CVaR Optimization of Real Estate Portfolios in an ALM Context

Bert Kramer

ORTEC
FINANCE

CVaR Optimization of Real Estate Portfolios in an ALM Context

Bert Kramer[1,2]

Ortec Finance Research Center

P.O. Box 4074, 3006 AB Rotterdam

Boompjes 40, The Netherlands, www.ortec-finance.com

ABSTRACT

In this paper we show the result of an optimization study of sector and region allocation within real estate portfolios. The optimizations are performed in an ALM context. That is, we try to determine optimal allocations based on a funding ratio risk measure for Dutch pension funds. The risk measure we use is the Conditional Value at Risk (CVaR). In our optimizations, we assume a fixed portfolio for the non-real estate part and optimize the composition of the part allocated to non-listed real estate. We conclude that within Europe, the number of countries and sectors that appear in the optimal portfolio is limited. So the diversification gain from investing in a large number of sectors and countries is relatively limited within Europe. For countries outside Europe, the results are sensitive to changes in the input and assumptions. Furthermore, we conclude that high leverage is only acceptable when the underlying real estate market is very stable. Finally, although the optimal weights differ per type of pension fund, the countries and sectors that appear in the optimal portfolios are quite stable.

Keywords: Conditional Value at Risk, Asset Liability Management, Optimization, Real estate

JEL Classification: G11, G23, C61

Table of Contents

1 Introduction

In this paper we show the result of an optimization study of sector and region allocation within real estate portfolios. Optimization methods are an essential tool in the process of determining appropriate asset allocations, in an asset-only or ALM context. On the one hand they can ensure that the solution found is optimal according to a specified set of criteria and on the other hand they allow one to find the allocation solution in a time efficient manner reducing the need for manual searches and/or sensitivity analyses.

The optimizations are performed in an ALM context. That is, we try to determine optimal allocations based on a funding ratio risk measure for Dutch pension funds. Therefore, the characteristics of the pension fund are important. We will analyze the sensitivity of the optimal allocations for the type of fund. We will analyze four pension fund types: an average fund, a young fund (i.e., the average age of the participants is relatively low), a mature fund (i.e., the average age of the participants is relatively high), and a so-called sleeping fund (i.e., no active participants anymore, no new premium receipts, only outgoing pension payments).

Given the strategic asset allocation resulting from an ALM study, and given the resulting allocation to non-listed real estate, we want to determine the optimal allocation within this asset class towards regions and sectors.

A downside risk measure better matches the risk perception of institutional investors like pension funds than a volatility risk measure (like variance) that also takes the upside into account. The downside risk measure we use is the Conditional Value at Risk (CVaR). During the last 10 years, CVaR optimization has become the method of choice for dealing with non-normal distributions in the academic literature and is also spreading quickly among practitioners.

As far as we know, this is the first study that applies CVaR optimization combined with scenario analysis in an ALM context to real estate portfolios.

In section 2 we will introduce the CVaR optimization approach and we explain why it is preferred above mean variance optimization and Value at Risk. In section 3 we will discuss the data that is used in our analyses. Section 4 shows the optimization results and section 5 concludes.

2 CVaR optimization

2.1 Tail Risk Measures: VaR and CVaR

The primary goal of the Conditional Value at Risk (CVaR) optimization is to combine the portfolio choice problem with modern risk management techniques based on tail risk measurement. The first tail risk measure which has gained large industry acceptance is the Value at Risk (VaR). Introduced by J.P. Morgan more than 15 years ago, the VaR has spread quickly to become one of the most widely used measures for quantifying risk. For a given portfolio, confidence level and time horizon, VaR is defined as a threshold value such that the probability that the loss on the portfolio over the given time horizon exceeds this value is equal to one minus the specified confidence. Its main advantages compared to the use of volatility (variance) are that:

- it takes into account the non-normality of the expected returns distribution
- it encapsulates all the relevant information about non-normality in a single value
- it puts a dollar figure on the possible loss given a confidence level, and this seems to be a more intuitive way of expressing the risk rather than using volatility

Due to these reasons it is now very frequently used by practitioners, as well as regulators, and has also drawn significant attention in the academic research.

Naturally, the idea of constructing portfolios using VaR instead of variance in a MV-like approach has been very appealing. However, the problem which arises is that VaR is not a coherent risk measure (Artzner *et al*, 1999). More explicitly, in certain cases, it can happen that the VaR of a portfolio with two instruments may be greater than the sum of individual VaRs of the two instruments. When calculated on a set of scenarios, this translates into a non-convex measure (Mauser and Rosen, 1998), and therefore can present several different local minima. These drawbacks make VaR less desirable as a portfolio selection tool, unless the return distributions are non-normal to a limited degree only.
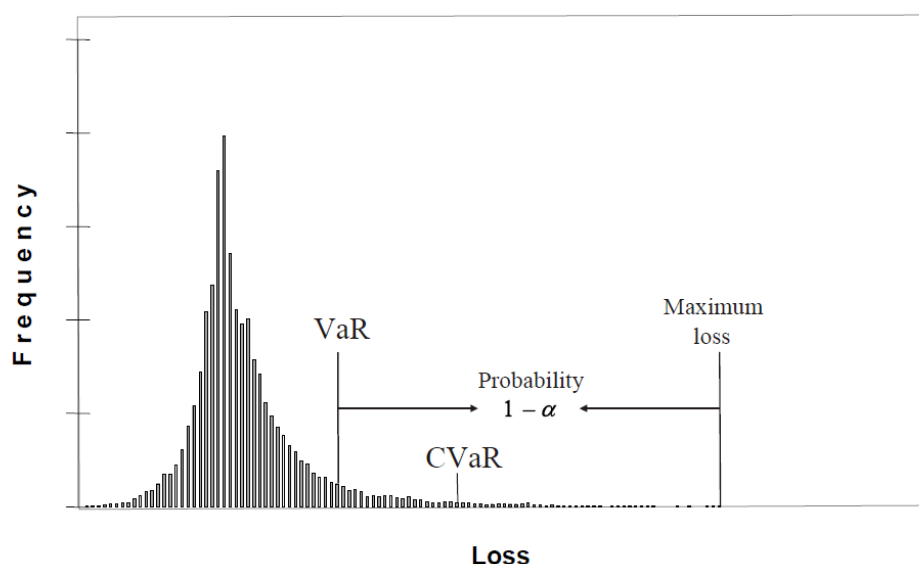
Figure 2.1 VaR and CVaR. Extract from Uryasev (2004)

A different downside risk measure has emerged as a viable alternative to VaR for portfolio optimization problems. Portfolio Optimization based on Conditional Value at Risk (CVaR) was introduced by Rockafellar and Uryasev (2000). As opposed to VaR which neglects the tail of the distribution beyond the confidence level threshold, CVaR averages over all the instances when the loss is larger than the VaR value. Figure 2.1 illustrates the definitions of VaR and CVaR with respect to the loss probability distribution. Due to the full account of the distribution tail, CVaR is a sub-additive risk measure and therefore satisfies the requirements set by Artzner *et al* (1999) for coherent risk measures. CVaR has the same advantages as those listed above for VaR. In addition, when evaluated over a set of scenarios-based samplings, the expected return distribution:

- is a convex risk measure
- can by approximated by a linear formula for which the optimization problem can be solved very efficiently

Although CVaR did not (yet) replace VaR as the preferred risk measure for risk management purposes, it has established itself as a viable alternative for portfolio optimization purposes.

