Dynamic asset allocation

Guaranteed defined contribution pension schemes are faced with the tough challenge of maximising pensions while remaining solvent at all times. Aggressive investment strategies can be counter productive and destabilise fund operation, so a balance must be struck between operational stability and investment return.

BY MICHAEL PREISEL, SØREN F.JARNER AND RUNE ELIASEN

Introduction

For defined-contribution (DC) pension schemes which combine a minimal return guarantee with a bonus option – sometimes denoted with-profits pension schemes – the objective of the company is to generate excess returns while staying funded at all times.

Excess returns are not necessarily to the immediate benefit of clients, though, since surplus can be kept within the fund as a collective surplus. It is when a bonus is first declared that funds are actually made available to clients. The stability and efficiency of the conversion of investment outcomes to clients' consumption is therefore very important. It is therefore of paramount importance that pension funds operate close to some equilibrium state. Otherwise some generations of clients might directly finance other generations' pensions or vice versa.

The analysis of the long-term stability of pension companies is done in three steps. First we develop a model for the operation of a pension company and discuss the tools available to management to control the company. We then explore designs of dynamic investment strategies – which we denote dynamic rules – as a tool of control. Finally we demonstrate that long-term stability is not guaranteed by "just" avoiding insolvency. This is because DC schemes basically operate in one of two modes: A 'troubled' mode, where investment risk must be curtailed and reserves rebuilt; and a 'surplus' mode, where investment returns can safely be spent as bonus. To maximise the time spent in 'surplus' mode, pension companies should consider putting an absolute cap on equity allocation.

Model Company

We will assume our model company to operate under fair value accounting standards where assets are marked-to-market and liabilities are reported at their discounted value using a market term structure. Hence the company is solvent when its funding ratio (the asset/liability ratio) is 100% or above. The initial funding ratio is set to 115%.

The company has run a defined contribution scheme in which each client bought with-profits, whole-life annuities [1] (or guaranteed annuity obligations (GAOs)) at a 5% guaranteed rate for €1,000 a year from age 20 till retirement at age 65. Each year a new cohort of 1,000 20-year-olds has entered the company. The scheme has been in operation for so long that the client population has reached its steady state[1] [2]. For simplicity we assume that realised and expected mortality are the same and we neglect effects of longevity, taxation, and inflation which are all assumed to be zero. Adding some – or all – of these effects would put more pressure on the company to produce (higher) returns to cover the increased cost, but would not change the dynamics of the model fundamentally.

To emphasise the irreplaceable nature of free reserves we will assume that the scheme is now closed for new contributions. Thus there is neither a sponsor nor new contributions to restore free reserves once lost. Further, we will assume that free reserves belong to the clients and no further (equity) capital is required.

Simulations

The management of the company has to make decisions under the uncertainty of a capital market. For simplicity we have reduced the market to three asset classes only: a swap market, cash and equity (see the appendix). Liabilities are marked-to-market off the swap curve and a (swap) hedge programme is implemented to match the duration of liabilities at all times, cf. [3] and references therein.

The computations below are based on 10,000 simulated capital market scenarios reaching 10 or 30 years into the future. In each scenario, the evolution of the company is computed. This includes the computation of a full balance in each node (quarter) of each scenario.

Tools of Control

A main objective of management is to ensure the company stays funded at all times. Thus the company must decide an acceptable level of long-term– or strategic – risk at which the company should operate. To quantify long-term risk, we introduce the probability of (path-wise) insolvency on a 10-year horizon as a key strategic risk measure. For each scenario it is checked whether the company has remained solvent or not. If the company at some point – maybe just a single term – has been insolvent, the scenario is marked. The probability of path-wise insolvency is then computed as the fraction of marked scenarios to the total number of scenarios.

The process leading to this decision is highly individual and must express the essence of each company's values. It is of utmost importance, though, since it sets the overall boundaries for the operation of the company and therefore – ultimately – the obtainable levels of pensions. In the following we shall assume the model company has decided to accept a maximum of 1% strategic risk. The tools to control this risk are the investment policy and the bonus policy.

Investment Policy

To be of operational value, an investment policy must be expressed as a day-to-day risk budget. We shall use the short-term probability, p, of loosing free reserves within the next 3 months as our short-term risk measure. Thus, the short-term risk measure involves a simultaneous simulation of assets and liabilities one single time-step ahead.

