Master Thesis

A Data-Driven Analysis of the Solvency II Review: Bond Portfolio Optimization under the Revised Volatility Adjustment for a Hypothetical Dutch Life Insurer

Author: Tijmen Tans *2745239*

1st supervisor: Ger Koole
Company supervisor: Marleen Tuijl
2nd reader: Ronald Meester

A thesis submitted in fulfillment of the requirements for
VU Master of Science degree in Business Analytics (Financial Risk Management)

Management Summary

The Solvency II 2020 Review changes the way in which the Volatility Adjustment (VA) is determined. Under the revised framework, the new Credit Spread Sensitivity Ratio (CSSR) creates a direct link between the spread sensitivity of an insurer's fixed-income assets and the VA applied to its liabilities. Consequently, asset allocation can affect not only the asset side and Solvency Capital Requirement (SCR), but also the Best Estimate Liability (BEL). This creates a new consideration for insurers when determining their strategic asset allocation.

This research investigates how these changes affect the capital-optimal bond portfolio of a hypothetical Dutch life insurer. A proxy Solvency II balance-sheet model is used to optimise a portfolio of government and corporate bonds by maximising capital surplus. Monthly optimisations from March 2016 to April 2025 are used to compare the current and revised regulatory frameworks and to examine the underlying mechanisms under different investment strategies and market and insurer conditions.

The main finding is that the revised VA changes the trade-off between the BEL and SCR. Portfolio changes can simultaneously affect both quantities, allowing their effects to offset each other. For duration allocation, shifts of up to approximately 20 percentage points were found to produce only limited changes in capital surplus and the Solvency Ratio. This creates a new steering possibility that gives the investment manager more freedom in investing while achieving the same regulatory balance-sheet outcomes.

The revised framework also results in a tendency towards longer-maturity bonds. Across the analysed period, the optimal portfolio shifts on average three percentage points from bonds with maturities below ten years towards bonds with maturities above ten years. This may result from both the positive effect of longer-duration bonds on the VA and the lower risk-free rates under the revised framework, which increase liability duration. Capital surplus increases by 1.21% on average, although this increase cannot be attributed exclusively to the revised VA. The magnitude of the observed effects also varies across risk appetites, market conditions, and liability profiles.

For life insurers, the main implication is that the revised VA introduces a new mechanism through which strategic asset allocation can influence the regulatory balance sheet. The results indicate additional flexibility to steer duration allocation and an increased attractiveness of longer-maturity bonds. However, the exact magnitude of these effects is insurer- and market-specific. The findings therefore provide stronger evidence for the existence and direction of these mechanisms than for their precise quantitative size.

Contents

1 Introduction

A defining characteristic of a life insurer is the horizon of its obligations. The products sold by life insurers involve payments that can extend over decades, meaning that a substantial share of future liabilities is known, or can at least be estimated more accurately. The cash flows arising from life insurance contracts are generally more predictable over long horizons than those of non-life contracts, which typically have a term of one year or less. Against the long-dated obligations stands an investment portfolio whose task is to fund them. The composition of the assets therefore largely reflects the liability profile they must support (van Bragt et al., 2010).

For European insurers, this asset–liability problem is governed by Solvency II (SII), the regulatory framework that prescribes how both sides of the balance sheet are valued and how much capital must be held against the risks they carry (Directive 2009/138/EC (Solvency II), 2009). The technical core of the framework is maintained by the European Insurance and Occupational Pensions Authority (EIOPA). They advise the European legislator on the rules and publish, on a monthly basis, the inputs that insurers must use in their valuation: the risk-free interest-rate term structure and the Volatility Adjustment (Thorendal, 2023). The framework itself is not static. It is periodically reviewed, and the most substantial revision since its introduction is now about to take effect.

That revision is the outcome of a process spanning more than half a decade. It was set in motion by EIOPA's 2020 opinion (EIOPA, 2020), enacted through Directive (EU) 2025/2 (Directive (EU) 2025/2, 2025), and given its technical detail in Commission Delegated Regulation (EU) 2026/269 in early 2026 (Commission Delegated Regulation (EU) 2026/269, 2026). The new rules are due to apply from 30 January 2027. Their substance has therefore been settled, although they are not yet in force. Insurers are currently assessing whether their asset allocations will remain capital-efficient under the regime with which they will shortly have to comply. This period, between regulatory certainty and practical application, provides a particularly relevant opportunity to analyse the effects of the revised rules.

The review affects several building blocks of the framework. From an investment perspective, one change is of particular interest: the revision of the Volatility Adjustment (VA). Under the current rules, the VA, the add-on to the rate used to discount liabilities, is exogenous to the individual insurer. It is derived from a currency-wide reference portfolio and applied uniformly, regardless of the assets held by the insurer itself. Under the revised rules, the adjustment is scaled by a new credit spread sensitivity ratio (CSSR), which compares the spread sensitivity of the insurer's own fixed-income assets with that of its liabilities. Asset allocation therefore helps determine the discount rate applied to the very liabilities that the portfolio must fund.

The revision therefore changes the economic structure of the allocation problem. The portfolio no longer affects only the asset side and the Solvency Capital Requirement (SCR); through the CSSR–VA channel, it also affects the Best Estimate Liability (BEL). A single allocation can consequently alter both the value of the liabilities and the capital requirement. Together with the value of the assets, these quantities determine the resulting capital surplus and Solvency Ratio. The portfolio is thus partly evaluated through consequences that it creates itself. Here, a capital-optimal portfolio is defined as the portfolio that maximises capital surplus, measured in this study as Own Funds minus the SCR, or equivalently, Assets minus the BEL minus the

SCR. The portfolio and the balance sheet it produces must consequently be determined jointly. This requires a model that constructs the full SII balance sheet for each candidate portfolio and searches directly over allocations using a decade of market data. This thesis examines that joint optimisation problem under the relevant changes introduced by the SII review, with particular emphasis on the allocation-sensitive VA mechanism.

1.1 Research questions

This leads to the central question of this thesis:

> *In what way, and through what mechanisms, do the Volatility Adjustment changes in the Solvency II 2020 revision change the capital-optimal bond portfolio of a Dutch life insurer?*

This question is addressed through four sub-questions, each examining a different dimension of the problem. The first establishes the effect itself: *in what way does the optimal portfolio under the revised regime differ from that under the current regime?* A direct comparison of the two optima would combine several effects. The revised rules may change the valuation of an unchanged portfolio, while the insurer may also re-optimise in response to those changes. The analysis therefore uses a stepwise decomposition in which the relevant changes are introduced in a controlled sequence. This makes it possible to distinguish the direct effect of the revised rules, the contribution of the allocation-sensitive VA mechanism, and the insurer's response through reallocation.

The second sub-question turns from what changes to why it changes: *how do particular asset allocations affect the BEL and SCR under the revised framework, and how can this trade-off be quantified?* The revised rules create a link from the spread sensitivity of the portfolio to the CSSR, from the CSSR to the VA, and from the VA to the BEL. At the same time, the allocation affects the SCR through the relevant market-risk charges. This sub-question therefore examines how fixed or deliberately varied allocations affect both quantities, before considering the insurer's response through optimisation.

The third sub-question asks whether the results depend on the insurer's risk appetite: *to what extent do different levels of risk appetite produce different outcomes under the revised regime?* Insurers subject to identical rules may nevertheless adopt very different investment strategies, ranging from cautious liability matching to a greater willingness to accept risk in pursuit of higher returns. The revision may affect these strategies differently.

The fourth sub-question asks whether the results depend on market and insurer conditions: *under which conditions is the impact of the VA revision on the capital-optimal bond portfolio strongest, and how robust is the central result to changes in these conditions?* The main analysis is based on one market history and one liability profile. Sensitivity analyses then vary the risk-free interest-rate level, the liability cash flows, and the risk-corrected spread environment. These experiments show where the effect is economically meaningful and where it has limited practical significance.

The operationalisation of these questions and the experiments used to answer them are set out in Chapter 5.

1.2 Approach and central goal

The study models a hypothetical Dutch life insurer with a balance sheet of EUR 15 million. Its investment portfolio consists exclusively of long positions in government and corporate bonds. The eligible bonds are limited to the rating classes AAA, AA, A, and BBB. No short positions or other asset classes are included.

The model is evaluated on a monthly basis from March 2016 to April 2025. For each month, it constructs the relevant SII balance sheet and identifies the bond allocation that maximises capital surplus. The portfolio may only contain bonds that are active in the market during that month. Each optimisation is performed independently. Positions, gains, and losses are not carried over from one month to the next; every optimisation starts with an empty portfolio. This design isolates the effect of the regulatory rules and of the market and insurer conditions observed in each month.

The central aim is to understand how the VA changes introduced by the SII revision affect the bond allocation and, consequently, the insurer's balance sheet. The main regulatory comparison includes the risk-free-rate and VA changes introduced by the review. The focus is on the revised Volatility Adjustment, because the new rules create a link between the asset allocation and the BEL through the CSSR. The BEL and SCR are the two balance sheet items that may have substantial impact on the Surplus.

The study therefore examines whether the BEL becomes a meaningful counterweight to the SCR under the revised regime. When an allocation increases spread exposure, it may increase the SCR while also increasing the VA, which reduces the BEL. The question is whether this reduction in the BEL partially offsets the higher SCR, offsets it to a similar extent, or remains too small for the SCR effect to dominate. The analysis also examines how the optimal allocation and its performance, measured through capital surplus and the Solvency Ratio, change under different regulations, investment strategies, and market and insurer conditions.

1.3 Scope

The scope of the model follows one simple rule: the model includes only the parts of the balance sheet that respond to the bond allocation. This keeps the analysis focused. If a component does not change when the allocation changes, it cannot explain why one portfolio produces a better result than another. It may still affect the absolute size of the balance sheet, but it does not affect the comparison between portfolios.

The model isolates two changes from the SII review: the revised Volatility Adjustment and the new risk-free interest-rate methodology. The VA is the main focus of the study. The new methodology changes the rates used in the model, while the revised VA creates a direct link between the bond allocation and the liability valuation. Both changes can affect the BEL and the capital surplus, but through different channels. To study the impact of the VA realistically, the model must therefore include the asset valuation, CSSR, VA, BEL, and the relevant parts of the SCR for every candidate allocation.

The SCR is calculated with the Standard Formula, but only the market-risk modules that can be affected by the bond allocation are included. These are spread risk, interest-rate risk, concentration risk, and currency risk. The risk margin is set to zero and life underwriting risk is given a constant value. This distinction is made to achieve a more accurate aggregation effect in the SCR calculation.

These choices also affect how the results should be interpreted. The model's balance sheet is validated against published figures for the Dutch life insurance market, which provides a realistic reference for the absolute values. The purpose of the study, however, is not to reproduce one real insurer or identify one universally optimal portfolio. It is to understand the response of the model: how the allocation, VA, capital surplus, and Solvency Ratio change when the rules, investment strategy, or market and insurer conditions change. The conclusions are therefore based mainly on differences between scenarios, rather than on the absolute levels produced by the model.

2 Financial Context and Optimisation Theory

This chapter provides the financial context and optimisation theory for the optimisation problem studied in this thesis. It first sets out the insurer's context, after which the optimisation context follows.

2.1 Financial and Insurance Context

European insurers operate within the Solvency framework, which was initially developed to reduce the risk of insurers becoming insolvent and to protect policyholders. The central objective of the framework is to ensure that insurers hold sufficient capital to meet their financial obligations.

Solvency II, introduced in 2009 and brought into force in 2016, was designed to address the shortcomings of the Solvency I framework. The system changed from a rules-based to a risk-based approach, replacing standard percentages with more detailed formulas. The framework also introduced multiple risk categories to reflect more closely the characteristics and choices of individual insurers (Höring, 2013).

The revision studied in this thesis, referred to as the 2020 Review, introduces changes throughout the framework. Among the changes affecting capital and quantitative requirements are a new method for calculating the risk-free rate, a new interest-rate risk formula, and a revised Volatility Adjustment structure (Alfonsi et al., 2020; Directive (EU) 2025/2, 2025).

Section 2.1.1 introduces the insurer's balance sheet and its main building blocks under the rules currently in force. Section 2.1.2 then focuses on the regulatory changes resulting from the 2020 Review.

2.1.1 Balance Sheet of the Insurer

The insurer's balance sheet consists of the following components:

Assets The asset side of the balance sheet, shown in Figure 1, consists of the insurer's investment portfolio. Its role is to generate sufficient returns to meet the guarantees embedded in the insurer's liabilities and, where possible, to generate additional profit through increases in market value (MV). Some investment products prioritise stable returns, whereas others aim to generate higher returns through capital appreciation.

Among the instruments available to insurers, bonds occupy a substantial position on the balance sheet of a typical life insurer, often accounting for a large share of total assets (Drenovak et al., 2021). Their relatively predictable cash flows and favourable treatment under SII make them natural building blocks for matching long-dated liabilities. They therefore warrant more detailed treatment before the broader portfolio is discussed.

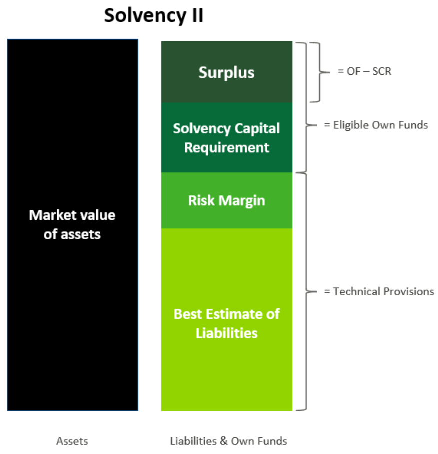

Figure 1: Balance sheet items

Products Bonds are fixed-income products that are particularly relevant to life insurers because of their relatively predictable future cash flows. However, bonds differ substantially in their characteristics: maturity and coupon structure affect their duration and therefore their sensitivity to interest-rate movements, while their credit rating reflects credit risk and influences the return investors require (Alfonsi et al., 2020).

Other return-seeking products include residential mortgages, which can generate relatively stable returns, and real estate, which is often expected to preserve or increase its value over time. These asset classes expose the insurer to various types of risk, including interest-rate risk. Hedging instruments can be used to reduce such exposures.

Interest-rate swaps are one example of a hedging instrument. In an interest-rate swap, two parties exchange interest payments, typically by exchanging fixed-rate payments for floating-rate payments. This allows an insurer to reduce its exposure to changes in interest rates.

Strategic trade-offs The composition of the asset portfolio is ultimately the result of strategic choices that reflect the insurer's risk appetite, liability profile, and view on the market. There is no single optimal portfolio: insurers weigh expected return, capital efficiency, liquidity, and matching quality against one another (Kouwenberg, 2017). A conservative liability matcher will tilt heavily towards government bonds and use derivatives to lock in duration, accepting modest returns in exchange for stability of the Solvency position. A more return-oriented insurer will allocate more to corporate credit, mortgages, real estate, and private markets, accepting higher capital requirements and greater sensitivity to market shocks in pursuit of a higher long-term return (Drenovak et al., 2021). Therefore, as mentioned in Subsection 1.3, the focus is not on creating one optimal portfolio, but on giving understanding of how the allocation, VA, Surplus and SR change due to different financial environment changes.

Best Estimate Liability Usually the largest item in size on the right side is the best estimate liability (BEL). Unlike the asset side, the BEL is not driven by investment choices but by product choices: the type of contracts the insurer has sold and the characteristics of the policyholder base together determine the future obligations the insurer will need to meet (Curry, 2019). Formally, the BEL is the expected present value of all future cash flows arising from these obligations:

$$\text{BEL} = \sum_{t=1}^{T} \frac{\mathbb{E}[CF_t]}{(1 + r_t)^t} \tag{1}$$

where $\mathbb{E}[CF_t]$ is the expected cash flow at time t, r_t is the relevant discount rate, in this case the sum of the risk-free interest rate and the Volatility Adjustment. This sum runs over all future time points until the portfolio has fully run off. As the name suggests, the BEL is by construction an estimate. The actual future obligations will deviate from it (Thorendal, 2023).

Intuition on BEL profile curve From a management perspective, the value lies in the interaction between the liability profile and the asset side. Long-dated liabilities are particularly sensitive to changes in interest rates, creating an incentive to match this exposure through long-duration assets or interest-rate swaps. The liability profile, reflected in the BEL and its cash flow structure, therefore provides an important starting point for asset–liability decisions.

However, this relationship can also run in the opposite direction. Management may influence the liability profile itself through product design or the composition of its policyholder base. For example, targeting older policyholders could reduce the expected duration of the liabilities if mortality assumptions imply earlier benefit payments. This would also change the cash flow profile of the BEL.

Risk-free rate term structure The risk-free interest rate plays a central role in financial modelling. It is used to discount future cash flows, value liabilities, and price financial instruments. In practice, different proxies are used for the risk-free rate, such as nearly risk-free government bonds or swaps rates. For European insurers, however, a single risk-free rate curve is prescribed by EIOPA, which acts as the technical advisor on insurance matters to the European Commission and ensures the consistent application of SII across member states (Thorendal, 2023).

The RFR serves as the base rate for discounting liability cash flows and determining their present value (PV). The effect of a change in the RFR depends strongly on both the size and the duration of the liability profile. Duration measures the first-order sensitivity of the PV to changes in the RFR.

For example, for a liability cash flow profile with a PV of approximately EUR 15 million, a 1 bp (0.01%) change in the RFR corresponds to a change in value of approximately EUR 18,000. A median change of 10 bps corresponds to a value change of approximately EUR 122,000 to EUR 180,000. In a tail-stress scenario involving a change of 50–90 bps, the resulting change in PV can approach 10% of the total liability value.

andere

Volatility Adjustment The Volatility Adjustment (VA) is the add-on to the risk-free rate used to discount the BEL, published by EIOPA and the same for all companies within each country. Its purpose is to limit the impact of short-term fluctuations in credit spreads on the value of liabilities and Own Funds. When spreads widen during market stress, the VA increases and thereby cushions the impact on the Solvency position (Barucci et al., 2025). It is applied only up to the Last Liquid Point (LLP) of the curve, the first 20 years of maturities. Above the LLP, the VA and RFR are extrapolated as one towards the Ultimate Forward Rate. The general formula for the interest rate used to discount the BEL is (Delegated Regulation (EU) 2015/35, 2025):

$$r(t) = \text{RFR}(t) + \text{VA}, \tag{2}$$

A higher VA raises the discount rate and lowers the PV of the BEL.

In the current legislation, the size of the VA is derived from a reference portfolio that EIOPA specifies per currency. EIOPA publishes the VA monthly, while the underlying reference portfolios are refreshed annually (Thorendal, 2023). The currency VA is

$$\text{VA} = \text{GAR} \cdot \text{RCS}, \qquad \text{GAR} = 65\%, \tag{3}$$

where GAR is the General Application Ratio and RCS, the Risk-Corrected Spread on the reference portfolio. The risk-corrected spread is the market spread, stripped of the part that compensates for default and downgrade:

$$\text{RCS} = S - \text{RC} = \sum_i w_i \, (S_i - \text{RC}_i), \tag{4}$$

summed over the asset categories i in the reference portfolio, with weights w_i, market spreads S_i, and risk corrections RC_i. Only this risk-corrected portion enters the VA calculation, thereby excluding the part of the spread attributed to default risk.

