
Investment risk-sharing

A State-of-the-Art Report

Rami Chehab* and Catherine Donnelly†

<https://risk-insight-lab.com/>

*Risk Insight Lab,
Department of Actuarial Mathematics and Statistics,
Heriot-Watt University, Edinburgh, Scotland EH14 4AS*

29 May 2018

This work was funded by the Actuarial Research Centre (ARC) of The Institute and Faculty of Actuaries (IFoA), UK.

Disclaimer: The views expressed in this publication are those of the authors and not necessarily those of the IFoA. The IFoA do not endorse any of the views stated, nor any claims or representations made in this publication and accept no responsibility or liability to any person for loss or damage suffered as a consequence of their placing reliance upon any view, claim or representation made in this publication. The information and expressions of opinion contained in this publication are not intended to be a comprehensive study, nor to provide actuarial advice or advice of any nature and should not be treated as a substitute for specific advice concerning individual situations.

<p>Please cite this report as: Chehab, R. and Donnelly, C. (2018). Investment risk-sharing: A State-of-the-Art Report. Technical report number 2, Risk Insight Lab, Heriot-Watt University, UK. https://risk-insight-lab.com/outputs/</p>
--

*R.Chehab@hw.ac.uk

†C.Donnelly@hw.ac.uk

Foreword

The authors are part of the ARC research project “Minimising Longevity and Investment Risk while Optimising Future Pension Plans”. The goal is to develop new pension product designs that keep the customers’ needs at the forefront. As a first step, this report was written to familiarize the project team with the existing knowledge on decumulation strategies.

The key question that we sought to answer is: “What structures in the pension arena have been proposed for investment risk-sharing?”. The structures cover participating policies (or with-profits contracts), defined benefit pension plans and collective defined contribution plans as well as theoretical structures which have proposed in the literature.

Contents

1	Executive summary	3
2	Introduction	4
3	Participating policies	5
3.1	Terminology	5
3.2	Financial market models	6
3.3	Contracts with a minimum guarantee	7
3.3.1	(Investment) return-based distribution	7
3.3.2	Reserve-based distribution	9
3.3.3	An inter-generational contract with a minimum guarantee	12
3.3.4	Different frameworks for defining bankruptcy	13
3.4	Contracts without a minimum guarantee	15
4	Funded, collective pension plans	23
4.1	Defined benefit (DB) plans	23
4.2	Collective defined contribution (CDC) plans	24
4.3	Analyses of inter-generational risk-sharing in private pension plans	24
5	Theoretical risk-sharing frameworks	27
5.1	Financial fairness and Pareto-efficiency	27
5.1.1	One-period model	27
5.1.2	Multi-period model	28
6	Conclusion	34

1 Executive summary

Investment risk-sharing is a means of sharing money among a collection of people, companies and institutions. It can happen at a point in time or over a period of time. Here the focus is on investment risk-sharing in the context of saving and spending money for retirement. The motivation for risk-sharing is to improve the welfare of the participants.

There is relatively little in the literature on different risk-sharing mechanisms. Much of it focuses on the analysis of collective schemes where the risk-sharing is enabled by a shared fund. Funded pension schemes, where the shared fund is the entire fund, and participating policies, where the shared fund is a buffer account separate to individual customer accounts, fall under this category.

Participating policies are investment products which share investment returns among policyholders. Many of them incorporate a minimum guaranteed return. They may include inter-generational risk-sharing between policyholders via a buffer account. Policyholders have individual accounts and, in the case of inter-temporal risk-sharing, may have a claim on a shared account called a buffer. However, many policies – as described in the literature – do not and it is the insurance company who is responsible for the buffer.

Collective defined contribution schemes allow for inter-generational risk-sharing. Defined benefit pension plans may include a small element of it. There are no individual account values or buffer account in funded pension schemes. Instead, the entire fund represents the totality of the members' claims on the fund. Collective defined contribution schemes are shown to be welfare-improving in several papers, compared to the members investing individually.

Several authors have solved for risk-sharing mechanism that share money among the participants in a financially-fair and Pareto-efficient way. The motivations and set-ups of the authors can be different, and they don't all translate easily into the retirement context. The mechanisms are currently too abstract to be considered as useful in a practical setting. The rules for allocating money are calculated from algorithms, which make them difficult to communicate. They rely on assigning a utility function to each agent in the risk-sharing system, which immediately makes them quite abstract.

2 Introduction

Investment risk-sharing is a means of sharing money among a collection of people, companies and institutions. It can happen at a point in time or over a period of time. Here the focus is on investment risk-sharing in the context of saving and spending money for retirement.

There are many structures for investment risk-sharing which are observed in practice. For example, there are products sold by insurance companies, called with-profits or participating policies. There are funded private pension schemes which are generally only open to employees of certain companies or industries. These pension schemes and products may share investment values indirectly over time among a group of people. However, they don't all utilize investment risk-sharing. For many, the insurance company or the sponsoring employers alone assume the risk of meeting any guarantees. These structures are presented in Sections 3 and 4.

Where there is investment risk-sharing, a common feature is the existence of a buffer. For new entrants to a pension scheme, or new customers, the value of the buffer is a distorted reflection of the previously experienced investment history. If returns have been good, the buffer should be large. If returns have been poor, the buffer should be small and perhaps even negative. The buffer allows participants to access historical returns, for the time before they were participants, as well as for the time during which they are participants.

The buffer is used to smooth the investment values of, or the income paid out to, the participants. If current investment returns are poor, money may be withdrawn from the buffer and distributed among the current participants. If current returns are high, some of the return may be deposited in the buffer rather than being allocated to the current participants.

Inter-generational fairness is often mentioned in the context of risk-sharing schemes. It is often not considered explicitly when setting up a scheme. The rules for allocating values among members do not necessarily involve any consideration of inter-generational fairness.

One way of defining fairness is to have a contract which is financially fair at the outset, in that the expected discounted value of benefits received equals the discounted value of contributions paid by the individual at the start of a contract or pension plan membership. On top of this, fairness could also mean that the risk faced by the individual of receiving less than their expected or promised benefits is equal across all generations.

However, despite investment risk-sharing being relatively common, few methods of sharing returns incorporate explicitly a notion of inter-generational fairness. The closest are the Pareto-efficient and financially fair methods outlined in Section 5. However, it is difficult to see how the methods could be communicated and applied easily in practice.

Lots of economics papers have focused on studying whether risk-sharing is welfare-improving to society. For example, Beetsma et al. (2012) studied the feasibility and welfare consequences of a funded pension scheme, which has inter-generational risk-sharing, when participation is voluntary. Gollier (2008) also consider funded pension schemes which involved inter-generational risk-sharing. Both papers concluded that risk-sharing is indeed welfare-enhancing.

3 Participating policies

Participating policies, also called with-profits policies, are the focus of this section. They are primarily a long-term investment contract, as opposed to an insurance contract. The policyholder pays either a single premium at the initiation of the contract, or they pay regular premiums. The premiums are invested by the insurance company. From the policyholder's perspective, the premiums accumulate as a function of the underlying investment returns. Participating policies are a common product in many countries.

There are two common features which are often seen in these contracts. One is a minimum guaranteed interest rate at which the policyholder's premium accumulates. The second is a smoothing of the returns granted on the premiums, which is enabled by a buffer.

The focus here is on the underlying mechanisms which allocate the invested monies among policyholders. After introducing some basic terminology, participating policies which have been studied in the literature are described. Most of these policies are based on real-life contracts. A notable exception is the contract proposed by Goecke (2013). Although the policyholder can usually surrender the policy before its maturity date in exchange for a payment from the insurance company, this option has been ignored in the presentation. Similarly, any insurance benefits paid upon death or ill-health have also been excluded.

Not all of these contracts feature inter-generational risk-sharing. In spite of this, their structure is presented because they may help in proposing new risk-sharing structures which have similar, attractive features.

Few papers analyze specifically either the inter-generational risk-sharing aspect of participating policies. Most papers consider the relationship between a single policyholder and the company. How much should the company charge the policyholder for the policy so that it is fair, i.e. the expected benefits received by the policyholder are equal to the premiums paid? How much of the buffer account and the investment returns should be received by the policyholder and how much by the company?

3.1 Terminology

Customer account with value P

The customer account represents the minimum value paid to the policyholder at the maturity date of the contract. The policyholder may receive exactly the value of the customer account, or they may get an additional amount paid to them. The customer account is found on the liabilities side of a balance sheet of an insurance company. Early surrender is ignored in this report.

The customer may receive a larger value than the customer account value, upon the maturity of their policy. The value of what the customer receives at the maturity time T is denoted $L(T)$.

Buffer with value B

The primary purpose of the buffer – also called the smoothing account, bonus account or collective reserve – is to smooth investment returns for customers. It does not belong to a single customer but it belongs to the customers collectively and usually also the company. When investment returns are poor, money is transferred from the buffer to the individual customer accounts. The opposite may occur when investment returns are good: money is transferred from the individual customer accounts to the buffer. Hence the buffer helps to reduce the fluctuations of the financial market returns on the customer accounts.

The buffer is the account that may enable risk-sharing between the customer and the company, and between the different generations of policyholders. Or it may be the responsibility of the company alone to make good any deficit in the buffer and there is no risk-sharing whatsoever.

Company account with value C

The company account or shareholders' account, as its name suggests, is the account of the shareholders' or owners' of the insurance company that is providing the insurance contract. It is found on the liability side of the insurance company balance sheet.

Not all of the contracts studied in the literature include a company account. However, whenever the insurance company provides guarantees and has a claim on the buffer, it should be understood that the buffer account incorporates the company account. Similarly, some of the contracts do not include a buffer but do have a company account.

Reference portfolio with value A

Investment return values used for crediting returns to the different accounts are calculated based on the reference portfolio. The portfolio may represent the actual portfolio holdings backing the policy or it may not.

Unless stated otherwise, the initial value $A(0) = P(0)$.

3.2 Financial market models

The most common continuous-time model applied in the literature, the Black-Scholes model, models the price S_0 of a risk-free bond as

$$S_0(t) = S_0(0)e^{rt} \quad (1)$$

and the price S_1 of a risky stock as

$$S_1(t) = S_1(0)e^{(\mu - \sigma^2/2)t + \sigma W(t)}, \quad (2)$$

at time $t \geq 0$, in which the constant r is the risk-free rate of interest, the constant μ is the mean rate of return on the risky stock, the constant $\sigma > 0$ is the volatility of return on the risky stock and W is a standard Brownian motion. It is usually assumed that $\mu > r$. The initial prices $S_0(0)$ and $S_1(0)$ are strictly positive constants.

The Black-Scholes model can easily be extended to include more assets and it is usually a straightforward extension of a derived result to allow for more assets. It is also possible to allow the parameters r , μ and σ to become deterministic functions or non-trivial stochastic processes.

Due to its structure and properties, the Black-Scholes model is one of the simplest models under which closed-form solutions can be found. Closed-form solutions are very useful. The sensitivity of the solution to the various parameters can generally be much better understood, and understood more quickly, when dealing with a closed-form solution. Moreover, the solution in a more complicated model has often the same broad structure.

Contrast this to trying to understand a solution that is the numerical output of a very complicated econometric model, without having the guidance of a solution derived in a simple model like the Black-Scholes model. There is no help as to what are the important parameters for the solution. Doing sensitivity testing becomes a difficult task.

The Black-Scholes model is often criticized for not capturing well enough some observed features of actual stock-market returns. For example, the probability of extreme returns under the Black-Scholes model is too low, compared to the observed chance of very high or very low returns. It

assumes a constant volatility of returns and that returns are independent in non-overlapping time periods, which are not observed in practice.

Unless stated otherwise, the Black-Scholes model has been used as the financial market model in the papers described.

3.3 Contracts with a minimum guarantee

Here participating policies which include a minimum guarantee are described. Effectively, this means that the policyholder's premiums are guaranteed to accumulate at a minimum interest rate by the insurer.

To describe the contracts, the following three accounts are used for the descriptions: a reference portfolio with value A , the customer account with value P and the buffer with value B . There may also be a company account with value C included, or it may be implicitly part of the value of the buffer. The reference portfolio may be simply a reference index or it may represent the underlying assets. The policyholder's contract begins at time 0 and ends at integer time $T > 0$. Mortality and surrender are ignored here.