As we demonstrate in the next section, a short-term risk budget is sufficient to control strategic risk, but as every period's P/L directly feeds into next period's risk budget, this leads to excessive trading if unconstrained. It is therefore advisable to supplement the short-term risk measure by a drag on trading.

We shall further introduce an equity cap to put an absolute upper limit to the equity allocation – even when free reserves allow a larger allocation – to control absolute volatility of free reserves. As we demonstrate later, this is a vital parameter to control long-term stability.

Bonus Policy

The bonus decision can basically be broken down in two decisions: when to attribute bonus and how much to attribute. In our model company, bonus is attributed once a year if the funding ratio, F, at that date is above a given

1 We use a unisex Makeham mortality law (truncated at age 120) with parameters $A=5\cdot10^{-4}$, $B = 7.5858\cdot10^{-5}$ and $c=1.09144$. The parameters correspond to those used in the Danish G82M mortality table used by the industry till recently.

Figure 1. Strategic risk vs. short-term risk

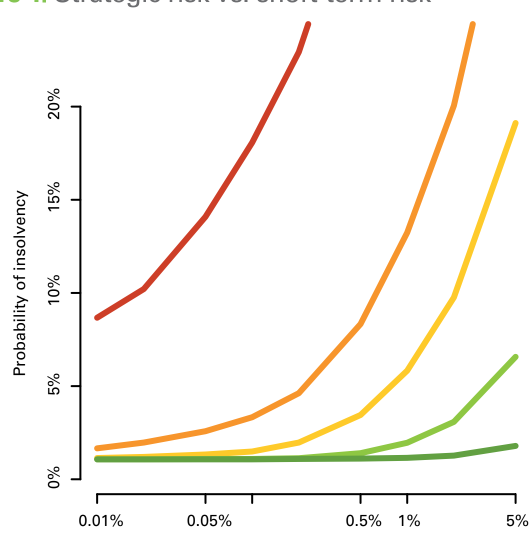

Worst-case estimates of 10Y insolvency risk for no risk buffer (red) and risk buffers of 1% (orange), 2% (yellow),5% (light green) and 10% (dark green) as a function of short-term risk, p.

Figure 2. Reversion risk

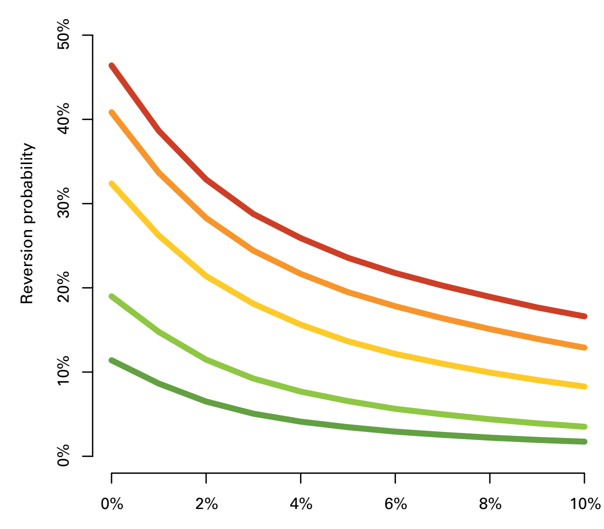

Probability of reverting a decision to increase equity exposure the following term without a buy limit (red),and with buy limits of 10% (orange), 5% (yellow), 2% (light green) and 1% (dark green) as a function of equity headroom.

Figure 3. Dynamic rule

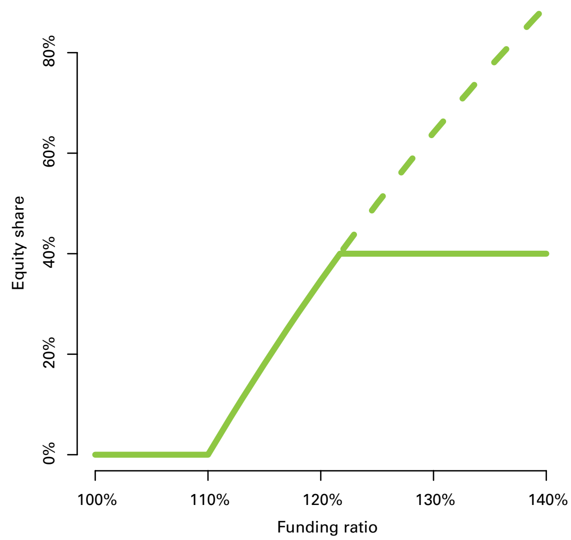

Maximal equity allocation under the dynamic rule with a risk buffer of 10%, short-term risk tolerance p = 0.5% and an equity cap at 40% (solid line), respectively, no cap on equity (dashed line).