On top of this currency-wide figure, a country-specific increase can apply when the risk-corrected spread on a country's reference portfolio is elevated relative to the currency spread, so that localised stress in a single market is recognised (Barucci et al., 2025).

Risk Margin The risk margin is the second component of the technical provisions, held on top of the BEL. It represents the amount a third party would require to take over the insurer's obligations. This compensates for the cost of holding capital against non-hedgeable risks until the portfolio has run off. Together, BEL and risk margin make up the full value of the obligations on the balance sheet. Like the BEL, the risk margin is driven by the products an insurer sells, since only non-hedgeable risks such as longevity, lapse, expense, and operational risk are included in the calculation of the risk margin (Curry, 2019).

Solvency Capital Requirement

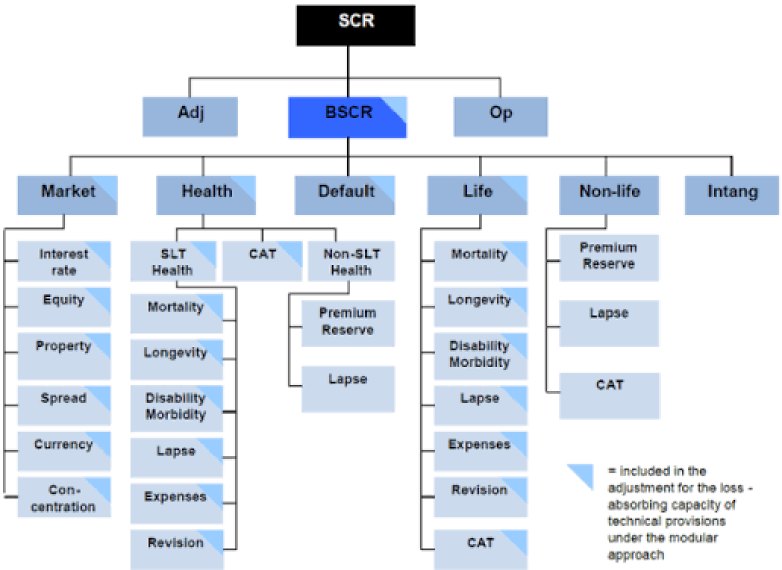

Figure 2: Complete module structure of the SCR (EIOPA)

General structure of the SCR The Solvency Capital Requirement is the amount of Own Funds an insurer must hold to absorb unexpected losses over a one-year horizon, calibrated to a Value-at-Risk at a 99.5% confidence level (Directive 2009/138/EC (Solvency II), 2009). In other words, it is the capital that should still leave the insurer solvent after a one-in-two-hundred-year adverse outcome (Boonen, 2017). Under SII the SCR may be computed with the Standard Formula or with a firm-specific Internal Model. The Standard Formula has a modular structure: each main risk module contains a number of submodules, whose capital charges are aggregated through prescribed correlation matrices so that diversification between risks is recognised (Delegated Regulation (EU) 2015/35, 2025, art. 87):

where the SCR aggregates the main risk modules,

$$SCR = \sqrt{\sum_{i,j} Corr_{i,j} \cdot SCR_i \cdot SCR_j} + SCR_{\text{intangibles}}. \tag{5}$$

The full modular tree is shown in Figure 2. The full correlation matrix is provided in Appendix A.

Main risk modules The Basic SCR is built from several main risk modules, of which only a subset is relevant for a typical life insurer. For this thesis the modules that are taken into account are (Drenovak et al., 2021):

- *Market risk*: captures the risk arising from the investment portfolio, and is split into the interest rate, equity, property, spread, currency and concentration submodules discussed below.
- *Counterparty default risk*: captures the risk from the default of counterparties to risk-mitigating contracts, such as reinsurance arrangements, derivatives and deposits.
- *Life underwriting risk*: captures the risk in the life insurance liabilities themselves, including mortality, longevity, lapse, expense, revision and catastrophe risk.

For a bond-only model, market risk is the only module that changes due to investment decisions and therefore handled in more detail.

Market risk submodules The market risk capital charge aggregates the submodules through the market-risk correlation matrix. As this thesis handles a bond-only universe (Delegated Regulation (EU) 2015/35, 2025) (Table 9, Appendix A):

$$SCR_{\text{market}} = \sqrt{\sum_{i,j} Corr_{i,j} \cdot SCR_i \cdot SCR_j}. \tag{6}$$

Each submodule is defined as the loss of basic Own Funds under a prescribed stress (Kouwenberg, 2017). As this thesis deals with a bond-only universe, the following modules are considered:

- *Interest rate risk*: captures losses when assets and liabilities move unevenly as rates shift. A downward shift is typically worse for insurers, since long-dated liabilities lose more value than the shorter assets. A rate increase hurts insurers with the opposite mismatch. The module stresses both directions (Tables 11, 12) and charges the worse outcome (Delegated Regulation (EU) 2015/35, 2025, art. 165–167).
- *Spread risk*: captures losses from credit spreads widening on bonds, loans, securitisations and derivatives. The charge rises with duration and falls with creditworthiness. This means that longer, lower-rated exposures are hit hardest, since their value is most sensitive to a spread move (Delegated Regulation (EU) 2015/35, 2025, art. 175, 176).
- *Currency risk*: captures losses from exchange-rate movements on positions held in a currency other than the reporting currency. Both a 25% appreciation and depreciation are tested per currency, with the worse outcome retained. The charge rises with the size of the net open position and the number of currencies involved (Delegated Regulation (EU) 2015/35, 2025, art. 188).
- *Concentration risk*: captures the extra risk from relying too heavily on a single counterparty. The charge kicks in once an exposure exceeds a credit-quality-dependent threshold, and rises with both the size of the excess and the counterparty's credit quality. High-quality domestic sovereigns are exempt entirely (Delegated Regulation (EU) 2015/35, 2025, art. 182–186).

Resulting components The components introduced in the previous paragraphs combine into a small set of headline figures that summarise the financial position of the insurer under SII. The market value of the assets, minus the technical provisions (BEL plus risk margin), gives the Own Funds of the insurer. Own Funds represent the economic value available to absorb unexpected losses. Subtracting the SCR from the Own Funds yields the Surplus: the buffer the insurer holds above the regulatory minimum. The ratio of Own Funds to SCR is the Solvency Ratio, expressed as a percentage:

$$\text{Solvency Ratio} = \frac{\text{Own Funds}}{\text{SCR}} \tag{7}$$

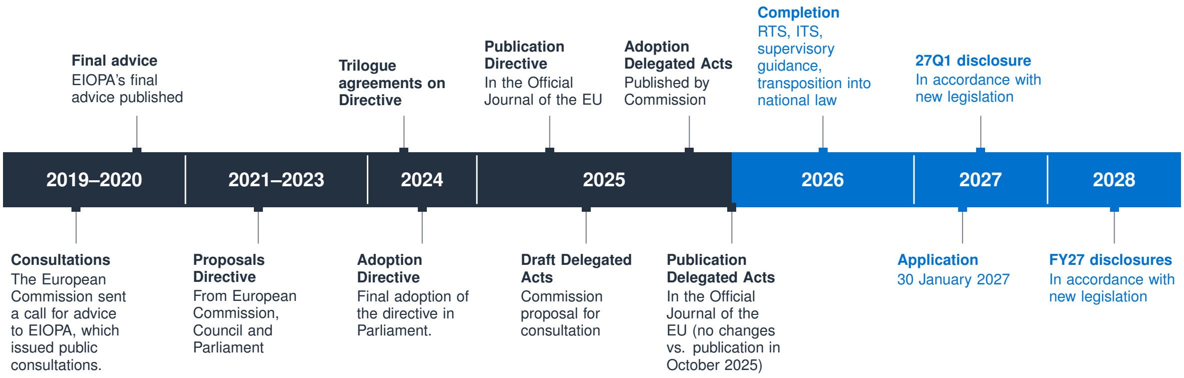

Figure 3: Timeline of the Solvency II Review. (Milliman, 2026)

A ratio of 100% means the insurer just meets the regulatory capital requirement. In practice, insurers maintain a buffer well above this level, in recent years between 175% and 220%, to absorb adverse market movements without breaching the requirement (Kouwenberg, 2017). These metrics are leading in the reflection of the solvency position within the market. Stakeholders and the market gain confidence in the insurer when their margins are solid.

2.1.2 Regulatory Changes

Section 2.1.1 described the Solvency II balance sheet under the rules currently in force. The Solvency II Review revises several of these building blocks. This section first provides an overview of the review process and then explains the changes that are relevant to this research. The second part examines the changes to the Volatility Adjustment in greater detail, focusing on their implications for the balance sheet and the allocation problem.

Timeline The starting point of the review is a clause in the original Solvency II framework that required the European Commission to reassess a number of technical elements, including the long-term guarantee measures (Directive 2009/138/EC (Solvency II), 2009). By the time this reassessment came due, a prolonged period of low and negative interest rates had exposed weaknesses in several of those elements. The Commission framed the review around four aims: supporting the long-term financing of the European economy, increasing the risk sensitivity of the framework, easing reporting and proportionality requirements for smaller undertakings, and improving group supervision. To prepare it, the Commission requested technical advice from EIOPA, which, following public consultation, issued its final opinion in 2020 (EIOPA, 2020). That advice underpinned the Commission's legislative proposal of 2021, and the subsequent negotiation between the Commission, the Council, and the Parliament produced a political agreement in 2024. The review was enacted through amending Directive (EU) 2025/2, published in the Official Journal in 2025 (Directive (EU) 2025/2, 2025), with the supporting detail set out in amending Delegated Regulation (EU) 2026/269 in early 2026.

What matters for this thesis is where the framework stands now. Both the Directive and the Delegated Regulation have been published, so the substance of the rules is fixed. Only the remaining technical standards, supervisory guidance, and national transposition are still being completed. The new rules do not apply until 30 January 2027, with the first disclosures under the regime reported thereafter (Directive (EU) 2025/2, 2025). The rules can therefore already be modelled ahead of their application date, which is what makes studying them now both possible and timely.

Overview of the changes The Review touches a number of these building blocks (Directive (EU) 2025/2, 2025), spanning the long-term guarantee measures, the technical provisions, and the SCR:

- **Yield curve extrapolation** The point beyond which the risk-free curve is extrapolated rather than read from the market is renamed the FSP, and the method used to extrapolate beyond it is replaced (Directive (EU) 2025/2, 2025, (39)). Both changes serve the same aim: to allow the curve to draw on market information further along the maturity spectrum, so that it reflects the deepest and most actively traded part of the market more faithfully than before. The new regulations are expected to produce lower RFR values, thereby increasing the BEL.

- **Risk margin** The cost-of-capital rate falls from 6% to 4.75%, and a time-dependent factor (a floored lambda term) is introduced that reduces the contribution of capital requirements far into the future (Commission Delegated Regulation (EU) 2026/269, 2026, (9, 51)). Together these substantially lower the risk margin and reduce its interest-rate sensitivity. Because the risk margin is driven by non-hedgeable risks, this is a change in the level of the technical provisions and not modelled in this thesis.

- **Interest rate risk** The stress applied to the risk-free curve is recalibrated. Under the current rules, the shock is purely relative to the base rate. This means that it becomes almost negligible as rates approach zero, understating the risk in a low-rate environment. The revised approach adds a fixed absolute component alongside the relative component, ensuring that the stress remains meaningful regardless of the rate level. The floor that previously prevented the downward shock from producing negative rates is also removed, allowing rates to fall below zero when this reflects market conditions. Since interest-rate risk is typically the dominant capital charge for insurers holding long-dated bonds against long-dated liabilities, this recalibration directly affects the capital charge relevant to asset–liability management (Commission Delegated Regulation (EU) 2026/269, 2026, pp. 43, 44).

Beyond these, the Review introduces a broad set of changes to proportionality, governance, and reporting: a lighter regime for small and non-complex undertakings, a liquidity risk management plan and other qualitative obligations, and revisions to the disclosure templates. These concern supervision, governance, and reporting, and therefore fall outside the scope of this thesis (PwC, 2025).

The revised Volatility Adjustment This is the central change examined in this thesis and the reason why the optimisation problem needs to be revisited. Under the current rules, the

VA is exogenous to the insurer's portfolio: it is derived from a country-wide reference portfolio. The revised VA is partly endogenous. Through a new credit spread sensitivity ratio (CSSR), it depends on the spread sensitivity of the insurer's own assets relative to that of its liabilities. This creates the chain of variables shown in Equation 8, from bond allocation to BEL, a link that did not exist under the current rules.

$$\text{Bond allocation} \rightarrow \text{PVBP (Assets)} \rightarrow \text{CSSR} \rightarrow \text{VA} \rightarrow \text{BEL} \tag{8}$$

From the actuarial perspective, this turns a fixed input into a feedback loop: the capital-optimal portfolio and the VA it earns are now jointly determined. That joint determination is precisely what this thesis sets out to examine.

Formulaic changes The revised VA is built from two parts, a permanent currency component and a macroeconomic component, so that the total adjustment is their sum.

$$\text{VA} = \text{VA}_{\text{cu}} + \text{VA}_{\text{macro}}. \tag{9}$$

The permanent component extends the currency VA of Equation (3).

$$\text{VA}_{\text{cu}} = \text{GAR} \cdot \text{CSSR} \cdot \text{RCS}_{\text{cu}}, \tag{10}$$

where the general application ratio GAR rises from 65% to 85%. The CSSR is new, and RCS_{cu} is the risk-corrected currency spread of the previous paragraph (Directive (EU) 2025/2, 2025, (41)). The macroeconomic component replaces the former country add-on,

$$\text{VA}_{\text{macro}} = \text{GAR} \cdot \text{CSSR} \cdot \max(\text{RCS}_{\text{co}} - 1.3\,\text{RCS}_{\text{cu}},\ 0) \cdot \omega_{\text{co}}, \tag{11}$$

where RCS_{co} is the risk-corrected spread on the country reference portfolio, the factor 1.3 sets the level at which a country's spread counts as elevated relative to the currency, and ω_{co} is a country activation factor (Directive (EU) 2025/2, 2025, (41)). The macro-component is generally much smaller than the permanent component.

The CSSR is the heart of the change. It compares how strongly the insurer's own fixed-income assets move when credit spreads shift to how strongly its liabilities move through the VA,

$$\text{CSSR} = \max\left(0,\ \min\left(\frac{\text{PVBP}_{\text{A}}}{\text{PVBP}_{\text{L}}},\ 1\right)\right). \tag{12}$$

The numerator is the spread sensitivity of the fixed-income assets. The denominator is the sensitivity of the technical provisions to the VA. The PVBP is the Price Value of a Basis Point, the price sensitivity of an item to interest rate changes. The CSSR equals one when the two match and falls below one when the assets are less spread-sensitive than the liabilities. Because it is capped at one, the CSSR can only scale the VA down, never amplify it (Directive (EU) 2025/2, 2025).

Both sensitivities are obtained by revaluation under a common shock, the notional VA, which is $\text{VA}_1 = \text{GAR}{\cdot}\text{RCS}_{\text{cu}}$, that is the currency VA with the CSSR set to one. Because numerator and denominator share this shock, its size cancels in the ratio. On the asset side, each fixed-income holding j is repriced at its yield y_j and at that yield raised by VA_0, and the per-instrument changes, scaled by the shift, are summed across the portfolio,

$$\text{PVBP}_\text{A} = \sum_j n_j \frac{P_j(y_j) - P_j(y_j + \text{VA}_0)}{\text{VA}_0}, \tag{13}$$

with n_j the number of instruments one holds of j and P_j its price. On the liability side, the best-estimate cash flows are discounted twice, on the risk-free curve and on the same curve shifted by VA_0, and the difference is divided by the shift,

$$\text{PVBP}_\text{L} = \frac{\text{BEL(RFR)} - \text{BEL(RFR} + \text{VA}_0)}{\text{VA}_0}. \tag{14}$$

The risk-corrected spread enters the VA in more than one form. The currency version RCS_cu and the country version RCS_co remain market-wide quantities: EIOPA continues to derive them from representative reference portfolios, which it updates annually and publishes for the VA calculation. The construction of the risk-corrected spread is also refined. The risk correction, previously almost constant, becomes a piecewise-linear function of the spread level, so that it grows as the spreads widen and dampens the VA more in stressed markets (Directive (EU) 2025/2, 2025, (45)).

2.2 Optimisation theory

The theory underlying the optimisation is relevant to understanding the methods used in this thesis. The following subsections first introduce the mathematical foundations of optimisation and then discuss the different optimisation algorithms considered for this problem.

2.2.1 Mathematical optimisation theory

Mathematical optimisation is concerned with selecting from a set of admissible alternatives, the one that performs best according to a stated criterion (Nocedal & Wright, 2006). A problem is described by three ingredients: a vector of decision variables $x = (x_1, \ldots, x_n)$ that the analyst controls; an objective function $f(x)$ that assigns a single value to every choice of those variables; and a feasible set $\mathcal{X}$ that collects the choices regarded as admissible. The canonical form is

$$\min_{x \in \mathcal{X}} f(x), \tag{15}$$

where a maximisation problem is recovered by minimising $-f$, so it suffices to treat one of the two.

The feasible set rarely includes all of $\mathbb{R}^n$; it is carved out by constraints. Three types recur: equality constraints $g(x) = 0$, which fix a quantity at an exact value: inequality constraints $h(x) \leq 0$, which keep it on one side of a threshold, and bound (or box) constraints $l \leq x \leq u$, which confine each variable to an interval. Together they define

$$\mathcal{X} = \{ x \in \mathbb{R}^n : g(x) = 0,\ h(x) \leq 0,\ l \leq x \leq u \}, \tag{16}$$

and a point satisfying all of them is called feasible. An optimum, if it exists, is a feasible point for which no other feasible point yields a better objective value.

How difficult a problem is to solve and which methods can be applied, depends less on its size than on the shape of the relationship between the decision variables and the objective value. Examples of structural characteristics that describe the shape:

- *Linearity* Whether the objective and constraints are linear in the decision variables. A linear objective improves at a constant rate in every direction, whereas a nonlinear one curves, so that the rate of improvement itself changes with position.
- *Smoothness* Whether the objective varies smoothly enough to possess a well-defined gradient everywhere. A smooth function has, at each point, a local direction of steepest improvement. At a kink, a point where the function stays continuous but its slope changes abruptly, no gradient is defined.
- *Convexity* Whether the feasible set and the objective are convex. A set is convex if the segment between any two of its points stays inside it, and a function is convex if its graph never bulges above the chord joining two of its points (Boyd & Vandenberghe, 2004).

These characteristics matter because they decide what it means to have solved a problem. A local optimum is a feasible point that is best within its immediate neighbourhood, with no nearby feasible point performing better. A global optimum is best over the entire feasible set. When the problem is convex, the two coincide: every local optimum is automatically global, and it is frequently unique, so locating any local optimum settles the question. When the problem is non-convex, that guarantee is lost. The landscape may then hold several local optima of differing quality, and a search that only moves towards nearby improvement can come to rest on one of them while a better solution exists elsewhere. Smoothness and convexity together therefore determine both what "solved" means and how a solution can be reached.