The customer account value P increases at the random policy interest rate r_P . The evolution of r_P depends on which contract is being considered. Note that the value of the customer account at the maturity time T is not necessarily what is paid to the policyholder. To emphasize this point, denote by $L(T)$ the value of what is paid to the policyholder at the maturity time T .

The buffer is determined residually as

$$B(n) = A(n) - P(n) - C(n), \quad \text{for } n = 1, 2, \dots, T - 1,$$

and, to allow for $L(T) > P(T)$, set $B(T) = A(T) - L(T) - C(T)$.

3.3.1 (Investment) return-based distribution

Return-based contracts allocate a return r_P to the customer account that is calculated using the annual return of the reference portfolio and a minimum guaranteed interest rate r_g . The general idea is that a percentage of the investment return on the reference portfolio is granted on the customer account, subject to a minimum return equal to r_g .

Haberman et al. (2003) and Bacinello (2001)

Haberman et al. (2003) consider three smoothing mechanisms in use in the UK.

- Arithmetic average of the last M time period returns, for a constant integer $M \geq 1$. For a fixed distribution ratio $\alpha \geq 0$, the policy interest rate at time n is calculated over the last $m = \min\{n, M\}$ time periods as

$$r_P(n) = \max \left\{ r_g, \alpha \left(\frac{1}{m} \left[\frac{A(n)}{A(n-1)} + \dots + \frac{A(n-m+1)}{A(n-m)} \right] - 1 \right) \right\}$$

for $n = 1, 2, \dots, T$. The customer account value at integer time $n \geq 1$ is

$$P(n) = P_0 \prod_{k=1}^n (1 + r_P(k)).$$

Bacinello (2001) studies a similar Italian contract which has $M := 1$.

- Geometric average of the last M time period returns, for a constant integer $M \geq 1$. For a fixed distribution ratio $\alpha \geq 0$, the policy interest rate at time n is calculated over the last $m = \min\{n, M\}$ time periods as

$$r_P(n) = \max \left\{ r_g, \alpha \left(\left(\frac{A(n)}{A(n-m)} \right)^{1/m} - 1 \right) \right\}$$

for $n = 1, 2, \dots, T$. The customer account value at integer time n is

$$P(n) = P_0 \prod_{k=1}^n (1 + r_P(k)).$$

- Smoothed asset share. For a fixed distribution ratio $\alpha \geq 0$, the policy interest rate at time n is

$$r_P(n) = \max \left\{ r_g, \alpha \left(\frac{A(n)}{A(n-1)} - 1 \right) \right\},$$

for $n = 1, 2, \dots, T$. Here we are obliged to define the unsmoothed asset share (AS), which accumulates as

$$(AS)_n = (AS)_{n-1} (1 + r_P(n)).$$

The customer account value at integer time $n \geq 1$ is

$$P(n) = \beta(AS)_n + (1 - \beta)P(n-1),$$

for a constant weighting factor $\beta \in (0, 1)$.

For the three contracts considered by Haberman et al. (2003), the terminal payout to the policyholder is

$$L(T) = P(T) + \gamma \max \{A(T) - P(T), 0\}$$

for a participation constant $\gamma \in (0, 1)$. If we could set $\gamma = 1$ then the policyholder would receive $\max \{A(T), P(T)\}$ at the maturity date. In that case, the customer account value acts as a minimum payout. However, as $\gamma < 1$, the policyholder obtains less than the excess value of the reference portfolio over the customer account value. To whom does the residual value go? In Haberman et al. (2003), it is the company. The company has guaranteed the minimum annual return r_g on the policy interest rate, and so it is the company who benefits from any asset out-performance over the customer account value.

Haberman et al. (2003) set the company's account value to

$$C(T) = A(T) - L(T),$$

which implies that $B(T) = 0$. Haberman et al. (2003) do not describe an explicit buffer. This suggests that there is no inter-generational risk-sharing and it is the company who bears any investment risk resulting from the guarantees and smoothing mechanism, through the company account.

The fair value of the contract is the single premium equal to the expected value of the discounted terminal payout $P(T) + \gamma \max \{A(T) - P(T), 0\}$, calculated under a risk-neutral measure. Haberman et al. (2003) study the parameter sensitivity of the fair value and other aspects of the three types of the contracts.

Miltersen and Persson (2003)

Miltersen and Persson (2003) analyze a policy which can model a Norwegian product launched in 1998. Their analysis indicates that the Norwegian product is over-priced and they suggest that may be the reason the product was not a success. There are similar Danish and Dutch contracts. There is no risk-sharing in this contract as it is described. Instead, the company bears the risk of the investment return being insufficient to pay the minimum interest rate guaranteed on the policyholder's premium.

For the analyzed policy, define the annual investment return process $R(n) := A(n)/A(n-1) - 1$. This may not be the actual investment return achieved by the underlying assets. The customer account value is calculated at integer time n is

$$P(n) = P(n-1) (1 + r_g + \alpha \max\{R(n) - r_g, 0\}) \times P(n-1),$$

with $A_0 = P_0$, the initial premium paid by the policyholder. If the investment return in year n satisfies $R(n) > r_g$, there is an excess monetary return $(R(n) - r_g)P(n-1)$. A specified fraction α of the excess return is credited to the customer account. Otherwise, if the return satisfies $R(n) \leq r_g$, there is no excess and the customer account is credited by the amount $r_g P(n-1)$ only.

At the maturity time of the contract, the customer gets $P(T)$ and the value of the buffer account if the latter is positive, i.e.

$$L(T) = P(T) + \max\{B(T), 0\}.$$

The company's account value increases as

$$C(n) = C(n-1) + \beta \max\{R(n) - r_g, 0\} \times P(n-1),$$

with $C(0) = 0$. If there is an excess monetary return $(R(n) - r_g)P(n-1)$, then the company account is given a fraction β of this amount. Otherwise, if $R(n) \leq r_g$, there is no excess return and nothing is credited to or deducted from the company account.

The bonus account is defined residually as $B(n) = A(n) - P(n) - C(n)$. In particular, $B(0) = 0$. It is possible for the buffer account to have a negative value. For this policy, the company is responsible for the buffer. The company must cover any deficit in the buffer account at the maturity date of the contract.

Miltersen and Persson (2003) calculate possible values for the triples (r_g, α, β) which mean that a single premium is the fair price, under a Black-Scholes model like the one described in Section 3.2.

3.3.2 Reserve-based distribution

For reserve-based contracts, returns are distributed to policyholders based on the security of the reserve, as represented by the buffer ratio. The buffer ratio is the size of the buffer B divided by the customer account value P . There may be a target buffer ratio which must be exceeded before returns are paid to a customer account.

The general idea is that a percentage of the buffer ratio above the target buffer ratio is granted on the customer account, subject to a minimum return equal to r_g .

Grosen and Jørgensen (2000)

Grosen and Jørgensen (2000) analyze a reserve-based participating policy. For a target buffer ratio $\gamma \geq 0$, the policy interest rate at integer time n is calculated as

$$r^p(n) = \max\left\{r_g, \alpha \left(\frac{B(n-1)}{P(n-1)} - \gamma\right)\right\},$$

Simulated Market Returns and Policy Interest Rates
 $r_g=4.5\%$, $r=8\%$, $\sigma=15\%$, $\alpha=30\%$, $\gamma=10\%$, $B(0)=10$, $P(0)=100$




Figure 1: Taken from Grosen and Jørgensen (2000, Figure 2). The policy interest rate $r^P(n)$ is plotted against n , for $n = 1, \dots, 20$ for one simulated future. On the same chart, the underlying investment return $A(n)/A(n-1) - 1$ is shown. The minimum guaranteed interest rate $r_g = 0.045$. The target buffer ratio is $\gamma = 0.1$ and the initial value of the buffer is 10% of the initial value of the policyholder’s account. The parameter $\alpha = 0.3$. For the Black-Scholes model of the financial market, the risk-free interest rate $r = 0.08$, and the volatility of the risky stock price is 0.15.

for $n = 1, 2, \dots, T$. They suggest that a realistic range for the target buffer ratio is $\gamma \in [0.1, 0.15]$ and for the parameter $\alpha \in [0.2, 0.3]$. Figure 1 shows a possible development of the policy interest rate compared to the investment return. As a comparison, Figure 2 shows the same items from 1983 to 1998 for one of Denmark’s life insurance company.

The customer account value at time n is

$$P(n) = P_0 \prod_{k=1}^n (1 + r_P(k)).$$

In this contract, the maturity payment to the policyholder is $L(T) = P(T)$. The buffer value is determined as $B(n) = A(n) - P(n)$. Grosen and Jørgensen (2000) do not model an inter-generational aspect of their reserve-based participating policy. Instead, they focus on an individual policy.

Hansen and Miltersen (2002)

Hansen and Miltersen (2002) apply a similar smoothing method to Grosen and Jørgensen (2000), and also use the Danish market as the backdrop. The main difference is that Hansen and Miltersen (2002) allow for fees to cover the minimum return guarantee. They consider two different ways of collecting the fees: a percentage of the customer’s savings or as a share of the buffer which is above the target buffer. In their model, the terminal payout to the customer is the sum of the customer account and the maximum of zero and the bonus account. Once again, a single cohort is studied.

Historical Asset Returns and Policy Interest Rates in Danica 1983-1998




Figure 2: Taken from Grosen and Jørgensen (2000, Figure 3). The chart shows the actual policy interest rates given by the largest Danish life insurance company, Danica, in the period 1983–1998. Also plotted is the tax-adjusted historical market return on a typical institutional portfolio, which is invested 30% in the Danish stock market portfolio and 70% in an index of liquid, Danish mortgage-backed bonds.

Hansen and Miltersen (2002)

Hansen and Miltersen (2002) consider two customers, with the older customer having a higher minimum return guarantee than the younger customer. They find that, when the two policies share a bonus account, there is a re-distribution of wealth from the younger customer to the older customer. They suggest that separate bonus accounts would eliminate the re-distribution.

Kling et al. (2007)

Kling et al. (2007) model a German participating policy, which is somewhat similar to Grosen and Jørgensen (2000). Again, the size of buffer ratio determines the bonus granted. However, it is calculated as $B(n)/P(n)$ at time n , rather than based on the previous time period's values, $B(n-1)/P(n-1)$. As long as the buffer ratio stays within a specified range, the customer account is credited with a fixed interest rate above the minimum guaranteed interest rate. If the buffer ratio falls below the range, then the credited interest rate is reduced using a specified formula. If it rises above the range, the credited interest rate is increased using a slightly different formula.

The focus of Kling et al. (2007) is on the risk to the insurer. There is no company account in the model, but the insurance company is in default if $B(n) < 0$ at any time n . A single cohort is studied and there is no linkage with subsequent cohorts.

Analysis of some participating policies with a minimum guarantee

The default put option is defined as the expected risk-neutral discounted value of $P(T) - A(T)$. Zemp (2011) calculates it for five contracts: Bacinello (2001), Grosen and Jørgensen (2000), Hansen

and Miltersen (2002), Kling et al. (2007) and the arithmetic average contract of Haberman et al. (2003). The default put option is a measure of risk for the insurance company. She finds that the contract analyzed by Bacinello (2001) is the most sensitive to changes in the underlying asset volatility σ and the initial reserve $B(0)$. The least sensitive to changes in the underlying asset volatility are the contracts of Hansen and Miltersen (2002) and Kling et al. (2007), and the least sensitive to changes in the initial reserve is the one studied by Grosen and Jørgensen (2000).

3.3.3 An inter-generational contract with a minimum guarantee

Døskeland and Nordahl (2008) investigate the financial impact of inter-generational risk-sharing on overlapping generations of policyholders. It is one of the few papers in the actuarial or financial literature to do so. The risk-sharing is done via a shared buffer, which is empty when the first generation joins. A cohort receives a fixed fraction of the buffer when their contract matures. The contract is motivated by the one studied by Miltersen and Persson (2003).

