Aalto University

School of Science

MS-E2177 Seminar on Case Studies in Operations Research

Project Report

Modelling the Risks of Collateralized Debt Obligations

Terho Mutikainen
Aaro Niinistö
Tuukka Pallonen (Project Manager)
Panu Salonen

May 25, 2026

Contents

1 Introduction

Ilmarinen is a Finnish pension insurance company whose investment portfolio contains a wide range of asset classes, including fixed income instruments and more complex structured-credit products. Because pension assets must be invested profitably but prudently, risk measurement is a central part of the investment process. Finnish pension insurers are governed by a solvency framework that requires risk values and expected returns to be calculated and monitored across different asset classes [5, 14]. However, the standard regulatory framework is relatively simple compared with the risk structure of collateralized debt obligations (CDOs) and collateralized loan obligations (CLOs), where losses depend not only on the credit quality of the underlying assets but also on tranche seniority, maturity, default dependence and recovery assumptions.

The purpose of this project is to develop a simplified but scientifically defensible model for assessing the risk and return of CDO/CLO tranches for Ilmarinen. The client-requested outputs are a one-year 97% Value-at-Risk loss assumption and an expected return estimate. These outputs are intended to be analogous to the risk and return inputs used in solvency-style calculations, although the model is not intended to calculate Ilmarinen's full regulatory capital requirement directly. Instead, the objective is to produce CDO-specific risk and return estimates that can be used as inputs to a broader internal risk-management framework.

The main modelling challenge is that CDO tranche risk differs substantially from ordinary corporate bond risk. In a CDO structure, the underlying portfolio losses are allocated through a waterfall defined by attachment and detachment points. Junior tranches absorb losses first, while mezzanine and senior tranches are protected by the subordination below them. This creates nonlinear risk behavior, where a moderate change in the underlying portfolio loss distribution can have a large effect on a specific tranche. Previous studies have emphasized that CDO risk measures are sensitive to credit fundamentals, ratings, correlation assumptions and modeling choices [16, 7].

Our baseline model is a one-factor Gaussian copula model, following the classical default-correlation approach introduced by David X. Li[13]. The model links obligor defaults through a common systematic factor while preserving interpretability. This makes it suitable for a project whose goal is not to build a full trading-desk pricing system, but a transparent risk model that can be explained, audited and eventually approximated by a simple proxy tool. Factor-copula models are also widely used as benchmark models in the CDO literature because they provide a tractable way to connect marginal default probabilities, default dependence and tranche-level loss allocation [12, 4].

The model developed extends the simplest large homogeneous pool setting in three important ways. First, the underlying portfolio is treated as a finite pool, allowing the model to reflect different portfolio sizes, rating buckets and exposure assumptions. Second, recovery rates are allowed to be stochastic, so that default losses are not forced to be identical across all scenarios. This is motivated by extensions of the Gaussian copula framework that incorporate random recovery and random factor loadings [1]. Third, the model includes conditional repricing after defaults. This means that after losses have occurred by the one-year observation horizon, the remaining tranche cash flows are revalued using the updated state of the pool. This feature is essential because a mezzanine or senior tranche can lose market value even before it suffers realized principal loss, if junior protection below it has been eroded.

The practical implementation consists of a Python-based model engine and an Excel-based proxy interface. Python performs the simulation of defaults and recoveries, tranche loss allocation, conditional repricing and output generation. Excel is used as a client-facing interface, where the user enters rating, maturity, attachment and detachment, and the workbook retrieves or interpolates outputs from a pre-computed scenario table. This structure keeps the computationally demanding model in Python while producing a tool that can be used in a form closer to the simple tables and rule sets used in regulatory risk frameworks.

The report is structured as follows. Section 2 summarizes the solvency-law motivation and explains why a CDO-specific model is needed. Sections 3 and 4 introduce CDO risk drivers, including tranche structure, credit quality, duration and spread risk. Section 5 presents the mathematical model, including the Gaussian copula framework, conditional repricing, finite heterogeneous pools and stochastic recovery. Section 6 describes the implementation architecture, while Section 7 records the main parameter choices and calibration assumptions. Section 8 explains the Excel proxy deliverable. Sections 9 and 10 describe validation and numerical results. Finally, Sections 11 and 12 discuss limitations, future work and conclusions.

2 Solvency legislation

Due to the nature of managed funds, the investing strategies are heavily regulated by Finnish legislation, more precisely Finnish Solvency Law (315/2015). This legislation sets various guidelines and limits for the investment portfolio regarding diversification, risk management and liquidity. The goal of the legislation is to ensure that pension insurance funds are able to pay out pensions even in unfavorable market conditions as well as prevent excess risk taking by requiring certain diversification within the portfolio. The ability to pay out pensions is ensured by solvency limit, which is a requirement to keep a certain amount liquid funds based on the corresponding investments risk and expected profit.

2.1 Diversification

Before examining the risk management framework and different risk classes presented by the Solvency Law, let us look into the required diversification of investments - which slightly differ between investment types while we focus on the bonds and other fixed income investments. Chapter 2 of Solvency Law (315/2015) states that maximum 5 percent of managed funds can be invested in the shares, bonds and other instruments issued by the same entity. It is important to note that the calculation is only applied to the amount of debt surpassing the collateral value. On the other hand, subsection 7 gives more specific requirements to bonds depending on their credit rating:

- **Credit rating 1:** the maximum share of investment portfolio is 10 percent
- **Credit rating 2:** the maximum share of investment portfolio is 5 percent

Subsection 8 states that if the limit is surpassed, 15 percent of the surpassing amount is added to the solvency limit. On the other hand, if the limit is surpassed by more than 15 percent, the surpassing amount is added to the solvency limit at 100 percent.

2.2 Risk management

Next, let us look into how legislation governs risk management and how the actual solvency limit is determined for certain investment. The Finnish Solvency Law (315/2015) recognizes 18 different risk classes, from which the most relevant to CDOs are:

- Interest rate risk.
- Credit spread risk in government bonds and debt instruments - credit rating 1.
- Credit spread risk in secured premium loans and investment loans, and in non-government bonds/debt instruments rated at credit class 1.
- Credit spread risk in unsecured premium loans and investment loans, and in government or non-government bonds/debt instruments rated at credit class 2.
- Credit spread risk in subordinated loans under the Limited Liability Companies Act (624/2006) Chapter 12, and in government or non-government bonds/debt instruments rated at credit class 3.

The division of credit spread risks based fixed income investment type provides a strong foundation for determining the quality of underlying loan mass in our risk model. Furthermore, we can now examine how the risk value, expected profit and the solvency limit (for a single investment or to pension insurance company as a whole) is calculated. For each risk class j, the risk value is determined by

$$V_j = \sum_i A_i \min[(1 + \tau L_i)S_j; 1] \tag{1}$$

where A_i is the amount of investment i exposed to the risk in risk class j, L_i is the debt portion included in that investment, and S_j is the loss assumption defined for risk class j. The loss assumption S_j and the constant τ representing risk related to debt are set by the authorities. In addition, the expected profit is determined by

$$\mu_j = \sum_i (m_j + L_i(m_j - m_6))A_i, \tag{2}$$

where A_i is the amount of investment i exposed to the risk in risk class j, L_i is the debt portion included in that investment, m_j is the expected profit in risk class j and m_6 is the expected profit in interest rate risk class. The risk class -specific expected profit m_j is set by the authorities.

As we examine risk within the classes related to interest rate risk and credit spread risk, subsection 14 of Finnish Solvency Law (315/2015) provides additional requirements for calculating the risk value and expected profit in these risk classes:

- When calculating the expected profit for interest rate risk class, m_j is to be replaced by $m_j D_j^{\gamma}$ where D_j is the duration of the investment facing interest rate risk and γ is a constant representing yield curve shape, set by the authorities.
- When calculating the risk value of interest rate risk, S_j is to be replaced by $S_j D_j - \mu_i$ where D_i is the duration of the investment and μ_i is its expected profit when A_i is set to 1.
- When calculating the risk value of credit spread risk, S_j is to be replaced by $S_j D_j - \mu_i$ where D_i is the duration of the investment and μ_i is its expected profit when A_i is set to 1.

Based on these additional requirements, the interest rate risk classes expected profit and risk value are now:

$$V_j = \sum_i A_i \min[(1 + \tau L_i)(S_j D_j - \sum_i (m_j D_j^{\gamma} + L_i(m_j D_j^{\gamma} - m_6))); 1] \tag{3}$$

$$\mu_j = \sum_i (m_j D_j^{\gamma} + L_i(m_j D_j^{\gamma} - m_6))A_i. \tag{4}$$

For credit spread risk, the corresponding equations are:

$$V_j = \sum_i A_i \min[(1 + \tau L_i)(S_j D_j - \sum_i (m_j + L_i(m_j - m_6))); 1] \tag{5}$$

$$\mu_j = \sum_i (m_j + L_i(m_j - m_6))A_i. \tag{6}$$

Now that we have established the risk class specific calculations of risk value and expected profit, subsection 23 of Finnish Solvency Law (315/2015) states how to combine these into a solvency limit. The amount of required solvency capital is set by

$$V_{total} = -\sum_j \mu_j + \sqrt{\sum_i \sum_j \rho_{ij}(V_i + \mu_i) + (V_j + \mu_j) + \sum_j \beta_j^2 B_j^2} + \sum_k K_k, \tag{7}$$

where V_j is the risk value of risk class j, μ_j is the expected return of risk class j, ρ_{ij} is the correlation between risk classes i and j, B_j is the smaller of the sum of long positions and the sum of short positions in risk class j, and K_k is the amount arising from the excess exposure to counterparty k as provided in subsection 8. However, the solvency limit is always at least 5 percent of investments. Also, B_j and ρ_{ij} are constants set by the authorities.