2.2.2 Optimisation algorithms

Once the mathematical properties of a problem have been established, the next step is to select an appropriate optimisation algorithm. This choice can affect the efficiency with which a solution is found. Selecting an algorithm often involves a trade-off between computational speed and solution quality. A suitable algorithm should provide an appropriate balance between these two objectives. The algorithms discussed below are those relevant to this research. The SLSQP was chosen as it is a basic algorithm, suited for many problems. The algorithms mentioned afterwards try to logically solve the shortcoming of the previous algorithms.

Sequential Least Squares Quadratic Programming (SLSQP) looks for a better point by using the slope of the objective to move downhill. At each step it replaces the real problem, which is hard, by a simpler one it can solve quickly: a bowl-shaped approximation of the objective around the current point, with the constraints replaced by straight lines. Solving that gives a direction d to step in,

$$\min_{d} \ \tfrac{1}{2}\, d^\top B_k\, d + \nabla f(x_k)^\top d \quad \text{s.t.} \quad \nabla c_j(x_k)^\top d + c_j(x_k)\ \{= 0,\ \geq 0\}, \tag{17}$$

where $\nabla f(x_k)$ is the slope and B_k stands for the curvature, which the method estimates as it goes. Then it takes a step and repeats (Kraft, 1988). An advantage of SLSQP is that equality and inequality constraints can be incorporated directly into these local subproblems, making it suitable for constrained nonlinear optimisation. This approach can be computationally efficient when the objective and constraints are smooth. However, because it relies on local gradient information, it is less suitable for the problem in this study, where several relationships contain kinks and are therefore non-smooth.

Differential Evolution (DE) uses no slope at all, to overcome the problem of having a flat solution space. It keeps a whole group of candidate solutions and improves them over many rounds. To propose a new candidate it takes one member of the group and shifts it by the difference between two other, randomly chosen members,

$$v_i = x_{r_1} + F\,(x_{r_2} - x_{r_3}), \qquad F \in (0, 1], \tag{18}$$

a step that is automatically large when the group is still spread out and small once it has clustered around a good region. The shifted candidate is then mixed with the original, governed by a probability CR, and kept only if it gives a better objective value (Storn & Price, 1997). Because DE evaluates candidate solutions based only on their objective-function values and does not require gradient information, it can handle non-smooth functions and irregular search spaces. However, the absence of gradient information can also make the search less directed and result in slower convergence.

Covariance Matrix Adaptation Evolution Strategy (CMA-ES) CMA-ES also works without gradients and with a group of candidate solutions, allowing it to guide the search more effectively. It maintains a "cloud" describing where promising solutions are likely to lie, characterised by a centre, an overall size, and a shape. In each iteration, it draws a batch of candidate solutions from this cloud (Ha, 2017):

$$x_k \sim m + \sigma\,\mathcal{N}(0, C), \qquad k = 1, \ldots, \lambda, \tag{19}$$

with m the centre, σ the size, and C the shape. After scoring the candidates it adjusts all three: it moves the centre towards the best ones, $m \leftarrow \sum_i w_i\, x_{i:\lambda}$; it grows or shrinks the size depending on whether it keeps making progress in one direction or merely zig-zags around; and it reshapes the cloud to stretch along the directions that have been paying off. In effect it learns the shape of the landscape while searching, if the good region is a long, narrow valley, the cloud becomes a matching long, narrow ellipse, so its candidates land usefully (Hansen & Ostermeier, 2001). This is what lets it handle irregular, awkwardly shaped objectives efficiently.

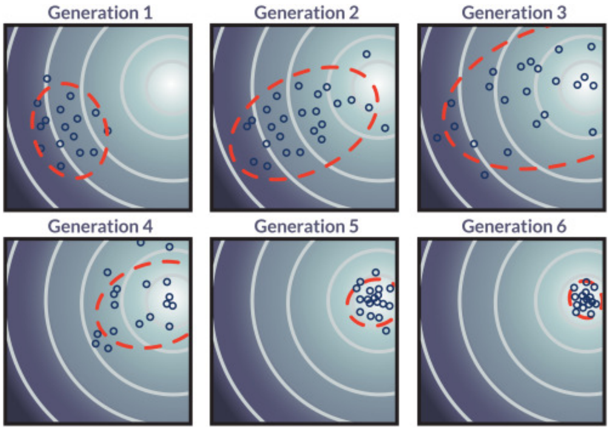

Figure 4: Covariance matrix adaptation evolution strategy (CMA-ES)

3 The model

The following chapter explains what the model used in this research can do and why. It first links the mathematical theory of optimisation to the problem at hand, then describes the pipeline of the programmed model in more practical terms. Finally, it discusses a few relevant modelling decisions.

3.1 Mathematical formulation of problem

It is useful to first connect the mathematical theory to the problem. At each optimisation, the insurer allocates a fixed budget across the available instruments to optimise its regulatory capital position. The optimisation performed for a moment in time is therefore independent from other optimisations.

The decision vector is a set of portfolio weights $w = (w_1, \ldots, w_n)$, one per investable individual bond. Each weight is bounded, $0 \leq w_i \leq 1$, and the weights exhaust the budget, $\sum_i w_i = 1$, so that w lies on the unit simplex. A weight maps to a holding through

$$\text{units}_i = \frac{w_i V}{P_i}, \tag{20}$$

with V the total portfolio value and P_i the unit's price.

The capital position is optimised through the surplus $\text{OF}(w) - \text{SCR}(w)$, maximised to enlarge the absolute buffer of Own Funds over the requirement. OF denotes Own Funds. The feasible set is fixed by three constraints: the budget $\sum_i w_i = 1$, the box bounds $0 \leq w_i \leq 1$, and a Solvency-ratio floor $\text{SR}(w) \geq x$, a management target above the regulatory minimum of 1.00.

How the objective depends on w is not obvious, and tracing it is what fixes the class of problem to be solved. On the budget axis the asset value is constant, $\sum_i \text{units}_i P_i = V$, so the weights act on the objective through three other channels. The liabilities are non-linear: the CSSR is a clipped quantity,

$$\text{CSSR}(w) = \min\big(\max(\frac{\text{PVBP}_A(w)}{\text{PVBP}_L}, 0), 1\big), \tag{21}$$

linear inside its bounds but kinked at 0 and 1. The Volatility Adjustment scales with it, $\text{VA}(w) = \text{GAR} \cdot \text{CSSR}(w) \cdot \text{RCS}$ (plus the macroeconomic term, which adds a further maximum). The best estimate discounts the liability cash flows at a rate that contains it,

$$\text{BEL}(w) = \sum_t CF_t \left(1 + r_t + \text{VA}(w)\right)^{-t}, \tag{22}$$

non-linear in w and, through the CSSR, non-smooth. Own funds are then $\text{OF}(w) = V + \text{carry}(w) - \text{BEL}(w)$.

The capital requirement compounds these features. The interest-rate charge is a worst case over shock directions,

$$SCR_{\text{interest rate}} = \max\left(\sum_{\text{currency}} SCR_{\text{rate up}}, \sum_{\text{currency}} SCR_{\text{rate down}}\right) \tag{23}$$

non-differentiable where the binding direction switches. The concentration charge is a sum of hinges over issuer thresholds.

$$SCR_{\text{conc}}(w) = \sum_j RF_j \max\left(E_j(w) - T_j,\, 0\right). \tag{24}$$

Where $E_j(\omega)$ is the exposure to bond j in percentages, T_j is the concentration threshold dependent on the credit score and RF_j is the risk factor. The modules are aggregated through a square root of a quadratic form,

$$\text{SCR}(w) = \sqrt{\mathbf{s}(w)^\top C\, \mathbf{s}(w)}, \tag{25}$$

with $\mathbf{s}(w)$ the vector of sub-module requirements and C is the correlation matrix 10. The objective then closes the construction, $\text{Surplus}(w) = \text{OF}(w) - \text{SCR}(w)$.

Three properties follow. Together they determine which optimiser can be used. The problem is non-linear: the discounting $\text{BEL}(w)$ and the square-root aggregation $\text{SCR}(w)$ are non-linear in w. It is non-smooth: the clip in the CSSR, the maxima in the interest-rate charge, and the hinges in the concentration charge each introduce kinks. These are points at which the objective remains continuous but is not differentiable, so that no gradient is defined there. Finally, it is non-convex: the objective is the difference between two w-dependent quantities, a form that is not convex in general. Convexity would be required for a local optimum to be global, and differentiability for gradient-based search; the problem offers neither, which directs the choice of method towards the derivative-free, population-based approach.

Combining this information leads to the following optimisation problem.

Objective	$\max_w \ \text{Surplus}(w)$	
Constraints (Hard)	s.t. $\sum_{i=1}^{n} w_i = 1,$	*(Portfolio allocation constraints)*
	$0 \le w_i \le 1,$	
(Soft)	Solvency Ratio ≥ 1.5	
	$\dfrac{\sum_t (aCF_t - lCF_t)^2}{\sum_t (lCF_t)^2}$	*(Cashflow matching penalty)*
	$\dfrac{1}{T-1}\sum_{t=1}^{T-1}(aCF_{t+1} - aCF_t)^2$	*(Cashflow difference penalty)*
Decision variables	$W = (w_1, w_2, \ldots, w_n) \in \mathbb{R}^n,$	n = Amount of available bonds at t

The soft constraints are added to the objective in the model itself, however are presented in separate lines for the sake of clarity. The soft constraint of the Solvency Ratio is implemented via Big-M. The SR constraint is soft to be able to track performance when the constraint may not be satisfied. The cash flow penalties make sure that the cash flow of the assets are realistic. In Section 3.3 this will be discussed more extensively.

3.2 Model architecture

The model developed to solve this problem is described in the following sections. First, the overall model structure is presented in Figure 5.

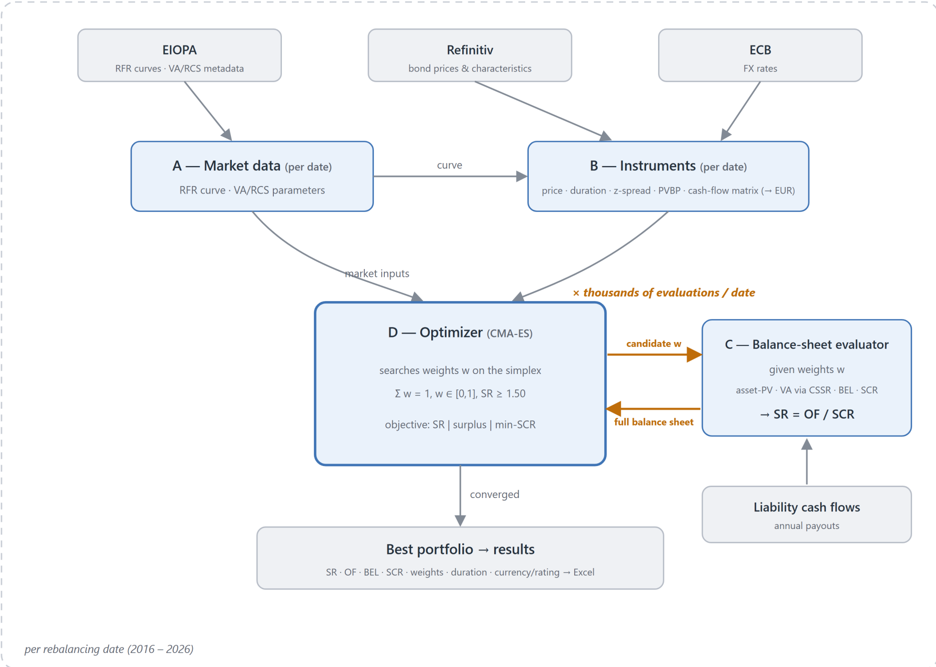

Figure 5: Conceptual model pipeline with four layers

Market data For each date, the model extracts the EIOPA risk-free curve and its associated parameters. The risk-corrected spread underlying the VA is not published directly. It is therefore inferred from the published VA by dividing it by the old GAR (65%), RCS = VA/GAR. This proxy is a pragmatic choice.

Instruments For each bond, the model combines a weekly market price and YTM time-series data with its static characteristics, including issuer, country, rating, coupon, maturity, and currency. Only the price and yield are obtained from market data. The cash flow schedule is derived from the static bond characteristics, while the Z-Spread, duration measures, and PVBP are calculated using these characteristics and the relevant curve. These quantities are computed once per date because none of them depend on the portfolio weights. Pre-computing them keeps the inner-optimisation loop computationally efficient.

Optimiser The optimiser receives the input data and proposes an initial candidate portfolio. Using its weights, the balance-sheet evaluator constructs the corresponding balance sheet. The optimiser then uses this result to propose a new candidate portfolio based on the previous one.

Balance-sheet evaluator Given a set of weights, the evaluator builds the balance sheet in a fixed order: total asset value, Volatility Adjustment, best-estimate liability (discounted at the curve raised by the VA), SCR and its sub-modules, and the Solvency Ratio.

3.3 Modelling choices

During the development of the model, several modelling decisions were made to ensure appropriate model behaviour and realistic results. The most relevant decisions are discussed in chronological order.

Proxy of RCS The risk-corrected spread (RCS) is not published directly by EIOPA in the same way as the VA. Instead, its value is determined through a series of formulas. The RCS used in the model is therefore reverse-engineered from the published VA value. Since $VA = GAR * RCS$ and the VA is published, the RCS can be calculated as $RCS = VA/GAR$.

Liability cash flow creation The liability cash flows are based on values reported in the literature (Curry, 2019; Thorendal, 2023). They are implemented as fixed absolute cash flow amounts for each projected year. This is a deliberate high-level representation, so the cash flows cannot be traced back to specific characteristics of the insurer. The underlying characteristics would have to be assumed as well and do not give more insight to the insurance company. It would only give a longer trace of assumptions to reason back to their own company.

Decision variables: bonds or buckets The choice between individual bonds and bond buckets reflects a trade-off between precision and computational efficiency. The buckets were defined based on a combination of rating, time to maturity, and bond type, namely government or corporate. Experiments showed that optimisation over buckets was approximately twice as fast, but did not produce sufficiently realistic results. The research therefore uses individual bonds as the decision variables.

Static versus dynamic portfolio The choice between a static and a dynamic portfolio reflects a trade-off between independence and realism. Under the static approach, each optimisation is performed independently and the amount available for investment is newly determined from the BEL at that date. The portfolio selected in the previous optimisation therefore has no effect on the next optimisation. Under a dynamic approach, the previous portfolio would instead form the starting point for the next optimisation. Its value would evolve through changes in the market value of the held bonds, coupon payments, and principal repayments at maturity. Consequently, the Assets-to-BEL ratio would depend partly on the performance of portfolios selected in earlier periods, introducing carry effects and path dependence into the results.

The aim of this thesis is not to develop an optimisation algorithm that performs optimally or to give a direct example of the optimal allocation for every insurer in real-world investment management. Instead, it is to investigate the effects of regulatory changes within the Solvency II framework and measure how that impacts the choice for an optimal portfolio and its performance. Independence is therefore prioritised, and a static portfolio is used.

Cash flow matching penalties Initial results showed that the asset cash flows were highly volatile across the projected years and could not be considered sufficiently realistic. While this is not necessarily problematic given the use of a proxy model, it does indicate that an adjustment is required. To improve the realism of the cash flow profile, two penalty terms were therefore added to the objective function. The first penalises differences between asset and liability cash flows in each year. The second penalises differences between asset cash flows in consecutive years.

These penalties affected the results by reducing Surplus. Without the penalties, the model achieves a Solvency Ratio of approximately 1.90. Introducing them therefore trades part of this higher Solvency Ratio for a more realistic cash flow profile. In the resulting setup, the Solvency Ratio generally remains at or below 1.55.

Disabling EU spread exemption Under the Standard Formula, government bonds issued by EEA countries are exempt from the spread shock. In this model, the spread-risk sub-module makes the largest contribution to the total SCR. The model contains enough eligible government bonds to invest almost the entire portfolio in this type of bond. This reduces the total SCR from more than EUR 1 million to approximately EUR 70,000, resulting in a Solvency Ratio of around 70 rather than approximately 1.5 and a Surplus that is almost as large as the Own Funds.

To prevent this exemption from dominating the optimisation, it is disabled in the model. This is a deliberate deviation from the unmodified Standard Formula and should be considered when interpreting the SCR results. Also different forms where one would turn the spread exemption off during optimisation but turn it on for calculating the final balance sheet would decrease the Spread SCR with the ratio of government/total investments.

Choice of objective variable Three candidate objectives were considered: minimising SCR, maximising Surplus, and maximising the Solvency Ratio. Minimising SCR alone ignores the return side of the balance sheet, while maximising the Solvency Ratio is already enforced as a soft constraint and would therefore be redundant as an objective. Maximising Surplus is chosen instead, as it directly measures the capital remaining after accounting for both liabilities and the SCR.

Addition of excess-return An excess return term was added to the Asset PV, also increasing the Surplus, to account for the potential additional return associated with taking more investment risk. It is calculated by mutliplying the amount invested to the Z-Spread, thereby assigning a higher potential return to bonds with higher spreads. As this is a simplified and deterministic approximation that does not account for factors such as the probability of default, the excess return is deliberately excluded from the optimisation objective. It is, however, included when calculating the resulting balance sheet and Surplus, as it represents potential additional income generated by future payments of the fixed income products. These will be higher for more risky investments.

4 Data

The model is driven by three input streams: bond-level market data from Refinitiv (Refinitiv, 2026), the monthly risk-free interest-rate term structures published by EIOPA (European Insurance and Occupational Pensions Authority, 2026), and a set of foreign-exchange rates from the ECB (European Central Bank, 2026). As the results of the model are completely dependent on the data going into the model, the validity of any result rests on those inputs. An unrepresentative sample would render the output uninformative regardless of the model's correctness. A biased sample would not directly create a problem, however it is important to know how it is biased. In that way the biases in the results can be better understood. This section sets out what the data is, how much of it there is, and how it was selected, ahead of the exploratory analysis in Section 4.2.

4.1 Data framework

The core of the dataset is a panel of 513 bonds observed weekly over 2016 - 2025, yielding on the order of 170,000 weekly observations, comprising 403 corporate and 110 government issues. All are investment grade, spread across four rating buckets (BBB 146, A 143, AA 123, AAA 101). By currency the panel is split almost evenly between the euro (256 bonds, 49.9%) and the US dollar (255, 49.7%), with sterling marginal (2 bonds, 0.4%). Issuance dates span 1997 to April 2025, 177 bonds issued before 2015 and 336 from 2015 onward, while maturities run from late 2015 to 2071, so that the panel deliberately mixes seasoned and recently issued bonds across a wide maturity spectrum. Issue sizes range from 1.0 to 71 billion USD, with a median of 1.75 billion. By sector the corporate book is led by a broad services bucket (191 bonds, 47%), after which manufacturing and energy (55 each, 14%), with consumer goods (35), telecommunications (29), electric utilities (27), transport (9) and gas distribution (2) making up the remainder, spanning eight industry sectors in total. The government issues form a separate sovereign sector.