Their main result is that later cohorts have a higher expected risk-adjusted return than earlier cohorts. The later generations benefit from a high difference between the risk-free interest rate and the minimum interest rate, high allocations to the bonus reserves and a conservative asset allocation. The opposite is true for earlier generations.

The liability side of the balance sheet is the sum of the company account (C), the buffer (B) and the customer account (P). The company maintains the proportion of total liabilities in the company account (the equity) at a constant $(1 - \alpha)$ at the end of each year, by either paying out dividends or injecting capital into the company account.

The h th generation of policyholders joins at integer time h and leaves at integer time $h + T$. The h th generation has a customer account value denoted by $P^{(h)}$. For $n \notin \{h, h + 1, \dots, h + T\}$, $P^{(h)}(n) = 0$ since either the h th generation has not bought a policy yet or their policy has matured. The sum of the in-force policies is given by P . The customer account value of the h th generation at integer time $n \in \{h, h + 1, \dots, h + T\}$ develops from the previous time period's account values as follows.

$$P^{(h)}(n) = \begin{cases} 0, & \\ \text{for } A(n) \leq \sum_{j=1}^{h-1} P^{(j)}(n-1) \cdot (1 + r_g) & \\ A(n) - \sum_{j=1}^{h-1} P^{(j)}(n-1) \cdot (1 + r_g), & \\ \text{for } 0 < A(n) - \sum_{j=1}^{h-1} P^{(j)}(n-1) \cdot (1 + r_g) \leq P^{(h)}(n-1) \cdot (1 + r_g) & \\ P^{(h)}(n-1) \cdot (1 + r_g), & \\ \text{for } \sum_{j=1}^h P^{(j)}(n-1) \cdot (1 + r_g) < A(n) \leq (P(n-1) + C(n-1))(1 + r_g) + B(n-1) & \\ P^{(h)}(n-1) \cdot (1 + r_g) & \\ + \frac{P^{(h)}(n-1)}{P(n-1)} \alpha \delta (1 - b) [A(n) - (P(n-1) + C(n-1))(1 + r_g) - B(n-1)], & \\ \text{for } A(n) > (P(n-1) + C(n-1))(1 + r_g) + B(n-1). & \end{cases}$$

The insurance company wishes to allocate a fixed annual rate r_g first to in-force customer

accounts, next to the company account and finally, if there are sufficient assets, to give an additional return on those two accounts. The earlier generations have priority.

For each generation h , there are four possible scenarios during the time that generation h 's policies are in-force. Here, "increased" means "increased at the fixed annual rate r_g ". The asset value is: (1) not enough to cover generation h 's customer account value, having been exhausted by covering the increased value of earlier generations' customer account values; (2) enough to cover some but not all of generation h 's increased customer account values; (3) enough to cover all of generation h 's increased customer account values but no more; (4) enough to cover generation h 's increased customer account value and add additional money into the customer account. The last case occurs because the asset value covers all in-force generation's increased customer account values, the increased company account account and last period's buffer account.

In the scenario of case (4) above, a fixed proportion b of the excess assets are allocated to the buffer account. The remaining proportion is divided among the company account and the customer accounts. The company account gets $\alpha\delta$, to compensate the shareholders for the risk of injecting capital into the company account.

Note that the company would be declared bankrupt if the assets were not enough to cover the fixed rate increase on the customer account values. Døskeland and Nordahl (2008) assume that the company would then be re-capitalized, with the liabilities equal to the asset value just prior to re-capitalization.

At the maturity time $h + T$ of generation h 's contract, the payment to that generation is

$$L^{(h)}(h + T) = P^{(h)}(h + T) + pB(h + T) \frac{P^{(h)}(h + T)}{\sum_{j=\max\{h-T,1\}}^h P^{(j)}(h + T)}.$$

Thus, in addition to their individual customer account, the generation gets a proportion $p \times P^{(h)}(h + T) / \sum_{j=\max\{h-T,1\}}^h P^{(j)}(h + T)$ of the buffer on their policy's maturity date.

In the contract, earlier generations receive the full value of their customer account, including the increase at the fixed rate r_g , before later generations. The benefit security of the longest in-force policyholders is prioritized over the others. Despite this, the results of Døskeland and Nordahl (2008) show that the more recent generations have a higher risk-adjusted return than the older generations. This suggests that a bankruptcy event is unlikely to occur.

3.3.4 Different frameworks for defining bankruptcy

Orozco-Garcia and Schmeiser (ress) study inter-generational subsidies in participating policies. They do this in two frameworks: an accounting framework and a run-off framework. There are different definitions of bankruptcy in the two frameworks. All policyholders have the same contract duration but they join at different times. The aim of Orozco-Garcia and Schmeiser (ress) is to see if it is possible to have both fair premiums and the same level of default risk for every cohort of policyholders. The single premium paid by a policyholder is considered to be fair if the difference between the present value of what the policyholder could receive from their contract (allowing for default events) and what they actually pay as their single premium is zero. The contract studied is an extension of Grosen and Jørgensen (2002) (described in Section 3.3.2), extended from an individual to a multi-cohort setting.

Broadly, in the accounting framework there is a bankruptcy event if the present value of liabilities exceeds that of the assets. This captures bankruptcy for most insurance company. The run-off framework would not declare bankruptcy in such a situation, but would continue paying out benefits as long as the asset value was large enough to pay the benefits. It is closer to how public pension systems work. However, the company is not allowed to issue new policies during a run-off event.

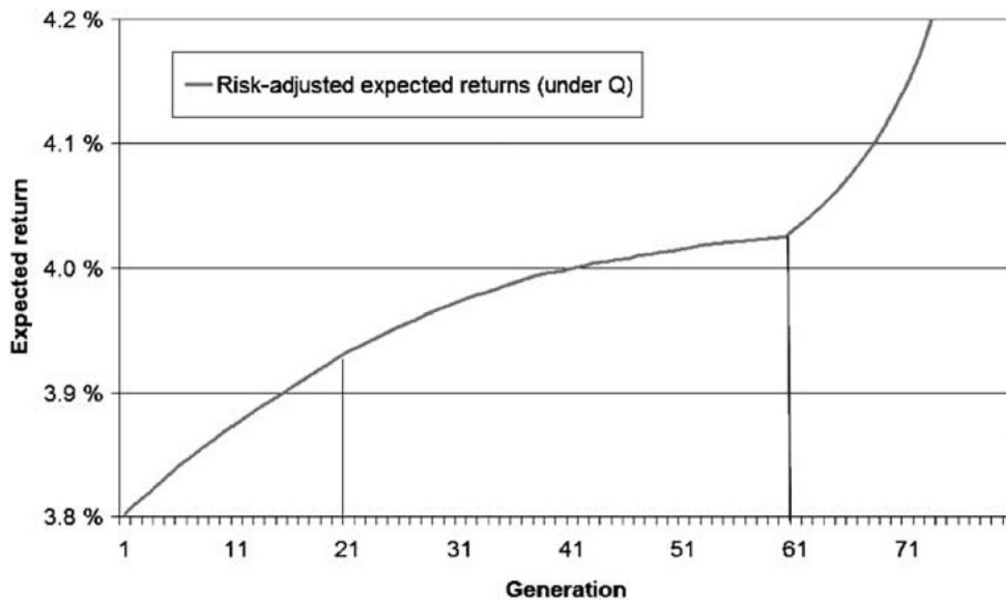


Figure 3: Taken from Døskeland and Nordahl (2008, Figure 3). The plot shows the risk-adjusted return, namely the continuously-compounded annual return on each generation's initial premium, $\ln(L^{(h)}(h+T)/P^{(h)}(h))/T$, over 80 generations who each purchase a 20-year policy. The risk-free interest rate is 4%, the mean return on the risky stock is 8% and its volatility is 16%. The fixed annual rate $r_g = 0.02$. The investment strategy is to invest 20% of asset value in the risky stock and the remainder in a bond whose price accumulates at the risk-free interest rate. For the other parameters, $\alpha = 0.9$, $b = 0.3$, $p = 0.36$ and $\delta = 0.9711$ is chosen so that the contract is fairly priced. The values of the full calibration are shown in Døskeland and Nordahl (2008, Table 4).

If a default event occurs, each policyholder receives a share of the assets which is proportional to their customer account value, with the amount capped at their customer account value. The loss experienced by a policyholder upon a default event occurring is then the difference between their customer account value and that capped proportional share of the assets. The default risk of a policyholder is the value of a “default put option”, equal to the expected discounted value of the possible losses.

Orozco-Garcia and Schmeiser (ress) show that the run-off framework is advantageous for the longer-in-force policyholders as they continue to be paid their benefits at the expense of the benefit security of the shorter-in-force policyholders. The accounting framework penalizes all in-force policyholders at an earlier stage. Orozco-Garcia and Schmeiser (ress) are able to determine risk management strategies so that all contracts operating under the accounting framework face the same amount of default risk. This is not possible under the run-off framework as the shorter-in-force policyholders always have a higher default risk than the longer-in-force policyholders. On the downside for the accounting framework approach, the amount of default risk may be high and fair pricing may not be possible.

3.4 Contracts without a minimum guarantee

Here two contracts without a minimum guarantee are described. The contract proposed by Goecke (2013) is a novel one and not based on an real-world contract. The second is the accumulation phase of a Danish contract that is scrutinized by Guillén et al. (2006). Although the Danish contract has no inter-generational risk-sharing and neither does it need to pose any financial risk to the company, it is included because it looks superficially like it may do: there is a buffer defined as part of the contract. However, there is no need for a buffer since a replicating strategy exists to pay each policy’s maturity value.

Goecke (2013)

Goecke (2013) introduces a return-smoothing mechanism without any guarantees. There are only three accounts: the customer account and the buffer on the liability side, and the reference portfolio on the asset side. In this contract, the reference portfolio value represents the true value of the underlying assets. He specifies rules to decide two things: how much to invest in equities and how much to declare as a return to the customers. He illustrates his rules in a simple setting and in continuous time.

The broad idea is to set a long-term, strategic exposure to equity risk. The proportion of assets that is invested in equities is in line with the strategic exposure, but adjusted to allow for the size of the collective reserve relative to the customer account value. The return declared to customers is the expected return on the investments, but similarly adjusted for the relative size of the collective reserve.

Goecke (2013) applies the same Black-Scholes model that is detailed in Section 3.2. There is a single investment strategy for the entire fund. Goecke (2013) studies a closed scheme; the customer account value changes only because of declared returns.

Suppose the pension fund invests a fraction $\beta(t)$ in the risky stock at time t , for all $t \geq 0$. Defining the volatility of return on assets

$$\sigma(t) := \beta(t)\sigma,$$

the usual market price of risk $h = \frac{\mu-r}{\sigma}$ and the mean return on the asset portfolio

$$\mu(t) := r + h\sigma(t) - \frac{1}{2}\sigma^2(t),$$

the asset value at time t obtained by following the strategy implied by β can be expressed as

$$A(t) = A(0) \exp\left(\int_0^t \mu(s)ds + \int_0^t \sigma(s)dW(s)\right).$$

The return on the customer accounts is denoted by $\eta(t)$ at time t . The total value V of the customer accounts is

$$V(t) = V(0) \exp\left(\int_0^t \eta(s)ds\right)$$

at time t . Thus the return η declared to the customers is a continuously-compounded rate. It is a deterministic function, so that it changes with time.

The reserve ratio, which is a measure of the relative difference of the buffer to the customer account value, is

$$\rho(t) := \ln\left(\frac{A(t)}{V(t)}\right).$$

A target reserve ratio ρ_{target} , a constant value, is fixed. The rules of the scheme adjust the investment strategy and the declared return on the customer accounts to try to get the reserve ratio to the target reserve ratio. The dynamics of the reserve ratio are

$$d\rho(t) = (\mu(t) - \eta(t))dt + \sigma(t)dW(t).$$

Thus the instantaneous expected growth rate of the reserve ratio is the difference between the expected return on the assets (albeit adjusted for risk) and the return declared on the customer accounts. The volatility of the reserve ratio is the same volatility as that of the asset return.