3 Modelling CDO risk

In evaluating the risk profile of Collateralized Debt Obligations (CDOs), drawing direct parallels to standard corporate credit ratings can be highly misleading. Tarashev and Zhu [16] highlight that the non-linear nature of the CDO waterfall structure inherently amplifies risk; a single-notch downgrade in the underlying collateral pool can trigger multi-notch downgrades in the derivative tranches. Furthermore, tranche ratings and Value-at-Risk (VaR) metrics are acutely sensitive to underlying correlation assumptions and sudden probability of default (PD) shocks. To calculate expected loss and default probabilities, the authors model defaults using strict correlation matrices, but importantly, they assume the Loss Given Default (LGD) to be an independent variable across the portfolio, isolating the dependency purely to the default event itself.

Fender and Mitchell [7] further explore the mechanics and limitations of various risk models, notably the Binomial Expansion Technique (BET) traditionally favored by rating agencies. The BET simplifies a complex, correlated portfolio into a smaller "diversity score" (DS) of entirely independent, homogeneous assets. This mathematical simplification yields $DS + 1$ potential default scenarios, allowing default probabilities to be modeled using a standard binomial distribution. While computationally elegant, the BET fundamentally underestimates the fat-tailed nature of real-world default correlations, which is a flaw that rating agencies historically patched by artificially inflating the underlying default probabilities. The authors contrast this aggregate approach with structural credit risk models, such as the Merton model, which link default directly to the financial health of the obligor. Ultimately, Fender and Mitchell emphasize the concept of model risk, noting that the chosen mathematical methodology fundamentally dictates the meaning, severity, and reliability of the resulting risk metrics.

To capture the complex dependency structures missed by the BET, modern risk management relies heavily on copula functions. Hwang and Kim [11] evaluate various copula structures for calculating the VaR of multivariate portfolios, including Archimedean and Hierarchical models. While Hierarchical copulas offer high flexibility by allowing different correlation structures for different subgroups, they introduce a significant number of parameters to estimate. For simpler, non-synthetic CDOs lacking highly distinct, sector-based asset classes, this can lead to severe overfitting.

4 Variables affecting risk of CDO

Ultimately the risk of a CDO will depend greatly on three variables. One is the individual risks of the underlying assets which is dependent on the credit risk rating of these assets. Then second variable will be the duration of the underlying bonds. Third, the specific CDO tranche affects the risk of the investment into the CDO.

4.1 Credit risk ratings and issuer ratings

To derive the risk of complete CDO we will need understanding of the underlying assets. Two variables will be needed: 1. the annual (marginal) probability of default (PD) and 2. loss given default (LGD). The probability of default will tell us the probability of default per year given that the entity has not yet defaulted. Then LGD gives us the percentage of the loan that is lost in the case of default. Together these two variables define the risks associated with default of the underlying individual assets. These variables are strongly associated with credit ratings.

4.2 CDO tranche structure and loss allocation

A CDO redistributes the credit risk of an underlying portfolio into several layers, or tranches, that differ in seniority and therefore in risk. In the standard structure, the most junior tranche is the equity tranche, followed by one or more mezzanine tranches, while the senior tranche is protected by all lower layers. This means that the same portfolio loss has different consequences for different investors depending on the tranche they hold. The equity tranche absorbs losses first, the mezzanine tranche principal absorbs losses only after the equity tranche has been exhausted, and the senior tranche is affected only after both junior layers have been substantially eroded [17].

The structure of a single tranche is formalized through attachment and detachment points. If a tranche j is defined by lower boundary K_L^j and upper boundary K_U^j, then it is exposed only to that part of cumulative portfolio loss which falls between these two thresholds. In other words, losses below K_L^j do not affect the tranche, while losses above K_U^j have already fully written it down. This threshold structure is the key reason why tranche risk is highly nonlinear and thus small changes in the portfolio loss distribution can produce much larger changes in mezzanine or senior tranche risk than in the underlying collateral pool itself [4].

Let L_t denote cumulative portfolio loss at time t. Then the cumulative loss allocated to tranche j can be written as

$$L_{j,t} = \left(L_t - K_L^j\right)^+ - \left(L_t - K_U^j\right)^+,$$

where $(x)^+ = \max(x, 0)$ for more concise notation. Thus, the tranche loss of j is zero when portfolio loss is below attachment, increases linearly while portfolio loss lies between attachment and detachment, and is capped once the tranche is exhausted. This clipped-loss structure is a standard representation in the CDO literature and makes clear that tranche cash flows depend on the realized path of cumulative portfolio losses rather than only on the number of defaults [4].

The same waterfall logic determines the cash flows of the tranche. The protection seller compensates the protection buyer for losses hitting the tranche, while the protection buyer pays a periodic premium on the outstanding tranche notional, here $\Gamma_{j,t}$, denotes the outstanding tranche notional of tranche j at time t. If tranche j has attachment point K_L^j, detachment point K_U^j, original portfolio notional M, and cumulative tranche loss $L_{j,t}$, then

$$\Gamma_{j,t} = (K_U^j - K_L^j)M - L_{j,t}.$$

Each default that affects the tranche reduces the principal on which future spread payments are made. If portfolio losses exceed the detachment point, no tranche notional remains and premium payments stop. Due to this, tranche valuation depends simultaneously on the expected loss path and on the remaining notional available to receive future premium cash flows [9, 4].

This structure also explains the implicit leverage of CDO tranches. We will also see that mezzanine and senior tranches are not only sensitive to realized defaults, but also to the erosion of subordination below them. This observation is central for our project, because it motivates the later extension from a static loss engine toward conditional repricing after defaults have already occurred.

4.3 Duration and spread risk

Duration affects the risk of a CDO tranche in two ways. The first is the regulatory way, because in the Finnish solvency framework, duration directly scales both interest-rate risk and credit-spread risk, so that otherwise similar fixed-income positions generate different risk values depending on their duration. In our current framework this is already visible from the solvency formulas, where for interest-rate and credit-spread risk classes the risk term is adjusted by duration D_i. Consequently, duration must be understood as one of the main determinants of the investment's regulatory risk contribution [5].

The second way is the economic valuation of the CDO tranche. A tranche is valued from its expected future cash flows, which consist of a premium leg and a protection leg. In standard tranche-pricing formulas, the premium leg is paid on the outstanding tranche notional, while the protection leg reflects expected future losses allocated to the tranche. Therefore, the value of a tranche depends not only on whether defaults occur, but also on how long the tranche continues to receive spread payments before maturity. Thus, maturity and duration determine how much present value is exposed to possible unknown future changes in credit conditions [4].

A convenient way to summarize this effect is through the notion of credit-spread duration. If V_j denotes the value of tranche j and s_j its market spread, then a first-order approximation of mark-to-market sensitivity is

$$D_j^{\mathrm{cs}} = -\frac{1}{V_j}\frac{\partial V_j}{\partial s_j}, \qquad \Delta V_j \approx -D_j^{\mathrm{cs}} V_j \Delta s_j.$$

Thus, for a given spread widening $\Delta s_j > 0$, a tranche with larger credit-spread duration experiences a larger negative mark-to-market move. This means that two tranches with identical current expected loss may still have different risk profiles if one has substantially longer remaining maturity or more premium cash flows still outstanding.

This effect is particularly important for mezzanine and senior tranches. Junior tranches are exposed first to realized defaults, so their risk is often dominated by immediate principal losses. By contrast, mezzanine and senior tranches can suffer substantial mark-to-market losses even before they are directly hit by defaults. The reason is that once junior protection is eroded, the expected future loss of more senior tranches increases and the fair spread required by the market rises. Since the premium leg of these tranches is often relatively long-dated, the present value impact of that repricing can be large. This is consistent with the broader factor-model literature, which emphasizes that tranche pricing depends on expected losses over multiple horizons rather than only on immediate realized defaults [4, 10].

Duration also interacts with the timing of losses. In dynamic CDO settings, forward tranche spreads tend to increase when protection is moved further into the future, because later periods are associated with higher cumulative default risk and less subordination protection remaining. Hull and White [10] show that more advanced dynamic models are needed when one wants to track how portfolio credit risk evolves over time and how tranche values respond to changing future credit conditions. This provides an intuition for why longer effective duration increases tranche risk as the longer the horizon over which the tranche is exposed to deterioration in the credit environment, the larger the repricing effect.

For our project, this implies that duration should not be interpreted only as classical interest-rate duration. For structured credit, the more relevant notion is the combination of maturity, credit-spread duration, and the remaining premium-paying notional of the tranche. Altogether, a longer-duration tranche is more exposed to worsening credit spreads, more vulnerable to mark-to-market losses after subordination erosion, and more heavily penalized in the solvency framework. This interpretation is also consistent with evidence that CDO valuation based only on expected cash flows can miss important mark-to-market effects, especially for more senior tranches [3].