Selection applied a common filter to both classes and a class-specific one on top. Throughout, only plain-vanilla fixed-coupon bonds were retained, all investment grade, with coupons capped at 8% (realised maxima of 6.95% for corporates and 6.50% for governments) and a minimum issue size of approximately one billion USD (the smallest realised issues are 1.0 billion for corporates and 1.18 billion for governments). This size floor sits well above the minimum issue sizes used for inclusion in major fixed-income benchmark indices (Bloomberg Index Services Limited, 2026), so that every retained bond is liquid and market-representative by a comfortable margin. Corporate issues were drawn through their TRBC industry classification and government issues through their sovereign instrument types (Dutch State Loans, Bundesanleihen, Obligations Assimilables du Trésor, and similar) and country of issue, covering thirteen countries in total. Financial-sector issuers, banks and insurers were excluded by design: risk in this sector as a whole is argued to be too large. The exclusion is therefore a deliberate simplification. The panel was constructed with a roll-over design across rating and maturity buckets, so that each bucket stays populated as bonds mature and are replaced.

For each bond the dataset records two kinds of information. Static attributes: issuer, rating, sector or instrument type, currency, issue and maturity dates, coupon, day-count convention and issue size. These fix the bond's identity and risk classification, while all static information is assumed to stay the same. Time-varying attributes, weekly price and yield, track its market behaviour across the 2016 - 2026 window. Together they feed the balance sheet construction and the exploratory analysis after that.

The two remaining streams support the bond data. Discounting rests on the EIOPA-published risk-free interest-rate term structures, which supply the curves applied to the liability and bond cash flows. Foreign-exchange rates convert the multi-currency positions to a common basis.

4.2 Exploratory data analysis

Before the data enters the model it is worth characterising it along two dimensions: the cross-sectional composition of the bond panel, and the time-series behaviour of the market inputs that drive the balance sheet. The aim is twofold, to confirm the inputs are sound, and to surface the features that will later shape the results.

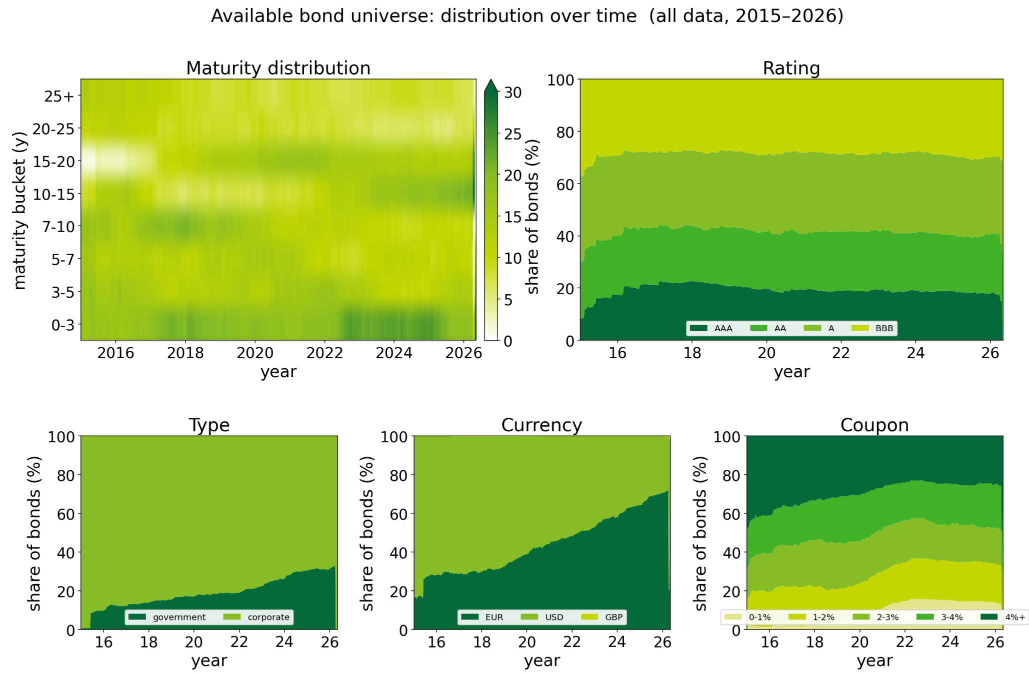

Figure 6: Cross-sectional decomposition of bonds available per type

The distributions of rating and coupon remain relatively stable over time. By contrast, the distributions of bond type, namely government or corporate, and currency change over the observation period. The distribution of time to maturity is generally well spread, although several lighter and darker areas are visible. This type of figure is used again later because it provides a clear overview of the portfolio allocation.

Time to maturity and rating are considered to have the greatest influence on the model results. Both characteristics affect the calculation of the spread shock. Time to maturity also determines when the bond's notional is repaid, making it important for cash flow matching and interest-rate risk calculations. Several dark bands are visible in the figure, while the BBB and A categories appear darker in the second half of the observation period.

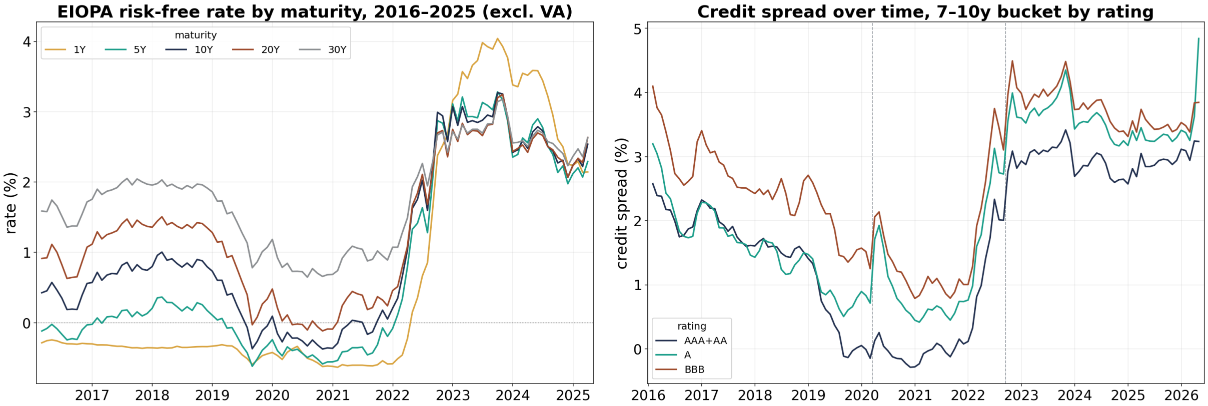

Figure 7: (left) RFR over time & (right) Spread over time

The Yield-to-Maturity and the Risk-Free Rate together form the discount base for assets and liabilities, and both display a similar three-part pattern over the sample period. Figure 7 (left) shows the EIOPA risk-free rate across maturities. From 2016 to 2019 rates were low and dispersed across maturities; from 2019 to 2022 they settled into a uniformly low, near-zero-to-negative range, with one- and five-year rates around −0.5%. From 2022 onward rates repriced sharply upward, reaching 3–4% within a year, with the curve briefly inverting in 2023–2024 before easing through 2025. This regime shift matters directly, as it drives both liability discounting and the interest-rate stress in the capital requirement.

A comparable pattern emerges in the YTM data, although the transition between the first two sections is smoother than in the RFR data. Figure 7 (right) shows credit spreads for the 7–10-year bucket by rating, with a clear and monotonic ordering throughout: BBB above A, followed by AAA and AA. The same stress episodes are visible. Spreads spike during COVID-19 and widen sharply through 2022, rising from well below 1% in 2021 to 3–4% thereafter. High-grade spreads even dip slightly below zero in 2019–2021, consistent with top-rated bonds trading through the swap curve in a negative-rate environment. Across the sample, maturity shifts yields across a wider range than rating quality within a given maturity. This indicates that the term structure, rather than credit quality, is the dominant driver of yield levels, although this difference narrows once the curve flattens after 2022.

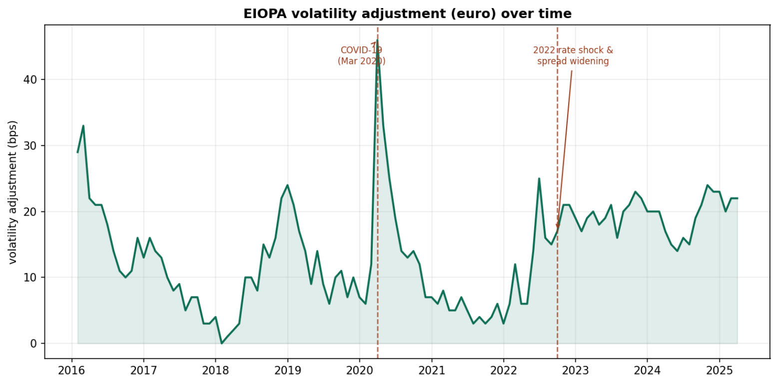

Figure 8: Volatility adjustment over time

Figure 8 plots the euro Volatility Adjustment over the same period. It is small in calm markets, near zero in 2018 and 2021, and spikes precisely when markets are stressed: it reaches roughly 46 basis points at the onset of COVID-19 in March 2020 and stays elevated around 20 basis points through the 2022 rate shock and spread widening. This counter-cyclical behaviour is exactly what the volatility adjustment is designed for. As the mechanism at the centre of this study, its historical path is worth establishing here.

5 Experimental Design

5.1 Research questions

The main research question of this thesis is as follows:

Main research question: *In what way, and through what mechanisms, do the Volatility Adjustment changes in the Solvency II 2020 revision change the capital-optimal bond portfolio of a Dutch life insurer?*

The central question is broad and cannot be answered through a single experiment. The sub-questions are therefore designed to examine the optimisation of a bond portfolio under the revised regulatory framework from four different perspectives. Each sub-question changes one main element and examines how the outcome responds. Across Questions 1 to 4, these elements are regulation, portfolio allocation, investment strategy, and market or insurer conditions. The four sub-questions are introduced below.

Sub-question 1 Comparison *In what way does the optimal portfolio under the new regime differ from the current one?*

This is assessed by comparing bond-portfolio optimisations under the current and revised regulatory frameworks. Two intermediate scenarios are added to provide more insight into which changes drive the observed effects. In these scenarios, the optimal portfolio under the current framework is held fixed and evaluated under a subset of the revised regulations. The comparison consists of the following four scenarios:

A. Optimise under the old regime.

B. Hold the old optimal portfolio under the new regime, but with the old Volatility Adjustment.

C. Hold the old optimal portfolio under the full new regime, including the new Volatility Adjustment.

D. Optimise under the full new regime.

This comparison between the different scenarios aim to understand better:

- *A vs B:* The impact on performance the other changing regulations.
- *B vs C:* The impact on performance of the changing VA.
- *C vs D:* The change in portfolio allocation through changing mechanisms.
- *A vs D:* the full difference in performance and allocation.

Splitting the results in this way allows the total change to be decomposed into three contributions: the effect of the remaining regulatory changes, the additional effect of the VA change, and the effect of the insurer's own reallocation. In the present analysis, the remaining changes are introduced first and the VA change second. Reversing this order would provide an alternative attribution of the VA effect and show how sensitive the decomposition is to the order in which the changes are introduced.

Sub-question 2 Mechanism *In what way do particular asset allocations translate into the long-term guarantee (LTG) measures under the new framework, and how can these differences be quantified?*

Where Sub-question 1 shows *that* the portfolio changes, this question asks *why*, by examining the link created by the revised rules. It seeks to isolate and understand the trade-off between the BEL and the SCR. The portfolio allocation is changed deterministically along one bond-characteristic axis at a time, such as rating or duration. Measuring the resulting changes in, and sensitivities of, the BEL and SCR provides insight into how these quantities drive changes in Surplus and the Solvency Ratio.

For this question, the analysis starts from an optimal portfolio. Five percentage points of the allocation to AAA-rated bonds are exchanged for an equal allocation to BBB-rated bonds. The same procedure is then repeated for 10%, 15%, 20%, and 25%, and in the opposite direction by shifting allocation from BBB-rated bonds to AAA-rated bonds. The resulting changes in the BEL, SCR, Surplus, and Solvency Ratio are compared to identify the key drivers of the optimal portfolio. This procedure is repeated for multiple dates and along multiple bond-characteristic axes.

Sub-question 3 Risk appetite *To what extent do different levels of risk appetite produce different outcomes under the new regime?*

Risk appetite is divided into three levels: risk-seeking, medium, and risk-averse. This distinction is implemented by changing the Solvency Ratio constraint. The constraint levels are set at 1.4, 1.6, and 1.8 for the risk-seeking, medium, and risk-averse strategies, respectively. This analysis aims to examine how changes in investment strategy affect the optimal portfolio weights and resulting performance.

The cash flow matching penalties in the objective function generally push the Solvency Ratio down. It is therefore relevant to assess whether the model can still achieve a Solvency Ratio of 1.8 under the risk-averse constraint.

Sub-question 4 Sensitivity analysis *In which region of the parameter space is the impact of the VA revision on the capital-optimal bond portfolio strongest, and how robust is the central VA-impact result to changes in these parameters?*

The preceding questions are evaluated using a single market state and one liability profile. This question examines the extent to which the findings depend on these inputs. For this sensitivity analysis, the following variables are shocked:

- The interest-rate level is shocked with an absolute shock of plus and minus 50 and 100 bps
- The Risk Corrected Spread is set to be 80 bps for each optimisation. This is the other term next to the CSSR determining the VA.
- The liability cash flows are shocked to change the duration of the liabilities.

This question assesses the robustness of the model in two ways. First, it examines whether the optimisation produces a feasible solution. Second, it examines how the resulting allocation and performance change.

5.2 Common experimental setup

All four sub-questions are based on the same baseline setup, allowing their results to be compared and the analysis to be reproduced. The focus differs between the sub-questions, however, and the variables and characteristic axes examined may differ accordingly.

Base configuration

- **Computational implementation:** The model and experiments were implemented in Python, using NumPy and pandas for numerical computations and data processing. Portfolio optimisation was performed using the CMA-ES implementation from the cma package. Matplotlib and openpyxl were used for visualisation and Excel-based output, respectively.
- **Decision variables:** individual bonds, numbering approximately 500. Optimising at the individual-bond level allows the model to favour or exclude each instrument based on its own characteristics. This gives the allocation a high degree of flexibility, but also creates a larger search problem with approximately 500 decision variables.
- **Portfolio sizing:** static, with the Asset side reset to 1.2× BEL at every date. The BEL is approximately EUR 15 million. Each date is solved independently, so neither portfolio allocations nor investment results are carried over from earlier dates. As a result, comparisons between Scenarios A, B, C, and D are made at a specific point in time and are not affected by previous allocations. The 1.2 multiple keeps the insurer above the regulatory minimum while leaving room for the Solvency Ratio constraint to influence the solution.
- **Objective:** Surplus, defined as Own Funds minus the SCR. Surplus represents the insurer's capital buffer in monetary terms and can be compared directly across solutions. It also avoids the unstable behaviour that the Solvency Ratio may show when the SCR is small. In addition, it provides a natural basis for incorporating the risk-appetite parameter used in Sub-question 3. The Solvency Ratio and a minimum-SCR objective remain available as additional checks.
- **Solvency Ratio constraint:** at least 1.50. This is an arbitrary management target above the 100% regulatory minimum. Market practice is between 1.8 and 2.2 (Table 17). It keeps the solutions away from the regulatory boundary and within a more realistic range.
- **Evaluation and optimisation frequency:** monthly, around the fifteenth day of each month, from March 2016 to April 2025. This frequency corresponds to the monthly publication of the EIOPA risk-free curve and VA. Each date therefore represents a new set of regulatory and market inputs, while also capturing stress periods such as the COVID-19 crisis and the 2022 repricing.

6 Results

The results of the experiment will be shown and discussed in the following chapter. First, it will be shown that the model behaves accordingly, after which the sub-questions will be answered one-by-one. This will lead to the conclusion, where the main question will be answered.

6.1 Validation

Before the model's results can be used to draw conclusions, it must first be established that its outputs are realistic and stable. This is assessed using three forms of evidence: whether the allocation for a single date is realistic, whether the output is consistent with the relevant market conditions, and whether the optimiser produces stable results. To keep the evidence concrete and comparable, the analysis focuses on three reference dates. These dates were selected to represent the different conditions under which the model operates. The main differences between these conditions are the RFR, credit-spread levels, and the number of bonds available:

Table 1: Overview of the three reference dates: market and portfolio inputs.

	2016-12-16	2020-04-17	2024-06-14
Regime	low rate	COVID stress	normalised
RFR 1y / 10y / 20y	−0.30 / 0.58 / 1.11%	−0.38 / −0.19 / 0.02%	3.51 / 2.76 / 2.68%
VA (EIOPA)	14.5 bps	38.6 bps	14.9 bps
Active bonds (universe 513)	181	206	287

The distribution of portfolio weights is an important criterion for assessing whether the solution is realistic. The graph shows a smooth and gradual decline in portfolio weights when the bonds are ranked by their invested weights, with no abrupt jumps or clear clustering. The differences between the weights are larger in 2016 than in 2024, as can be seen more clearly in Table 18 in Appendix B. In 2016, the largest position represents 2.1% of the portfolio, compared with 0.8% for the fiftieth-largest position. In 2024, these weights are 0.7% and 0.6%, respectively. This difference may be related to the number of bonds available in each year. Overall, the observed weight distributions are considered realistic.

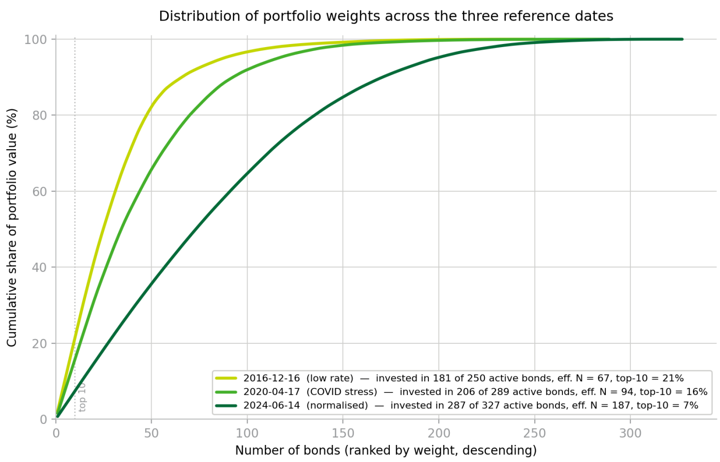

Figure 9: Distribution of portfolio weights

During the model construction phase, the cash flow profiles of the assets and liabilities were examined regularly to assess whether the model behaved appropriately. Figure 10 shows a clear difference between the three profiles. The first profile displays considerably more volatility than the other two. It has peaks at the beginning of the maturity range and at ten years to maturity, indicating that the portfolio is concentrated in bonds maturing at these points. Beyond these peaks, the cash flow changes considerably from year to year.