Figure 4. Funding ratio

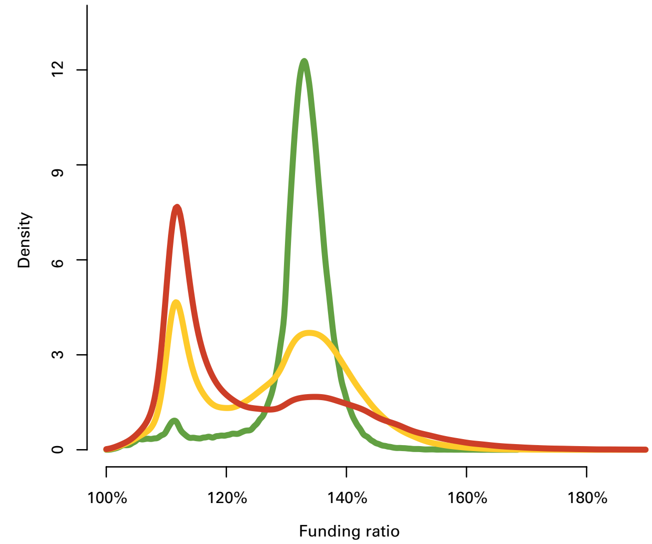

Distribution of funding ratio in steady state with a bonus threshold of 130% and equity caps of 10% (green), 30% (yellow) and 50% (red).

bonus threshold, F_T. The bonus threshold therefore is the 'when' of the bonus policy.

All pension rights are increased by the same bonus percentage, b, such that the funding ratio after bonus attribution becomes $F/(1 + b)$. The maximal bonus percentage is the one that brings the funding ratio down to the bonus threshold. However, to amortise bonus over time the actual bonus percentage is set to a fraction, α, of the maximal percentage:

1 $$b = \alpha\left(\frac{F}{F_T} - 1\right) \text{ for } F > F_T.$$

In other words, the bonus fraction, α, is the 'how much' of the bonus policy. Both parameters are very significant to the development of the pension fund but we will restrict the analysis to only vary the bonus threshold keeping α=1/2 fixed.

Dynamic Rule

As introduced in the previous section, an investment strategy must satisfy a number of criteria to effectively control a pension company. We will discuss such criteria in terms of a hierarchy of rules which together regulate the maximum equity allocation at any point in time. We shall denote the collection of these sub-rules the *dynamic rule*.

Risk reduction rule

From a practical point of view, the absolutely necessary risk controlling parameter to be decided is the maximum tolerable short-term risk, p, since it is the critical link between the day-to-day risk budget and acceptable (long-term) strategic risk. At any point in time the maximum allowed financial risk is therefore given as the equity allocation for which the probability of the funding ratio next quarter falling below 100% is at most p. If – at the end of a quarter –the funding ratio and current equity allocation lead to a probability higher than p, the equity allocation must be reduced, ultimately to 0%, until the probability gets below p.

To demonstrate the isolated effect of this rule, we show in Figure 1 strategic risk (10Y path-wise insolvency) versus short-term risk tolerance, p, disregarding bonus (and additional sub-rules described below).

It is quite clear the objective of at most 1% strategic risk (red line) cannot be satisfied for the percentiles, p, evaluated. An even lower percentile of p would therefore be necessary to satisfy the strategic risk tolerance.

However, using very small probabilities as the basis for risk control is inadvisable. It is therefore better to deduct a risk buffer, R, from free reserves such that the short-term risk measure becomes the probability of the funding ratio falling below 100%+R on a three month horizon.

Figure 1 therefore also compares strategic risk for risk buffers of R=1%, 2%, 5% and 10%. Introduction of the buffer clearly reduces strategic risk markedly. There is no universal guide as to the choice of the pair (p,R) (short-term risk tolerance and risk buffer, respectively) and in practice a company must decide from implementational ease and computational efficiency. In the following we shall assume a short-term risk tolerance of p=0.5% and a risk buffer of R=10% satisfying the requirement of a strategic risk about 1%.