5 Model Description

5.1 Single-factor Gaussian copula model

To simulate the risk and loss distribution of the CDO portfolio, we employ a Single-Factor Gaussian Copula model, evaluated via Monte Carlo simulation. This framework, rooted in Merton's structural model of default, provides a robust and computationally efficient methodology to model the joint default probability of a large pool of correlated assets.

The core assumption of the model is that the default of any individual loan or asset i is driven by an unobservable, continuous latent variable A_i. This variable is typically interpreted as the normalized asset value of the obligor. To capture the correlation between different obligors in the portfolio, we decompose this latent variable into two mutually independent, standard normally distributed components: a systematic market factor M and an idiosyncratic, obligor-specific factor ϵ_i.

The structural equation for the latent asset value is defined as

$$A_i = \sqrt{\rho}M + \sqrt{1-\rho}\epsilon_i, \tag{8}$$

where $M \sim \mathcal{N}(0,1)$ captures macroeconomic shocks affecting the entire portfolio, and $\epsilon_i \sim \mathcal{N}(0,1)$ represents the unique, company-specific risks of obligor i. The parameter $\rho \in [0,1]$ represents the asset correlation to the market factor. Because both the systematic and idiosyncratic components are standard normal, the resulting latent variable A_i is also standard normally distributed.

A default event is triggered if the latent variable A_i falls below a specific default threshold X_i within the evaluation period. The default indicator D_i can be expressed as

$$D_i = \mathbf{1}_{A_i < X_i}. \tag{9}$$

To align this structural mechanism with empirical risk data, the threshold X_i is calibrated directly to the obligor's marginal probability of default (PD_i). Using the inverse cumulative distribution function (quantile function) of the standard normal distribution, Φ^{-1}, the threshold is established as

$$X_i = \Phi^{-1}(PD_i). \tag{10}$$

This ensures that the expected default rate in the model matches the individual credit ratings provided by external rating agencies. Upon determining the default states for all assets in a given simulation path, the model calculates the raw portfolio losses. This is achieved by multiplying the default indicators by their respective loss severities, represented by $(1 - \text{Recovery Rate})$, and then weighting them according to the portfolio composition. To reflect real-world uncertainty, the recovery rate (or Loss Given Default, LGD) can be modeled either as a fixed percentage or sampled stochastically, for example from a Beta distribution with wanted mean and concentration parameters.

Finally, these aggregated portfolio losses are passed through the CDO waterfall structure. The losses absorbed by a specific investment tranche are calculated by clipping the aggregate portfolio loss against the tranche's predefined attachment and detachment points. We run this process over thousands of simulated paths to construct an approximated loss distribution, allowing us to derive key risk metrics such as Expected Tranche Loss and Value-at-Risk (VaR).

5.2 Determining the correlation with common market factor

As established in Equation (8), our one-period default model relies heavily on the asset correlation parameter ρ. It is important to estimate this correlation accordingly since it has a significant effect on joint default probabilities. Basel 2 provides a framework [2] on how to estimate ρ depending on the underlying loan type to capture systematic risk. This Basel 2 states that the correlation between market factor and a corporate loan can be determined by formula

$$R = 0.12\frac{1-e^{-50PD}}{1-e^{-50}} + 0.24(1 - \frac{1-e^{-50PD}}{1-e^{-50}}), \tag{11}$$

which is dependent on the probability of defaults (credit rating) of the underlying loan. Basel 2 uses R instead of ρ for denoting correlation. The formula is based on the principle that higher credit rating (lower probability of default) pushes the correlation towards its minimum value 0.12 while lower credit rating (higher probability of default) increases it. According to Basel 2 [2], this formula is same for corporate, sovereign and bank loans but for SME (small and medium enterprises whose annual sales are less than 50 million euros) it is adjusted by

$$0.04(1 - (S - 5)/45), \tag{12}$$

where S is the annual sales in millions of euros, ranging between 5 and 50 million euros. On the other hand, we should use a fixed correlation of 15 percent with residential mortgage exposures and 4 percent with revolving retail exposures [2]. Finally, ”other retail exposures” is likely a major type of exposure in our CDOs, which has similar formula than corporate/sovereign/bank loans but uses different parameters

$$R = 0.03\frac{1 - e^{-35PD}}{1 - e^{-35}} + 0.16(1 - \frac{1 - e^{-35PD}}{1 - e^{-35}}) \tag{13}$$

All in all, to estimate correlation ρ in our model we could take in the distribution of underlying exposure in each class, the distribution of credit rating (probability of default) within these classes and estimate their correlation to systematic market factor with

$$\rho = a\frac{1 - e^{-cPD}}{1 - e^{-c}} + b(1 - \frac{1 - e^{-cPD}}{1 - e^{-c}}) \tag{14}$$

where parameters a, b, c can be retrieved from

Exposure	a	b	c
Corporate	0.12	0.24	-50
Sovereign	0.12	0.24	-50
Bank	0.12	0.24	-50
Other retail	0.03	0.15	-35

For retail mortgage $\rho = 0.15$ and for revolving retail $\rho = 0.04$.

5.3 Probability of default

The PD of individual assets affects the latent asset value A_i through ρ, as we can see from equation (14). In addition to that, as defined in our structural model, we assert a threshold to the latent asset value to determine default events. If the threshold is breached then the loan has defaulted. As previously shown in equation (10), this threshold depends strictly on the PD of the individual assets. These default probabilities are published by S&P Global annually. Table 1 shows the statistics of these 1 year default probabilities over time

Table 1: Descriptive statistics on one-year global default rates (%) [15]

	AAA	AA	A	BBB	BB	B	CCC/C
Minimum	0.00	0.00	0.00	0.00	0.00	0.25	0.00
Maximum	0.00	0.38	0.38	1.00	4.24	13.84	49.46
Weighted long-term average	0.00	0.02	0.04	0.13	0.55	2.87	26.12
Median	0.00	0.00	0.00	0.05	0.45	3.12	25.90
Standard deviation	0.00	0.06	0.10	0.24	0.95	3.20	11.47
2008 default rates	0.00	0.38	0.38	0.49	0.82	4.09	27.27
Latest four quarters (Q1 2025–Q4 2025)	0.00	0.00	0.00	0.00	0.08	1.18	25.90

We weighted long-term averages for the default probabilities since yearly statistics have relatively low absolute default amounts and thus have large error boundaries. Table 2 shows the yearly default rates per rating class and their 95 % error boundaries.

The error in delta has obviously direct affect on the threshold of the defaults in the model. Additionally, it has affect on the market correlation ρ with the following relationship.

$$\frac{d\rho}{dPD} = -\frac{(a - b)\, c\, e^{c\,PD + c}}{e^{c} - 1} \tag{15}$$

Table 2: Estimated one-year PD and 95% confidence bounds (%)

	AAA	AA	A	BBB	BB	B	CCC/C
PD (long-term avg)	0.00	0.02	0.04	0.13	0.55	2.87	26.12
PD min (95%)	0.00	0.00	0.00	0.00	0.00	0.00	3.64
PD max (95%)	0.00	0.14	0.24	0.60	2.41	9.14	48.60
Δ PD (95%)	0.00	0.12	0.20	0.47	1.86	6.27	22.48

5.4 Conditional repricing after defaults

The one-factor Gaussian copula model described above is a natural starting point for simulating the loss distribution of a CDO portfolio. However, for the purposes of this project, a pure hold-to-maturity loss engine is not sufficient. In particular, if defaults have already occurred by an intermediate observation date u, the value of a tranche must be updated to reflect both the realized losses up to u and the expected future losses from u to maturity. This is especially important for mezzanine and senior tranches as discussed previously [4, 10].

Let M denote the initial portfolio notional and let tranche j be defined by attachment and detachment points K_L^j and K_U^j. The cumulative portfolio loss at time t is denoted by L_t, and the cumulative loss allocated to tranche j is

$$L_{j,t} = \min(L_t, K_U^j M) - \min(L_t, K_L^j M).$$

Equivalently, this was written in another form in chapter 4 as

$$L_{j,t} = (L_t - K_L^j M)^+ - (L_t - K_U^j M)^+,$$

i.e. the tranche absorbs only the part of the portfolio loss lying between its attachment and detachment levels [4].

To value the tranche after defaults have already occurred, we use the outstanding tranche notional

$$\Gamma_{j,t} = (K_U^j - K_L^j)M - L_{j,t}.$$

This was the remaining principal of tranche j after losses up to time t, as stated earlier. Since future premium payments are made only on the outstanding tranche notional, $\Gamma_{j,t}$ is a state variable linking realized losses to future tranche value.

Let u denote the observation date, for example one year from initial valuation. At time u, the realized portfolio loss L_u, the realized tranche loss $L_{j,u}$, and the remaining tranche notional $\Gamma_{j,u}$ are known. The mark-to-market value of tranche j at time u is then defined as the conditional present value of the remaining premium leg minus the conditional present value of the remaining protection leg

$$V_j(u) = PV_{\text{prem}}(u) - PV_{\text{prot}}(u).$$

Assuming future payment dates $u < t_{m+1} < \cdots < t_K = T$, this can be written as

$$PV_{\text{prem}}(u) = s_j \sum_{k=m+1}^{K} \Delta t_k \, D(u, t_k) \, \mathbb{E}[\Gamma_{j,t_k} \mid \mathcal{F}_u] \,,$$

$$PV_{\text{prot}}(u) = \sum_{k=m+1}^{K} D(u, t_k) \, \mathbb{E}\big[L_{j,t_k} - L_{j,t_{k-1}} \mid \mathcal{F}_u\big] \,,$$

where s_j is the contractual tranche spread, Δt_k is the accrual fraction between payment dates, $D(u, t_k)$ is the discount factor from t_k back to u, and $\mathcal{F}_u$ denotes the information available at time u.