The 2020 and 2024 cash flow profiles are more similar. They show less volatility and no extreme peaks. In both cases, the assets generate higher cash flows at shorter maturities and lower cash flows at longer maturities, which is considered realistic. Overall, the observed cash flow behaviour is considered reasonably realistic.

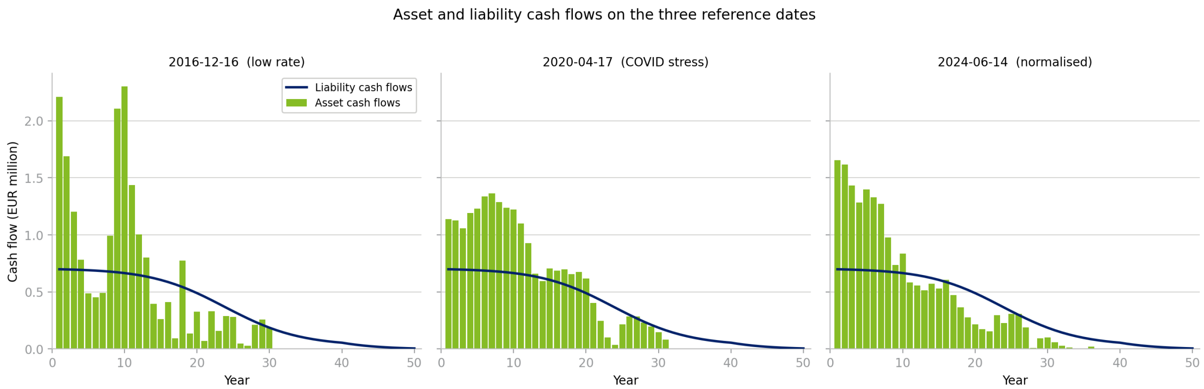

Figure 10: cash flow profiles of the assets and liabilities

The most important constraint in the optimisation, the Solvency Ratio constraint, is satisfied for all three reference dates. The margin above the constraint is, however, remarkably small. This suggests that the optimisation operates close to the constraint boundary. The cash flow matching penalties included in the objective function are a plausible explanation for this behaviour. Nevertheless, the results confirm that feasible solutions can be found for all three reference dates.

Table 2: Solvency Ratio on the three reference dates.

	2016-12-16	2020-04-17	2024-06-14
Regime	low rate	COVID stress	normalised
Solvency Ratio	1.505	1.560	1.542

Comparison with the market Looking at the size of the balance sheet items alone would not establish whether the model behaves appropriately. Therefore the results are compared to the largest Dutch life insurers, who also are main clients for the Insurance team of the Financial Risk department of Deloitte.

Table 3: Position of the model relative to the Dutch solo life market (SFCR 2022 - 2025).

Metric	Market (range, mean)	Model (3 dates)	Assessment
BEL / Assets	0.67 - 0.87 (0.76)	0.83 (0.83)	comparable (anchor)
SCR / Assets	3.3 - 4.7% (4.0%)	11.3 - 11.9% (11.7%)	∼3× higher
Own funds / Assets	6.1 - 9.3% (7.6%)	17.4 - 18.6% (18.0%)	∼2.4× higher
Solvency Ratio	1.69 - 2.31 (1.92)	1.50 - 1.56 (1.54)	acceptable, conservative

Table 3 shows that the BEL-to-Assets ratio of the hypothetical insurer is comparable to that of the real insurers and, in some cases, slightly higher. This represents a deliberately conservative calibration. However, the comparison should be interpreted with care because the model does not include the Risk Margin, which forms part of the balance sheets of real insurers.

The SCR of the hypothetical insurer is approximately three times as high as that of the real insurers, making it more difficult for the modelled insurer to remain solvent. One important driver of this relatively high SCR is the spread-risk sub-module. The average spread shock calculated by this sub-module is approximately 8%. By comparison, real insurers generally have lower exposure to assets subject to spread risk and may apply a lower or zero spread shock to European government bonds.

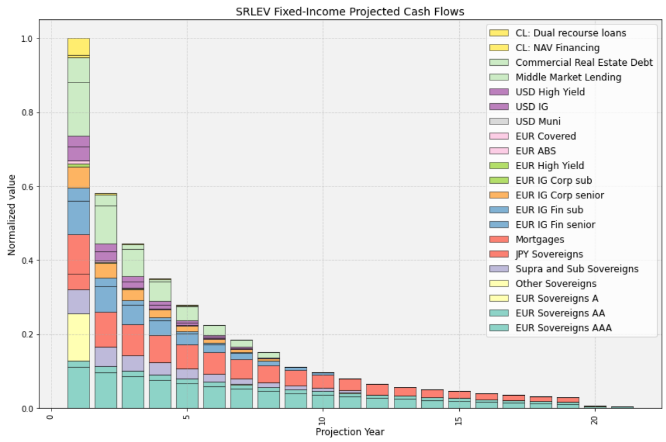

Figure 11: Fixed-income cash flow profile of SRLEV N.V., life insurance company of Athora (Thorendal, 2023)

Figure 11 shows the projected cash flow from SRLEV N.V., the life insurance entity of Athora responsible for brands as Zwitserleven and Reaal. All but 5 categories are types of bonds that SRLEV holds. This gives an idea of the form of a possible cash flow profile of fixed-income products. Comparing this Figure 11 to Figure 10, it can be concluded that the profiles resemble more as we move forward in time. We should keep in mind that the liability cash flow profile of SRLEV N.V. is not known. The picture also gives insight in the different types of bonds that such a life insurance company has.

Stability of the optimiser Given that the optimiser is stochastic, it is valid to test whether the solutions are constant. This is tested by optimising 5 times over the reference dates, keeping the data exactly the same.

Table 4: Optimizer stability across five identical runs, at three levels of aggregation

	2016-12-16	2020-04-17	2024-06-14
Single product			
Weight overlap, same bonds	77.6%	72.2%	70.6%
Category (rating × duration)			
Weight overlap, same category	93.9%	94.3%	93.8%
Result variables			
cash flow profile agreement	96.7%	97.5%	95.7%
Surplus (coeff. of variation)	0.17%	1.28%	0.63%
SCR (coeff. of variation)	0.23%	0.52%	0.45%
BEL (coeff. of variation)	0.02%	0.10%	0.02%

Table 4 shows that even though the data is the same, the weights on the level of individual stocks does change. This is expected as different products do have similar characteristics and therefore the system may be indifferent between them. This is backed by the second line, where more stability is seen on the level of product categories. The resulting cash flow as a third level, has an even higher overlap between different runs. This means that even though products may differ, the cash flows created by them are relatively constant. Finally, the last layer is the performance of the objective value, which is even more constant. This suggests that the optimisation method performs constant, however there are multiple product ensembles leading to approximately the same result. The date within the COVID peak has the most change.

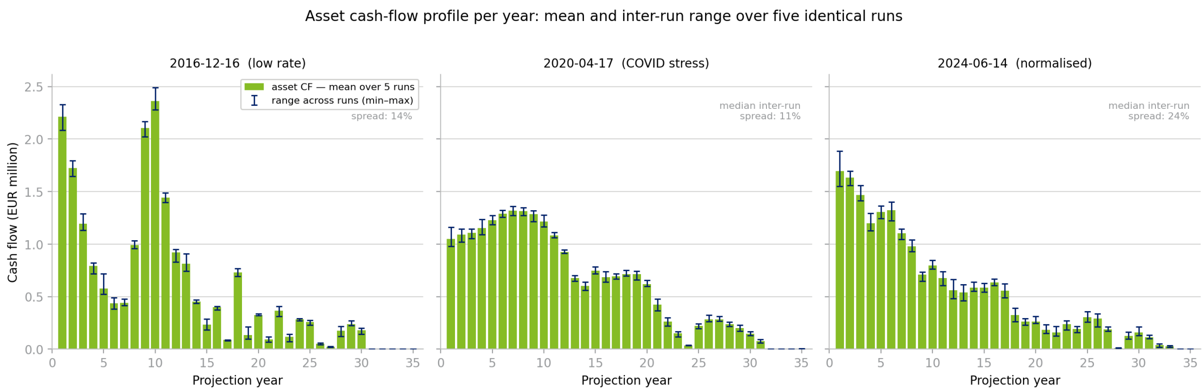

Figure 12: Asset cash flow profile per projection year

Figure 12 shows an expected structure, considering the overlap from Table 4. One observation that stands out, is within the picture from 2016. The diagram looks like a relatively random behaviour with a volatile cash flow after the second peak. However, it appears that this cash flow is found to be optimal by the model in all the runs.

Taking all this evidence together, the realism of the allocations, the comparison with the market, and the consistency across runs, it can be concluded that the model performs in a generally valid way. The volatility of the cash flow profile of 2016 is not preferred, and therefore we want to understand what creates this volatility and see how far in the timeline the volatility perseveres.

6.2 Sub-question 1 Regulation change

Before investigating the decomposition between the different scenarios, a starting point must be given. The allocation of the portfolio is as follows:

Figure 13: Portfolio allocation old regime

Several observations can be made from Figure 13. In the maturity plot, one bucket dominates the first half of the observation period and accounts for roughly 40% of the portfolio weights. During this period, the maturity buckets between three and seven years remain remarkably small.

The portfolio is concentrated in corporate bonds and euro-denominated bonds. This can be explained by the way these assets are treated in the model. Government and corporate bonds are subject to the same mechanisms, while corporate bonds with the same rating typically offer higher spreads. Euro-denominated bonds do not incur a currency SCR, whereas USD-denominated bonds do. The coupon distribution changes over time, but the changes are small and relatively spread out across the portfolio.

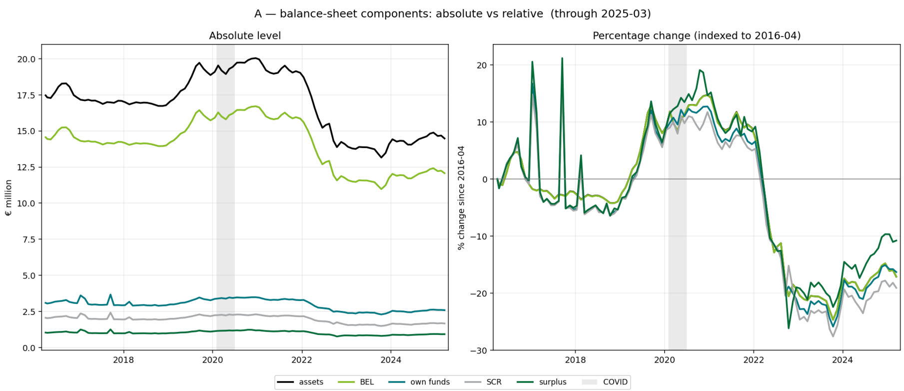

Figure 14: Balance sheet items over time

Figure 14 shows that movements in the BEL are closely reflected in the other balance-sheet values, with the RFR and VA acting as the main drivers of its changes. In the left-hand panel, the pattern observed in the assets and liabilities is also visible in OF, the SCR, and Surplus, although it appears less pronounced. The right-hand panel shows percentage changes and makes clear that this weaker appearance results from the different scales of the values rather than from a genuine difference in their relative movements. The relative changes across all five lines are closely aligned. The main exception is a set of three early spikes in the SCR, OF, and Surplus, which are caused by a sharp increase in the Z-Spread of bonds approaching maturity.

Table 5: A bond within days of maturity distorting the expected-return proxy (valuation date 15 Sep 2017).

Bond	Maturity	ttm (y)	RFR	spread	Z-Spread	mkt. value	Δexp. ret.
FR0011321447	2017-09-18	0.008	−0.36%	−11.1%	24,096%	335K	664K

The bond matures in three days leading to unrealistically low duration numbers. The duration is used within the formula of the Z-Spread in the denominator and the Z-Spread is used to calculate the excess-return. It reaches the Own Funds via this route. The SCR spikes through a spike in concentration and interest rate risk. Both signs that the optimiser has picked up on the opportunity. From the resulting measures side, it can be concluded that the RFR and VA seem a dominant driver in changes within the balance sheet. These spikes therefore have limited impact on the overall interpretation, as the balance-sheet components move largely in parallel.

Scenario $A \to$ Scenario B: The extrapolation This section compares how the balance-sheet numbers react to the new method of extrapolation. While we keep in mind that over the analysed time period, the RFR will be lower within the new method.

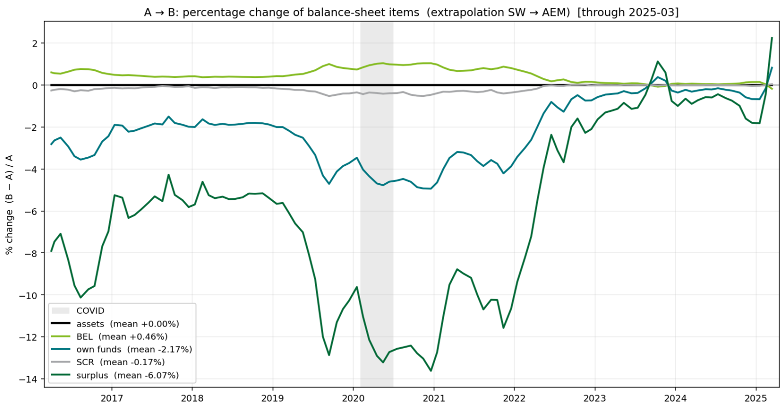

Figure 15: Percentage change of balance sheet due to new extrapolation

Figure 15 shows that the BEL is higher due to lower RFR, the SCR is lower due to a lower SCR IR, as it depends on the level of the RFR. Own Funds is lower because the BEL is higher. The decrease in the SCR has a positive effect on Surplus, but is insufficient to offset the impact of the changing BEL and Own Funds. As the relative change to the Surplus is higher than to the SCR, we can reason that the Solvency Ratio is lower as well. Also we may recognise the shape of the Figure 7, showing the change of the RFR.

Scenario $B \rightarrow$ Scenario C: The Volatility Adjustment The changes within the VA formula are mainly a changing GAR ratio from 65% to 85% and the addition of the CSSR term. The change in GAR makes the VA higher, where the CSSR only has the ability to make the VA lower. Firstly, we want to understand how this impacts our timeline.

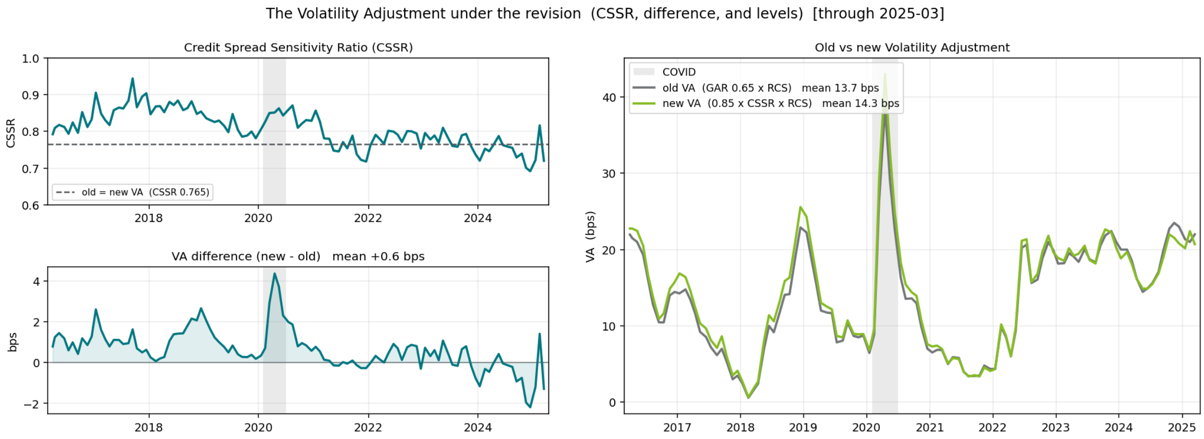

Figure 16: Difference in VA, decomposed

Figure 16 shows the difference of VA in percentage and in absolute numbers, with the CSSR as item to decompose its difference. We observe that whether the new VA is smaller or larger, depends on whether the CSSR is at or above 0.765. This is the rate for which the $GAR_{old} = GAR_{new} * CSSR$. The magnitude of the old VA together with the distance to the 0.765 of the CSSR, determines how much the VA changes. This can be seen during the COVID period. In the end, the difference between the old and new VA is relatively small at a mean of 0.6 bps.

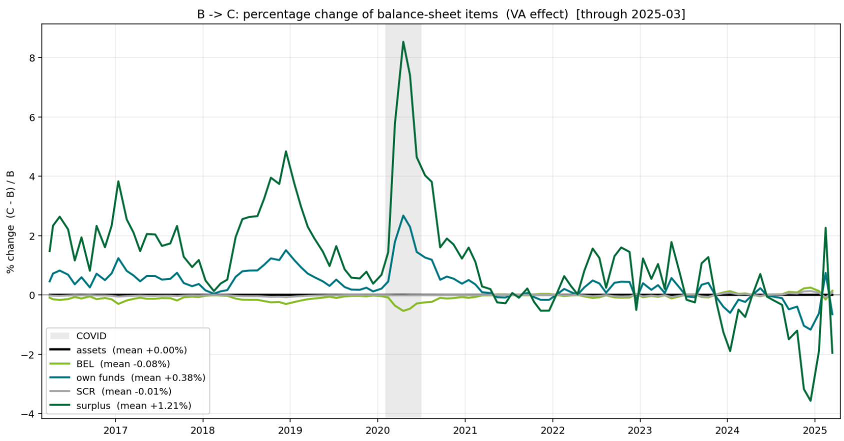

Figure 17: Impact of new VA on the balance sheet

A higher VA means a lower BEL and higher Own Funds. The SCR is also lower, meaning that all effects on Surplus are positive. The shape of the change in Surplus, resembles the change in VA strongly. Concluding that the overall changes are relatively small in this instance.

Scenario $C \rightarrow$ Scenario D: The response The last change will be the re-optimisation within the same set of rules.

From Figure 18 we observe that overall the portfolio changes to be more invested into longer

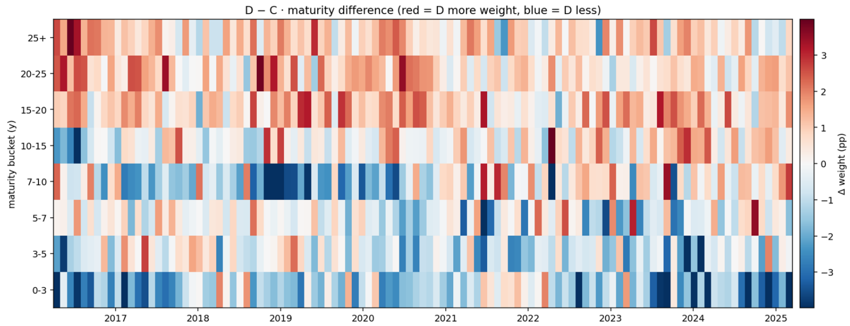

Figure 18: Change of portfolio allocation

maturity products. On average 3% of the weight shifted from maturity under 10 years, to maturities of over 10 years. This may have something to do with the decrease in RFR in longer maturities, leaving a higher present value of the liabilities for those years.