Sufficient notation and definitions have now been introduced to define the operational rules of the scheme in Goecke (2013). A long-term, strategic volatility of return on assets $\hat{\sigma} > 0$ is set. It represents the desired, long-term exposure to investment risk. A risk exposure adjustment parameter, a constant value a , is fixed. It adjusts the strategic volatility of return to allow for the deviation of the actual reserve ratio from the target reserve ratio. The investment strategy rule is defined by setting the volatility of return on assets to

$$\sigma(t) = \hat{\sigma} + a(\rho(t) - \rho_{\text{target}}).$$

The volatility of return on assets is called the risk exposure. From the above equation, the proportion of the reference portfolio value to invest in the risky stock at time t is

$$\beta(t) = \frac{\hat{\sigma} + a(\rho(t) - \rho_{\text{target}})}{\sigma}.$$

The size of the risk exposure adjustment parameter a affects the proportion invested in the risky stock. The higher the value of a , the greater the investment risk taken and hence the more that is invested in the risky stock.

For the return on the customer accounts, a declaration adjustment parameter, a constant value $\theta \geq 0$, is fixed. It can be interpreted as how quickly the reserve ratio “surpluses” or “deficits” are rectified. The value of θ adjusts the strategic volatility of return to allow for the deviation of the actual reserve ratio from the target reserve ratio.

The return on the customer accounts is determined as

$$\eta(t) = \mu(t) + \theta(\rho(t) - \rho_{\text{target}}).$$

In words, the customer accounts are continuously credited by the mean return on the assets with an adjustment for how different the actual reserve ratio is to the target reserve ratio. The speed of the adjustment is governed by the parameter θ . The higher the value of θ , the more volatile are the adjustments. For two different strategies, Figure 4 illustrates the values of the risk exposure, $\sigma(t)$ against the return on the customer accounts, $\eta(t)$, for five different possible values of the reserve gap, $\hat{\rho}(t) := \rho(t) - \rho_{\text{target}}$. A strategy corresponds to a triple $(\hat{\sigma}, a, \theta)$.

Goecke (2013) finds that, compared to a constant-mix benchmark strategy, his proposed mechanism gives a higher expected return on the customer accounts for a fixed standard deviation of return, upon setting the risk exposure adjustment parameter $a = 0$ (Figures 5 and 6). A constant-mix investment strategy is one that is continuously rebalanced back to a constant proportion of assets in the risky stock. As the time horizon increases, the expected return on the customer accounts falls towards the constant-mix benchmark expected return, for a fixed standard deviation. Similarly, as the adjustment coefficient b increases, the expected return falls towards the constant-mix benchmark expected return.

Setting the parameter a to a range of strictly positive values, he finds a similar result; his proposed mechanism gives a higher expected return on the customer accounts for a fixed standard deviation of return (Figure 7). As he notes, the distribution of the return on the customer accounts is no longer normal when $a \neq 0$.

Guillén et al. (2006)

Guillén et al. (2006) analyze the accumulation phase of a product sold in the Danish market. It looks similar to the return-based distribution contracts of Section 3.3.1. The customer account is first increased at a fixed interest rate r_p (which can vary in time, but here it is assumed to be constant for simplicity). Then a proportion $\alpha \in [0, 1]$ of the value of an investment index in excess of the increased customer account is added to the customer account. As before, $P(n)$ denotes the customer account value and $A(n)$ denotes the value of a specified reference portfolio at time n , with initial condition $A(0) = P(0)$ a.s. The reference portfolio value may be different to the value of the underlying investments.

The customer account at time n is

$$P(n) = (1 + r_p)P(n-1) + \alpha [A(n) - (1 + r_p)P(n-1)],$$

assuming that a single premium $P(0)$ has been paid at time 0 by the customer. The customer is paid $L(T) = P(T)$ at the maturity time T . No surrender is allowed before this time. Indeed, the value $P(n)$ does not represent the market value at time n of the customer's terminal payout.

Considering the evolution of the customer account above, the value of the last term on the right-hand side, $\alpha [A(n) - (1 + r_p)P(n-1)]$, can be negative, so the fixed interest rate r_p is not a guaranteed return on the customer account. If the value is negative, this does not necessarily represent a bankruptcy situation (unlike the contract of Døskeland and Nordahl 2008). The underlying asset value at time n is not necessarily equal to $A(n)$; the company can invest to get a different return than the one implied by $\{A(n)\}_{n \in \{0,1,\dots,T\}}$. Indeed, Guillén et al. (2006) show that there exists a replicating investment strategy which exactly replicates the customer's terminal payout $L(T)$. The value of the replicating strategy is not given by $A(n)$ at time n .

Although a buffer $B(n) = A(n) - P(n)$ at time n is defined, it is redundant since a replicating strategy exists for each policyholder. While the contract looks superficially to be a risk-sharing contract, it is not a risk-sharing contract due to the existence of a replicating strategy.




Figure 4: Taken from Goecke (2013, Figure 2). The figure illustrates the bonus declaration, $\eta(t)$, and desired volatility of return, $\sigma(t)$, for two different strategies. For both strategies, $r := 0.03$, $h = 0.25$ and the long-term strategic risk exposure $\hat{\sigma} := 0.05$. The reserve gap $\hat{\rho}(t) = \rho(t) - \rho_{\text{target}}$. Strategy 1 is a pure liability management strategy, since $a := 0$. Therefore $\sigma(t) = \hat{\sigma} = 0.05$. It does not adjust the strategic risk exposure as the reserve gap varies. Only the return $\eta(t)$ on the customer accounts changes with $\hat{\rho}(t)$. For Strategy 1, $\eta(t) = 0.04125 + 0.4\hat{\rho}(t)$, for the chosen parameter $\theta := 0.4$. The vertical line for Strategy 1 is a plot of the return on the customer accounts $\eta = 0.04125 + 0.4\hat{\rho}$ against the risk exposure $\sigma = 0.05$, for values $\hat{\rho} \in \{-0.1, -0.05, 0, +0.05, +0.1\}$. Strategy 2 has the same, long-term strategic risk exposure, $\hat{\sigma} = 0.05$. The parameters $a := 0.5$ and $\theta := 0.2$. The “day-to-day” risk exposure varies with the reserve gap: $\sigma(t) = 0.05 + 0.5\hat{\rho}(t)$. The return on the customer accounts is less sensitive to changes in the reserve gap than Strategy 1. For Strategy 2, $\eta(t) = 0.03 + 0.25\sigma(t) - 0.5\sigma^2(t) + 0.2\hat{\rho}(t)$. The line for Strategy 2 is a plot of the return on the customer accounts $\eta = 0.03 + 0.25\sigma - 0.5\sigma^2 + 0.2\hat{\rho}$ against the risk exposure $\sigma = 0.05 + 0.5\hat{\rho}$, for values $\hat{\rho} \in \{-0.1, -0.05, 0, +0.05, +0.1\}$.

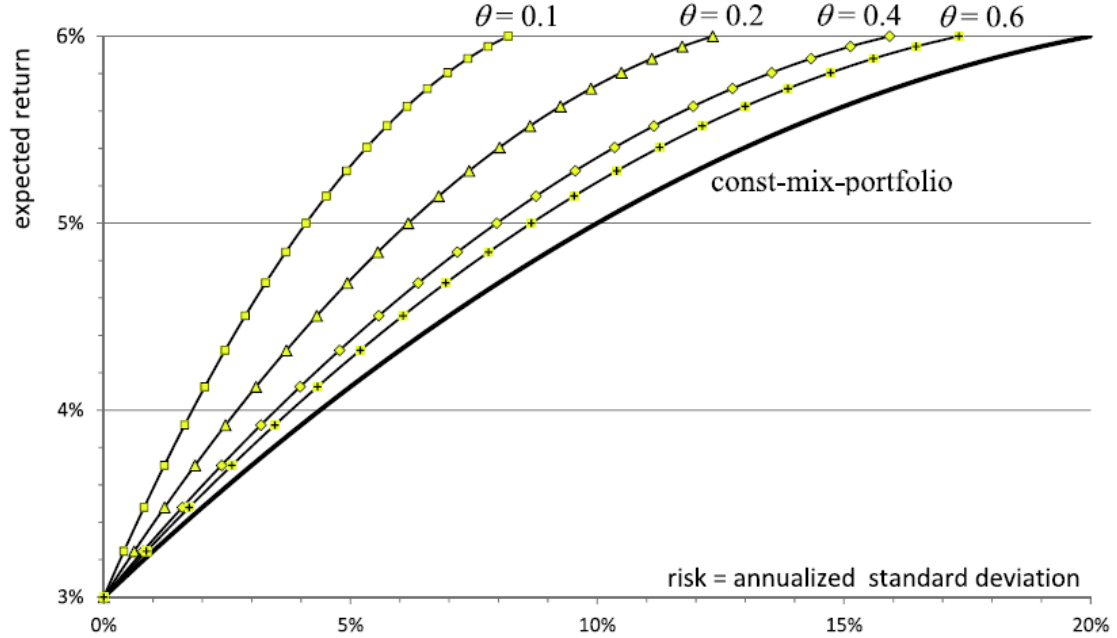


Figure 5: Taken from Goecke (2013, Figure 6A). For all strategies, the risk-free interest rate $r := 0.03$, the market price of risk $h = 0.25$ and the time horizon is 10 years. For the risk-sharing strategies which are represented by the non-solid lines, the parameters are $a := 0$, the value of θ is as marked beside each line and the initial reserve gap is zero, i.e. $\rho(0) = \rho_{\text{target}}$. The value of $\hat{\sigma}$ is a variable. A constant-mix strategy, with a constant percentage invested in the risky stock, is represented by the solid line. For the 10-year annualized return on the customer accounts, $\bar{\eta}(10) := 1/10 \int_0^{10} \eta(t) dt$, the non-solid lines show how its expected value $\mathbb{E}(\bar{\eta}(10)) = 0.03 + 0.25\hat{\sigma} - 0.5\hat{\sigma}^2$ varies against a function of its standard deviation, for four different values of θ . Although its standard deviation is $SD(\bar{\eta}(10)) := \hat{\sigma}/\sqrt{10} \cdot \sqrt{1 - (1 - e^{-10\theta})(3 - e^{-10\theta})/(20\theta)}$, the horizontal axis shows $\sqrt{10}SD(\bar{\eta}(10))$ in order to eliminate the time dependence. The investment strategy underlying the ‘const-mix-portfolio’ solid line has a standard deviation of return equal to $\hat{\sigma}/\sqrt{10}$. The solid line shows how the annualized return on the underlying strategy, $0.03 + 0.25\hat{\sigma} - 0.5\hat{\sigma}^2$, varies with the choice of $\hat{\sigma}$ (again, in order to eliminate the time dependence so that there is no need to choose a value for T for the constant-mix strategy). Note that the expected returns on the customer account (the four risk-sharing strategies) and the constant-mix investment strategy (solid line) are the same. It is the standard deviations which are different. The risk-sharing strategies have a lower volatility for each expected return value. For example, the line with squares representing $\theta = 0.1$ is $\mathbb{E}(\bar{\eta}(10))$ plotted against (approximately) $0.41\hat{\sigma}$ whereas the constant-mix solid line shows $\mathbb{E}(\bar{\eta}(10))$ plotted against $\hat{\sigma}$.