This repricing can also be expressed in terms of the remaining loss capacity of the tranche after time u. Defining the currency attachment and detachment levels

$$A_j = K_L^j M, \qquad B_j = K_U^j M.$$

After realized portfolio loss L_u, the effective future attachment and detachment levels become

$$\widetilde{A}_j(u) = \max(0, A_j - L_u), \qquad \widetilde{B}_j(u) = \max(0, B_j - L_u).$$

If $\Delta L_{u,t}$ denotes the future portfolio loss from u to t, then the future loss allocated to tranche j may be written as

$$\Delta L_{j,t|u} = (\Delta L_{u,t} - \widetilde{A}_j(u))^+ - (\Delta L_{u,t} - \widetilde{B}_j(u))^+.$$

This formulation makes the economic effect of subordination erosion more visible. If equity losses have already consumed most of the lower protection, then $\widetilde{A}_j(u)$ becomes small. As a result, even moderate additional future losses are much more likely to hit the mezzanine tranche, which increases the expected future protection payment and decreases the expected future premium-paying notional. Consequently, the mark-to-market value of the tranche can decline sharply even when the realized tranche loss at time u is still zero.

For our project, the conditional repricing framework provides a practical bridge between static tranche-loss simulation and economically meaningful 1-year risk measurement for holder of the tranche. Rather than only recording whether the tranche has already suffered principal loss by time u, the model also accounts for the deterioration in the value of the remaining cash flows. This is the feature needed to capture the client's requirement that defaults in junior tranches should also reduce the value of more senior holdings through erosion of subordination and higher expected future losses [4, 10].

5.5 Realized loss, mark-to-market change and expected return

Once conditional repricing has been introduced, it becomes important to separate three related but distinct quantities: realized tranche loss, mark-to-market value change, and expected return. They capture different dimensions of CDO risk and investment performance. As noted before, a tranche may suffer a large mark-to-market loss even when its realized principal loss is still zero. This distinction is essential in structured credit, where the waterfall structure amplifies changes in expected future losses and spread levels [4, 3].

The first quantity is the realized tranche loss. At observation date u, this is simply the cumulative loss already allocated to tranche j through the waterfall, namely $L_{j,u}$ and is the part of the loss that has already materialized through defaults and recoveries. It represents an actual reduction in tranche principal and is therefore a realized cash loss from the investor's point of view. If $L_{j,u} = 0$, then the tranche has not yet suffered direct principal write-down.

The second quantity is the mark-to-market value change. From Section 5.d, the tranche value at time u is

$$V_j(u) = PV_{\text{prem}}(u) - PV_{\text{prot}}(u),$$

that is, the conditional present value of the remaining premium leg minus the conditional present value of the remaining protection leg. The corresponding mark-to-market change over the period $[0, u]$ is

$$\Delta V_j^{\text{MTM}}(u) = V_j(u) - V_j(0).$$

This term reflects the repricing of the remaining cash flows after the market has observed defaults, recoveries, and erosion of subordination by time u. It is therefore possible that

$$L_{j,u} = 0 \qquad \text{but} \qquad \Delta V_j^{\text{MTM}}(u) < 0.$$

This is an important case where, for example, equity losses reduce the protection available to mezzanine, causing the mezzanine tranche to reprice downward even before any realized mezzanine principal loss occurs [3, 10].

In addition to realized tranche losses, the model separates two mark-to-market effects. The first is the endogenous repricing effect, which is generated by the simulated state of the portfolio itself. This component reflects realized defaults, realized recoveries, erosion of subordination, changes in the surviving pool and the realized systematic factor. The second is an exogenous spread-driven mark-to-market effect, which represents broader market spread movements not generated directly by the simulated defaults inside the pool.

For each rating, maturity and tranche scenario, the model computes a scenario-specific fair running spread. This is the spread that makes the tranche value equal to zero at inception, meaning that the present value of the premium leg equals the present value of the protection leg. This is the standard in synthetic CDO pricing: the protection seller receives periodic spread payments on the outstanding tranche notional and, in exchange, compensates the protection buyer for losses allocated to the tranche [4].

Let $PV^{\mathbb{Q}}_{\text{prot},j}(0)$ denote the risk-neutral present value of the protection leg of tranche j. This term measures the discounted expected losses that the protection seller must pay when portfolio losses hit the tranche. Let

$$PV01^{\mathbb{Q}}_{\text{prem},j}(0) = \sum_{k=1}^{K} \Delta t_k D(0,t_k)\mathbb{E}^{\mathbb{Q}}[\Gamma_{j,t_k}]$$

denote the risky premium annuity of the same tranche. Here Δt_k is the accrual fraction, $D(0,t_k)$ is the discount factor and Γ_{j,t_k} is the outstanding tranche notional at payment date t_k. The premium annuity is therefore the present value of one unit of running spread paid on the expected surviving tranche notional.

If the running spread is s_j, the premium leg is

$$PV^{\mathbb{Q}}_{\text{prem},j}(0;s_j) = s_j PV01^{\mathbb{Q}}_{\text{prem},j}(0).$$

The fair spread s_j^{fair} is obtained from the break-even condition

$$PV^{\mathbb{Q}}_{\text{prem},j}(0;s_j^{\text{fair}}) = PV^{\mathbb{Q}}_{\text{prot},j}(0),$$

which gives

$$s_j^{\text{fair}} = \frac{PV^{\mathbb{Q}}_{\text{prot},j}(0)}{PV01^{\mathbb{Q}}_{\text{prem},j}(0)}.$$

The numerator is the expected discounted compensation for tranche losses, while the denominator is the expected discounted premium-paying notional. A tranche with higher expected losses has a larger protection leg. A tranche that is expected to amortize or be written down quickly has a smaller premium annuity. Both effects increase the fair spread. This also explains why distressed tranches can produce very large fair spreads. Consider for example if the expected surviving tranche notional is close to zero, the premium annuity becomes very small, so the break-even spread can become large even for a moderate protection leg.

The exogenous spread-driven mark-to-market component is then approximated by

$$\Delta V^{(s)}_{j,\text{spread}} \approx -D_j^{\text{spread}} V_j(0)\Delta s_j^{(s)},$$

where D_j^{spread} is the spread duration of the tranche, $V_j(0)$ is the initial tranche value and $\Delta s_j^{(s)}$ is the one-year spread shock in simulation path s. Following the project simplification, the spread duration is set equal to the tranche maturity

$$D_j^{\text{spread}} = T.$$

The one-year holding-period return in simulation path s is therefore written as

$$R_j^{(s)} = C_j^{(s)} - L_{j,u}^{(s)} + \Delta V^{(s)}_{j,\text{endo}} + \Delta V^{(s)}_{j,\text{spread}},$$

where $C_j^{(s)}$ is premium carry over $[0,u]$, $L_{j,u}^{(s)}$ is realized tranche loss, $\Delta V^{(s)}_{j,\text{endo}}$ is the endogenous conditional repricing effect and $\Delta V^{(s)}_{j,\text{spread}}$ is the exogenous spread-driven mark-to-market effect. The premium carry is computed using the model-implied fair spread s_j^{fair}.

The spread shock in our model is modeled as

$$\Delta s_r^{(s)} \sim N(0, \sigma_s(r)^2),$$

where $\sigma_s(r)$ is the one-year spread volatility for rating bucket r. A horizon profit-and-loss measure is therefore

$$\Pi_j(0,u) = C_j^{\text{prem}}(0,u) - L_{j,u} + \Delta V_j^{\text{endo}}(u) + \Delta V_j^{\text{spread}}(u).$$

The corresponding expected holding-period return can then be defined as

$$ER_j(0,u) = \frac{\mathbb{E}[\Pi_j(0,u)]}{V_j(0)},$$

provided that the tranche is valued at market value $V_j(0)$ at inception. If instead the tranche is normalized by its principal amount, one may replace the denominator by the initial tranche notional $(K_U^j - K_L^j)M$. In either case, it is important that the expected return combines both realized cash flow effects and repricing of the remaining position.

This distinction is important because each quantity answers a different question. Realized loss answers how much principal has actually been lost so far. Mark-to-market change answers how has the value of the remaining tranche changed given the updated credit environment. Expected return answers what is the total economic gain or loss of holding the tranche over the horizon. In a simpler hold-to-maturity loss engine, only the first question is addressed. In our project, however, the client's requirement extends to the second and third questions as well, since subordination erosion changes the value of the holding even before the tranche suffers direct principal write-down.

For this reason, the 1-year risk analysis of a tranche should not be based only on the distribution of $L_{j,u}$. It should also incorporate the distribution of $\Delta V_j^{\mathrm{MTM}}(u)$ and, the distribution of the full holding-period profit and loss $\Pi_j(0,u)$. This is also consistent with the broader structured-credit literature, where tranche valuation depends on expected losses at multiple dates and not just on immediate realized defaults [4, 10].