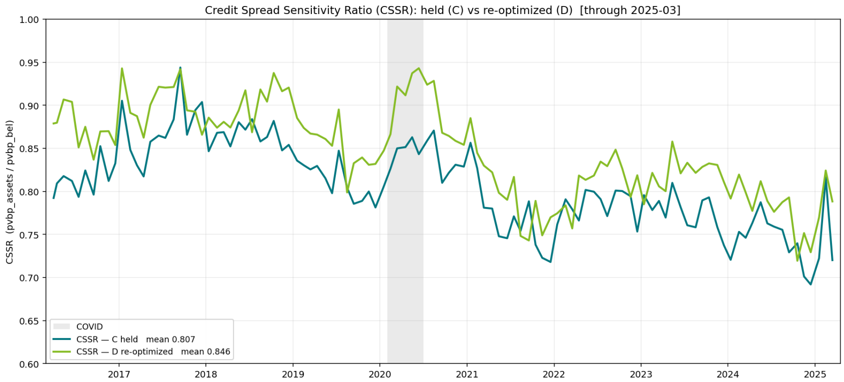

Figure 19: difference in CSSR

Figure 19 shows the CSSR before and after re-optimisation. The line of the new CSSR is generally above the old one, demonstrating that the optimiser takes the opportunity to use the CSSR to increase the VA and decrease the BEL.

Scenario $A \rightarrow$ Scenario D: The total effect The combined effect is examined next. How does the result within the two scenarios differ?

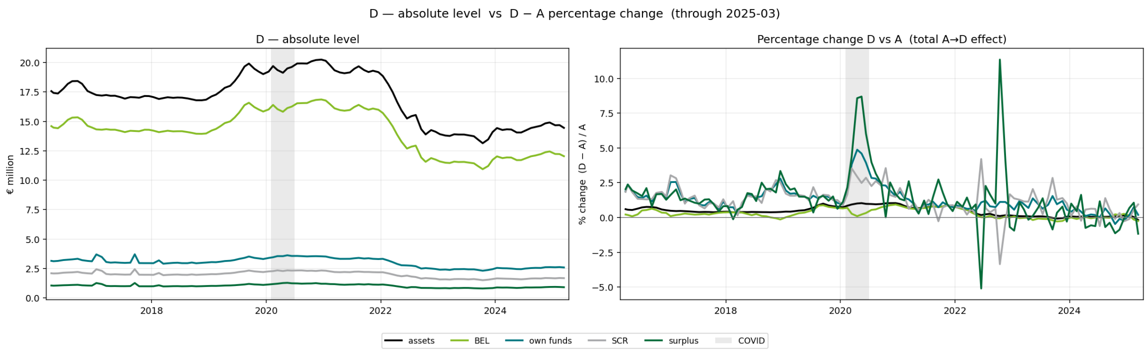

Figure 20: The resulting components of the optimisation within new regulations and percentage difference between A and D

The absolute balance-sheet values in Scenario D do not show any clear differences. They move in the same direction, broadly following changes in interest rates, and the relative distances between the balance-sheet items remain similar. The percentage changes show that all items generally increase in value. Assets increase slightly more than the BEL, although the difference is small. Own Funds, the SCR, and Surplus follow approximately the same path, with a few pronounced spikes in Surplus during the COVID period and the low-interest-rate period in 2022. Overall, Scenarios A and D appear to behave in almost the same way, with a slight advantage for Scenario D.

6.3 Sub-question 2 Deterministic portfolio change

To examine the mechanism through which portfolio allocation affects the balance sheet, the optimal portfolio is deterministically shifted along different bond characteristics. The centre of each figure represents the optimal portfolio, while movements along the x-axis represent percentage-point reallocations between two groups of bonds. For example, a shift of -10 in the first figure reallocates 10 percentage points from AAA–AA bonds towards A-rated bonds. By evaluating BEL, SCR, Surplus and the Solvency Ratio after each deterministic shift, the sensitivity of the balance sheet to portfolio allocation can be examined. This provides insight into why particular portfolio allocations are preferred by the model.

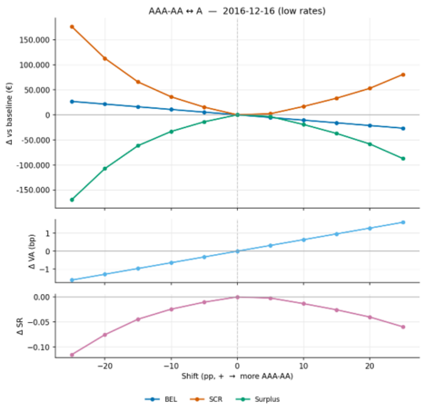

Figure 21: Portfolio change along bond rating axis

Figure 21 shows the effect of reallocating portfolio weight between A-rated and AAA–AA-rated bonds. While the BEL responds approximately linearly to the allocation change, the SCR shows a considerably stronger and non-linear response. The only deviation from the linear BEL relationship occurs when the Credit Spread Sensitivity Ratio (CSSR) reaches its upper bound of one. Around the optimal portfolio, changes in SCR therefore outweigh the opposing changes in BEL, causing Surplus to deteriorate when moving in either direction. This indicates that, for changes in credit rating, the SCR is the dominant balance-sheet mechanism determining the preferred allocation. Similar behaviour is observed for most of the other rating reallocations.

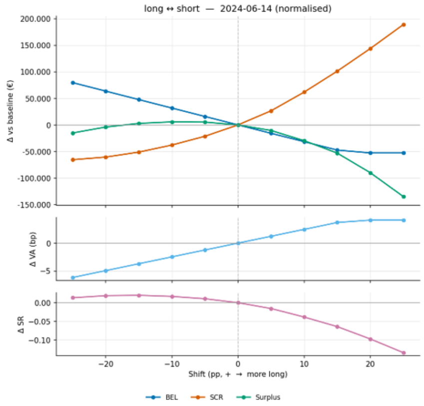

Figure 22: Portfolio change along duration axis

A different interaction emerges when portfolio weight is shifted between short- and long-duration bonds. Between a shift of approximately -20 percentage points and the optimal portfolio, substantial changes in the allocation result in only limited changes in Surplus and the Solvency Ratio. Over this range, the change in BEL largely counterbalances the change in SCR, creating a relatively flat balance-sheet outcome despite the considerable change in portfolio composition. Consequently, multiple substantially different duration allocations can produce almost equivalent balance-sheet performance. This creates a steering possibility, as the insurer can alter the duration composition of the portfolio within this range without materially affecting Surplus or the Solvency Ratio.

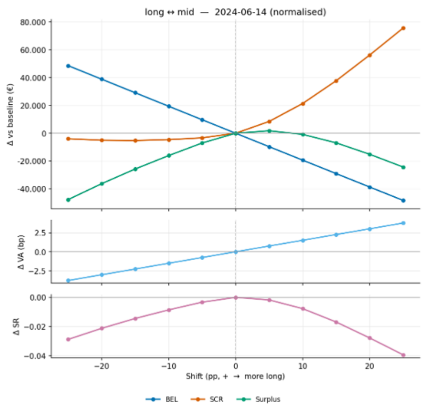

Figure 23: Portfolio change along duration axis

Figure 23 illustrates a different balance between the same two mechanisms. The SCR curve is relatively flat around its minimum, meaning that SCR alone provides only a weak preference for a particular allocation. In contrast, the BEL changes more strongly and approximately linearly as the portfolio is shifted towards longer-duration bonds. The BEL effect can therefore outweigh the relatively small SCR differences around the minimum, shifting the preferred balance towards longer maturities. This demonstrates that BEL can materially affect the portfolio allocation when the SCR is sufficiently insensitive to changes in duration. The resulting shift towards longer-duration bonds is consistent with the allocation change observed after the regulatory revision in Sub-question 1.

6.4 Sub-question 3 Risk appetite change

This sub-question examines how different levels of risk appetite lead to different portfolio choices and how those portfolios perform.

Figure 24 shows a clear, stepwise difference between the different levels of risk appetite. The trends in the maturity-distribution heatmaps become more pronounced as risk aversion increases. The concentration appears to be in the same group of products, which mature around 2026. Their time to maturity decreases at roughly the same rate as calendar time advances.

The differences between the risk-appetite levels are also visible in the rating-mix and currency-mix diagrams. As observed earlier, euro-denominated bonds become more attractive when stricter Solvency Ratio constraints must be met.

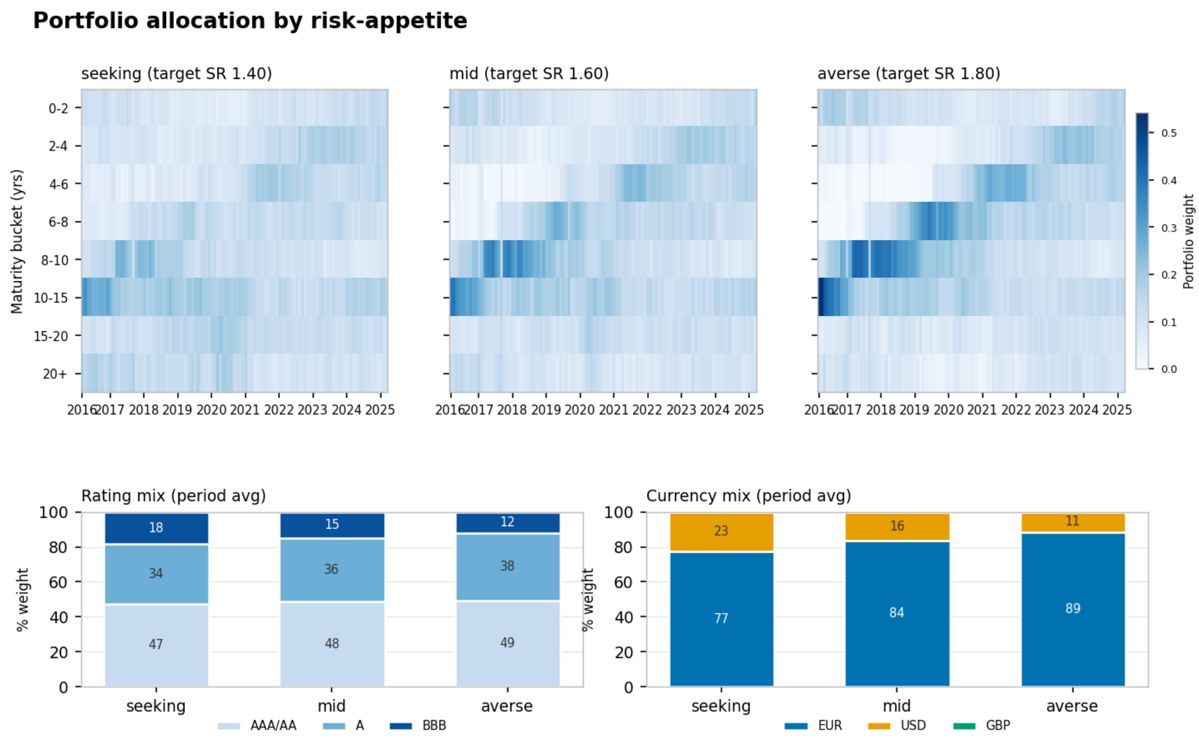

Figure 24: The allocation choices of the different levels of risk appetite

There is a clear difference in the size of the SCR across the three strategies. The total market SCR decreases from €3.62 million for the risk-seeking strategy to €3.10 million for the mid strategy and €2.73 million for the risk-averse strategy. Although the risk-averse strategy is more concentrated in particular groups of bonds, the concentration risk submodule decreases from €0.28 million to €0.18 million. The largest relative decrease is observed for currency risk, which approximately halves from €0.97 million to €0.49 million. The other submodules decrease as well, with spread risk falling from €1.51 million to €1.34 million and interest-rate risk from €0.46 million to €0.32 million.

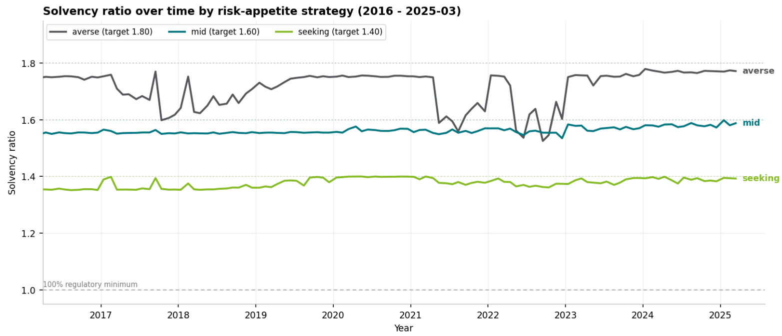

Figure 25: SR level of different strategies over time

Figure 25 shows that for the seeking and mid strategy, it was possible to satisfy the SR constraint over the whole timeline. For the risk-averse strategy this was not the case for almost half of the optimisations (47/109). An important driver is the spread risk module, which is highly dependent on the available products. When there are products available with certain maturity and spread in a particular combination, the shock is too high to still satisfy the constraint. Overall, the mid strategy has a consistent performance on the SR.

6.5 Sub-question 4 Sensitivity analysis

Sub-question 1 varies the regulation and Sub-question 3 the insurer, Sub-question 4 varies the market and insurer conditions. Everything runs under the new regime with the baseline risk appetite fixed. Only the environment moves, the interest-rate level, the BEL duration, and the RCS are manipulated. Each is a one-at-a-time shock around the baseline, re-optimised over the monthly series.

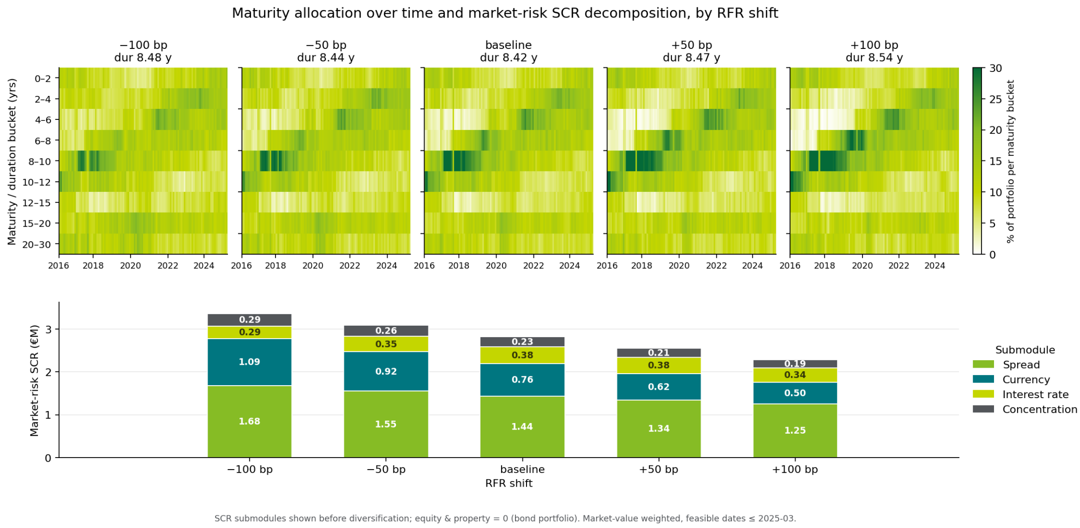

Figure 26: Allocation differences over RFR manipulations

Figure 26 shows that the maturity distribution over time follows a pattern similar to that of the different risk-appetite strategies in Section 6.4. The concentration in the dominant buckets becomes more pronounced as the RFR increases. Similar patterns can also be observed in the rating distribution and the SCR decomposition. The allocation under the higher-RFR scenario resembles that of the risk-averse strategy.

One difference is that the interest-rate SCR does not decline monotonically. This may be related to the relaxation of the cash flow matching constraint, which creates larger differences between the asset and liability cash flows.

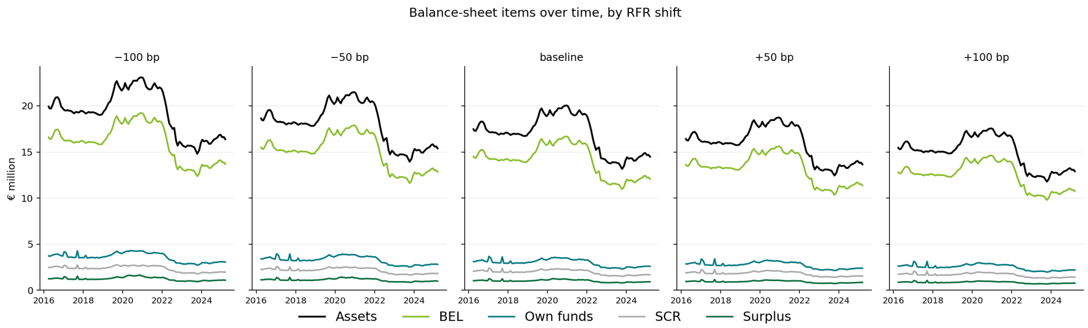

Figure 27: Balance sheet items over the manipulations

Figure 27 illustrates how the different RFR scenarios affect the overall balance sheet. Changes in the RFR affect the BEL through the discounting of the liability cash flows. Since the Asset side is approximately 1.2 times the size of the BEL, these changes also affect the scale of the balance sheet as a whole.

The change in balance-sheet scale affects both the portfolio allocation and its performance. One possible explanation for the changes in allocation is the difference between the regular spread and the Z-Spread, which alters the relative attractiveness of the bonds under the different RFR scenarios. The discount rate also has a larger effect on longer-dated liabilities. Consequently, the impact of the RFR changes is not uniform across the entire time period.

Table 6: Share of feasible optimisations per RFR manipulation

	RFR −100 bp	RFR −50 bp	Baseline	RFR +50 bp	RFR +100 bp
Feasible (clean, /109)	108	109	108	108	107
% feasible (clean)	99.1	100.0	99.1	99.1	98.2

Table 6 shows that the feasibility is not significantly impacted by the RFR manipulation.

Table 7: Feasibility of the reshaped-liability optimum per target duration

Target duration (y)	6	8	10	11	12	13	14	16	20
2016-04-01 (low rates)	×	×	✓	✓	✓	×	×	×	×
2016-12-16 (low rates)	×	×	×	✓	✓	×	×	×	×
2019-08-02 (low rates)	×	×	✓	✓	✓	✓	×	×	×
2020-04-17 (COVID)	×	×	✓	✓	✓	✓	×	×	×
2022-12-02 (high rates)	×	×	×	×	×	×	×	×	×
2024-06-14 (normalised)	×	×	×	×	✓	×	×	×	×
% feasible	0	0	50	67	83	33	0	0	0

It appeared that the manipulations on the liability cash flow caused a relatively high percentage of infeasible results. Table 7 shows which optimisations got feasible solutions. There were no feasible solutions found for the date with the high rates, even though there was a feasible solution within the baseline run. The duration of the liabilities in that run was 12.5. The duration of 12 has the best result, with a corridor of at least one feasible solution between 10 and 13. Outside of that area there are no feasible solutions found. It can be concluded that the duration and form of the cash flow profile has a significant impact on the performance of the model.