In a recent survey, 51% of the European investment management practitioners indicated that they set absolute tail risk objectives (measured as VaR or CVaR) at the beginning of their portfolio optimization process, while only 46% uses volatility (Amenc *et al*, 2011). However, when implementing the portfolio optimization most investment managers stick to Normal distribution VaR, neglecting the extreme risks following from fat tails. CVaR is the most popular approach that takes into account fat tails; it is used by 22% of the practitioners (Amenc *et al*, 2011).

Note that in order to achieve all the desirable properties of CVaR including the cases of probability distributions with discontinuities, Rockafellar and Uryasev (2002) have taken a particular definition of CVaR out of several, closely related, risk measures. One can define $CVaR^+$ as the expected losses strictly exceeding VaR (a.k.a. Expected Shortfall), and $CVaR^-$ as expected losses which are equal to or exceed VaR (a.k.a. Tail VaR). Neither VaR, nor $CVaR^+$ or $CVaR^-$ is convex (see Figure 2.2). However the following linear combination is convex (Rockafellar and Uryasev, 2002):

$$CVaR = \lambda \, VaR + (1-\lambda) \, CVaR^+,$$

where

$$\lambda = (\Psi(VaR) - \alpha) / (1-\alpha) \, , \, (\, 0 \leq \lambda \leq 1 \,)$$

and $\Psi(VaR)$ is one minus the probability that losses exceed the VaR.

For continuous distributions, all the above CVaR measures are equivalent. In case of discontinuities arising due to discrete sampling, like in the case of a scenario based approach, or due to the presence of derivatives with discontinuous payoffs, the distinction is however very important.
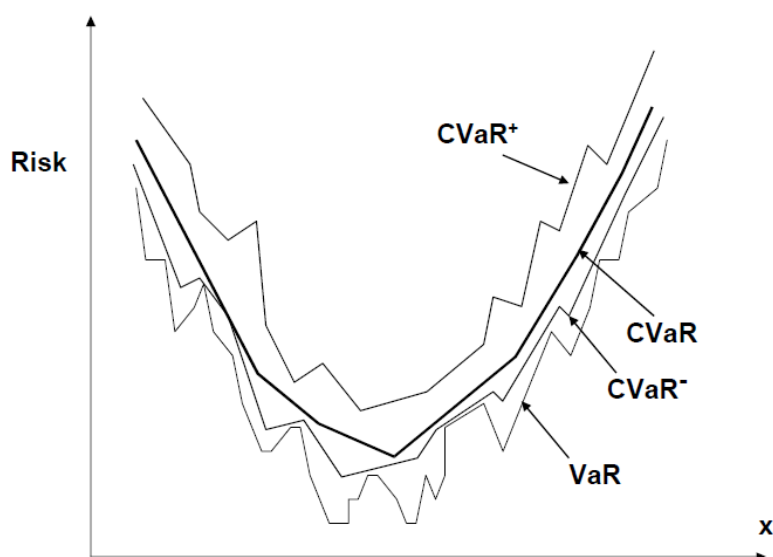

Figure 2.2 VaR, $CVaR^+$, $CVaR^-$ and CVaR. Extract from Uryasev (2004).

2.2 CVaR Optimization vs. MV Optimization

The most important advantage of Mean-CVaR over MV optimization is the fact that it accounts for the non-normality of the expected return distribution. The parameter which controls which part of the return distribution tail is considered when calculating the CVaR, is the Confidence level (e.g., 95% or 99%). Different types of results can be investigated by varying this parameter. A lower value, for example 70% or 60% for the confidence, means that a large part of the distribution is considered. Except for very special situations, one normally sees that in these cases the resulting asset allocations are very similar to the solutions of the MV optimization. Very strong departures from normality are required in order to change the solutions at this confidence level. On the other hand, at the other extreme, when the confidence level is close to unity (e.g., 99% or 99.9%) the results are most sensitive to the non-normal features of the distributions and the resulting allocations could be very different from the MV ones.

Xiong and Idzorek (2011) analyzed the impact of skewness and fat tails on the asset allocation decision in a MV and a Mean-CVaR optimization. They find that compared to MV optimization, Mean-CVaR optimization monotonically underweighs global high yield, U.S. REITs, and commodities because of their more extreme negative skewness and higher kurtosis, and it overweighs non-U.S. government bonds, U.S. nominal bonds, and non-U.S. REITs because of their more attractive combined skewness and kurtosis. The Mean-CVaR optimization tends to pick positively skewed (or less negatively skewed) and thin-tailed assets, whereas the MV optimization ignores the information from skewness and kurtosis. An out-of-sample test showed that the Mean-CVaR outperformed the MV optimization in the financial crisis of 2008.

The mean-CVaR portfolio optimization problem can be formulated in three equivalent variants:

1. Minimize CVaR with a minimum expected return constraint
2. Maximize Expected return with a maximum CVaR constraint
3. Maximize/minimize a linear combination of expected return and CVaR using a constant in order to adjust the relative importance of different terms of the optimization objective.

Using any of the above three methods, one can draw a Mean-CVaR efficient frontier. Note that this efficient frontier will be generally sub-optimal when evaluated in a purely MV framework. Exception is the case when the return distribution is normal, when the two coincide. However, in the case of non-normal return distributions, the MV efficient frontier will be sub-optimal when evaluated using the Mean-CVaR framework and only by calculating the Mean-CVaR efficient frontier can we assess how far it is from the true optimal solution. In our application, we will use the first variant. That is, we set the expected return and calculate the portfolio that minimizes the CVaR given the expected return. We will replicate this procedure for varying expected return levels.