Figure 6: Taken from Goecke (2013, Figure 6B). For all strategies, the risk-free interest rate $r := 0.03$ and the market price of risk $h = 0.25$. For the risk-sharing strategies which are represented by the non-solid lines, the parameters are $a := 0$, the value of $\theta = 0.4$, the initial reserve gap is zero, i.e. $\rho(0) = \rho_{\text{target}}$ and the time horizon is marked beside each line. The value of $\hat{\sigma}$ is a variable. A constant-mix strategy, with a constant percentage invested in the risky stock, is represented by the solid line. Its time horizon is irrelevant as, in this model, the annualized return and the square root of the time horizon multiplied by the standard deviation of return do not depend on the time horizon. For the T -year annualized return on the customer accounts, $\bar{\eta}(T) := 1/T \int_0^T \eta(t) dt$, the non-solid lines show how its expected value $\mathbb{E}(\bar{\eta}(T)) = 0.03 + 0.25\hat{\sigma} - 0.5\hat{\sigma}^2$ varies against a function of its standard deviation, for four different values of T . Although its standard deviation is $SD(\bar{\eta}(T)) := \hat{\sigma}/\sqrt{T} \cdot \sqrt{1 - (1 - e^{-0.4T})(3 - e^{-0.4T})/(0.8T)}$, the horizontal axis shows $\sqrt{T}SD(\bar{\eta}(T))$ in order to eliminate the time dependence. The investment strategy underlying the ‘const-mix-portfolio’ solid line has a standard deviation of return equal to $\hat{\sigma}/\sqrt{T}$. The solid line shows how the annualized T -year return on the underlying strategy, $0.03 + 0.25\hat{\sigma} - 0.5\hat{\sigma}^2$, varies with the choice of $\hat{\sigma}$ (again, in order to eliminate the time dependence so that there is no need to choose a value for T for the constant-mix strategy). Note that the expected returns on the customer account (the four risk-sharing strategies) and the constant-mix investment strategy (solid line) are the same. It is the standard deviations which are different. The risk-sharing strategies have a lower volatility for each expected return value. For example, the line with squares representing $T = 3$ is $\mathbb{E}(\bar{\eta}(T))$ plotted against (approximately) $0.46\hat{\sigma}$ whereas the constant-mix solid line shows $\mathbb{E}(\bar{\eta}(T))$ plotted against $\hat{\sigma}$. For the longest time horizon, $T = 30$ years, the risk-return profile of the risk-sharing strategy is very similar to the constant-mix profile. As the time horizon shortens, the risk-sharing strategies show a lower volatility for a given value of expected return.

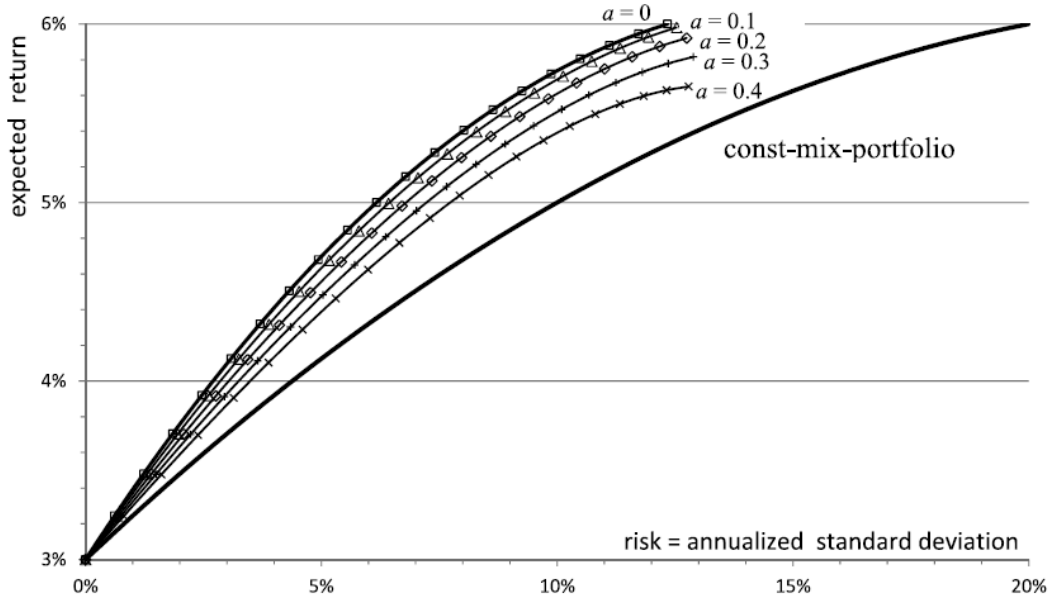


Figure 7: Taken from Goecke (2013, Figure 7). For all strategies, the risk-free interest rate $r := 0.03$ and the market price of risk $h = 0.25$. For the risk-sharing strategies which are represented by the non-solid lines, the parameters are $\theta = 0.4$, the initial reserve gap is zero, i.e. $\rho(0) = \rho_{\text{target}}$, the time horizon $T = 5$ years and the value of a is marked beside each line. The value of $\hat{\sigma}$ is a variable. A constant-mix strategy, with a constant percentage invested in the risky stock, is represented by the solid line. Its time horizon is irrelevant as, in this model, the annualized return and the square root of the time horizon multiplied by the standard deviation of return do not depend on the time horizon. For the 5-year annualized return on the customer accounts, $\bar{\eta}(5) := 1/5 \int_0^5 \eta(t) dt$, the non-solid lines show how its expected value $\mathbb{E}(\bar{\eta}(5)) = 0.03 + 0.25\hat{\sigma} - 0.5\hat{\sigma}^2 f(a)$, where $f(a)$ can be deduced from Goecke (2013, Equation (18)), varies against a function of its standard deviation, for four different values of a . Although its standard deviation is $SD(\bar{\eta}(5)) := \hat{\sigma}/\sqrt{5} \cdot g(a)$, where $g(a)$ can be derived from Goecke (2013, Equation (19)), the horizontal axis shows $\sqrt{5}SD(\bar{\eta}(5))$ in order to eliminate the time dependence. The investment strategy underlying the ‘const-mix-portfolio’ solid line has a standard deviation of return equal to $\hat{\sigma}/\sqrt{T}$. The solid line shows how the annualized T -year return on the underlying strategy, $0.03 + 0.25\hat{\sigma} - 0.5\hat{\sigma}^2$, varies with the choice of $\hat{\sigma}$ (again, in order to eliminate the time dependence so that there is no need to choose a value for T for the constant-mix strategy). For the non-zero choices of a , the distribution of $\bar{\eta}(5)$ is not normal. The risk-sharing strategies have a lower volatility for each expected return value compared to the constant-mix strategy.

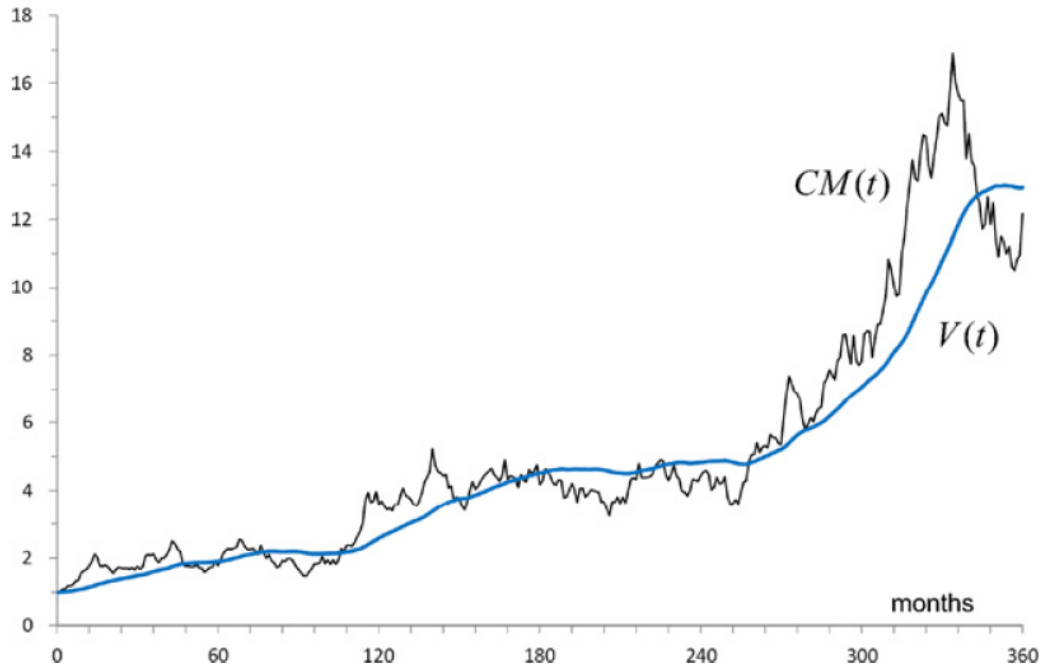


Figure 8: Taken from Goecke (2013, Figure 8). For all strategies, the risk-free interest rate $r := 0.03$ and the market price of risk $h = 0.25$. For the risk-sharing strategy which is represented by the smoother, blue line labelled “ $V(t)$ ”, the parameters are $a := 0$, $\theta = 0.4$, the initial reserve gap is zero, i.e. $\rho(0) = \rho_{\text{target}}$, the time horizon $T = 30$ years and $\hat{\sigma} = 0.2$. A constant-mix strategy, with a constant percentage invested in the risky stock so that the volatility of return is the same as the risk-sharing strategy’s $\hat{\sigma} = 0.2$, is represented by the black line labelled “ $CM(t)$ ”. The lines show one particular future state of the world over the 30 year time horizon. Although the risk-return profile of these strategies is very similar (as evidenced in Figure 6), the daily customer account values ($V(t)$ line) fluctuate less than the daily asset values of the constant-mix strategy ($CM(t)$ line).

4 Funded, collective pension plans

4.1 Defined benefit (DB) plans

Defined benefit (DB) plans are designed to pay a pension income and/or a lump-sum to its members when they retire. They are usually offered to individuals through their workplace. Their employer sponsors the pension plan, which is overseen by trustees. While both the active (i.e. non-retired and haven't left the scheme) members and the employer contribute to the scheme, the employer usually makes larger contributions per member. The interplay of a defined benefit pension plan with a statutory or industry-wide fund which may secure part or all of the benefits of insolvent pension funds (for example, the UK's Pension Protection Fund) is not considered here.

The surplus or deficit is calculated as the difference between the value of the assets and the value of the liabilities. Periodically, the plan is valued and the contributions are adjusted to allow for any change in the cost of accruing benefits and the plan's funding level. If there is a deficit then the employer is ultimately responsible for rectifying it. If the employer is unable to make good a deficit then the scheme will fail and the benefits accrued by the members may be reduced.

In the typical DB plan, the employer and the active members pay contributions towards both the cost of the accruing benefits and to rectify any deficit in the plan. Neither pensioners nor deferred pensioners – members who no longer work for the employer and who have not yet retired – are liable for further contributions. Individual members are typically asked to contribute the same percentage of their salary, regardless of their age. Usually, the actual cost of accruing benefits is lower for younger members than for older members. From the younger individual's point-of-view, they are subsidizing the older member's benefit if they are asked to contribute the same percentage of their salary. Nevertheless, since the employer usually pays the bulk of the contributions of each member, this is rarely a point of controversy for defined benefit members. The fact that only contributions can be adjusted to rectify deficits means that the employer and perhaps also the active members bear most¹ of the downside risk of scheme deficits.

Benefits are paid out to retired members. There may also be benefits paid to the partners and dependants of deceased members. A typical benefit is a lifetime income to the retirees. The annual income paid is generally based on a formula, such as the number of years of service multiplied by 2% multiplied by the salary of the member before retirement. It is common for the income to be increased each year by the increase in price inflation. Another typical benefit is a lump-sum benefit paid to each member upon their retirement, which may be in addition to or instead of a lifetime income.

In DB plans, the nominal income paid to retirees is not adjusted if there is a funding deficit or surplus. However, if pension increases are discretionary then they may be withheld if, for example, the scheme is in deficit. Whether pensions increases are mandatory or discretionary will depend on the plan rules and the jurisdiction.

If the active members' contributions are adjusted to allow for surpluses or deficits then there is inter-generational risk-sharing in the DB plan. Otherwise, the sponsoring employer is entirely responsible for funding risks. There is no explicit buffer account as there is with participating policies. Instead, the entire fund is equivalent to the totality of the customer accounts and the buffer account of participating policies. The ratio of the value of the assets to the value of the liabilities – the funding ratio or funding level – may be taken as a measure of risk. There may be a target of keeping the funding ratio above some level to maintain the security of accrued benefits

¹Since it is possible that the employer is unable to support the plan, there always remains a benefit risk to the deferred pensioners and pensioners.

and provide a cushion for future accruing benefits.

4.2 Collective defined contribution (CDC) plans

Collective defined contribution (CDC) plans can be thought of as DB plans without the guarantees. They promise but do not guarantee a benefit. If there is a deficit, then it is not the legal responsibility of the sponsoring employer to eliminate the deficit (although the employer could choose to make additional contributions). Instead, the benefits can be reduced to try to restore the funding ratio to a desired level. For example, pensions increases may not be granted. In extreme cases, the nominal level of benefits may be cut, including benefits already in payment.