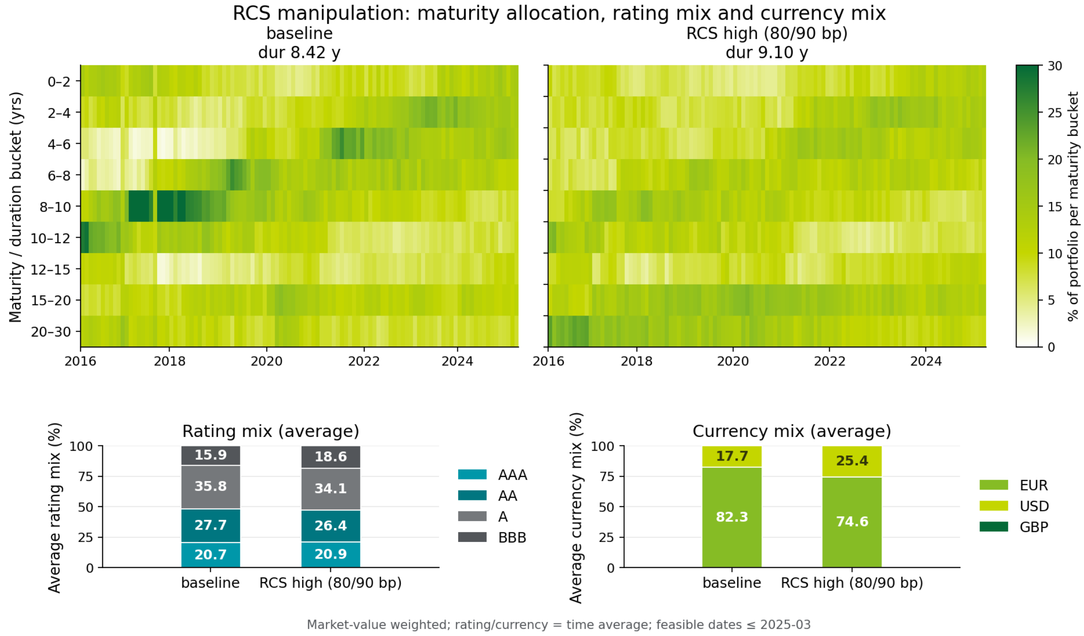

Figure 28: Allocation difference baseline vs RCS

The maturity allocation shown in Figure 28 is relatively well-diversified. This confirms the earlier pattern: under conditions with less 'pressure' on the algorithm to reach a certain SR, the maturity distribution tends to be more diversified. Also the rating mix and currency mix follow this pattern. The rating leans more to worse ratings and more investments into USD bonds. That the higher RCS has more space is an expected consequence, as the higher RCS only has impact on the VA, making sure it is a higher number.

Table 8: Market-risk SCR submodule decomposition, baseline vs. RCS high

Scenario	Spread	Currency	Interest rate	Concentration	Σ (undiv.)	Market (div.)
Baseline	1.44	0.76	0.38	0.23	2.81	1.84
RCS high (80/90 bp)	1.57	1.08	0.46	0.33	3.44	2.17

The riskier investment profile is also visible in the decomposition of the SCR Market module. All submodules get a higher value, which is also expected for most of them, however the maturity distribution was more diversified, therefore it is an increased concentration risk is against expectations. The larger Own Funds show to have more room to cherry pick more towards bonds that are good for the system.

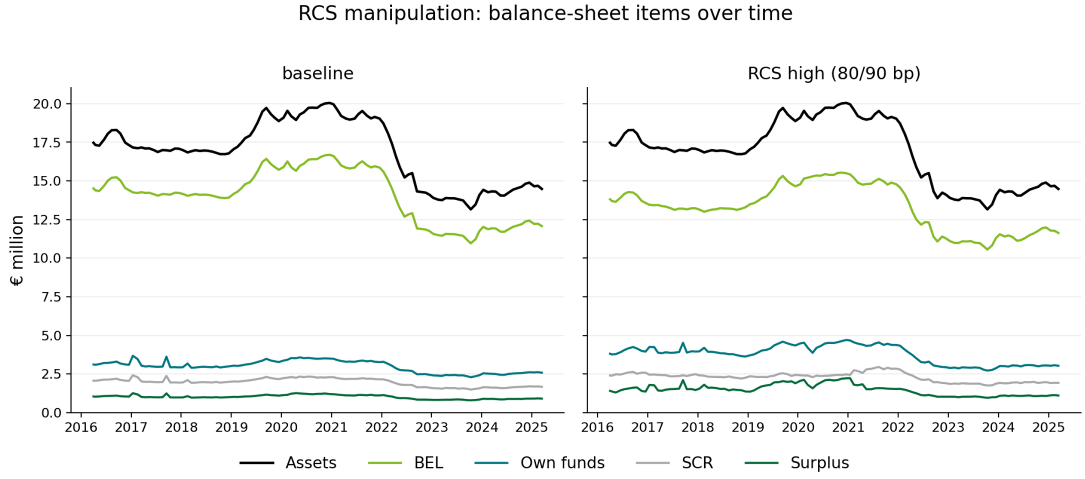

Figure 29: Balance sheet items over the manipulations

Figure 29 shows clearly that the assets stay the same and the gain comes from the decreased BEL. It is the first time that we observe that the behaviour of the SCR is not parallel to the OF and Surplus. When the model has more liberty to choose it, as it is not hard to satisfy the SR constraint, the SR is able to go up. On its high-point the SR almost reaches 200%.

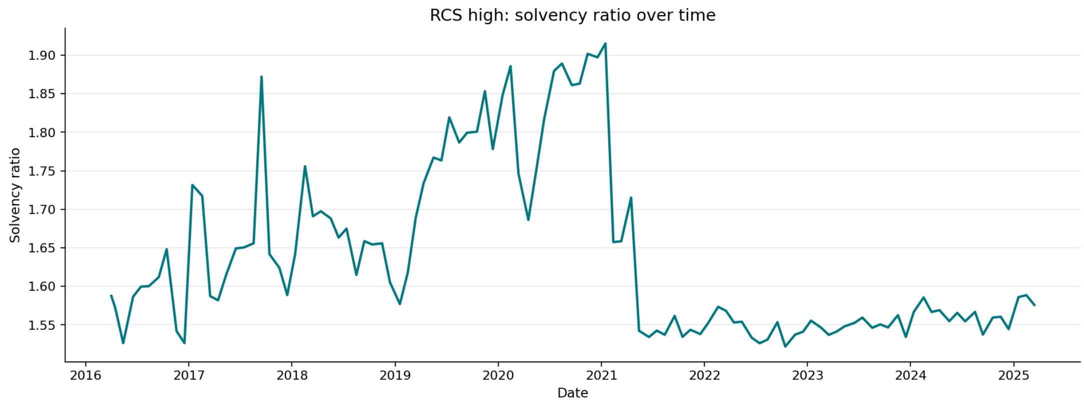

Figure 30: Solvency Ratio in RCS manipulation

Figure 30 underlines the similar behaviour of the model. During periods where the model has the possibility to rise, there is high volatility within on the SR and has the ability to rise significantly. If not, it will stick close to the 1.5 threshold.

7 Conclusion

For this research, the objective was to investigate how the changes to the VA introduced by the SII 2020 Review affect the capital-optimal bond portfolio of a Dutch life insurer. The effect of the revised VA is assessed in the context of the other modelled regulatory change, the new RFR, and under different investment strategies and market and insurer conditions. The analysis focuses on changes in the optimal portfolio allocation and in its performance, measured through capital surplus and the Solvency Ratio. Related balance sheet items, including the BEL and the SCR, are also examined to explain these changes.

The results show that the revised VA affects both the composition of the optimal portfolio and the way in which portfolio allocation translates into balance sheet performance. In particular, the allocation-sensitive VA creates a new interaction between the BEL and the SCR. Changes in the allocation can therefore alter the BEL in a way that partly offsets changes in the SCR: shifts of up to 20 percentage points in the duration allocation were found to leave Surplus and the Solvency Ratio largely unchanged, allowing different portfolios to produce similar performance. This effect is not equally strong in all scenarios.

The results also show a tendency towards longer-maturity bonds, on average a 3 percentage-point shift from maturities under 10 years to maturities over 10 years. The Surplus increased on average by 1.21%. The shift towards longer maturities may reflect the positive effect of these bonds on the VA. It may also be related to the lower RFR, which can increase the duration of the liabilities and make longer-duration assets more relevant for asset–liability management. Therefore, the observed changes should be understood as the combined result of the revised VA, the new RFR, and the prevailing market conditions. The size of this effect also depends on risk appetite: total market SCR fell from €3.62 million under a risk-seeking strategy to €2.73 million under a risk-averse one, though the latter failed to meet the Solvency Ratio constraint in 47 of 109 months.

The primary contribution of this research for life insurers is the insight that the revised VA introduces opportunities to steer portfolio allocation and that longer-maturity bonds are generally rewarded under the new framework. The existence of this steering effect is relevant for all life insurers, although its magnitude depends on the current portfolio composition and the prevailing market conditions, as it was not consistently observed across all scenarios. A similar conclusion applies to the shift towards longer bond maturities. Over time, little consistency was observed regarding the maturity groups that were replaced by others, indicating that these changes remain dependent on the existing portfolio allocation and the market environment.

The increase in Surplus is a more general finding. Based solely on this study, it cannot be determined whether this increase is primarily driven by the revised regulation, the composition of the portfolio, or the modelling choices made within the research. Nevertheless, given the extent of the changes introduced by the SII 2020 Review, it is highly likely that both the Surplus and the Solvency Ratio will change when the same portfolio is held under the revised regulatory framework.

Overall, the results demonstrate that the SII 2020 Review affects not only the optimal asset allocation of life insurers, but also the relationship between allocation decisions and balance sheet performance. Although the exact magnitude of these effects depends on the composition of the portfolio and market conditions, the findings indicate that the revised VA introduces new opportunities and considerations that should be taken into account in strategic asset allocation decisions.

8 Discussion

The results of this research should be interpreted in light of the model used to represent the insurer's balance sheet being a proxy. A proxy model requires assumptions, modelling choices, and simplifications. The model therefore includes only those balance sheet components that are expected to respond to changes in bond allocation. This keeps the analysis focused on the mechanism under investigation, but it also means that not every relationship between balance sheet components is captured.

Each monthly optimisation is performed independently. No portfolio positions or investment results are carried over from one month to the next. This reduces path dependence and prevents biases from accumulating across consecutive optimisations. At the same time, it means that the model does not represent a dynamic investment strategy in which an insurer holds assets over time and benefits from coupon payments, maturities, or other carry effects. The independent starting point is therefore both a strength and a limitation.

A particularly notable omission is the risk margin. Within the model, the risk margin would not respond to the bond allocation and was therefore excluded under the scope rule. However, the SII review also changes the rules governing the risk margin. These changes reduce its level and its sensitivity to interest-rate movements (Commission Delegated Regulation (EU) 2026/269, 2026; Directive (EU) 2025/2, 2025). A lower RFR can increase the BEL and reduce Own Funds, while the revised risk margin can partly offset this effect. By excluding the risk margin, the model may therefore underestimate Own Funds and Surplus under the revised framework. This creates a possible negative bias in the estimated effect of the RFR and the other relevant SII changes.

There may also be an opposing bias. The model assumes that the insurer invests entirely in bonds, while real insurers generally hold a broader range of assets. Because the CSSR is linked to the spread sensitivity of fixed-income investments, the effect of bond allocation on the VA may appear stronger in a portfolio that consists entirely of fixed-income assets. The model may therefore overstate the importance of bond allocation for the overall balance sheet and the size of the VA effect. These two sources of bias work in opposite directions, but there is no reason to assume that they cancel each other out.

It is useful to make a distinction between the existence of a finding and the exact size of that finding. The possibility of steering the portfolio while keeping Surplus and the Solvency Ratio relatively stable is likely to follow from the revised regulatory mechanism. The exact amount of reallocation and the resulting change in Surplus are more dependent on the modelling choices. In Sub-question 2, the portfolio was adjusted using a greedy approach in which the weights were changed by an equal multiplier. Larger existing positions therefore received larger absolute changes, which may have increased concentration risk disproportionately. A different allocation method could produce a different change in the BEL and a different size of the steering effect. The model therefore provides stronger evidence for the existence of the mechanism than for the precise size of its effect.

The increase in Surplus is more difficult to attribute to one specific cause. Numerically, it appears in Sub-question 1 to be driven mainly by the difference between the change in asset values and the change in the BEL. This may reflect the effect of the portfolio allocation on the BEL through the revised VA. It may also result from higher VA values under the revised rules or from the new RFR used in the valuation of the assets. The omission of the risk margin adds another source of uncertainty to this interpretation. These effects cannot be separated completely in the current model and may interact with one another.

The experiments that varied risk appetite and the RFR produced a similar pattern in portfolio allocation. The level of Surplus, however, moved in the opposite direction across the two experiments. Surplus was higher when the Solvency-ratio constraint was set at a higher level, but it was also higher under the negative RFR shock than under the positive shock.

This pattern can largely be explained by the modelling choices. Risk appetite was represented by setting the Solvency-ratio constraint within a predefined corridor. A higher constraint therefore makes a higher Solvency Ratio more likely and, in the modelled results, is associated with a higher Surplus. The RFR shock affects the scale of the balance sheet. The negative shock produces a larger balance sheet than the positive shock, so the resulting Own Funds and Surplus are also larger in absolute terms. The higher Surplus under the negative RFR shock should therefore not be interpreted as evidence that this shock is intrinsically more favourable. It mainly reflects the balance sheet scaling created by the way the shock was modelled.

Further research The proxy model offers a controlled way to study the relevant SII changes, but several extensions could improve its realism. These extensions differ in both their expected benefit and their modelling complexity.

The simplest extension would be to increase the number of bonds in the current asset universe. This would improve the balance and availability of the data without requiring additional model components. Several experiments indicated that a shortage of suitable bonds contributed to infeasible optimisation problems. This was particularly visible in Sub-question 3, where different levels of risk appetite were represented through changes in the Solvency-ratio constraint. The more risk-averse strategy became infeasible during longer periods of the dataset, especially between 2017 and mid-2019, when interest rates were low or negative. These conditions increased the duration of the liabilities and created a greater demand for long-duration bonds. Because few such bonds were available, concentration risk increased and feasible solutions became harder to obtain.

A further extension would be to include additional asset classes. This would broaden the scope of the model, but it would also make the balance sheet more representative of an actual insurer. Different asset classes respond differently to market conditions and may provide alternatives to bonds when managing interest-rate or spread risk, which would make it possible to model hedging strategies. However, this would require additional assumptions and further SCR modules, such as Counterparty Default Risk. The resulting balance sheet calculations would become more complex and computationally more demanding.

Computational efficiency is therefore an important trade-off. The current model uses a sequence of relatively simple calculations, which makes it suitable for repeated optimisation experiments. Adding asset classes, hedging instruments, and additional SCR modules would increase the realism of the model, but could make large-scale optimisation much more difficult.

Another extension would be to replace the historical deterministic data with stochastic data. This would make it possible to test whether the findings remain robust across a wider range of market conditions. It would also create the basis for a dynamic portfolio model. In such a model, portfolio positions would be carried over time and the performance of purchased assets would be tracked through coupon payments, maturities, and changes in market value. A major challenge would be to generate data that is sufficiently realistic. If this could be achieved, the model could provide a stronger basis for analysing long-term investment strategies.

Company-specific application The model could also be applied using company-specific data. An insurer's own projected liability cash flows could be incorporated so that interest-rate movements on the liability side become directly relevant to its balance sheet. The insurer could also evaluate its historical investment strategy within the model. This would make it possible to assess whether changes in the allocation to particular asset classes could have improved its Surplus or Solvency Ratio under the relevant historical conditions. Such an application would move the model from a general research framework towards a potential decision-support tool.

Taken together, these limitations do not remove the relevance of the steering effect, but they limit the extent to which its exact size can be generalised. The model is best understood as a controlled framework for studying how the relevant SII changes alter the relationship between bond allocation, BEL, SCR, and balance sheet performance. Future extensions can test whether the same mechanism remains visible with a broader asset universe, a more complete risk-margin treatment, and a dynamic portfolio. The main result is therefore not a universal preference for one type of bond, but the need to evaluate portfolio allocation and regulatory balance sheet effects jointly.

Glossary

- **Basis point (bp):** one hundredth of a percentage point (0.01%).
- **Best Estimate Liability (BEL):** the expected present value of all future cash flows arising from an insurer's obligations to policyholders.
- **Bond:** a debt instrument in which an investor lends a principal amount to an issuer in exchange for periodic interest payments and repayment of the principal at maturity.
- **Capital surplus:** the amount of Own Funds remaining after deducting the Solvency Capital Requirement.
- **Carry:** the income or return earned from holding a position over time.
- **Concentration risk:** the risk arising from excessive exposure to a single issuer or group of related exposures.
- **Cost-of-capital rate:** the rate used to calculate the cost of holding capital against non-hedgeable risks.
- **Coupon:** the periodic interest payment a bond makes to its holder.
- **Credit Quality Step (CQS):** a regulatory scale grouping credit ratings into risk categories.
- **Credit Spread Sensitivity Ratio (CSSR):** a ratio introduced in the revised Volatility Adjustment that relates the spread sensitivity of an insurer's assets to that of its liabilities.
- **Currency risk:** the risk of losses resulting from changes in foreign-exchange rates.
- **Day-count convention:** the method used to determine the fraction of a year between two dates for interest calculations.
- **Discount rate:** the rate used to convert future cash flows into their present value.
- **Duration:** a measure of the sensitivity of the value of a financial instrument to changes in interest rates.
- **First Smoothing Point (FSP):** the point from which the revised Solvency II risk-free curve begins to converge towards the Ultimate Forward Rate.
- **General Application Ratio (GAR):** the regulatory factor determining what proportion of the risk-corrected spread enters the Volatility Adjustment.
- **Hedging:** the use of financial instruments to reduce exposure to a specific risk.
- **Interest-rate risk:** the risk of changes in the value of assets or liabilities resulting from movements in interest rates.
- **Investment grade:** a rating category indicating relatively low credit risk.

- **Last Liquid Point (LLP):** the maturity up to which the risk-free yield curve is based directly on sufficiently reliable market information under the current methodology.
- **Market value (MV):** the price at which an asset can be valued in the market at a given point in time.
- **Maturity:** the date on which a bond reaches the end of its contractual term and its principal is repaid.
- **Notional:** the principal amount of a bond that the issuer is contractually required to repay at maturity.
- **Own Funds (OF):** the capital available to an insurer to absorb losses under Solvency II.
- **Percentage point (p.p.):** a unit for the arithmetic difference between two percentages.
- **Present value (PV):** the current value of future cash flows after discounting them at an appropriate discount rate.
- **Present Value of a Basis Point (PVBP):** the change in the present value of an asset or liability resulting from a one-basis-point change in the relevant discount rate or spread.
- **Rating:** a credit-agency assessment of an issuer's creditworthiness and likelihood of default.
- **Risk-Corrected Spread (RCS):** the market spread after deducting the portion attributed to expected default and downgrade risk.
- **Risk-free rate (RFR):** the interest-rate term structure used under Solvency II as the basis for discounting liability cash flows.
- **Risk margin:** the additional component of technical provisions representing the cost a third party would require to take over and support the insurer's non-hedgeable obligations.
- **Solvency Capital Requirement (SCR):** the capital an insurer is required to hold to absorb unexpected losses over a one-year horizon at the prescribed Solvency II confidence level.
- **Solvency Ratio (SR):** the ratio of an insurer's eligible Own Funds to its Solvency Capital Requirement.
- **Spread:** the additional yield of a bond relative to a reference risk-free rate or curve.
- **Spread risk:** the risk of losses resulting from changes in credit spreads.
- **Technical Provisions:** the sum of the Best Estimate Liability and the risk margin.
- **Time to maturity (TTM):** the remaining time until a bond reaches its maturity date.
- **TRBC:** a classification system used to categorise companies by industry sector.