The main drawbacks of the CVaR optimization compared to the MV are twofold: computation time increases and errors due to the finite sampling of the return distributions. The original approach described by Rockafellar and Uryasev (2000), also used in our application, transforms the Mean-CVaR optimization process into a linear optimization problem which can then be solved using general-purpose linear solvers. Note that the number of asset classes is irrelevant, as the problem size increases linearly with N - the number of scenarios used. On the other hand, for MV optimization, the most demanding part is the calculation of the covariance matrix based on the used scenario set, which also scales linearly with N. Once the means vector and the variance matrix are calculated, the generation of the efficient frontier only requires operations which scale according to the number of asset classes, much smaller than N. In conclusion, for a single optimization, the computation effort is similar between CVaR and MV if one uses up to 5,000 scenarios, but the difference increases when calculating an entire efficient frontier.

One of the problems of the CVaR optimization is that the parameters entering the calculation, the asset class expected means and the covariance matrix, need to be calculated using available data obtained through a finite sampling of the presumed true distributions of the asset class returns. This becomes an even larger problem when one uses only the tail of the distribution to define the risk. The use of tail risk measures can be regarded as accounting for higher order moments of the distributions. However, higher order moments are even more difficult to estimate, which makes the error due to finite sampling larger. Alternatively, we can consider the fact that the tail risk is estimated using only a small fraction of the overall distribution sample (1% to 10% typically), which again leads to the conclusion that finite sample errors have a larger effect in this case compared to MV. However, note that the increased estimation errors do not affect CVaR optimization only, they are already present in

the moment when the tail-risk measures are calculated for the purpose of policy evaluation. Therefore, if one decides to accept them for the purpose of tail risk evaluation, they should be equally acceptable for the purpose of optimization based on tail risk.

For a discussion of the way estimation errors affect VaR and CVaR see Acerbi (2004). His study compares the errors affecting the two tail measures and concludes that, even if it is difficult to correctly account the tail of the distributions beyond VaR, as CVaR does, still it is better to do it rather than to ignore them completely as in the case of VaR. In addition, he advocates the use of Extreme Value Theory (EVT) in order to better estimate the parameters which are truly relevant for the accurate description of the distribution tail and for a more proper estimation of CVaR - see Gilli and Këllezi (2006) and references therein for EVT application for VaR and CVaR calculation.

2.3   CVaR in a scenario approach – numerical optimization

CVaR optimization is a very flexible method of portfolio construction. The CVaR portfolio optimization methodology does not make any assumptions on the specific form of the return distributions of the assets which enter the optimization. The problem is numerically feasible as long as a sampling of the expected return distribution is available. The Rockafellar and Uryasev (2000) formulation provides a mathematically sound and computationally efficient way of solving tail-risk based optimization problems. The computational efficiency is particularly related to the use of a scenario based sampling of the expected return distributions. Once a scenario set has been generated it is straight forward to extract asset and liabilities distributions for any time horizons and construct the associated linear optimization problem, including the asset class constraints. Afterwards, this can be solved by a general purpose linear optimization solver. It is possible to extract the distribution at several time horizons and impose maximal CVaR constraints for each of them simultaneously. The dimension of the optimization problems to be solved is equal to N+S, where N is the number of scenarios considered and S is the number of asset classes. Therefore, in normal conditions, the number of asset classes does not impose significant restrictions on the size of the portfolios to be optimized.

In the optimization we use 1,000 scenarios over the time period 2010 - 2030 for all relevant investment categories, interest rates and inflations obtained from the Ortec Finance Dynamic Scenario Generator. For a short explanation of how scenarios are generated in this scenario generator, see the Appendix. Volatilities and correlations in the scenario set are based on historical data. We have used Matlab in combination with Excel to run the optimizations. The Matlab code is based on Rockafellar and Uryasev (2000) and Uryasev (2000).

2.4 CVaR Optimization in an ALM Context

The CVaR risk measure and Mean-CVaR optimization output can be calculated using only the expected return distributions of different asset classes. No assumptions on the exact nature of the distribution or on the underlying economic factors which generate those distributions are required. For this reason, CVaR Optimization can be translated to ALM problems just as easily as the MV optimization. The distribution on which the Mean-CVaR optimization is applied can be set to either:

- Return on assets
- Cumulative return on assets
- Surplus (asset minus liabilities)
- Funding ratio return

where the distribution can be obtained from the set of scenario paths for any desired time horizon.

Bogentoft *et al.* (2001) have applied CVaR optimization to the problem of ALM for pension funds. They have used a path-scenario set and have formulated a multi-stage decision problem. Using a grouping procedure for the different paths, they calculate different allocations for different groups of paths. Asset allocation policies were described either as fixed number of shares or fixed dollar amounts. Another example of the use of CVaR optimization in an ALM context is that of Claessens and Kreuser (2004), in the context of currency reserves problems for central banks. They use a tree based scenario approach and use CVaR in order to impose several risk constraints on the probability density function of the portfolio returns or liquidity.

In the current application, we will apply CVaR optimization to the funding ratio return of a pension fund. The funding ratio is defined as the market value of the assets divided by the market value of the liabilities. The funding ratio return is calculated as:

(return on assets – growth rate liabilities) / (1 + growth rate liabilities)

So, in each year in each scenario, the return on assets is corrected for the growth rate of the (market value of the) liabilities. This is a very relevant measure as pension funds need a positive funding ratio return to be able to meet their liabilities in the long run. Especially when pension premiums cannot rise much further, and when the current funding ratio is close to 100% (so no large buffer). Dutch pension funds currently have low buffers and high premiums. We will apply the CVaR optimization to four types of Dutch pension funds: an average fund, a young fund (i.e., the average age of the participants is relatively low), a mature fund (i.e., the average age of the participants is relatively high), and a so-called sleeping fund (i.e., no active participants anymore, no new premium receipts, only outgoing pension payments).

3 Data

In this study, we want to optimize the sector, region and country allocation within the real estate portfolio of a pension fund. We limit ourselves to the regions Europe, North America and South East Asia. Furthermore, to be included in the analysis the real estate market of a country has to be (reasonably) transparent, the local real estate market should be sufficiently large, and a benchmark with sufficient historical data should be available.

3.1 Europe

For Europe, historical returns are obtained from IPD. Capital growth rates are unsmoothed. On average, unsmoothing increases the standard deviation of the returns with 70%. The countries and sectors that are included are reported in Table 3.1.

Table 3.1: European countries and sectors included in the analysis (marked with a "V")

Country	Retail	Office	Industrial	Residential	Start year
Netherlands	V	V		V	1977 / 1973 (resid)
UK	V	V	V	V	1981
Denmark	V	V		V	2000
Finland	V	V	V	V	1998
France	V	V	V	V	1998
Germany	V	V		V	1996
Ireland	V	V	V		1984
Norway	V	V			2000
Portugal	V	V			2000
Sweden	V	V		V	1997

The IPD reports returns for more European countries and sectors. However, countries and sectors not included in the analysis either are too small, or the benchmarks are based on too few objects or have a too short history. We require at least 10 years of historical returns. Relatively short historical time series of returns is a major problem in South European countries and in upcoming markets in Central and Eastern Europe.