5.6 Finite heterogeneous pool

The single-factor Gaussian copula model introduced earlier in this chapter provides a baseline for modeling default dependence in a CDO portfolio. However, one of its most restrictive simplifications is the large homogeneous pool assumption, under which the underlying assets are treated as if they had identical notionals, default probabilities, recovery rates, and factor loadings. This approximation is analytically convenient, but it is not fully appropriate for the type of structured credit portfolio considered in our project. In practice, a CDO or CLO pool is finite, and the names within it typically differ in exposure size, rating quality, loss given default, and sensitivity to the common market factor [4, 12].

To account for this, we extend the baseline model to a finite heterogeneous pool with n underlying assets. Let N_i denote the notional of exposure i, PD_i its marginal probability of default, R_i its recovery rate, and ρ_i its correlation with the common market factor. The latent asset variable of name i is then written as

$$A_i = \sqrt{\rho_i}\, M + \sqrt{1-\rho_i}\, \epsilon_i,$$

where $M \sim N(0,1)$ is the common market factor and $\epsilon_i \sim N(0,1)$ is the idiosyncratic shock of obligor i. A default occurs when

$$A_i < X_i, \qquad X_i = \Phi^{-1}(PD_i),$$

so that the threshold is calibrated name by name to the corresponding marginal default probability.

Conditional on a realization $M = y$, the defaults are independent and each name has conditional default probability

$$p_i(y) = \Phi\left(\frac{\Phi^{-1}(PD_i) - \sqrt{\rho_i}\, y}{\sqrt{1-\rho_i}} \right).$$

This preserves the main computational benefit of the factor approach as even though the pool is heterogeneous, the dependence structure is still driven by a single systematic factor, and the conditional loss distribution can be obtained either by Monte Carlo simulation or by semi-analytical numerical methods [12, 4]. We will discuss a semi-analytical numerical method implemented in chapter 6.

The cumulative portfolio loss at time t is now written as

$$L_t = \sum_{i=1}^{n} N_i(1-R_i)\, \mathbf{1}_{\{\tau_i \le t\}},$$

where τ_i denotes the default time of obligor i. In contrast to the homogeneous-pool setting, each default contributes a name-specific loss size $N_i(1-R_i)$, so the aggregate loss distribution is no longer determined only by the number of defaults. Instead, it depends on which names default. This difference is important because two scenarios with the same number of defaults can produce very different tranche outcomes if the underlying names have different notionals or recovery rates.

The tranche waterfall is then applied exactly as before. If tranche j has attachment point K_L^j and detachment point K_U^j, its cumulative loss is

$$L_{j,t} = (L_t - K_L^j M)^+ - (L_t - K_U^j M)^+,$$

and its outstanding tranche notional is

$$\Gamma_{j,t} = (K_U^j - K_L^j)M - L_{j,t}.$$

Thus, the finite heterogeneous extension does not change the tranche-loss formula itself, but it changes the distribution of the underlying portfolio loss L_t that enters the waterfall.

This extension is particularly relevant for conditional repricing after defaults. In a homogeneous large-pool approximation, the state of the tranche can often be summarized by aggregate loss alone. In a finite heterogeneous pool, this is no longer generally sufficient. After some names have defaulted, the future loss distribution depends not only on current loss L_u, but also on the remaining composition of the pool, since the surviving names may have systematically different risk characteristics from the names that have already defaulted. In this sense, defaulted name heterogeneity makes the future loss process depend on the current portfolio structure and not only on the current loss level [4].

For our project, the finite heterogeneous pool should therefore be understood as a extension of the baseline Vasicek model rather than a different framework. The common-factor structure is retained, but the underlying pool is allowed to contain different name sizes, different default probabilities, and different loss severities. If full look-through data are not available, the same idea can be implemented at the bucket level by grouping assets according to rating class, exposure type, or maturity. This still preserves the economic effect of the the portfolio being finite and heterogeneous, so tranche risk depends on the composition of the collateral pool rather than on a single representative obligor.

5.7 Stochastic recovery

Previously, recovery has entered the model through the term $(1 - R_i)$, where R_i is the recovery rate of obligor i. A common simplification in Gaussian copula models is to assume that each obligor has a fixed recovery rate. While this is convenient, it is not fully realistic for structured credit. Empirical evidence suggests that recovery rates are uncertain and tend to be lower in adverse default environments. For this reason and due to client demands, we retain stochastic recovery in our model rather than treating recovery as a deterministic constant [1, 4].

In the finite heterogeneous pool introduced above, the cumulative portfolio loss is

$$L_t = \sum_{i=1}^{n} N_i(1 - R_i)\mathbf{1}_{\{\tau_i \le t\}},$$

where N_i is the notional of exposure i, τ_i is its default time, and $R_i \in [0, 1]$ is now a random recovery rate. Thus, each default contributes a random loss amount $N_i(1-R_i)$. Compared with the deterministic-recovery case, stochastic recovery broadens the portfolio loss distribution even if the mean recovery is unchanged. This matters for tranche valuation because the tranche-loss mapping is highly nonlinear and additional dispersion in portfolio loss can affect mezzanine and senior tranche risk.

To preserve tractability, we model the recovery rate of obligor i with a Beta distribution,

$$R_i \sim \mathrm{Beta}(\alpha_i, \beta_i),$$

which is naturally supported on the interval $[0, 1]$. It is convenient to parameterize this distribution by a mean recovery μ_i and a concentration parameter ν_i, so that

$$\alpha_i = \mu_i \nu_i, \qquad \beta_i = (1 - \mu_i)\nu_i.$$

Under this parameterization,

$$\mathbb{E}[R_i] = \mu_i, \qquad \mathrm{Var}(R_i) = \frac{\mu_i(1 - \mu_i)}{\nu_i + 1}.$$

Hence, μ_i determines the expected recovery level while ν_i controls the dispersion around that mean. A higher value of ν_i implies lower recovery uncertainty, whereas a lower value implies a wider range of possible loss-given-default outcomes.

Conditional on a realization of the common market factor $M = y$, the expected portfolio loss becomes

$$\mathbb{E}[L_t \mid M = y] = \sum_{i=1}^{n} N_i\big(1 - \mu_i\big)\, p_i(y, t),$$

where $p_i(y, t)$ is the conditional default probability of obligor i by time t. Thus, if only the mean recovery changes, stochastic recovery does not alter the conditional expected loss. Its main effect is instead on the dispersion of losses. If defaults and recoveries are independent conditional on the factor, then the conditional variance of the loss contribution of obligor i is

$$\mathrm{Var}\big(N_i(1-R_i)\mathbf{1}_{\{\tau_i\le t\}} \mid M=y\big) = N_i^2\left[p_i(y,t)\sigma_{R_i}^2 + p_i(y,t)\big(1-p_i(y,t)\big)(1-\mu_i)^2\right],$$

where $\sigma_{R_i}^2 = \mathrm{Var}(R_i)$. The first term is absent under deterministic recovery and shows how recovery uncertainty increases loss variability.

This is important for tranche risk because different tranches react differently to wider loss distributions. Junior tranches are already exposed to early losses, but for mezzanine and senior tranches the additional right-tail mass generated by stochastic recovery can increase expected tranche loss and worsen mark-to-market outcomes. In particular, after subordination has been eroded, recovery uncertainty in the surviving pool can heavily impact the conditional repricing formulas from Section 5.d, since future portfolio loss $\Delta L_{u,t}$ is then driven by both default uncertainty and recovery uncertainty.

For our project, the stochastic-recovery specification should be understood as a controlled extension of the baseline Gaussian copula model. The dependence structure of defaults still comes from the common market factor, but the severity of default losses is no longer fixed ex ante. In the implementation, the mean recovery μ_i and concentration ν_i can be assigned either name by name or at the bucket level by rating class, exposure type, or seniority. This is consistent with the project's broader objective as the model should remain tractable and explainable, but it should also be realistic enough to reflect the fact that defaults in a structured credit portfolio do not all generate the same loss severity [1, 4].

6 Implementation architecture

This section describes how the mathematical framework above is implemented in the model. The implementation is designed to remain consistent with the scope of the project. It should be sufficiently realistic to capture the main economic mechanisms of CDO risk, but also transparent and computationally light enough to be explainable and simplifiable into a client-facing proxy tool.

The model is implemented as a two-layer architecture. The first layer is a physical simulation of defaults and recoveries up to the observation horizon u, which in our baseline case is one year. The second layer is a conditional repricing step from u to maturity, where the remaining tranche cash flows are valued given the realized state of the portfolio at time u. This architecture allows us to move beyond the static hold-to-maturity loss framework and incorporate the mark-to-market effect of subordination erosion discussed in Section 5.4 and 5.5 [4, 10].

6.1 Input layer

The implementation begins from a configuration layer that specifies the main model inputs. These include:

- the number of loans or exposure buckets in the underlying pool,
- the one-year default probabilities, either directly or through credit ratings and rating mixtures,
- the common-factor correlation parameter ρ,
- the tranche attachment and detachment points,
- the recovery specification, either fixed or stochastic,
- the return-model assumptions such as spread, spread duration, and repricing settings.