- **Ultimate Forward Rate (UFR):** the long-term rate towards which the Solvency II risk-free curve converges.
- **Volatility Adjustment (VA):** an adjustment added to the risk-free rate used to discount the BEL, intended to reduce the effect of short-term spread movements on insurers' balance sheets.
- **Yield:** the return implied by a bond's price and contractual cash flows.
- **Yield to Maturity (YTM):** the annualised return implied by holding a bond until maturity, assuming its contractual payments are made as scheduled.
- **Z-spread:** the constant spread added to the risk-free spot-rate curve such that the discounted value of a bond's contractual cash flows equals its market price.

References

Alfonsi, A., Cherchali, A., & Infante Acevedo, J. A. (2020). A synthetic model for asset-liability management in life insurance, and analysis of the SCR with the standard formula. *European Actuarial Journal*, *10*(2), 457–498. https://doi.org/10.1007/s13385-020-00240-3

Barucci, E., Marazzina, D., & Rroji, E. (2025). An investigation of the volatility adjustment. *Decisions in Economics and Finance*, *48*(2), 1337–1367. https://doi.org/10.1007/s10203-023-00416-y

Bloomberg Index Services Limited. (2026, January 8). *Bloomberg fixed income index methodology*. Retrieved June 17, 2026, from https://assets.bbhub.io/professional/sites/10/Bloomberg-Index-Publications-Fixed-Income-Index-Methodology.pdf

Boonen, T. J. (2017). Solvency II solvency capital requirement for life insurance companies based on expected shortfall. *European Actuarial Journal*, *7*(2), 405–434. https://doi.org/10.1007/s13385-017-0160-4

Boyd, S., & Vandenberghe, L. (2004). *Convex optimization*. Cambridge University Press.

Commission Delegated Regulation (EU) 2015/35 of 10 October 2014 supplementing Directive 2009/138/EC of the European Parliament and of the Council on the taking-up and pursuit of the business of Insurance and Reinsurance (Solvency II) (2025). https://eur-lex.europa.eu/eli/reg_del/2015/35/oj

Commission Delegated Regulation (EU) 2026/269 of 29 October 2025 amending Delegated Regulation (EU) 2015/35 as regards technical provisions, long-term guarantee measures, own funds, equity risk, spread risk on securitisation positions, other standard formula capital requirements, reporting and disclosure, proportionality and group solvency (2026). Retrieved June 17, 2026, from http://data.europa.eu/eli/reg_del/2026/269/oj

Curry, B. (2019). The stochastic full balance sheet model. *British Actuarial Journal*, *24*, e21. https://doi.org/10.1017/S1357321719000126

Directive (EU) 2025/2 of the European Parliament and of the Council of 27 November 2024 amending Directive 2009/138/EC as regards proportionality, quality of supervision, reporting, long-term guarantee measures, macro-prudential tools, sustainability risks, and group and cross-border supervision (2025, January 8). https://eur-lex.europa.eu/eli/dir/2025/2/oj

Directive 2009/138/EC of the European Parliament and of the Council of 25 November 2009 on the taking-up and pursuit of the business of Insurance and Reinsurance (Solvency II) (2009, November 25). https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:02009L0138-20190113

Drenovak, M., Ranković, V., Urošević, B., & Jelic, R. (2021). Bond portfolio management under solvency II regulation. *The European Journal of Finance*, *27*(9), 857–879. https://doi.org/10.1080/1351847X.2020.1850499

EIOPA. (2020). *Opinion on the 2020 review of solvency II*. Retrieved January 6, 2026, from https://www.eiopa.europa.eu/publications/opinion-2020-review-solvency-ii_en

European Central Bank. (2026). *Euro foreign exchange reference rates* [Daily reference rates, series EXR.D.<currency>.EUR.SP00.A; coverage January 2016–December 2025]. ECB Data Portal, dataset EXR. Retrieved June 17, 2026, from https://www.ecb.europa.eu/

stats/policy_and_exchange_rates/euro_reference_exchange_rates/html/index.en.html
European Insurance and Occupational Pensions Authority. (2026). *Risk-free interest rate term structures* [Monthly technical information on the relevant risk-free interest rate term structures, published under the Solvency II Directive]. Retrieved June 17, 2026, from https://www.eiopa.europa.eu/tools-and-data/risk-free-interest-rate-term-structures_en

Ha, D. (2017, October 29). *A visual guide to evolution strategies*. Retrieved June 9, 2026, from https://blog.otoro.net/2017/10/29/visual-evolution-strategies/

Hansen, N., & Ostermeier, A. (2001). Completely derandomized self-adaptation in evolution strategies. *Evolutionary Computation*, *9*(2), 159–195.

Höring, D. (2013). Will solvency II market risk requirements bite? the impact of solvency II on insurers' asset allocation. *The Geneva Papers on Risk and Insurance - Issues and Practice*, *38*(2), 250–273. https://doi.org/10.1057/gpp.2012.31

Kouwenberg, R. (2017, January 16). Strategic asset allocation and risk budgeting for insurers under solvency II. https://doi.org/10.2139/ssrn.2894809

Kraft, D. (1988). *A software package for sequential quadratic programming* (tech. rep. No. DFVLR-FB 88-28). DLR German Aerospace Center.

Nocedal, J., & Wright, S. J. (2006). *Numerical optimization* (2nd ed.). Springer. https://doi.org/10.1007/978-0-387-40065-5

PwC. (2025, January 17). *Uitkomsten van de Solvency II 2020-review*. Retrieved June 17, 2026, from https://www.pwc.nl/nl/dienstverlening/audit-assurance/pwc-accountancy-insights/verslaggeving/uitkomsten-van-de-solvency-ii-2020-review.html

Refinitiv. (2026). Fixed income pricing and reference data [Weekly bond prices, yields and static attributes for 513 corporate and government bonds; Refinitiv is part of the London Stock Exchange Group (LSEG)].

Storn, R., & Price, K. (1997). Differential evolution - a simple and efficient heuristic for global optimization over continuous spaces. *Journal of Global Optimization*, *11*(4), 341–359.

Thorendal, M. (2023). Navigating the volatility adjustment in solvency II.

van Bragt, D., Steehouwer, H., & Waalwijk, B. (2010). Market consistent ALM for life insurers—steps toward solvency II [Number: EI-1536]. *The Geneva Papers on Risk and Insurance - Issues and Practice*, *35*(1), 92–109. https://doi.org/10.1057/gpp.2009.34

A Appendix A: Correlation Matrices

Table 9: Market-risk correlation matrix under the standard formula. Order: interest rate, equity, property, spread, currency, concentration

	IR	Equity	Property	Spread	Currency	Conc.
IR	1.00	0.50	0.50	0.25	0.00	0.25
Equity	0.50	1.00	0.75	0.75	0.00	0.25
Property	0.50	0.75	1.00	0.50	0.00	0.25
Spread	0.25	0.75	0.50	1.00	0.00	0.25
Currency	0.00	0.00	0.00	0.00	1.00	0.00
Conc.	0.25	0.25	0.25	0.25	0.00	1.00

Table 10: Top-level SCR correlation matrix under the standard formula. Order: market, counterparty default (CDR), life, health, non-life

	Market	CDR	Life	Health	Non-life
Market	1.00	0.25	0.25	0.25	0.25
CDR	0.25	1.00	0.25	0.25	0.50
Life	0.25	0.25	1.00	0.25	0.00
Health	0.25	0.25	0.25	1.00	0.00
Non-life	0.25	0.50	0.00	0.00	1.00

B Appendix B: Additional Data Tables

Table 11: Relative upward stress on the risk-free term structure under the current regime (Art. 166, Delegated Regulation (EU) 2015/35).

Maturity (years)	Upward stress (%)
1	70
2	70
3	64
4	59
5	55
6	52
7	49
8	47
9	44
10	42
11	39
12	37
13	35
14	34
15	33
16	31
17	30
18	29
19	27
20	26
90	20

Table 12: Relative downward stress on the risk-free term structure under the current regime (Art. 167, Delegated Regulation (EU) 2015/35).

Maturity (years)	Downward stress (%)
1	75
2	65
3	56
4	50
5	46
6	42
7	39
8	36
9	33
10	31
11	30
12	29
13	28
14	28
15	27
16	28
17	28
18	28
19	29
20	29
90	20

Table 13: Spread-risk stress parameters a_i and b_i by credit-quality step and modified duration dur_i (Art. 176, Delegated Regulation (EU) 2015/35). The stress is $a_i + b_i\,(\text{dur}_i - \text{offset})$, floored/capped as indicated. CQS 5 and 6 share the final column.

Duration dur_i	stress_i	CQS 0		CQS 1		CQS 2		CQS 3		CQS 4		CQS 5 & 6	
		a_i	b_i	a_i	b_i	a_i	b_i	a_i	b_i	a_i	b_i	a_i	b_i
up to 5	$b_i \cdot \text{dur}_i$	–	0.9	–	1.1	–	1.4	–	2.5	–	4.5	–	7.5
5–10	$a_i + b_i(\text{dur}_i - 5)$	4.5	0.5	5.5	0.6	7.0	0.7	12.5	1.5	22.5	2.5	37.5	4.2
10–15	$a_i + b_i(\text{dur}_i - 10)$	7.0	0.5	8.5	0.5	10.5	0.5	20.0	1.0	35.0	1.8	58.5	0.5
15–20	$a_i + b_i(\text{dur}_i - 15)$	9.5	0.5	11.0	0.5	13.0	0.5	25.0	1.0	44.0	0.5	61.0	0.5
over 20	$\min[a_i + b_i(\text{dur}_i - 20);\, 1]$	12.0	0.5	13.5	0.5	15.5	0.5	30.0	0.5	46.6	0.5	63.5	0.5

Table 14: Concentration-risk parameters by weighted-average credit-quality step: relative excess-exposure threshold CT_i (Art. 185) and market risk-concentration factor g_i (Art. 186), Delegated Regulation (EU) 2015/35.

Credit-quality step	0	1	2	3	4	5	6
Excess threshold CT_i (%)	3	3	3	1.5	1.5	1.5	1.5
Risk factor g_i (%)	12	12	21	27	73	73	73

Table 15: Revised upward interest-rate stress under Delegated Regulation (EU) 2026/269. For each maturity m the absolute component $a_m^{\uparrow}$ (pp) and the relative component $b_m^{\uparrow}$ (%) enter $r^{\uparrow}(t) = r(t)\,[1 + b_m^{\uparrow}] + a_m^{\uparrow}$.

m	$a_m^{\uparrow}$ (%)	$b_m^{\uparrow}$ (%)	m	$a_m^{\uparrow}$ (%)	$b_m^{\uparrow}$ (%)
1.00	61.00	2.14	26.00	22.00	0.76
2.00	53.00	1.86	27.00	21.00	0.74
3.00	49.00	1.72	28.00	21.00	0.72
4.00	46.00	1.61	29.00	20.00	0.70
5.00	45.00	1.58	30.00	20.00	0.69
6.00	41.00	1.44	31.00	20.00	0.70
7.00	37.00	1.30	32.00	20.00	0.71
8.00	34.00	1.19	33.00	20.00	0.71
9.00	32.00	1.12	34.00	20.00	0.71
10.00	30.00	1.05	35.00	20.00	0.71
11.00	30.00	1.05	36.00	20.00	0.72
12.00	30.00	1.05	37.00	21.00	0.72
13.00	30.00	1.05	38.00	21.00	0.72
14.00	29.00	1.02	39.00	21.00	0.73
15.00	28.00	0.98	40.00	21.00	0.73
16.00	28.00	0.98	41.00	21.00	0.74
17.00	27.00	0.95	42.00	21.00	0.74
18.00	26.00	0.91	43.00	21.00	0.75
19.00	26.00	0.91	44.00	21.00	0.75
20.00	25.00	0.88	45.00	21.00	0.75
21.00	25.00	0.87	46.00	21.00	0.75
22.00	24.00	0.85	47.00	21.00	0.75
23.00	24.00	0.82	48.00	21.00	0.74
24.00	23.00	0.80	49.00	21.00	0.74
25.00	22.00	0.78	50.00	21.00	0.73

Table 16: Revised downward interest-rate stress under Delegated Regulation (EU) 2026/269. For each maturity m the absolute component $a_m^{\downarrow}$ (pp) and the relative component $b_m^{\downarrow}$ (%) enter $r^{\downarrow}(t) = r(t)\,[1 - b_m^{\downarrow}] - a_m^{\downarrow}$.

m	$a_m^{\downarrow}$ (%)	$b_m^{\downarrow}$ (%)	m	$a_m^{\downarrow}$ (%)	$b_m^{\downarrow}$ (%)
1.00	58.00	1.16	26.00	52.00	0.46
2.00	51.00	0.99	27.00	53.00	0.45
3.00	44.00	0.83	28.00	53.00	0.44
4.00	40.00	0.74	29.00	53.00	0.42
5.00	40.00	0.71	30.00	53.00	0.41
6.00	38.00	0.67	31.00	53.00	0.40
7.00	37.00	0.63	32.00	53.00	0.39
8.00	38.00	0.62	33.00	54.00	0.37
9.00	39.00	0.61	34.00	54.00	0.36
10.00	40.00	0.61	35.00	54.00	0.35
11.00	41.00	0.60	36.00	54.00	0.34
12.00	42.00	0.60	37.00	55.00	0.33
13.00	43.00	0.59	38.00	55.00	0.32
14.00	44.00	0.58	39.00	56.00	0.31
15.00	45.00	0.57	40.00	57.00	0.30
16.00	47.00	0.56	41.00	57.00	0.29
17.00	48.00	0.55	42.00	58.00	0.28
18.00	49.00	0.54	43.00	59.00	0.27
19.00	49.00	0.52	44.00	61.00	0.26
20.00	50.00	0.50	45.00	62.00	0.25
21.00	49.00	0.49	46.00	62.00	0.23
22.00	50.00	0.49	47.00	63.00	0.22
23.00	51.00	0.48	48.00	64.00	0.21
24.00	51.00	0.48	49.00	64.00	0.19
25.00	52.00	0.47	50.00	65.00	0.18

Table 17: SFCR source data used for the market comparison: Dutch solo life entities, 2022 - 2025. Amounts in EUR billion; BEL is the Life + Unit-Linked best estimate. The last three columns are derived ratios. Source: published SFCR reports.

Insurer	Year	Assets	BEL	RM	SCR	OF	SII	BEL/A	SCR/A	OF/A
ASR Leven	2025	53.5	39.3	0.71	2.15	4.97	2.31	73.4%	4.0%	9.3%
	2024	52.5	38.7	1.33	2.44	4.31	1.77	73.7%	4.6%	8.2%
	2023	51.5	38.7	1.33	2.24	3.95	1.76	75.1%	4.4%	7.7%
	2022	50.3	36.1	1.23	2.38	4.01	1.69	71.8%	4.7%	8.0%
Aegon Leven	2025	69.1	51.3	0.74	2.28	4.62	2.02	74.2%	3.3%	6.7%
	2024	66.6	52.4	1.47	2.31	4.48	1.94	78.6%	3.5%	6.7%
	2023	66.1	51.7	1.49	2.36	4.02	1.70	78.2%	3.6%	6.1%
	2022	60.1	49.9	1.43	2.21	4.63	2.10	83.0%	3.7%	7.7%
NN Leven	2025	126.8	104.2	2.28	4.84	10.77	2.23	82.2%	3.8%	8.5%
	2024	132.1	92.9	2.88	4.87	9.13	1.87	70.3%	3.7%	6.9%
	2023	130.2	87.1	2.92	5.20	10.21	1.96	66.9%	4.0%	7.8%
	2022	124.5	107.9	3.73	5.36	10.23	1.91	86.7%	4.3%	8.2%
Achmea Leven	2025	40.9	30.1	1.09	1.89	3.55	1.87	73.5%	4.6%	8.7%
	2024	42.1	32.1	1.25	1.79	3.13	1.75	76.2%	4.2%	7.4%
	2023	42.9	32.3	1.23	1.48	2.61	1.77	75.2%	3.4%	6.1%
	2022	42.1	31.4	0.95	1.51	3.10	2.04	74.5%	3.6%	7.4%
Benchmark (n=16)	mean						1.92	75.9%	4.0%	7.6%
	min						1.69	66.9%	3.3%	6.1%
	max						2.31	86.7%	4.7%	9.3%

Table 18: The fifty largest portfolio weights on each reference date (percentage of portfolio value), ranked in descending order.

Rank	2016-12-16	2020-04-17	2024-06-14	Rank	2016-12-16	2020-04-17	2024-06-14
1	2.148	1.573	0.743	26	1.622	1.342	0.720
2	2.144	1.573	0.742	27	1.614	1.333	0.719
3	2.142	1.572	0.742	28	1.600	1.320	0.719
4	2.141	1.572	0.742	29	1.587	1.319	0.717
5	2.137	1.572	0.742	30	1.575	1.296	0.711
6	2.137	1.572	0.742	31	1.540	1.278	0.710
7	2.136	1.572	0.742	32	1.500	1.239	0.710
8	2.134	1.572	0.741	33	1.490	1.237	0.709
9	2.112	1.571	0.741	34	1.465	1.183	0.706
10	2.111	1.570	0.741	35	1.378	1.171	0.702
11	2.108	1.569	0.740	36	1.374	1.129	0.696
12	2.105	1.567	0.740	37	1.292	1.084	0.693
13	2.101	1.566	0.739	38	1.250	1.060	0.692
14	2.064	1.560	0.739	39	1.250	1.011	0.689
15	2.052	1.554	0.736	40	1.234	1.010	0.688
16	2.002	1.554	0.733	41	1.206	0.995	0.687
17	1.976	1.552	0.733	42	1.194	0.992	0.678
18	1.960	1.539	0.731	43	1.152	0.981	0.671
19	1.927	1.536	0.730	44	1.060	0.977	0.666
20	1.916	1.499	0.729	45	1.007	0.967	0.665
21	1.843	1.464	0.729	46	0.987	0.953	0.664
22	1.808	1.443	0.728	47	0.919	0.944	0.662
23	1.737	1.380	0.727	48	0.919	0.931	0.660
24	1.713	1.358	0.724	49	0.865	0.909	0.659
25	1.675	1.354	0.723	50	0.806	0.870	0.658