Next to return series per sector and country, we have also established return series per country and for Europe as a whole. The weights for the sectors within a country and for the countries within Europe are based on the IPD estimates of the size of the total market in 2009. The resulting weights are summarized in Table 3.2.

Table 3.2: Sector weights within European countries

Country	Retail	Office	Industrial	Residential
Netherlands	0.28	0.22		0.50
UK	0.52	0.31	0.16	0.01
Denmark	0.25	0.61		0.13
Finland	0.29	0.44	0.08	0.18
France	0.22	0.52	0.09	0.16
Germany	0.29	0.55		0.16
Ireland	0.41	0.51	0.09	
Norway	0.32	0.68		
Portugal	0.70	0.30		
Sweden	0.17	0.73		0.10

For the European region index we exclude the Netherlands. We apply CVaR optimization to Dutch pension funds. Being the home market, Dutch real estate has a dominant position in most real estate portfolios. As such, Dutch real estate is analyzed as a separate category. The region index for Europe excluding the Netherlands is calculated as:

0.035DEN+0.041FIN+0.234FRA+0.288GER+0.004IRE+0.036NOR+0.254UK+0.016POR+0.093SWE

For countries outside the Euro zone, we assume that the currency risk is completely hedged. So in the optimization we use returns in local currencies. Note that we do not use weighted historical returns or indices as input for the scenario model. Each country and sector is modeled separately in the scenario model, and economic scenarios are generated for each country and sector separately. Only at the end do we construct scenarios for the weighted country and region indices based on the scenarios for the underlying components.

Next to unleveraged return series, we also look at indirect investments in unlisted real estate funds with 25% and 50% leverage. In our analysis, we only look at leveraged country funds, and not at leveraged sector funds within countries. We assume that within these funds debt is financed in local currencies with 50% short term and 50% long term financing. For funds with 25% leverage we assume an average credit spread of 150 basis points with a standard deviation of 75 basis points. For funds with 50% leverage we assume an average credit spread of 200 basis points with a standard deviation of 100 basis points.

3.2 North America

For Canada we use IPD data with unsmoothed capital growth returns. For the US we use the MIT Transaction-Based Index for retail, office and industrial and the Case-Shiller index for residential. The sector weights for the country indices are summarized in Table 3.3

Table 3.2: Sector weights within North America

Country	Retail	Office	Industrial	Residential	Start year
Canada	0.37	0.42	0.15	0.06	1985
USA	0.22	0.39	0.22	0.18	1994 / 1890 (resid)

Based on the 2009 estimated market sizes, the North America region index is calculated as:

$$0.935 USA + 0.065 CAN.$$

Leverage is added analogously to Europe.

3.3 South East Asia

Few historical time series for the return on non-listed Asian real estate exist. Only Japan and South Korea are included in the IPD database, and even those for only a few years. For residential real estate in Hong Kong, the University of Hong Kong maintains a (repeat sales) transaction price index which starts in 1991. For other South East Asian real estate markets, we will derive characteristics from listed real estate indices. Table 3.4 summarizes the return series used for South East Asian real estate.

Table 3.4: Return series for SE Asian real estate

Country	Index name	Index short	Start year
Hong Kong	University of Hong Kong Real Estate Index Series	HKU-ARPI	1991
Hong Kong	EPRA/NAREIT Hong Kong index	ELHK Index	1990
China	CHINA SE SHANG PROP INDX	SHPROP Index	1996
Taiwan	TAIWAN TAIEX CONSTRUC IX	TWSECON Index	1994
Malaysia	KUALA LUMPUR PROP. INDEX	KLPRP Index	1993
Singapore	FTSE EPRA/NAREIT Singapore Index	RDSI Index	1994
Thailand	THAI PROPERTY DEV INDEX	SETPROP Index	1994

Apart from the HKU-ARPI (residential) index, no distinction is made between sectors. Furthermore, it is not unlikely that some of the funds that are included in an index also invest in other countries. We cannot correct for that, so the SE Asian indices should be seen as an approximation.

To obtain indices for non-listed unlevered real estate, we correct the listed real estate indices for leverage and stock market sentiment. To determine the influence of leverage and general stock market sentiment, we have estimated the relationship between NAREIT returns (listed, levered) and unsmoothed NCREIF Property Index returns (non-listed, unlevered) for the USA. The following relationship between listed (L) and non-listed (NL) real estate has been estimated for the USA and is applied to the listed SE Asian indices to obtain an estimate for non-listed SE Asian real estate:

$$NL(t) = 0.148 L(t) + 0.308 L(t-1) + 0.358 L(t-2) + 0.186 L(t-3)$$

The resulting standard deviations after correction are summarized in Table 3.5.

Table 3.5: Impact of correction on volatility of SE Asian real estate

Country	Index name	Std.dev. original	Std.dev. corrected
Hong Kong	EPRA/NAREIT Hong Kong index	49.4%	16.6%
China	CHINA SE SHANG PROP INDX	68.3%	37.5%
Taiwan	TAIWAN TAIEX CONSTRUC IX	56.3%	24.8%
Malaysia	KUALA LUMPUR PROP. INDEX	59.3%	22.8%
Singapore	FTSE EPRA/NAREIT Singapore Index	39.2%	17.1%
Thailand	THAI PROPERTY DEV INDEX	63.3%	41.5%

Note that we will not use the historical average return as a proxy for the future expected return. So only the impact of the corrections on the standard deviation and the correlations with other investment categories is relevant. Table 3.6 gives the correlations between the SE Asian real estate returns after correction.

Table 3.6: Correlation matrix for South East Asian real estate returns

	Thailand	Malaysia	Hong Kong	China	Taiwan
Malaysia	0.26				
Hong Kong	-0.43	0.46			
China	-0.48	0.46	0.67		
Taiwan	0.39	0.41	0.22	0.10	
Singapore	0.15	0.55	0.77	0.26	0.60

The country weights for the SE Asia region index are based on the relative size of the investible stock in 2007 (source: Invesco). This gives the following composition of the region index:

$$0.057Thai + 0.033Mal + 0.083HKG + 0.644CHN + 0.135TWN + 0.048SIN$$

3.4   Expected future returns

We assume that the average geometric annual return over the coming 20 years equals 5.5% for all real estate sectors and regions. Consequently, the expected arithmetic annual return depends on the volatility of the return series[3]. The expected arithmetic return is higher with more volatile return series. Consequently, for Germany, where historical volatilities have been low, the expected arithmetic return is only slightly above 5.5%. The highest expected arithmetic returns are for Thailand and China (about 12%).