Step-up rule

To reduce strategic trading we put a drag on increases in equity allocation by introducing a 'headroom' needed to increase equity exposure, that is, equity allocation is increased only if the short-term risk tolerance, p, allows an extra 5% in equity. Alternatively, a buy limit which caps the maximal increase in equity allocation each quarter would have the same effect.

To quantify the effect of these drags we introduce the probability of immediately reverting a buy decision to a sell decision the following term as a measure of robustness. This is illustrated in Figure 2 where reversion risk is

Table 1: Steady state properties of model company for combinations of equity cap and bonus threshold

Bonus threshold	120				130				140			
Equity cap	**10**	**30**	**50**	**70**	**10**	**30**	**50**	**70**	**10**	**30**	**50**	**70**
Mean equity	9	19	21	22	10	24	28	29	10	25	33	34
Prob. at cap	93	45	20	10	95	70	40	23	95	79	53	33
Mean bonus	1.2	1.5	1.6	1.6	1.2	1.7	1.8	1.9	1.1	1.7	1.9	2
Prob. bonus	77	43	31	28	81	52	34	25	80	56	38	28
Mean fund	121	119	118	118	132	128	124	122	141	137	131	128

Row 3) Mean allocation to equity, 4) probability of equity allocation being at the equity cap, 5) mean annual bonus percentage, 6) probability of attributing bonus, and 7) mean funding ratio. All numbers are percentages.

shown for combinations of headroom and buy limit (assuming no bonus). We first note that with neither mechanisms in place, the reversion probability is just below 50% reflecting the fact that although equity prices are generally increasing in any one term, the probability of a drop in equity prices is just below 50%. Reversion risk drops sharply when equity headroom is just a few percent, as this creates a buffer which can absorb the many small equity losses. The buy limit obviously also has a strong effect on reversion risk although introduced for a different purpose. For instance, even without an equity headroom, a buy limit of 2% point each term reduces the reversion risk to only 20%. In the following we shall only use a headroom of 5%.

Equity Cap

The rules for risk reduction and risk increase described so far control the timing and speed with which the company moves its funds in and out of equity but do not limit the overall exposure. As already indicated, the following section will demonstrate there is good reason also to put an absolute cap on the equity allocation. In Figure 3 we illustrate how an equity cap of 40% will limit equity allocation under the dynamic rule. The maximum allocation is reached for funding ratios just above 120%, indicating the bonus threshold should be set at or above this level since a lower threshold could trigger forced reductions in equity when the bonus is attributed.

Collecting it all

The effect of the dynamic rule can be demonstrated by the following example: Assume a term has just passed in which equity markets rallied. This has resulted in a current equity share of, say, 35%.

Risk-reduction: As a result of the successful investment outcome free reserves have risen and the funding ratio is now 125%. The short-term risk, that is, the probability of the funding ratio falling below 110% next quarter, is now only $p=0.001\%$. Conversely, the maximum allowed equity allocation can be computed to 50% since this corresponds to $p=0.5\%$.

Step-up rule: In principle, an increase in equity allocation of 15% is therefore allowed but trading is reduced[2] to 10% since a headroom of 5% has been decided.

Equity cap: Finally, despite the fulfillment of both the previous sub-rules the increase in equity allocation is reduced further to only 5% since an equity cap of 40% has been decided.

2 Notice, that the primary objective of the step-up rule is to prevent trading overhead of many small adjustments of the portfolio.
3 Notice, that not all quantities actually reach a steady state, the primary example being the balance of the company. Relative quantities, though, for example, the funding ratio (balance divided by liability), typically do become stationary.
4 The fairly conservative dynamic rule guarantees that the probability of insolvency is less than 3% over the 30-year period considered. Numbers presented are therefore conditional on the company staying solvent, that is, based only on non-insolvency scenarios.

Long-term stability

Analytically we define 'long-term' as the properties of a pension fund when the fund has been so long in operation that the opening balance 'is forgotten', that is, in a steady state[3]. Table 1 summarises a number of key steady state properties[4] for twelve selected combinations of investment and bonus policies. Perhaps the most striking feature is that a seven-fold increase in the equity target (from 10-70%) only improves the average bonus by a third (from 1.2–1.6%) at a 120 bonus threshold. Also the probability of attributing a bonus decreases. This shows that high average bonuses are achieved by bursts of very high bonuses followed by a number of years with no bonus at all.

