure, *Journal of Performance Measurement*, vol. 6, no. 3, pp. 59–84.

Sharfman, M. (1996) The Construct Validity of the Kinder, Lydenberg & Domini Social Performance Ratings Data, *Journal of Business Ethics*, vol. 15, pp. 287-296.

Sharpe, W.F. (1964) Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk, *The journal of finance*, vol. 19, no. 3, pp. 425-442.

Sharpe, W.F. (1966) Mutual Fund Performance, *Journal of business*, vol. 39, no. 1, pp. 119-138.

Sharpe, W.F. and Tint, L.G. (1990) Liabilities-A New Approach, *The journal of portfolio management*, vol. 16, no. 2, pp. 5-10.

SOA (2003) Professional Actuarial Specialty Guide: Asset-Liability Management.

Sortino, F.A. and Van Der Meer, R. (1991) Downside Risk, *The Journal of Portfolio Management*, vol. 17, no. 4, pp. 27-31.

Statman, M. (2000) Socially Responsible Mutual Funds, *Financial Analysts Journal*, vol. 56, pp. 30-39.

Statman, M. (2006) Socially Responsible Indexes, *Journal of portfolio Management*, vol. 32, pp. 100-109.

Sturm, J.F. (1999) Using SeDuMi 1.02 a MATLAB Toolbox for Optimization Over Symmetric Cones, *Optimization Methods and Software*, vol. 11-12, pp. 625-653.

Surz, R.J. (1981) Elaborations on the Tax Consequences of Long-Run Pension Policy, *Financial Analysts Journal*, vol. 37, no. 1, January-February, pp. 52-54 & 60.

Tepper, I. (1981) Taxation and Corporate Pension Policy, *Journal of Finance*, vol. 36, no. 1, March, pp. 1-13.

Tola, V., Lillo, F., Gallegati, M. and Mantegna, R.N. (2008) Cluster Analysis for Portfolio Optimization, *Journal of Economic Dynamics and Control*, vol. 32, no. 1, pp. 235-258.

Tu, J. and Zhou, G. (2011) Markowitz Meets Talmud: A Combination of Sophisticated and Naive Diversification Strategies, *Journal of Financial Economics*, vol. 99, no. 1, pp. 204-215.

USS (1999-2012) Scheme Funding Reports of the Actuarial Valuation as at 31 March 1999 - 2011, Universities Superannuation Scheme.

USS (2013) Report and Accounts for the Year Ended 31 March 2013, Universities Superannuation Scheme, 110 pages.

USS (2014) 2014 Actuarial Valuation: A Consultation on the Proposed Assumptions for the Scheme's Technical Provisions and Recovery Plan, USS, October, 39 pages.

Utz, S., Wimmer, M., Hirschberger, M. and Steuer, R.E. (2014) Tri-Criterion Inverse Portfolio Optimization with Application to Socially Responsible Mutual Funds, *European Journal of Operational Research*, vol. 234, pp. 491-498.

Van Rooij, M., Siegmann, A. and Vlaar, P.J. (2004) PALMNET: A Pension Asset and Liability Model for the Netherlands, *SSRN Working Paper*.

Von Arx, U., and Ziegler, A. (2014) The Effect of Corporate Social Responsibility on Stock Performance: New Evidence for the USA and Europe, *Quantitative Finance*, vol. 14, pp. 977-991.

Votaw, D. (1973) Genius Becomes Rare. The Corporate Dilemma-Traditional Values versus Contemporary Problems, Prentice Hall, Englewood Cliffs, NJ.

Waddock, S.A (2003) Myths and Realities of Social Investing, *Organization and Environment*, vol. 16, pp. 369-380.

Waddock, S. A. and Graves, S. B. (1997) The Corporate Social Performance-Financial Performance Link, *Strategic Management Journal*, vol. 18, no. 4, pp. 303—319.

Wood, D. J. and Jones, R. E. (1995) Stakeholder Mismatching: A Theoretical Problem in Empirical Research on Corporate Social Performance, *International Journal of Organizational Analysis*, vol. 3, pp. 229-267.

Xing, X., Hu, J. and Yang, Y. (2014) Robust Minimum Variance with L-Infinity Constraints, *Journal of Banking and Finance*, vol. 46, pp. 107-117.

Young, T. W. (1991) Calmar Ratio: A Smoother Tool, *Futures*, vol. 20, p. 40.

Zhu, L.L. (2015) The Performance of Asset Allocation Strategies Across Datasets and Over Time, Working paper, University of Edinburgh Business School, August, 42 pages.

Ziembra, W.T. and Mulvey, J.M. (1998) *Worldwide Asset and Liability Modeling*, Cambridge University Press.