---

[3] Without autocorrelation the arithmetic return is approximately equal to the geometric return plus half the variance.

4 Results

In our optimizations, we assume a fixed portfolio for the non-real estate part consisting of:

- 66.7% Eurozone government bonds (average duration 5);
- 13.3% Equity Europe;
- 13.3% Equity US hedged;
- 3.3% Equity Japan unhedged;
- 3.3% Equity emerging markets unhedged.

This is a representative portfolio for an average Dutch pension fund. For the average pension fund, this portfolio leads to an average funding ratio return (geometric) of 0.51% with a standard deviation of 0.92% over a 20-year horizon.

We have optimized the composition of the non-listed real estate portfolio with both a total allocation of 10% and 20% in real estate. As the results were similar, we will only show the 20% allocation results. We have optimized the CVaR of the 5% left tail of the funding ratio return distribution. We will first show the results when we limit the optimization to direct unlevered investments in European real estate. In the second subsection we look at indirect continent funds with leverage.

4.1 Direct investments in European real estate

We first report the efficient frontier when we limit the analysis to European country funds without leverage. The results are presented in Table 4.1. This table shows the 20-year mean-CVaR optimal allocation to countries within the real estate portfolio for several mean funding ratio return levels. To obtain the weights of individual countries in the total investment portfolio, the weights in the table have to be multiplied by 20%.

Only five countries appear in the optimal portfolio for all types of fund and for all mean funding ratio return levels. Apart from the most risky portfolios, the UK receives a high allocation in all cases. In the most defensive portfolios, Denmark also receives a high allocation. The Netherlands mainly appears in portfolios for the average and young pension fund; and Germany only in the defensive optimal portfolios for the average pension fund. Ireland receives significant allocations in the most risky parts of the efficient frontiers. Such high allocations in this small economy which currently also suffers from (amongst others) a major real estate crisis might be too much. Later in this section we will show the results when we restrict ourselves to the large European countries.

Table 4.2 shows the efficient allocations when we zoom in to sectors within European countries for the average Dutch pension fund. In this case, Denmark drops out of the efficient portfolio. The other countries present in Table 4.1 also appear in Table 4.2 with at least one sector. Interesting to note is that the optimal portfolios in Table 4.2 are less favorable than those in Table 4.1. So optimizing over individual European sectors leads to a lower mean return and a higher risk than optimizing over aggregate European country indices. The Irish office market receives high allocations in the more risky optimal portfolios. These allocations might not be realistic given the size and characteristics of the current Irish office market. Table 4.3 summarizes the results when we restrict ourselves to the larger European countries. In that case, only one portfolio remains optimal when we use a 0.01% step size for the mean funding ratio return: the lowest risk portfolio in Table 4.2; so Irish offices are not replaced by other sectors. Only for the young pension fund, Swedish offices enter the optimal portfolio. So again, a relatively limited number of sectors appear in the optimal portfolios. The weights of these sectors can, however, differ significantly between different pension funds.

Table 4.1: optimal mean-CVaR portfolios, 20-year horizon, unlevered European countries only.

Panel A: Average Dutch pension fund

Mean fr ret.	NL	UK	IRE	GER	DEN	Std.Dev.	5% CVaR	5% VaR
0.76%	0.0%	48.5%	0.0%	32.5%	19.0%	0.86%	-0.94%	-0.58%
0.77%	4.0%	52.5%	0.0%	18.5%	25.0%	0.87%	-0.95%	-0.59%
0.78%	17.0%	59.0%	0.0%	0.0%	24.0%	0.88%	-0.97%	-0.60%
0.79%	19.6%	73.9%	3.0%	0.0%	3.5%	0.90%	-1.00%	-0.65%
0.80%	5.5%	67.0%	27.5%	0.0%	0.0%	1.03%	-1.12%	-0.69%

Panel B: Young Dutch pension fund

Mean fr ret.	NL	UK	IRE	GER	DEN	Std.Dev.	5% CVaR	5% VaR
0.77%	0.1%	53.3%	0.0%	2.6%	44.1%	0.89%	-0.98%	-0.64%
0.78%	9.9%	68.5%	0.0%	0.0%	21.7%	0.90%	-0.99%	-0.65%
0.79%	15.9%	77.2%	7.0%	0.0%	0.0%	0.92%	-1.02%	-0.66%
0.80%	0.0%	59.9%	40.1%	0.0%	0.0%	0.95%	-1.06%	-0.70%

Panel C: Sleeping Dutch pension fund

Mean fr ret.	NL	UK	IRE	GER	DEN	Std.Dev.	5% CVaR	5% VaR
0.76%	0.0%	43.9%	0.9%	0.0%	55.2%	0.63%	-0.52%	-0.23%
0.77%	0.0%	48.6%	12.6%	0.0%	38.9%	0.64%	-0.54%	-0.25%
0.78%	0.0%	57.9%	20.6%	0.0%	21.6%	0.66%	-0.57%	-0.26%
0.79%	0.0%	68.0%	27.8%	0.0%	4.2%	0.69%	-0.60%	-0.29%
0.80%	0.0%	10.0%	90.1%	0.0%	0.0%	0.84%	-0.72%	-0.35%

Panel D: Mature Dutch pension fund

Mean fr ret.	NL	UK	IRE	GER	DEN	Std.Dev.	5% CVaR	5% VaR
0.77%	0.0%	47.2%	0.0%	0.0%	52.9%	0.72%	-0.68%	-0.39%
0.78%	2.7%	62.8%	2.1%	0.0%	32.3%	0.73%	-0.70%	-0.40%
0.79%	2.9%	71.7%	10.5%	0.0%	15.0%	0.75%	-0.73%	-0.41%
0.80%	1.3%	75.3%	23.4%	0.0%	0.0%	0.77%	-0.76%	-0.43%
0.81%	0.0%	1.5%	98.6%	0.0%	0.0%	0.95%	-0.90%	-0.52%

Table 4.2: optimal mean-CVaR portfolios, 20-year horizon, unlevered European sectors, average Dutch pension fund

Mean fr return	NL retail	UK resid.	UK retail	IRE office	GER office	GER retail	Std. Dev.	5% CVaR	5% VaR
0.71%	3.2%	3.9%	31.1%	0.0%	13.6%	48.3%	0.84%	-0.95%	-0.62%
0.72%	11.4%	17.3%	26.2%	25.2%	20.0%	0.0%	0.87%	-1.00%	-0.65%
0.73%	8.0%	12.4%	14.1%	65.5%	0.0%	0.0%	0.98%	-1.14%	-0.76%

Table 4.3: optimal mean-CVaR portfolios, 20-year horizon, large unlevered European countries only.