The model also resolves the portfolio weights of the underlying exposures. If explicit weights are not given, the pool is treated as equally weighted. If rating mixtures are used instead of name-level data, the model expands the rating buckets into a finite pool with the desired number of names. In this way, the same implementation can work in two data modes.

6.2 Outer simulation of physical-measure one-year loss

In the outer layer, the model simulates the state of the portfolio at the observation horizon u. The latent asset value of each name is generated under the single-factor Gaussian copula structure

$$A_i = \sqrt{\rho}\,M + \sqrt{1-\rho}\,\epsilon_i,$$

where M is the common systematic factor and ϵ_i is the idiosyncratic shock of obligor i. A default occurs when the latent asset value falls below the threshold calibrated from the obligor-specific probability of default. If stochastic recovery is enabled, the recovery rate is sampled at default time and otherwise a fixed recovery level is used.

From the simulated defaults and recoveries, the implementation computes the portfolio loss

$$L_u = \sum_{i=1}^{n} N_i(1 - R_i)\mathbf{1}_{\{\tau_i \le u\}},$$

and the corresponding tranche loss

$$L_{j,u} = (L_u - K_L^j M)^+ - (L_u - K_U^j M)^+.$$

At the same time, the model records the surviving loan pool and the realized value of the systematic factor. These objects form the state that is passed to the conditional repricing layer.

6.3 Initial valuation at time 0

Before running the horizon return analysis, the tranche is valued at time 0 using a risk-neutral finite-pool loss distribution. This valuation produces the initial premium leg, protection leg and risky premium annuity of the tranche.

As discussed in section 5.5 a running-spread tranche, the premium leg is linear in the contractual spread

$$PV^{\mathbb{Q}}_{\text{prem},j}(0; s) = s\, PV01^{\mathbb{Q}}_{\text{prem},j}(0).$$

The fair running spread is the spread that makes the initial tranche value equal to zero:

$$V_j(0; s_j^{\text{fair}}) = s_j^{\text{fair}} PV01^{\mathbb{Q}}_{\text{prem},j}(0) - PV^{\mathbb{Q}}_{\text{prot},j}(0) = 0.$$

The fair spread is reported as a model output. Very large fair spreads can occur when the risky premium annuity is close to zero. In those cases, the value is not to be read as a realistic coupon, but as an indication that the tranche is close to exhaustion or not quoteable with our model.

6.4 Conditional repricing from u to maturity

The second layer of the architecture is the conditional repricing step. For each outer scenario, the implementation revalues the remaining tranche cash flows from time u to maturity T. This step uses the state variables generated in the outer simulation:

$$(L_u,\ L_{j,u},\ \Gamma_{j,u},\ \mathcal{S}_u,\ M_u)\,,$$

where $\Gamma_{j,u}$ is the outstanding tranche notional, $\mathcal{S}_u$ is the surviving pool, and M_u denotes the realized systematic factor at the observation date.

The implementation updates the forward default probabilities of the surviving names conditionally on the factor realization at the observation time. Thus, bad outer states correspond to worse forward credit conditions in the repricing step rather than reverting mechanically back to unconditional average conditions. This is important for capturing the mark-to-market effect of a deteriorated default environment.

Conditional on these forward default probabilities, the future loss distribution of the surviving pool is constructed using a bucketed semi-analytical loss. If stochastic recovery is enabled, each obligor's loss-given-default is represented by a discretized recovery distribution. The resulting future portfolio loss distribution is then passed through the tranche waterfall to compute the expected future tranche losses and the expected future outstanding tranche notional. From these, the model evaluates the conditional premium and protection legs

$$V_j(u) = PV_{\text{prem}}(u) - PV_{\text{prot}}(u).$$

6.5 Bucketed loss distribution for conditional repricing

The conditional repricing step requires the model to estimate the future loss distribution of the surviving pool after the observation date u. A full nested Monte Carlo simulation would be computationally expensive, since every outer one-year scenario would require a separate inner simulation from u to maturity. To avoid this, the implementation uses a bucketed loss-distribution method inspired by the probability bucketing approach of Hull and White [8]. The idea is to construct the conditional portfolio loss distribution directly on a discrete loss grid, contrary to simulating a large number of future loss paths.

The method is suitable for our one-factor Gaussian copula setting as conditional on the common systematic factor, defaults can be shown to be independent. Thus, the future loss distribution of the surviving pool can be built recursively by adding one exposure at a time. This is consistent with factor-copula CDO models, where the common factor generates dependence but conditional independence is recovered once the factor is fixed [12, 4].

As earlier, $\mathcal{S}_u$ denotes the pool of names that have survived to the observation date u. For a future payment date $t_k > u$, let

$$p_i(u, t_k \mid \mathcal{F}_u)$$

denote the conditional forward default probability of surviving name i between u and t_k, given the information available at time u. This information includes the realized portfolio loss, the surviving pool and the observed common factor. The future portfolio loss from u to t_k is:

$$\Delta L_{u,t_k} = \sum_{i \in \mathcal{S}_u} N_i(1 - R_i)\mathbf{1}_{\{u < \tau_i \le t_k\}},$$

where N_i is the notional of exposure i, R_i is its recovery rate and τ_i is its default time. The objective of the bucketing algorithm is to approximate the distribution of $\Delta L_{u,t_k}$.

First, define a loss grid

$$0 = x_0 < x_1 < \cdots < x_B,$$

where each grid point represents a possible future portfolio loss bucket. Let $\pi_b^{(m)}$ denote the probability that the cumulative future loss lies in bucket b after the first m surviving exposures have been processed. The algorithm is initialized by placing all probability mass at zero loss:

$$\pi_0^{(0)} = 1, \qquad \pi_b^{(0)} = 0 \quad \text{for } b > 0.$$

If recovery were deterministic, exposure i would create a fixed loss amount

$$\ell_i = N_i(1 - R_i)$$

in the event of default. Then each existing loss bucket has two transitions: the name either survives, in which case the loss bucket is unchanged, or defaults, in which case probability mass is shifted by ℓ_i. The recursion is

$$\begin{aligned} \pi_{\text{new}}(x) \ &\mathrel{+}= \ (1 - p_i)\pi_{\text{old}}(x), \\ \pi_{\text{new}}(x + \ell_i) \ &\mathrel{+}= \ p_i \pi_{\text{old}}(x), \end{aligned}$$

where $p_i = p_i(u, t_k \mid \mathcal{F}_u)$.

In our implementation recovery is stochastic. Therefore a default does not correspond to a single loss amount. Instead, the recovery distribution is discretized into a small number of recovery states. If the recovery of name i has states

$$r_{i,1}, \ldots, r_{i,K}$$

with probabilities

$$w_{i,1}, \ldots, w_{i,K}, \qquad \sum_{k=1}^{K} w_{i,k} = 1,$$

then the corresponding default loss states are

$$\ell_{i,k} = N_i(1 - r_{i,k}).$$

The bucketing update becomes

$$\begin{aligned} \pi_{\text{new}}(x) \ &\mathrel{+}= \ (1 - p_i)\pi_{\text{old}}(x), \\ \pi_{\text{new}}(x + \ell_{i,k}) \ &\mathrel{+}= \ p_i w_{i,k} \pi_{\text{old}}(x), \qquad k = 1, \ldots, K. \end{aligned}$$

Thus, each surviving exposure contributes one no-default branch and several default branches corresponding to the discretized recovery states. This allows the conditional repricing to retain stochastic recovery without running an inner Monte Carlo simulation. The use of stochastic recovery is also consistent with extensions of the Gaussian copula framework that incorporate random recovery rates as discussed by Andersen and Sidenius [1].

After all surviving names have been processed, the algorithm produces an approximate conditional distribution

$$\mathbb{P}(\Delta L_{u,t_k} = x_b \mid \mathcal{F}_u) \approx \pi_b.$$

This distribution is then passed through the tranche waterfall. Let

$$A_j = K_L^j M, \qquad B_j = K_U^j M$$

denote the currency attachment and detachment levels of tranche j. For each future loss bucket x_b, the total portfolio loss at time t_k is

$$L_u + x_b,$$

and the corresponding cumulative tranche loss is

$$L_{j,t_k}(x_b) = (L_u + x_b - A_j)^+ - (L_u + x_b - B_j)^+.$$

The conditional expected tranche loss is then computed as

$$\mathbb{E}[L_{j,t_k} \mid \mathcal{F}_u] = \sum_{b=0}^{B} \pi_b L_{j,t_k}(x_b).$$

Similarly, the expected outstanding tranche notional is

$$\mathbb{E}[\Gamma_{j,t_k} \mid \mathcal{F}_u] = (B_j - A_j) - \mathbb{E}[L_{j,t_k} \mid \mathcal{F}_u].$$

These conditional expectations are the inputs to the repricing formulas from Section 5.d. The premium leg depends on the expected outstanding tranche notional, whereas the protection leg depends on expected future tranche-loss increments. Therefore, the bucketed loss distribution provides the computational bridge between the simulated state at time u and the mark-to-market value

$$V_j(u) = PV_{\text{prem}}(u) - PV_{\text{prot}}(u).$$

The main approximation in this method is the discretization of future losses onto a finite grid. A finer grid improves accuracy but increases runtime. In our prototype, the grid size and the number of recovery states are therefore treated as numerical implementation parameters rather than fundamental modeling assumptions. This is appropriate for the project scope, since the purpose is to obtain a tractable conditional repricing approximation rather than a full nested Monte Carlo valuation.