A glance at the table shows that companies can essentially choose between a steady small bonus or a volatile large bonus. Companies targeting a steady bonus attribution of, say, 1.6% on average, can increase the bonus frequency by a combination of increasing the bonus threshold and decreasing the equity cap.

For a given equity cap, high bonus threshold strategies have a higher probability of attributing bonus and a much higher probability of being at their maximal allocation since a larger free reserve is built up to cushion short-term losses. Similarly, for a given bonus threshold, lower equity caps lead to a more steady bonus attribution and higher average funding ratios.

It is thus clear that the equity cap's importance is way beyond short-term risk management. The joint choice of bonus threshold and equity cap will together define the level and stability of bonus attribution and is essential to the long-term characteristics of the pension fund.

Funding Ratio Distribution

In addition to the summary statistics in Table 1, it is also instructive to see the entire steady state distributions of the funding ratio. This is depicted in Figure 4 for a bonus threshold of 130%, where a kernel density estimate of the distribution of the funding ratio (after bonus) for equity caps of 10%, 30% and 50% have been computed. With an equity cap of 10% one gets a narrow distribution of the funding ratio around the bonus threshold for both bonus policies. In fact, the average funding ratio is just above the bonus threshold as shown in the last row of Table 1.

For equity caps above 10%, the distributions have a characteristic bimodal form, with one mode centered just above the bonus threshold and the other mode just above 110%, corresponding to the risk buffer of 10% used by the dynamic rule, cf. Figure 3. The left mode is caused by the risk reduction rule which sells all equity as the ratio level approaches 110%. This prevents insolvency, but it also makes it hard to improve the funding ratio again, causing an accumulation of probability in that region.

For high equity caps, the distribution is very skewed with a heavy left mode and a long tail. This is a consequence of an aggressive strategy which has a high equity exposure when the funding ratio allows so, making it quite likely to suffer a great loss and a relapse to the 'sticky' region around 110%. In fact, as already observed from Table 1, the mean funding ratio decreases with the equity cap.

Model and parameter sensitivity.

In the previous sections we have made a number of explicit choices for risk tolerance and bonus policy as well as parameter assumptions for the capital market model. They have all been chosen for expositional clarity but are all within reasonable range of 'realistic' values. In general, though, results will depend on these choices. In particular the strategic decision variables – investment policy, bonus policy, and risk tolerance – are very important. This is a major point in itself, since it demonstrates that the success (or failure) of pension companies to a large extent can be attributed to how they are run – and not to the erratic behavior of capital markets. Here we have chosen to focus on the distribution of funding ratio as a measure of the stability of a pension company. The relative size of the two peaks in the distribution of funding ratios will therefore change with both capital market parameters and strategic decision variables.

The qualitative result is robust, though, that DC pension funds have two-modes of operation: a mode where the funding ratio is rebuilt to some target level, and another mode where bonus is declared regularly. This separation is basically independent on the choice of model and parameters, but is the result of the institutional restriction that the funding ratio must always be above one.

Conclusion

The outcome of the analysis therefore is a mixed message. It is confirmed that higher equity allocations, that is, higher equity caps, lead to a higher expected bonus – but a severe penalty in the form of irregular bonus attribution has to be paid. Thus clients are at risk of not properly sampling bonus attributions, that is, at risk of wealth not being distributed fairly among clients.

There are two effects contributing to the destabilising effect of high equity caps. First, the fund is non-sponsored and thus the only way to improve a low funding ratio is through successful investments and/or patience. Strate-gies with a generally high equity exposure are more sensible to the sticky nature of low funding levels as they are more likely to end up there in the first place. Second, the bonus policy will at some point reduce the funding level by attributing bonus. Again this impacts the high equity strategies the most since it (partly) disallows positive equity returns to compensate negative returns. Positive returns will only partly improve the funding level when a bonus is being paid while negative returns will reduce the funding level in full.

There is therefore no universal answer to the question of how to optimally operate a pension fund. Even within the same long-term insolvency risk limitation there is ample room for pension companies to distinguish themselves from each other.