Country	Average	Sleeping	Mature	Young
NL retail	3.2%	0.0%	1.1%	11.2%
UK residential	3.9%	20.1%	25.0%	17.9%
UK retail	31.1%	11.8%	14.8%	33.9%
GER office	13.6%	0.0%	59.2%	21.9%
GER retail	48.3%	68.1%	0.0%	8.6%
SWE office	0.0%	0.0%	0.0%	6.6%
Mean funding ratio return	0.71%	0.71%	0.72%	0.72%
Standard deviation	0.84%	0.61%	0.70%	0.88%
CVaR 5.0%	-0.95%	-0.53%	-0.68%	-0.99%
VaR 5.0%	-0.62%	-0.27%	-0.41%	-0.68%

4.2 Continent level indirect

In our next analysis, we want to determine optimal portfolios of levered and unlevered continent funds. In an unrestricted optimization, the optimal portfolios contain European continent funds with 25% and 50% leverage and the SE Asian continent fund without leverage. The unlevered SE Asia fund dominates the portfolio with a 50% allocation in the most defensive profile to almost 100% in the most aggressive portfolio. However, investing in unlevered SE Asian funds might be very difficult when you are not a local player. The optimal portfolios when the minimum leverage for SE Asian funds is set to 25% are summarized in Figure 4.1. Note that North America does not appear in any of the optimal portfolios. When we further restrict the SE Asian funds to 50% leverage, only 25% and 50% levered European funds appear in the optimal portfolios.

Figure 4.1 only shows the results for the average Dutch pension fund. For the other types of fund, the more risky portfolios are very similar to those of the average fund. The least risky portfolios contain a somewhat lower allocation to European funds with 25% leverage and a somewhat higher allocation to European funds with 50% leverage.

Figure 4.1: optimal weights per region, average fund
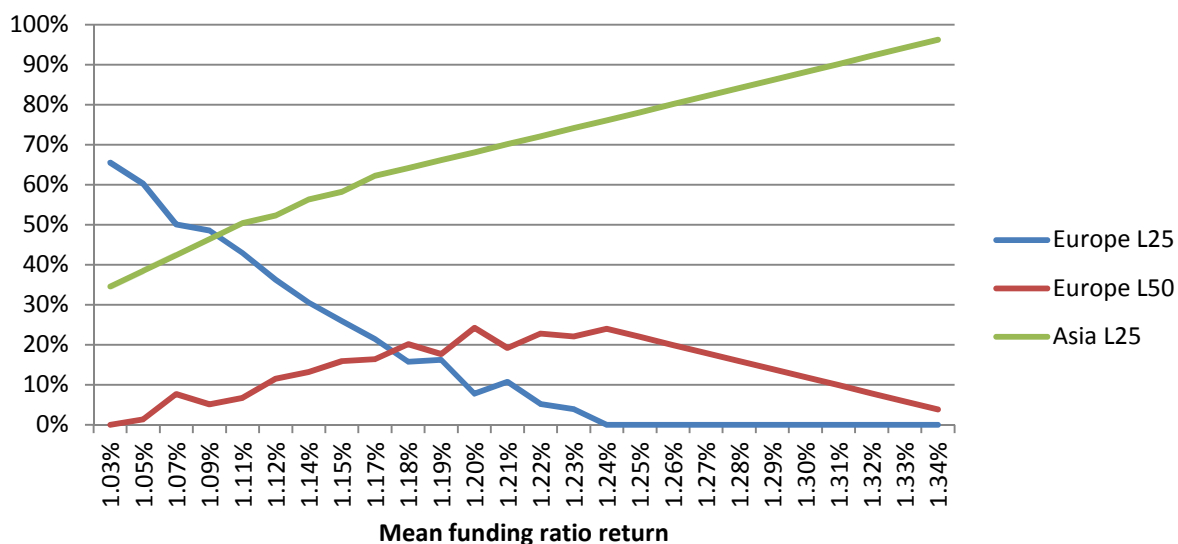
4.3 Country level indirect

As a final example we will show the optimization results when we directly optimize at the country level. Furthermore, apart from the Netherlands all investments will be in indirect, levered country funds. For the Netherlands, being the home country, we add the restriction of a minimum allocation of 12.5%. The results are summarized in Table 4.4. Our main conclusions are:

- The allocation to Dutch direct real estate never rises above the minimum of 12.5%.
- Within Europe, only Germany, Finland and Denmark enter the optimal portfolios.
- Apart from the Hong Kong residential sector, SE Asian countries do not enter the optimal portfolios when the minimum leverage is set to 25%.
- The USA does not enter the optimal portfolios and Canada only in one case.

Note that more European and SE Asian countries would enter the optimal portfolio when direct, unlevered investment is possible for all countries. And the allocation to Canadian real estate would be higher. But the USA would still not enter the optimal portfolios.

Table 4.4: optimal mean-CVaR portfolios, 20-year horizon, country level, only NL direct.

Country	Average	Sleeping	Mature	Young
Netherlands	12.5%	12.5%	12.5%	12.5%
Germany L25	0%	0%	14.5%	0%
Germany L50	0%	26.0%	0%	0%
Finland L25	38.5%	0%	0%	0%
Finland L50	25.0%	51.5%	47.0%	59.5%
Denmark L50	17.0%	9.5%	25.0%	24.0%
Canada L25	0%	0%	1.0%	0%
Hong Kong residential L25	7.0%	0.5%	0%	4.0%
Mean funding ratio return	0.71%	0.71%	0.71%	0.71%
Standard deviation	0.83%	0.57%	0.67%	0.84%
CVaR 5.0%	-0.92%	-0.44%	-0.64%	-0.93%

Germany, Finland and Denmark are the three countries in the data set with the smallest funding ratio return volatility. See Table 4.5. So high leverage is only acceptable when the underlying real estate market is very stable. For more volatile real estate markets, leverage should be kept to an absolute minimum. This is also why SE Asian real estate markets only enter the optimal portfolios with a significant allocation when it is possible to invest in unlevered funds or at least funds with very limited leverage.

The differences between optimal and non-optimal portfolios are often small. To analyze the sensitivity of the results for changes in the input, we have rerun the optimizations for a 15-year instead of a 20-year horizon. The results are summarized in Table 4.6.

Table 4.5: annualized geometric 20-year funding ratio return per country ordered by volatility, standard pension fund.