6.6 Tranche Return decomposition

Once both the realized one-year loss and the conditional mark-to-market value are available, the implementation decomposes the horizon performance into components. These are:

- premium carry over the period $[0, u]$,
- realized tranche loss $L_{j,u}$,
- endogenous repricing effect through $V_j(u) - V_j(0)$,
- exogenous spread-driven mark-to-market component.

This decomposition is consistent with the distinction introduced in Section 5.e between realized loss, mark-to-market change, and expected return.

6.7 Outputs

After repeating the two-layer procedure over all simulation paths, the model aggregates the pathwise results into the final outputs relevant for the client such as expected tranche loss, the 1-year loss distribution, the 1-year 97% Value-at-Risk, the distribution of mark-to-market value changes and the expected holding-period return. Thus, the final implementation does not only answer how much loss has already materialized by the observation date, but also how the remaining value of the tranche has changed in response to defaults, recoveries, and erosion of subordination.

6.8 Scope and limitations of the implementation

The implementation should be interpreted as a project-level structured-credit risk engine rather than a full trading-desk pricing framework. It is more realistic than a pure large homogeneous pool approximation because it allows finite pools, heterogeneous exposures, and stochastic recovery. At the same time, it remains efficient because the inner repricing step uses a semi-analytical bucketed loss distribution rather than a full nested Monte Carlo. The main simplifications are the use of a single common-factor structure, a fixed observation horizon, and a simplified mapping from the physical observation-time macro state to the repricing inputs. These simplifications are acceptable for the scope of the project, since the goal is not to replicate market base-correlation models, but to build a coherent and explainable tool for structured-credit risk assessment.

7 Parameter choices and calibration assumptions

The purpose of this section is to distinguish between parameters that are exposed to the end user and assumptions that are fixed in the model. This distinction is important because the final deliverable is intended to be a practical proxy tool. Therefore, the Excel interface contains only a small number of meaningful inputs, while the remaining modeling choices are fixed in the background.

7.1 User-controlled input parameters

Based on the client feedback, the main end-user inputs in the Excel proxy tool are the tranche attachment point, detachment point, rating and maturity. These variables define the structural position of the tranche and the baseline credit-risk profile used by the model.

Table 3: User-controlled input parameters in the Excel proxy tool.

Input	Description	Role in the model
Rating	Credit quality or rating bucket of the underlying exposure.	Determines the baseline one-year probability of default.
Maturity	Remaining maturity of the tranche in years.	Determines the horizon over which future losses and repricing effects are evaluated.
Attachment point	Lower loss threshold of the tranche.	Determines when the tranche starts absorbing portfolio losses.
Detachment point	Upper loss threshold of the tranche.	Determines when the tranche is fully exhausted.

Other parameters, such as recovery assumptions, correlation, simulation count and repricing settings, are fixed in the background and documented in the assumptions sheet.

7.2 Fixed structural model assumptions

The core dependence model is fixed to be the one-factor Gaussian copula. More complex market-implied base-correlation models are not included, as we only build an explainable risk tool for tranche-level loss and return assessment [13, 12, 4].

Table 4: Fixed structural model assumptions.

Assumption:	Baseline choice:
Copula family	One-factor Gaussian copula
Observation horizon u	1 year
Loss allocation	Attachment-detachment tranche waterfall
Loss metric	1-year 97% VaR
Conditional repricing	Enabled

7.3 Portfolio and pool assumptions

The model is intended to work both with limited summary data and with look-through data. In the baseline setting, the underlying portfolio is represented as a finite pool. If exposure weights are not available, the model assumes an equal-weighted pool. If look-through data are available, the same architecture can use name-level notionals, probabilities of default and recovery assumptions.

This means that the finite heterogeneous pool should be viewed as a realism-enhancing extension of the baseline model. The common-factor structure is retained, but the underlying pool can incorporate different ratings, exposures and recovery assumptions [12, 4].

7.4 Rating-to-PD calibration

In line with the client discussions, the rating input is interpreted as the average rating of the underlying collateral pool. For a selected rating bucket r, the baseline one-year default probability $PD_r(1)$ is taken from public historical default-rate data [15]. The copula model then maps this marginal default probability into a default threshold, following the standard default-correlation approach of David X. Li [13].

As only one-year default probabilities are available, we approximate the maturity-dependent cumulative probability using a constant-hazard survival model. Defining

$$\lambda_r = -\log(1 - PD_r(1)),$$

the survival probability to maturity T is

$$S_r(T) = e^{-\lambda_r T},$$

and the cumulative default probability is

$$PD_r(T) = 1 - S_r(T) = 1 - e^{-\lambda_r T}.$$

Equivalently,

$$PD_r(T) = 1 - (1 - PD_r(1))^T.$$

This approximation preserves monotonicity and ensures that cumulative default probabilities remain between zero and one.

For conditional repricing after the observation date u, the corresponding forward default probability from u to T, conditional on survival to u, is

$$PD_r(u, T) = \frac{PD_r(T) - PD_r(u)}{1 - PD_r(u)} = 1 - (1 - p_r)^{T-u}.$$

The conditional repricing layer adjusts future default probabilities using the realized common factor at the observation date. Let $p_r = PD_r(1)$ denote the one-year default probability associated with average portfolio rating r. The cumulative default probability over horizon t is computed using the constant-hazard approximation

$$PD_r(t) = 1 - (1 - p_r)^t.$$

At observation time u, the Monte Carlo path has a realized systematic factor $Y_u = y$. To condition future default probabilities on the macro state observed at time u, we use a time-consistent Gaussian factor approximation. Specifically, we assume that the systematic factor at horizon t can be written as

$$Y_t = \frac{W_t}{\sqrt{t}},$$

where W_t is a Brownian motion. This gives

$$\mathrm{Corr}(Y_u, Y_t) = \sqrt{\frac{u}{t}}.$$

Hence, observing $Y_u = y$ shifts the distribution of the future systematic factor Y_t, which in turn changes the conditional default probability to horizon t.

$$PD_r(t \mid Y_u = y) = \Phi\left(\frac{\Phi^{-1}(PD_r(t)) - \sqrt{\rho}\sqrt{u/t}\, y}{\sqrt{1 - \rho u/t}}\right).$$

The corresponding forward default probability from u to t, conditional on survival to u, is

$$PD_r(u, t \mid Y_u = y) = \frac{PD_r(t \mid Y_u = y) - PD_r(u \mid Y_u = y)}{1 - PD_r(u \mid Y_u = y)}.$$

This forward probability is used for the surviving names in the conditional repricing step. Negative realizations of Y_u increase the conditional forward default probabilities, while positive realizations decrease them. Thus, the model links bad one-year macro states to higher expected future losses and lower conditional tranche values.

7.5 Default correlation assumptions

The baseline model uses a single common-factor correlation parameter ρ. This means that all names are exposed to the same systematic credit factor, while idiosyncratic shocks remain independent conditional on the factor. For the Excel proxy model, ρ is also thus fixed. However, because tranche risk is sensitive to default correlation, sensitivity runs are produced for low, base and high correlation cases [13, 12, 4].

7.6 Recovery-rate assumptions

The implementation uses a Beta-distributed recovery rate that is naturally justified as discussed previously. The mean of 50 % controls the expected recovery level, while the concentration parameter (20) controls the dispersion around that mean [1, 4].

In the conditional repricing step, the continuous recovery distribution is discretized into five recovery states to keep the repricing engine stable. Recovery variance regularization is included as a modeling safeguard for highly distressed names after 85 % PD threshold.

7.7 Return and mark-to-market assumptions

The return model decomposes the one-year holding-period performance of tranche j into four components, which are premium carry, realized tranche loss, endogenous repricing and exogenous spread-driven mark-to-market change. The decomposition shows how a CDO tranche can lose value through different channels. A realized tranche loss corresponds to an actual allocation of portfolio losses through the waterfall, whereas mark-to-market changes reflect changes in the value of the remaining contractual cash flows.

The premium carry is based on a model-implied fair running spread. In a synthetic CDO, the protection seller receives periodic spread payments on the outstanding tranche notional and compensates the protection buyer for losses allocated to the tranche. Standard CDO pricing therefore determines the running premium by equating the present value of the premium leg and the present value of the protection, or default, leg under a risk-neutral measure [4, 6]. Hull and White express the same idea by writing the contract value in terms of an annuity factor for unit running payments and a payoff term arising from defaults; the breakeven spread is then the default payoff present value divided by the premium annuity, including any accrual adjustment [9].

Let $PV^{\mathbb{Q}}_{\text{prot},j}(0)$ denote the risk-neutral present value of the protection leg of tranche j. This is the discounted expectation of future tranche-loss increments:

$$PV^{\mathbb{Q}}_{\text{prot},j}(0) = \sum_{k=1}^{K} D(0, t_k)\mathbb{E}^{\mathbb{Q}}\left[L_{j,t_k} - L_{j,t_{k-1}}\right],$$

where $D(0, t_k)$ is the discount factor and L_{j,t_k} is the cumulative loss allocated to tranche j by payment date t_k.