From a steady state perspective, though, companies should reconsider targeting high equity allocations as the universal solution to provide large pensions and at least ensure they are willing to withhold bonuses to accumulate the free reserves necessary to support such strategies. Given the difficulty in maintaining high reserve levels, the solution to many pension companies would rather be to consider modest increases in their bonus thresholds and target mid-level equity allocations.

Appendix A: Capital markets

Simulations are done in three-month time-steps at horizons 10 or 30 years. Given a starting time, t_0, of the simulation, the state of the capital market is then known at discrete points in time, $t_i = t_0 + i \cdot 0.25$. The core of the capital markets model is a 'base' term structure, $R(t_i, T)$, defined for all maturities $T \geq 0$ and constructed from n reference zero rates, $\{r_1(t_i), \ldots, r_n(t_i)\}$ by linear interpolation between points and (constant) continuation outside their range. In this paper we use the 4 reference maturities $\{6M, 2Y, 10Y, 30Y\}$. Given an initial term structure, $R_{init}(T) = R(t_0, T)$, which in the simulations is set to points $\{2.5\%, 3\%, 3.5\%, 4\%\}$, the dynamics of the term structure is modelled as a mean-reverting vector-autoregressive model:

$$\begin{pmatrix} \log r_{i,1} \\ \\ \log r_{i,4} \end{pmatrix} = \begin{pmatrix} \log r_{i-1,1} \\ \\ \log r_{i-1,4} \end{pmatrix} + \Phi\left(\begin{pmatrix} \mu_1 \\ \\ \mu_4 \end{pmatrix} - \begin{pmatrix} \log r_{i-1,1} \\ \\ \log r_{i-1,4} \end{pmatrix}\right) + A\begin{pmatrix} z_{i,1} \\ \\ z_{i,4} \end{pmatrix} \tag{2}$$

where $\Phi = \mathrm{diag}(\varphi_1, \ldots, \varphi_4)$ is a diagonal matrix of mean reversion strengths, $(\mu_1, \ldots, \mu_4)$ are mean reversion levels, **A** is a lower-triangular matrix yielding the covariance structure of the term structure, and the $z_{i,j}$'s are iid. $N(0, 1)$-random variates. The model is calibrated to mean-revert around the initial term structure with variance-covariance matrix Σ in stationarity.

In the simulations we have used (stationary) volatilities vol(r_i) = (2.5%, 2.25%, 2.25%, 2%), correlations cor(r_i, r_j) = 1–0.05*|i–j|, and mean reversion strengths φ_i=0.007 from which the mean reversion levels μ_i can be computed.

Equity Market

Conceptually, the equity market (index) is modeled as giving a return equal to a generic 10Y (zero-coupon) bond plus an equity premium, γ, plus added volatility. The evolution of the equity market is modelled as an equity index, S, defined as:

$$S(t_i) = S(t_{i-1})\exp(e(t_i))$$

where $e(t_i)$ is the stochastic period return of equity index in the period from t_i–1 to t_i. Period returns are simulated as:

$$e_i = \eta r_{i-1}^{eq} + \gamma - (r_i^{eq} - r_{i-1}^{eq})\beta + Dz_i \quad (3)$$

where $r_i^{eq} = R(t_i, 10\Upsilon)$ is the 10-year zero rate at time t_i linking bond and equity returns (to control correlation), D is the one-period standard deviation of innovations, and η=0.25 is a scaling factor to correct zero-coupon bond returns to the sampling frequency of the model.

In the simulations we have used a period equity premium, γ= 0.5%, and the equity 'duration' β together with innovation volatility, D, calibrated to ensure a volatility of equity returns of 10% and 0% correlation between equity and r_i^{eq} zero-coupon bond return.

Michael Preisel is head of quantitative research, Søren Jarner is chief analyst and Rune Eliasen is quantitative analyst at ATP

References

H. U. Gerber.
Life Insurance Mathematics
3rd Ed. Springer-Verlag, Berlin, 1997.

C. W. Jordan.
Life Contingencies, 2nd Ed.
Society of Actuaries, Chicago, 1967.

M. Preisel, S. F. Jarner and R. Eliasen.
A fair-value enterprise.
Life & Pensions, October 2006.

M. Preisel, S. F. Jarner and R. Eliasen.
Investing for retirement through a with-profitspension scheme: A client's perspective. Preprint.
Submitted to Insurance: Mathematicsand Economics, December 2006.