Country	Mean	Std. Dev.	Country	Mean	Std. Dev.
Germany	1.5%	0.9%	USA	2.0%	1.8%
Finland	1.5%	0.9%	UK	1.9%	2.0%
Denmark	1.7%	1.0%	Ireland	1.9%	2.1%
Portugal	1.6%	1.1%	Hong Kong	1.6%	2.8%
Canada	1.6%	1.1%	Taiwan	1.6%	2.9%
Norway	1.6%	1.2%	Singapore	1.6%	3.0%
France	1.6%	1.2%	Malaysia	1.6%	3.4%
Sweden	1.6%	1.6%	China	1.7%	5.3%
Netherlands	1.7%	1.7%	Thailand	1.8%	7.1%

Table 4.6: optimal mean-CVaR portfolios, 15-year horizon, country level, only NL direct.

Country	Average		Sleeping		Mature		Young	
NL	12.5%	12.5%	12.5%	12.5%	12.5%	12.5%	12.5%	12.5%
USA L25	0%	53%	0%	50%	0%	51%	0%	53%
GER L50	10.5%	0%	6.5%	0%	15.5%	0%	22.0%	0%
FIN L50	77.0%	0%	81.0%	0%	72.0%	0%	65.5%	0%
DEN L50	0%	34.5%	0%	37.5%	0%	36.5%	0%	34.5%
Mean fr return	0.50%	0.60%	0.58%	0.68%	0.56%	0.66%	0.50%	0.60%
Std. Dev.	1.08%	1.19%	0.82%	0.94%	0.92%	1.04%	1.15%	1.26%
CVaR 5.0%	-1.61%	-1.87%	-1.01%	-1.30%	-1.24%	-1.51%	-1.73%	-2.00%

When we compare Table 4.6 with Table 4.4, we can see that the European countries in the optimal portfolios are unchanged. Canada and Hong Kong, however, have disappeared from the optimal portfolios. In the high risk portfolios, the USA fund with 25% leverage receives a significant allocation. The low risk portfolios show a high allocation to Finland and for the rest Germany and the Netherlands. The high risk portfolios contain Denmark, the USA and the Netherlands.

Based on this sensitivity analyses, we can conclude that the optimal portfolios are sensitive to the investment horizon. A number of "base" countries will always be in the optimal portfolios, but other countries will appear and disappear depending on the assumptions.

5 Concluding remarks

Conditional value at risk (CVaR) Optimization is a flexible method of portfolio construction which uses the average of the tail of the distribution function in order to characterize the risk, instead of distribution variance. Major advantages of this method are:

- It is particularly suited for the use with scenario based sampling of complicated expected return distributions.
- It can handle return distributions that are strongly non-normal and/or discontinuous.
- It allows for the combination in the same optimization problem of several CVaR restrictions with different time horizons and/or confidence levels.
- It translates well to an ALM context
- A downside risk measure like CVaR better matches the risk perception of institutional investors like pension funds than a volatility risk measure (like variance) that also takes the upside into account.

On average, pension funds pay out their pensions after 15 to 20 years. They, therefore, used to have a long-term investment horizon. The optimal portfolios presented in this paper are also based on this long investment horizon. However, as funding ratios are calculated based on market values nowadays, the short term volatility of funding ratios has increased dramatically. That is why pension funds also have to take shorter horizons into account. The relevance of steering on short horizons has increased further in recent years due to the extremely low current funding ratios caused by low interest rates and dissatisfactory (negative) equity returns. So both long and short horizons are of crucial importance. With a longer investment horizon, an investor can invest more in risky assets. Short term optimal allocations will be more defensive and with more focus on liquidity than long term optimal allocations. Under the current circumstances and with the currently low coverage ratios, short term restrictions will be extremely important for many pension funds. These short term restrictions are not taken into account in the current paper but are a topic for further research. The same holds for the impact of the illiquidity of non-listed real estate holding on optimal portfolios, both on a short and a long investment horizon.

For a long investment horizon, we conclude that within Europe the number of countries and sectors that appear in the CVaR optimal portfolio is limited. So the diversification gain from investing in a large number of sectors and countries is relatively limited within Europe. For countries outside Europe, the results are sensitive to changes in the input or assumptions. Outside Europe, investing in funds with 50% leverage is unwise. Inside Europe, 50% leverage is possible, but only for the most stable countries. So high leverage is only acceptable when the underlying real estate market is very stable. Finally, although the optimal weights differ per type of pension fund, the countries and sectors that appear in the optimal portfolios are quite stable.

References

Acerbi, C. (2004). Coherent representations of subjective risk aversion. in G. Szegö (Ed.) *Risk Measures for the 21st Century*. New York: Wiley, 147–207.

Amenc, N., F. Goltz and A. Lioui (2011). Practitioner Portfolio Construction and Performance Measurement: Evidence from Europe. *Financial Analysts Journal 67*, 39–50.

Artzner P., F. Delbaen, J.M. Eber, and D. Heath (1999). Coherent measures of risk. *Mathematical Finance 9*, 203–228.

Bogentoft E., H.E. Romeijn and S.Uryasev (2001). Asset/Liability Management for Pension Funds Using CVaR Constraints. *The Journal of Risk Finance 3*, 57–71.

Claessens, S. and J. Kreuser (2004). A Framework for Strategic Foreign Reserves Risk Management. In *Risk Management for Central Bank Foreign Reserves*. Frankfurt am Main: European Central Bank Publication.

Gilli, M. and E. Këllezi (2006). An application of extreme value theory for measuring financial risk. *Computational Economics 27*, 123.

Mausser H. and D. Rosen (1998). Beyond VaR: From Measuring Risk to Managing Risk. *ALGO Research Quarterly 1*, 5–20.

Rockafellar R.T. and S. Uryasev (2000). Optimization of conditional value-at-risk. *The Journal of Risk 2*, 21–41.

Rockafellar R.T. and S. Uryasev (2002). Conditional value-at-risk for general loss distributions. *Journal of Banking and Finance 26*, 1443–1471.

Steehouwer, H. (2010). A Frequency Domain Methodology for Time Series Modeling. In Berkelaar, A., J. Coche and K. Nyholm (Eds.), *Interest Rate Models, Asset Allocation and Quantitative Techniques for Central Banks and Sovereign Wealth Funds*, Palgrave Macmillan.

Uryasev, S. (2000). Conditional Value-at-Risk: Optimization Algorithms and Applications. *Financial Engineering News 14*, February,1–5

Uryasev, S. (2004). *Conditional Value at Risk, Methodology and Applications: Overview*. Presentation at the Integrated Risk-Return Management: New Approach to Management of Bank Portfolios Workshop, July 2004, New York.