The rules for granting pension increases or cutting nominal pensions may be defined via the funding ratio, via fixed rules. Pension increases may be granted on a sliding scale, from zero pension increases up to full indexation, depending on the funding ratio. Cutting nominal pension amounts may also depend on the funding ratio. If the funding ratio subsequently increases to a high enough level, nominal pensions are restored to their original value and additional pension increases may be granted to compensate for previously lost indexation. It may be that the rules are based on two funding ratios: one calculated using a liability value based on paying only the nominal pensions without any pension increases and a second calculated using a liability value based on paying only the nominal pensions with discretionary pension increases granted.

The inter-generational risk-sharing aspect of CDC plans comes from the possibility of both contributions and benefits being adjusted in response to funding risk. Bennett and Van Meerten (2018) provide a description of the legal and regulatory differences between UK DB plans, their Dutch equivalent and the Dutch CDC plans which have been in existence since 2004.

4.3 Analyses of inter-generational risk-sharing in private pension plans

Much of the academic literature shows that risk-sharing in funded pension plans has positive welfare implications for the members. Most papers use a welfare approach to measure the attractiveness of different pension plans, using a constant relative risk aversion (CRRA) utility function.

Gollier (2008) finds that collective risk-sharing increases the expected return for each cohort and does not reduce the standard deviation of return. He analyzes the lump-sum benefits paid to members of a pension scheme. It is assumed that members contribute annually for 40 years to the scheme and at the end of the 40 years, they receive a lump-sum payment. Each retiring member is replaced by a new one, so that there are always 40 cohorts actively contributing in the scheme. He considers two collective schemes, in which everyone contributes the same amount each year.

In the first collective risk-sharing scheme, Gollier (2008) maximizes the expected utility of the entire membership by optimally choosing how much to invest annually in a risky stock and how much to pay out to each retiring cohort. Compared to not being in a collective scheme but simply following the optimal investment strategy for an individual's utility function, there is an increase in the certainty-equivalent annual return for each cohort of around 0.8%.

In the second collective risk-sharing scheme of Gollier (2008), a minimum return of zero is required for each cohort. His motivation for the constraint is to make the risk-sharing scheme attractive to all generations. Compared to not being in a collective scheme, there is an increase in the certainty-equivalent annual return for each cohort of around 0.5%. This is slightly lower than the first collective scheme due to the solvency requirements of the minimum yield, which constrains the investment strategy.

Cui et al. (2011) find similar results: inter-generational risk-sharing through collective, funded pension plans is welfare-enhancing compared to investing without risk-sharing. Importantly, they

find that the anticipated expected welfare gain of a new cohort is not funded by existing or future cohorts. A collective plan that allows both contributions and benefits to be adjusted in response to the funding level produces more welfare gains than one which can adjust only contributions or only benefits.

Cui et al. (2011) point out that the welfare-enhancing aspect of collective pension plans must be considered, and not just the market values. In their words (Cui et al., 2011, page 3),

The market valuation of [inter-generational risk-sharing] shows that the risk-sharing arrangement is a fair deal when the fund is initially 100% funded. In other words, the market value of funding surplus equals the market value of the deficits in the generational account of the entry cohort. However, the schemes with [inter-generational risk-sharing] are potentially welfare enhancing, and thus a positive-sum game in welfare terms. An agent joining an underfunded collective scheme is not necessarily worse off in welfare terms. It is possible for well-structured pension schemes to absorb funding deficits up to 10-20% by a combination of higher contributions and benefit cuts over a number of years and still enhance the welfare for participants.

In the collective schemes analyzed by Cui et al. (2011), new members pay contributions which are expected to fund each member's pension in isolation. This avoids the plan becoming a Ponzi scheme. Members contribute for 40 years and receive a benefit for 15 years. Mortality and early withdrawals are ignored, as is the distinction between the sponsoring employer and the active members. They analyze:

- A standard DB scheme in which only the contributions are adjusted in response to funding risks. Thus only the the sponsoring employer and the active members bear the investment risk.
- A CDC scheme in which only the benefits-in-payment are adjusted in response to funding risks. Thus only the retirees bear the investment risk.
- A CDC scheme in which both the benefits in payment and the contributions are adjusted in response to funding risks. Thus the sponsoring employer, the active members and the retirees bear the investment risk.
- An individual defined contribution (DC) scheme. The active members and the retirees bear their own investment risk.

For the DB and CDC schemes, the adjustments to the contributions and benefits-in-payment are given by mathematical rules. The parameterization of the rules and the optimal investment and benefit-in-payment strategy are determined for the first cohort, based on maximizing their expected discounted utility of lifetime consumption. The optimal strategies and parameterizations are kept for the subsequent cohorts, modulo any adjustments made in line with the rules. The individual DC scheme members follow the optimal investment and consumption strategy that maximizes their own expected discounted utility of consumption.

Cui et al. (2011, Table 1) find that the CDC scheme in which both the benefits in payment and the contributions can be adjusted is the most welfare-enhancing under a range of investor and market models. Welfare is measured by the certainty equivalent consumption. Either the individual DC or the standard DB scheme are the next best, depending on the precise model used. The CDC scheme in which only the benefits-in-payment can be adjusted is the worst performing scheme.

Astrup Jensen and Nielsen (2016) consider how worse off is an investor if they invest collectively compared to investing on their own. They consider the double impact of being forced to follow a suboptimal investment strategy – namely, the collective’s investment strategy rather than their own, individually optimal investment strategy – and different terminal return-sharing mechanisms.

Specifically, they suppose that the members of a collective receive a share of the investment proceeds that is proportional to their initial contribution. This may be sub-optimal for the individual since they have different aversion to risk, as measured by different parameterizations of constant relative risk aversion utility functions.

Although their paper does not directly consider risk-sharing mechanisms, their results are worthy of consideration. Potentially, some members of a collective risk-sharing scheme could be better off investing individually. For example, suppose a risk-averse member does not contribute much initially relative to the majority of risk-seeking investors. Further assuming that the collective investment strategy is heavily weighted towards the optimal investment strategy of the risk-seekers, then the risk-averse investor may be – for a utility perspective – far from their optimal, individual strategy.

For a CDC plan with conditional indexation, similar to some Dutch CDC plans, Kleinow and Schumacher (2017) show that financial fairness is possible for each new cohort. They allow for a sponsoring employer as one of the plan’s stakeholders. When the nominal funding ratio is low, i.e. the ratio of the assets to the value of the nominal benefits (excluding pension increases), the fair contributions paid by the new cohort are less than when the nominal funding ratio is high. This is because indexation is less likely when the nominal funding ratio is low. However, even when the contributions from new cohorts are lower, the existing cohorts still benefit from risk-sharing.

Hoevenaars and Ponds (2008) value inter-generational wealth transfers as option prices in various DB and CDC pension plans. Their focus is on the impact of strategic or structural changes in the schemes. For example switching to a more risky investment strategy results in a loss of value for older members and a gain for younger members. A switch between a pension scheme targeting flexible contributions and fixed benefits to a scheme aiming for fixed contributions and flexible benefits. The latter switch involves a transfer of value from the older members of the scheme to the younger members. Chen et al. (2017) ask when members would leave a CDC scheme, if participation is voluntary. They use option pricing techniques to answer their question. They find that young members are more likely to stay in the scheme than older members; the value of the former’s ‘exit option’ is high due to the long time over which it may be exercised. However, assuming that each member has the same contribution rate regardless of their age, the value of each year’s pension accrual is much higher for older members than younger ones. This is an incentive for older members to stay in the scheme.

Donnelly (2017) adjusts the mechanism of Goecke (2013) to a pension plan setting, where members can leave and join, and benefits are paid out from the plan’s assets. It is assumed that contribution rates are fixed.

She finds that the pension scheme – which is a type of CDC plan – reduces the volatility of benefits paid out to members, compared to a benchmark Defined Contribution plan. However, members must be prepared to accept reductions in their benefits as a consequence of the lack of a sponsor; the asset-to-liability funding ratio is improved by benefit reductions and changes to the investment strategy rather than by additional contributions.

5 Theoretical risk-sharing frameworks

Inter-temporal risk-sharing is used to compensate for markets being incomplete. The incompleteness referred to in this context arises from the inability of individuals to be exposed to returns and risks outside of their own lifespan. By inter-temporal risk-sharing through a buffer or similar mechanism, a new generation of customers purchasing a contract are exposed to risks that occurred before they bought their contract. This addresses part of the incompleteness but not all of it; old generations whose contract has matured or who have died are not exposed to risks and returns which occur after their exit.

5.1 Financial fairness and Pareto-efficiency

A number of recent papers have targeted the question of operating a risk-sharing scheme in a financially fair and Pareto-efficient way. In these papers, a risk-sharing scheme is defined through a combination of a collective investment decision and a rule for the allocation of investment returns to the participants of the scheme.

Financial fairness means that the expected value of what an individual gets back from a contract should be equal to what they put in. Pareto-efficiency looks at the utility of all the scheme's members together. A group of contracts are Pareto-efficient if no individual's utility can be improved without reducing the utility of another individual in the scheme. One of the earliest papers to apply these two conditions in a one-period setting is Gale and Sobel (1979)².

The authors on this topic must devote a lot of effort to show that a solution exists and, if possible, that the solution is unique. Rarely are they able to find an explicit solution for a particular example. Sometimes they can only characterize a solution and sometimes they can find an iterative algorithm to calculate the solution.

The conclusions of these papers differ when it comes to the solution of their stated problem. This is because the problems are different. Some consider only one time-period. Others are multi-period set-ups, allowing agents to contribute to and receive payments from the risk-sharing scheme. At what times do agents contribute to the scheme? When do agents receive payments out from the scheme? Are investment returns allowed for in the framework? The answers to these questions are different in the models considered in the literature.

5.1.1 One-period model

Here the risk-sharing schemes of Pazdera et al. (2016) and Xia (2004) are presented. Both of them are defined in a one-period model. Therefore, there is no need for a buffer: money is paid in at time 0 and it is all paid out at time 1.

Pazdera et al. (2016)

Pazdera et al. (2016) consider a group of n agents who wish to divide up the proceeds of a collective investment. The value of the collective investment at the end of the period are modelled by a random variable X . The rules for the division are decided in advance by applying financial fairness and Pareto-efficiency.

The collective investment and the allocation rule are together considered a risk-sharing scheme. A risk-sharing scheme defines a contingent claim payable to the i th agent as a function $y_i : (0, \infty) \rightarrow$

²An earlier, now declassified, version by Gale can currently be found online at <http://www.dtic.mil/dtic/tr/fulltext/u2/a047596.pdf>.

$(0, \infty)$ of the end-period investment payoff, for each agent $i = 1, \dots, n$. In other words, agent 1 receives $y_1(X)$ at time 1, agent 2 receives $y_2(X)$ at time 1, and so on. To ensure that no more and no less than the investment proceeds are allocated to the agents, the functions must satisfy $\sum_{i=1}^n y_i(x) = x$ for all $x > 0$. Similarly, the collective investment proceeds must be positive, i.e. $X \geq 0$.

The agents contribute money to the scheme at the starting time, with the i th agent contributing a constant amount $w_i > 0$. Although agents assign the same probabilities to random events, they can have different utility functions.

To define financial fairness, suppose that the expected value of the collective investment X under a risk-neutral measure \mathbb{Q} equals the sum of the initial contributions, i.e. $\mathbb{E}_{\mathbb{Q}}(X) = \sum_{i=1}^n w_i$. Pazdera et al. (2016, Definition 2.1) define an allocation rule (y_1, \dots, y_n) to be financially fair if the expected value of what each agent receives from the collective investment proceeds equals the agent's initial contribution, i.e. $\mathbb{E}_{\mathbb{Q}}(y_i(X)) = w_i$ for $i = 1, \dots, n$.

To define Pareto-efficiency, some notation is required. For a fixed integer $M \geq 1$ and two vectors $\mathbf{a} = (a_1, a_2, \dots, a_M)$ and $\mathbf{b} = (b_1, b_2, \dots, b_M)$, write $\mathbf{a} \succeq \mathbf{b}$ if $a_m \geq b_m$ for every $m \in \{1, \dots, M\}$ and there exists at least one $m \in \{1, \dots, M\}$ such that $a_m > b_m$.

A collective risk X and an allocation rule $\mathbf{y} = (y_1, \dots, y_n)$ are together Pareto-efficient if there does not exist another collective investment proceeds \tilde{X} and allocation rule $\tilde{\mathbf{y}} = (\tilde{y}_1, \dots, \tilde{y}_n)$ such that

$$(\mathbb{E}_{\mathbb{P}}(u_1(\tilde{y}_1(\tilde{X}))), \dots, \mathbb{E}_{\mathbb{P}}(u_n(\tilde{y}_n(\tilde{X})))) \succeq (\mathbb{E}_{\mathbb{P}}(u_1(y_1(X))), \dots, \mathbb{E}_{\mathbb{P}}(u_n(y_n(X)))).$$

Pazdera et al. (2016, Theorem 3.4) are able to establish the existence and uniqueness of a Pareto efficient, financially fair risk-sharing scheme, as represented by the pair consisting of a collective risk X and an allocation rule $\mathbf{y} = (y_1, \dots, y_n)$. Importantly, they provide an iterative algorithm to compute the risk-sharing scheme. Their examples are relatively abstract.

Xia (2004)

Xia (2004) considers a similar setting to Pazdera et al. (2016). A group of n agents wish to pool their money together to make a collective investment. The payoff of the investment is then divided among the agents. There are traded assets available for investment in the market.

Agent i has initial wealth $w_i > 0$ and a utility function U_i . Agent i wishes to maximize the expected utility of terminal wealth by choosing a suitable investment strategy. Assuming that the optimal strategy exists, then the expected utility of terminal wealth is represented by $u_i(w_i)$.

Now suppose that Agent i joins instead the risk-sharing group, for some $i \in \{1, \dots, n\}$. They pay initially w_i and receive a random payoff $\xi_i \geq 0$ a.s. at the terminal time. In other words, the agent does not follow their own, individually optimal investment strategy. For Agent i to be at least as well off with risk-sharing,

$$\mathbb{E}(U_i(\xi_i)) \geq u_i(w_i).$$

For the group of n agents to benefit from risk-sharing, the above inequality should be strict for at least one agent. In other words, a Pareto-efficient solution is sought. Xia (2004) characterizes the solution. However, as no explicit solutions are calculated as examples, it is difficult to see how to apply the characterization to a practically-orientated situation.

5.1.2 Multi-period model

Bao et al. (2017)

Bao et al. (2017) propose a framework for risk-sharing within a multi-period model. There are a

finite number of agents in the discrete-time model, and the risk-sharing finishes at a finite time. All cashflows occur at fixed times. The set-up is quite different to that of Pazdera et al. (2016).

The agents wish to share their risks, which are cashflows paid to the agents. The risks are modelled by random variables which represent cashflows received by the system at the fixed times. There is a buffer which is used to spread the risks over the time horizon, through the payment of contingent claim to the agents at fixed times. Furthermore, returns are earned on the buffer. The total value of the buffer and the agents' cashflows received (the risks) is the money to be shared among the agents and the buffer.

The total contingent claims paid to the agents at each integer time and the value of the buffer at the terminal time are the decision variables. The random variable B_n represents the size of the buffer at the n th time point and the random variable P_n is the total contingent claim paid out at the n th time point as a consequence of the risk-sharing, for $n = 1, \dots, N$.

It is not required that the financial market is complete. Instead the agents simply have to agree on a real-world probability measure \mathbb{P} and a risk-neutral measure \mathbb{Q} .

Exogenous elements of the system are investment returns on the buffer and the risks brought into the system by the agents. The random variable X_n represents the risk from the $(n-1)$ th to the n th time point (money paid into the system), and the random variable R_n is the random accumulation factor of the buffer from the $(n-1)$ th to the n th time point, for $n = 1, \dots, N$. The pair (X_n, R_n) is independent of the pair (X_m, R_m) , for $n \geq m$, but the random variables X_n and R_n can be dependent with known joint distribution, for $n = 1, \dots, N$. The information available at the n th time point is the filtration $\mathcal{F}_n = \sigma\{(X_1, R_1), \dots, (X_n, R_n)\}$, for $n = 1, \dots, N$.

For the system to be in balance, it must hold that at the n th time point that the buffer size equals the $(n-1)$ th buffer size increased by the accumulation factor R_n plus the risk X_n paid into the system less the total contingent claim P_n paid out to the agents, i.e.

$$B_n + P_n = X_n + B_{n-1}R_n, \quad \text{for } n = 1, 2, \dots, N.$$

A risk-sharing rule is a random vector $(P_1, P_2, \dots, P_n, B_n)$ that satisfies the above constraint and some technical conditions.

To impose the Pareto-efficiency condition, a utility function u_n is used to evaluate the worth of the contingent claim P_n paid at the n th time point, for each n . A potentially different utility function u_p evaluates the worth of the final buffer size B_n .

According to Bao et al. (2017, Definition 3.1), a risk-sharing rule $(P_1, P_2, \dots, P_n, B_n)$ is called Pareto-efficient if there does not exist another risk-sharing rule $(\tilde{P}_1, \tilde{P}_2, \dots, \tilde{P}_N, \tilde{B}_N)$ such that

$$\begin{aligned} & (\mathbb{E}_{\mathbb{P}}(u_1(\tilde{P}_1)), \mathbb{E}_{\mathbb{P}}(u_2(\tilde{P}_2)), \dots, \mathbb{E}_{\mathbb{P}}(u_N(\tilde{P}_N)), \mathbb{E}_{\mathbb{P}}(u_p(\tilde{B}_N))) \\ & \succeq (\mathbb{E}_{\mathbb{P}}(u_1(P_1)), \mathbb{E}_{\mathbb{P}}(u_2(P_2)), \dots, \mathbb{E}_{\mathbb{P}}(u_N(P_n)), \mathbb{E}_{\mathbb{P}}(u_p(B_n))). \end{aligned}$$

Furthermore, Bao et al. (2017, Theorem 3.2) show that for a Pareto-efficient risk-sharing rule, there exists positive constants $\theta_1, \dots, \theta_N, \theta_p$ such that the rule maximizes the function

$$\mathbb{E}_{\mathbb{P}} \left(\sum_{n=1}^N \theta_n u_n(P_n) + \theta_p u_p(B_n) \right)$$

over all possible risk-sharing rules. For any vector of positive constants $\boldsymbol{\theta} = (\theta_1, \dots, \theta_N, \theta_p)$, they determine an explicit expression for the Pareto-efficient risk-sharing rule that corresponds to $\boldsymbol{\theta}$. There are infinitely many Pareto-efficient risk-sharing rules for $\{(X_n, R_n)\}_{n=1}^N$. However, by requiring financial fairness, the set can be reduced to a single risk-sharing rule.

Under financial fairness, the expected value of the payments to the agents at the n th time point is equal to a constant, with a different constant for each time point. For example, the constant

could equal the expected value of the agents' risk at the same time point or it could be set using some other method. Similarly, the expected value of the buffer at the last time point is equal to a constant. Bao et al. (2017, Section 4) define a *value profile* as a vector of constants, representing the expected value of the agents' risk at each time point. The expected values are calculated under the risk-neutral measure \mathbb{Q} . The numerical value of each component may be restricted due to technical conditions on the utility functions.

Given a value profile $\boldsymbol{\nu} = (\nu_1, \nu_2, \dots, \nu_N, \nu_p)$, a risk-sharing rule $(P_1, P_2, \dots, P_n, B_n)$ is financially fair if

$$\nu_n = \mathbb{E}_{\mathbb{Q}}(P_n), \quad \text{for } n = 1, 2, \dots, N, \quad \text{and} \quad \nu_p = \mathbb{E}_{\mathbb{Q}}(B_n),$$

where expectations are under the risk-neutral measure \mathbb{Q} . No discounting is applied.

Bao et al. (2017, Theorem 5.1) prove that, given a value profile, there exists a Pareto-efficient, financially fair risk-sharing rule which is unique. Furthermore, they have devised an iterative algorithm to calculate the risk-sharing rule to within a desired error bound. As the authors explain, the system decides at each time point on how much to pay out (P_n at the n th time point) and how much to keep until the next time point (i.e. how much to place in the buffer at time n).

To illustrate the result of Bao et al. (2017), one of their numerical examples is outlined next. Suppose within a three-period model that there are three agents. Agent n enters at time $n - 1$ and leaves at time n , for $n = 1, 2, 3$. At time n , Agent n pays the value of X_n into the system. The random variable X_n has the distribution

$$\begin{aligned} \mathbb{P}[X_n = 0.8] &= 0.4, & \mathbb{P}[X_n = 1.2] &= 0.6 \\ \mathbb{Q}[X_n = 0.8] &= 0.5, & \mathbb{Q}[X_n = 1.2] &= 0.5 \end{aligned}$$

and receives the contingent claim P_n from the system. Note that $\mathbb{E}_{\mathbb{Q}}(X_n) = 1$ and recall that the specified risk-neutral measure \mathbb{Q} is simply any risk-neutral measure that is agreed upon by the agents. The initial buffer size $B_0 = 1$ and the investment returns are zero, so that $R_n = 1$ for $n = 1, 2, 3$. Thus the budget constraint that must be satisfied by a risk-sharing rule is

$$B_n = B_{n-1} + X_n - P_n \quad \text{for } n = 1, 2, 3.$$

The financial fairness constraints are

$$\mathbb{E}_{\mathbb{Q}}(P_n) = \mathbb{E}_{\mathbb{Q}}(X_n) = 1 \quad \text{and} \quad \mathbb{E}_{\mathbb{Q}}(B_n) = B_0.$$

Thus each agent should expect to get back what they put into the system. Similarly, the expected terminal buffer should equal its initial value, so that – in terms of the expectation – the buffer is unchanged by the risk-sharing cashflows. Without requiring Pareto efficiency, it is clear that setting $X_n := P_n$ for $n = 1, 2, 3$ would satisfy the financial fairness constraints.

Imposing Pareto-efficiency, in addition to financial fairness, results in the system development as shown in Figure 9. Each path in the tree represents a possible future state of the world, within the chosen multi-period model. The utility functions are $u_n(x) := -\frac{1}{2}x^{-2}$ for $n = 1, 2, 3$ (and presumably also $u_p(x) := -\frac{1}{2}x^{-2}$, although the buffer's utility function does not appear to be defined in Bao et al. 2017, Section 7.1).

Statistics on the mean and standard deviation of the contingent claim payments P_n made to each agent is shown in Table 1. These are compared to the same statistics of the risks X_n which each agent would receive if they weren't in the risk-sharing system. The statistics are calculated under the real-world measure \mathbb{P} .

Under the risk-sharing system, the first two agents have the expected payments which are lower than what they would get without the system. The variability of the payment size is also lowered




Figure 9: The Open-Ended Buffer example in Bao et al. (2017, Section 7.1). The values of X_n are the payments into the system from the agents, the values P_n are the payments out of the system to the agents and the values B_n represent the size of the buffer, which stays in the system. Investment returns are zero and the agents have a power utility function. The Pareto efficient, financially fair risk-sharing rule is (P_1, P_2, P_3, B_3) . The values of B_3 are close to, but not exactly the same as, those shown in Bao et al. (2017, Table E.6). Here the intermediate and final buffer values are calculated from the budget equations $B_n = B_{n-1} + X_n - P_n$. Otherwise, the values are copied from Bao et al. (2017, Tables E.3-E.5).

by risk-sharing. The risk-sharing system reduces the volatility of payments for the agents and the price of this is a lower expected return, in absolute terms.

However, in relative terms, all of the agents get a higher expected return per unit of volatility of payment size. For example, for the contingent claim paid out at time 1, we have $\mathbb{E}_{\mathbb{P}}(P_1) / \text{sd}_{\mathbb{P}}(P_1) = 20.335$, compared to the same statistic for the payment into the system at time 1 (which is what the first agent would get if they didn't join the risk-sharing system): $\mathbb{E}_{\mathbb{P}}(X_1) / \text{sd}_{\mathbb{P}}(X_1) = 5.307$.