Xiong, J.X. and T.M. Idzorek (2011). The Impact of Skewness and Fat Tails on the Asset Allocation Decision. *Financial Analysts Journal 67*, 23–35.

Appendix: The scenario model

Scenarios are future trajectories modeling the external insecurities that decision makers must take into account in their policy determination and evaluation. Scenario analysis is gaining ground as the preferred method for gaining insight into risk and return of investments worldwide. In some countries it has reached a level of acceptance that has led the regulatory bodies to include it in their framework: the Dutch regulator (the Dutch Authority for the Financial Markets, AFM) for example has defined use of scenario analysis as best practice.

With the use of a corporate model of the company we can calculate, for every year and each scenario, what the consequences of the policy intentions are. The scenarios used should be realistic and plausible and therefore give a reliable representation of possible future outcomes. That is, scenarios generated by the model should contain all known real world features and dynamics. Examples of well known real world features as reported by the academic literature are:

- Term structure of risk and return: Risk and return properties such as means, volatilities, correlations and distributions vary with the investment horizon. For instance: the correlation between equity returns and inflation is around zero for short horizons, but it increases to over 0.5 on a 30 year horizon.
- Business cycle dynamics: For example, stock prices tend to lead in the business cycle while real estate returns typically lag in the business cycle.
- Volatility dynamics: Volatility of an asset class is not a fixed characteristic. It is dependent on return levels. Typically, the correlation between the actual return and volatility is negative for equity (high volatility and bad returns tend to happen together).
- Tail risk: correlations between asset classes increase in the left tails of the distribution. Consequently, the benefits of diversification disappear when it is needed most, in a crisis.
- Non-normal distributions: Distributions typically do not resemble the Normal distribution, but are skewed and have fatter tails.

Ignoring one or more of these real world features may impact the outcome of an analysis, resulting in incorrect conclusions and actions.

In order to consistently account for the economic risks, we separately assess the long term trend, medium term (business cycle) and a short term (monthly) components by using frequency domain techniques[4]. The central idea is that these components contain information about the long, medium and short term dynamics of, for instance, the stock price index (returns) or the interest rates. A correct representation of these dynamics in the scenarios is then achieved by modeling these components separately and through recombining the scenarios so generated (the stock price index or the interest rate). The underlying idea for the methodology is illustrated in Figure A.1.

In Figure A.1 a long interest rate series has been split up into a long term (low frequency), medium term (business cycle) and short term (monthly) component. The components are uncorrelated and add up to the original series. Amongst others, figure A.1 shows that high (low) interest rates and high (low) interest rate volatility often occur simultaneously. See for instance the patterns in the post-WWI period and in the 1970s to 1990s period. This time varying (level-dependent) volatility can also be taken into account in the scenario model. Our approach has additional advantages:

- There is no loss or suppression of long term information: both long and short term fluctuations are visible.
- Appropriate data can be used for the relevant aspects of capturing the behavior of the variables needed: for example, long term (annual) series for long term behavior and higher frequency (monthly, weekly or even daily) short term series for short term behavior. This way the data used is in line with the investment horizon.
- Empirical behavior in all horizons can be modeled simultaneously due to separate modeling of the various components. Given that the output is uncorrelated, separate models for the various components can be devised. In the end the results from the component models can then be reassembled again.

---

[4] The frequency domain methodology consists of a number of statistical and econometric techniques such as spectral analysis, frequency decomposition (filtering) and frequency restricted stochastic processes. The aim is to describe all the aspects of the time series behavior of economic variables at the same time, rather than focusing on a subset of aspects (e.g. only the long term properties). A more technical description of how this is done can be found in Steehouwer (2010).

- Time varying (i.e., stochastic) volatility can also integrated in this decomposition framework by using long term realized volatility time series. This enables high dimensional modeling of long and short term volatility dynamics including mutual and return correlations.
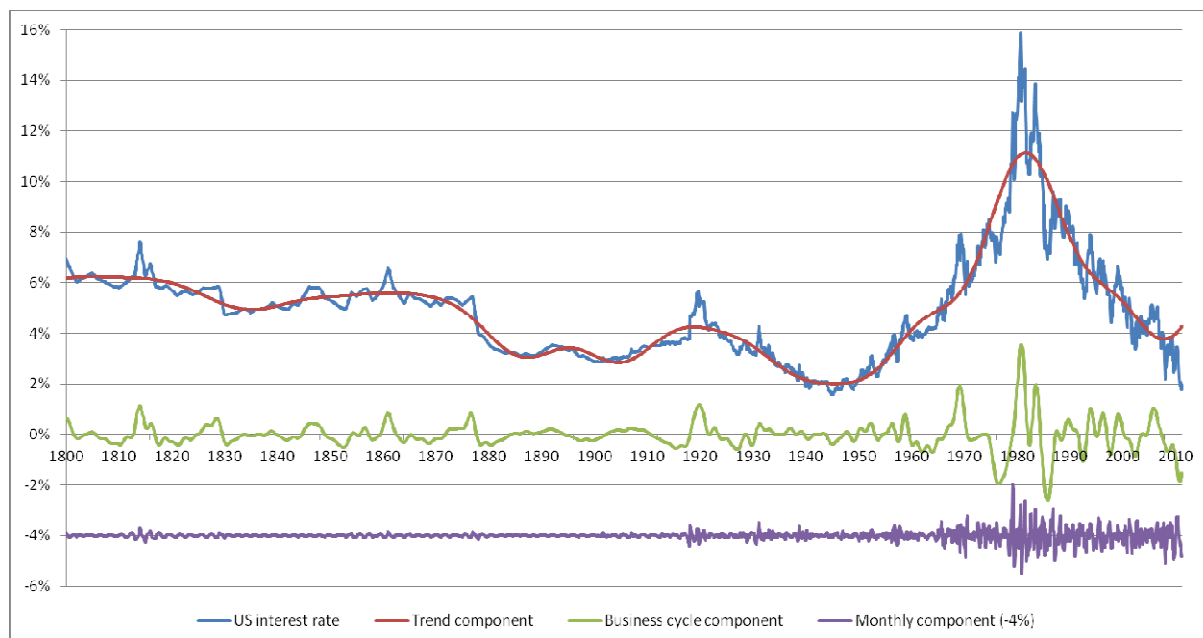

Figure A.1: Frequency decomposition of the US long interest rate

With this methodology it is possible to construct scenarios that provide a better joint description of the empirical behavior over time of economic and financial variables, for different horizons so providing insight in and understanding of the corresponding dynamic behavior. In this way, we can generate plausible scenarios for developing strategies to meet long, medium and short term goals.

ORTEC

FINANCE