The last agent has a higher expected payment with the risk-sharing system than without it. However, they have a lower volatility of payment size. The last agent takes more risk, something that is seen in other risk-sharing systems: the last generations have the highest expected return and the highest volatility of return (Donnelly, 2017; Døskeland and Nordahl, 2008).

n	$\mathbb{E}_{\mathbb{P}}(P_n)$	$\text{sd}_{\mathbb{P}}(P_n)$	$\frac{\mathbb{E}_{\mathbb{P}}(P_n)}{\text{sd}_{\mathbb{P}}(P_n)}$	$\mathbb{E}_{\mathbb{P}}(X_n)$	$\text{sd}_{\mathbb{P}}(X_n)$	$\frac{\mathbb{E}_{\mathbb{P}}(X_n)}{\text{sd}_{\mathbb{P}}(X_n)}$
1	1.010140	0.049676	20.335	1.040000	0.195959	5.307
2	1.023884	0.083940	12.198	1.040000	0.195959	5.307
3	1.043132	0.127110	8.207	1.040000	0.195959	5.307

Table 1: The Open-Ended Buffer example in Bao et al. (2017, Section 7.1). The values of X_n are the payments into the system from the agents and the values P_n are the payments out of the system to the agents. Investment returns are zero and the agents have a power utility function. The Pareto efficient, financially fair risk-sharing rule is (P_1, P_2, P_3, B_3) . The values of B_3 are close to, but not exactly the same as, those shown in Bao et al. (2017, Table E.6). Here the statistics are calculated using the values of X_n and P_n displayed in Figure 9. They differ slightly from the statistics in Bao et al. (2017, Table 1), due to rounding error differences.

Barrieu and Scandolo (2008)

Barrieu and Scandolo (2008) wish to determine the optimal risk exchange between exactly two agents over many time periods. The idea is that the agents will receive random cashflows at times $1, 2, \dots, N$ outside of the system. The agents will pay the cashflows received at time n into the risk-sharing system and, in return, get a cashflow from the system at the same time n . The rules for the allocation of the cashflows between the two agents are decided at time 0. Thus the total cashflows paid into the system at each time n by the two agents is re-distributed between them at time n , according to the rules decided at time 0. No buffer account is needed. Agents' preferences are modelled by quite general preference functionals.

Barrieu and Scandolo (2008) extend Pareto-efficiency from the one period to the multi-period setting, and determine necessary and sufficient conditions for the Pareto-efficiency of an allocation.

As an illustration of their results, consider Agent A and Agent B over a two-period model (Barrieu and Scandolo, 2008, Section 5.3.1). The two agents have exponential-type utility functions $u_k(x) = -\exp(-\gamma_k x)$, for $x \in \mathbb{R}$, where γ_k is the coefficient of risk aversion, for $k \in \{A, B\}$. Such a utility function implies constant absolute risk aversion. For example, an individual acting under exponential utility would invest a constant amount in a risky stock in order to maximize the expected value of the utility of their wealth at a fixed point in time.

Suppose that Agent k receives the random amount X_1 at time 1 and X_2 at time 2. The preference functional of Agent $k \in \{A, B\}$ is defined as

$$U(X_1, X_2) = \mathbb{E}(u_k(X_1)) + \beta_k \mathbb{E}(u_k(X_2)).$$

The constants $\beta_k > 0$ can represent a preference for having money sooner rather than later.

Suppose that the total amount of money received by both agents at time n is represented by the random variable \hat{X}_n , for $n = 1, 2$. The agents wish to share this money among themselves, so that Agent A receives X_n^A at time $n = 1, 2$. As all the money received at time n is re-allocated at time n , it follows that Agent B receives $X_n^B := \hat{X}_n - X_n^A$ at time $n = 1, 2$.

The result of Barrieu and Scandolo (2008, Section 5.3.1), here slightly paraphrased, is that an allocation $((X_1^A, X_1^B), (X_2^A, X_2^B))$ is Pareto-efficient if and only if Agent A receives the allocation

$$X_1^A = \frac{1}{\gamma_A + \gamma_B} \left(\gamma_B \hat{X}_1 + \ln \left(\frac{\gamma_B}{\gamma_A} \right) + \ln(\lambda) \right)$$

at time 1 and the allocation

$$X_2^A = \frac{1}{\gamma_A + \gamma_B} \left(\gamma_B \hat{X}_2 + \ln \left(\frac{\beta_A}{\beta_B} \right) + \ln \left(\frac{\gamma_B}{\gamma_A} \right) + \ln(\lambda) \right)$$

at time 2, for a constant $\lambda \in \mathbb{R}$. The constant λ is determined as the solution of

$$\begin{aligned} \gamma_A \exp(-\gamma_A X_1^A) &= \lambda \gamma_B \exp(-\gamma_B X_1^B) \\ \beta_A \gamma_A \exp(-\gamma_A X_2^A) &= \lambda \beta_B \gamma_B \exp(-\gamma_B X_2^B). \end{aligned}$$

The above forms of X_1^A and X_2^A imply that Agent B receives the allocation

$$X_1^B = \hat{X}_1 - X_1^A = \frac{1}{\gamma_A + \gamma_B} \left(\gamma_A \hat{X}_1 + \ln \left(\frac{\gamma_A}{\gamma_B} \right) - \ln(\lambda) \right)$$

at time 1 and the allocation

$$X_2^B = \hat{X}_2 - X_2^A = \frac{1}{\gamma_A + \gamma_B} \left(\gamma_A \hat{X}_2 + \ln \left(\frac{\beta_B}{\beta_A} \right) + \ln \left(\frac{\gamma_A}{\gamma_B} \right) - \ln(\lambda) \right)$$

at time 2.

6 Conclusion

Various mechanisms through which risk-sharing takes place have been described. These range from participating policies, to funded pension plans to theoretical frameworks. The risk-sharing mechanism is usually implicit: money is paid into the risk-sharing system, invested and then paid out following pre-determined rules which aren't explicitly fair to each generation. The theoretical frameworks which apply financial fairness and Pareto-efficiency are a notable exception.

In participating policies (Section 3) the risk-sharing is achieved through either a buffer account. The purpose of the buffer may be multi-purpose: in addition to risk-sharing across cohorts of policyholders, it may also smooth returns over a single cohort's policy duration and provide capital to support policy guarantees. Both the insurance company and the policyholders may have legitimate claims on the buffer. It is difficult to tease out solely the inter-generational risk-sharing aspect in those participating policies which include it. This is probably why there so few papers written on the topic.

In funded pension plans (Section 4), the funding ratio may be used as an indirect measure of a buffer. There is no explicit buffer account. Several papers consider the impact of CDC schemes on different generations, usually based on Dutch CDC schemes. The general consensus is that they are welfare-improving, compared to investing individually. The framework is generally the same: use a simple Black-Scholes market model and assume that equal numbers of generations enter and exit the scheme at the start of each time period.

The theoretical frameworks described in Section 5 come closest to risk-sharing mechanisms which are designed to be fair. However, it is hard to see how to apply them in practice in a pension saving and spending context. The rules for allocating money are calculated from algorithms, which make them difficult to communicate. They rely on assigning a utility function to each agent in the risk-sharing system, which immediately makes them quite abstract.

References

- Astrup Jensen, B. and Nielsen, J. (2016). How Suboptimal Are Linear Sharing Rules? *Annals of Finance*, 12(2):221–243.
- Bacinello, A. (2001). Fair Pricing of Life Insurance Participating Policies with a Minimum Interest Rate Guaranteed. *ASTIN Bulletin*, 31(02):275–297.
- Bao, H., Ponds, E., and Schumacher, J. (2017). Multi-period risk sharing under financial fairness. *Insurance: Mathematics and Economics*, 72:49–66.
- Barriau, P. and Scandolo, G. (2008). General pareto optimal allocations and applications to multi-period risks. *ASTIN Bulletin*, 38(1):105–136.
- Beetsma, R., Romp, W., and Vos, S. (2012). Voluntary participation and intergenerational risk sharing in a funded pension system. *European Economic Review*, 56(6):1310–1324.
- Bennett, P. and Van Meerten, H. (2018). Apples and oranges: a comparison of the key features of the legislative and regulatory framework for UK and Dutch defined benefit pension schemes (including Dutch CDC Schemes). Technical report, Pensions Institute, Cass Business School, City University London. Number PI-1803.
- Chen, D., Beetsma, R., Broeders, D., and Pelsser, A. (2017). Sustainability of participation in collective pension schemes: An option pricing approach. *Insurance: Mathematics and Economics*, 74:182–196.
- Cui, J., De Jong, F., and Ponds, E. (2011). Intergenerational risk sharing within funded pension schemes. *Journal of Pension Economics and Finance*, 10(1):1–29.
- Donnelly, C. (2017). A discussion of a risk-sharing pension plan. *Risks*, 5(1):1–20. Article number 12.
- Døskeland, T. and Nordahl, H. (2008). Intergenerational effects of guaranteed pension contracts. *The Geneva Risk and Insurance Review*, 33(1):19–46.
- Gale, D. and Sobel, J. (1979). Fair division of a random harvest: the finite case. In Green, J. R. and Scheinkman, J. A., editors, *General Equilibrium, Growth, and Trade: Essays in Honor of Lionel McKenzie*, pages 193–198. Academic Press, New York, USA.
- Goecke, O. (2013). Pension saving schemes with return smoothing mechanism. *Insurance: Mathematics and Economics*, 53(3):678–689.
- Gollier, C. (2008). Intergenerational risk-sharing and risk-taking of a pension fund. *Journal of Public Economics*, 92(5):1463–1485.
- Grosen, A. and Jørgensen, P. (2000). Fair valuation of life insurance liabilities: The impact of interest rate guarantees, surrender options, and bonus policies. *Insurance: Mathematics and Economics*, 26(1):37–57.
- Grosen, A. and Jørgensen, P. (2002). Life insurance liabilities at market value: An analysis of insolvency risk, bonus policy, and regulatory intervention rules in a barrier option framework. *Journal of Risk and Insurance*, 69(1):63–91.

- Guillén, M., Jørgensen, P., and Nielsen, J. (2006). Return smoothing mechanisms in life and pension insurance: Path-dependent contingent claims. *Insurance: Mathematics and Economics*, 38(2):229–252.
- Haberman, S., Ballotta, L., and Wang, N. (2003). Modelling and Valuation of Guarantees in With-Profit and Unitised With Profit Life Insurance Contracts. Technical report, Cass Business School, London. Working Paper Actuarial Research Paper No 146.
- Hansen, M. and Miltersen, K. (2002). Minimum rate of return guarantees: The danish case. *Scandinavian Actuarial Journal*, 2002(4):280–318.
- Hoevenaars, R. and Ponds, E. (2008). Valuation of intergenerational transfers in funded collective pension schemes. *Insurance: Mathematics and Economics*, 42(2):578–593.
- Kleinow, T. and Schumacher, J. (2017). Financial fairness and conditional indexation. *Scandinavian Actuarial Journal*, 8:651–669.
- Kling, A., Richter, A., and Ruß, J. (2007). The interaction of guarantees, surplus distribution, and asset allocation in with-profit life insurance policies. *Insurance: Mathematics and Economics*, 40(1):164–178.
- Miltersen, K. and Persson, S. (2003). Guaranteed investment contracts: Distributed and undistributed excess return. *Scandinavian Actuarial Journal*, 2003(4):257–279.
- Orozco-Garcia, C. and Schmeiser, H. (in press). Is Fair Pricing Possible? An Analysis of Participating Life Insurance Portfolios. *Journal of Risk and Insurance*.
- Pazdera, J., Schumacher, J., and Werker, B. (2016). Cooperative investment in incomplete markets under financial fairness. *Insurance: Mathematics and Economics*, 71:394–406.
- Xia, J. (2004). Multi-agent investment in incomplete markets. *Finance and Stochastics*, 8:241–259.
- Zemp, A. (2011). Risk comparison of different bonus distribution approaches in participating life insurance. *Insurance: Mathematics and Economics*, 49(2):249–264.