The premium leg is paid only on the outstanding tranche notional. If

$$\Gamma_{j,t} = (K^j_U - K^j_L)M - L_{j,t}$$

denotes the outstanding tranche notional at time t, then the risky premium annuity is

$$PV01^{\mathbb{Q}}_{\text{prem},j}(0) = \sum_{k=1}^{K} \Delta t_k D(0, t_k)\mathbb{E}^{\mathbb{Q}}[\Gamma_{j,t_k}].$$

This quantity is the present value of one unit of running spread paid on the expected surviving tranche notional. It is risky because the notional on which the premium is paid can decline when losses hit the tranche. Hence, for a running spread s_j, the premium leg is

$$PV^{\mathbb{Q}}_{\text{prem},j}(0; s_j) = s_j PV01^{\mathbb{Q}}_{\text{prem},j}(0).$$

The fair running spread is the value of s_j that makes the tranche value equal to zero at inception

$$PV^{\mathbb{Q}}_{\text{prem},j}(0; s^{\text{fair}}_j) = PV^{\mathbb{Q}}_{\text{prot},j}(0).$$

Since the premium leg is linear in the spread, this gives

$$s^{\text{fair}}_j = \frac{PV^{\mathbb{Q}}_{\text{prot},j}(0)}{PV01^{\mathbb{Q}}_{\text{prem},j}(0)}.$$

In addition to the fair-spread level, the model uses a separate rating-based spread volatility to represent exogenous market spread movements over the one-year horizon. The fair spread s^{fair}_j determines the premium carry of the tranche, whereas the spread-volatility input $\sigma_s(r)$ determines the size of the simulated market spread shock. The model therefore does not set the running spread from a broad rating table but it only uses rating information to assign a spread-shock volatility when tranche-specific spread-volatility data are not available.

For a simulation path s, the one-year spread shock is specified as

$$\Delta s^{(s)}_j \sim N(0, \sigma_s(r)^2),$$

where r denotes the rating bucket of the scenario. In the implementation, the annual spread-volatility inputs are given in basis points as

r	AAA	AA	A	BBB	BB	B	CCC
$\sigma_s(r)$ (bps)	51	72	97	103	117	110	475

The large volatility assigned to the CCC bucket reflects the fact that very weak credit-quality buckets are treated as highly unstable spread-risk states. In the simplified Excel proxy, the distressed C/Unrated category is treated consistently with this lower-credit-quality spread-risk convention. It is not used to determine the contractual running spread, which is instead obtained from the fair-spread calculation above.

7.8 Conditional repricing assumptions

The repricing layer values the remaining tranche cash flows after the observation date u, which is set to one year in the baseline specification. The remaining horizon is therefore equal to the tranche maturity minus one year. The model assumes annual payment dates in the conditional repricing step. The state passed to the repricing engine includes the realized portfolio loss, realized tranche loss, outstanding tranche notional, surviving pool and the realized common factor. The conditioning happens on surviving pool, realized loss and the realized common factor. This allows the model to capture both subordination erosion and deterioration in the credit environment after defaults have occurred.

The model uses a loss grid size of 256, but this can be increased in runs if runtime allows. The macro risk premium shift is not accounted for and should be treated as a simplifying assumption.

7.9 Numerical implementation parameters

These numerical parameters are implementation choices rather than fundamental model assumptions. The numerical implementation uses a fixed random seed of 123 to make the simulation results reproducible. The current batch size is 1000, which is an implementation parameter controlling how simulations are processed computationally. The Excel proxy interpolation is enabled in the current workbook version, allowing the workbook to return approximate outputs when the user inputs do not exactly match a precomputed Python scenario grid point.

7.10 Sensitivity parameters

Several model parameters are uncertain and should be tested through sensitivity analysis. In particular, tranche risk is expected to be sensitive to default correlation, recovery assumptions, spread volatility and pool granularity.

Table 5: Sensitivity parameters.

Parameter	Baseline	Sensible sensitivity cases
Default correlation ρ	20%	5%, 40%
Mean recovery	50%	40%, 60%
Beta concentration	20	10, 50

8 Excel proxy and deliverables

The final practical deliverable is a lightweight Excel-based proxy tool built on top of the Python model. The intended workflow is

$$\text{Python model} \to \text{scenario output table} \to \text{Excel proxy workbook.}$$

Python performs the simulation, loss allocation, fair-spread calculation, conditional repricing and final risk metric selection. Excel acts as a client-facing interface for retrieving and displaying the precomputed outputs.

8.1 Workbook structure

The workbook is divided into separate sheets, each with a specific role. The main sheets are summarized in Table 6.

Table 6: Structure of the Excel proxy workbook.

Sheet	Purpose
`README`	Explains the purpose of the workbook, the intended workflow, the input fields, the main outputs and the limitations of the proxy approach.
`Inputs`	Contains the user input fields: rating, spread duration or maturity, attachment point and detachment point.
`Outputs`	Presents the selected scenario output in a client-facing format, including final VaR, expected return and diagnostic model components.
`Assumptions`	Documents the fixed assumptions behind the Python run and Excel proxy, including rating buckets, tranche buckets, duration buckets and spread-risk treatment.
`Raw_Import`	Stores the raw scenario table exported from Python. This sheet is the only sheet that should be replaced when a new Python scenario run is imported.
`Model_Data`	Mirrors the imported Python data and adds helper fields used for lookup, interpolation and display.
`Lookup_Helper`	Performs the lookup calculation for the flexible input-output interface. It first checks for an exact match and otherwise uses nearby scenario rows for interpolation.
`Simplified_Table`	Presents the standardized 4-by-9 final VaR table indexed by rating bucket, tranche bucket and spread-duration bucket.

The workbook contains two complementary interfaces. The first is the flexible `Inputs/Outputs` interface, where the user selects a specific rating, maturity or spread duration, attachment point and detachment point. The second is the `Simplified_Table`, which gives a compact standardized table for the client's main rating, tranche and duration buckets.

8.2 Imported Python data and model-data layer

The Python model exports a long-form scenario table. This table is pasted into the `Raw_Import` sheet of the workbook and mirrored into the `Model_Data` sheet. The imported data include the rating bucket, spread duration, tranche attachment and detachment, fair spread, premium annuity, expected tranche loss, loss VaR, expected return, raw return VaR, final VaR, spread-driven mark-to-market components, endogenous repricing components and conditional remaining loss metrics.

The main client-facing risk metric is

`final_var_97_1y`.

The model also exports

`return_var_raw_97_1y`, `return_var_97_1y`, `final_var_source`, `final_var_fallback_reason`.

These fields make the reported VaR explainable. The final VaR is selected by the Python model, while Excel only retrieves the exported value. If a fallback rule is used for a distressed or weakly quoteable case, the source and fallback-reason fields explain the origin of the final reported value.

8.3 User inputs

The input interface uses four user-controlled inputs: rating bucket, spread duration, attachment point and detachment point. In the base model, spread duration is interpreted as the maturity bucket used in the Python scenario grid. These inputs were selected because they determine the main risk profile of a CDO/CLO tranche in the project setting. Other assumptions, such as default correlation, recovery model, simulation count and repricing settings, are fixed in the background and documented in the assumptions sheet and in the report.

8.4 Main outputs

The main workbook output is the one-year 97% final VaR,

`final_var_97_1y`.

This value is the primary proxy risk measure used in the simplified table. It is intended to represent the model-selected one-year downside risk of the tranche, taking into account realized tranche losses, return-based mark-to-market effects and any defined logic for distressed cases.

The workbook reports several diagnostic outputs:

- expected one-year return,
- expected one-year tranche loss,
- one-year tranche-loss VaR,
- raw return VaR,
- fair spread and premium annuity,
- final VaR source and fallback reason,
- expected spread-driven mark-to-market effect,
- spread-driven mark-to-market volatility,
- expected endogenous repricing effect,
- initial tranche value under the valuation measure,
- base remaining tranche expected loss,
- conditional remaining tranche expected loss after the one-year observation horizon.

8.5 Simplified 4-by-9 table

The `Simplified_Table` sheet provides a compact reporting table for the standardized scenarios. The table has four rating rows:

AA, BBB, B, C/Unrated,

and nine columns formed from three tranche buckets and three spread-duration buckets. The tranche buckets are

0–10%, 14–18%, 60–100%,

representing equity-style, mezzanine-style and senior-style tranche positions. The spread-duration buckets are

$$2.5, \qquad 5.5, \qquad 9.0$$

years.

Each cell in the simplified table reports `final_var_97_1y` from the Python scenario output. The table is therefore a standardized final VaR matrix rather than a separate Excel risk calculation.

8.6 Lookup and interpolation

The `Inputs/Outputs` interface first checks whether the selected input combination exactly matches one row in the imported Python scenario table. If an exact match exists, the workbook returns the corresponding model output directly. If the selected combination is not exactly available in the scenario grid, the workbook uses nearby rows to form an interpolated proxy value.

The simplified 4-by-9 table is more restrictive. It is designed around the standardized rating, tranche and duration buckets and therefore uses exact lookups from the precomputed scenario grid. This makes the simplified table easier to interpret and reduces the risk that the reported values are affected by interpolation assumptions.

The Excel workbook should therefore be interpreted as a proxy interface and values outside the generated scenario grid should be treated as extrapolations. For material decisions or unusual tranche structures, the Python model should be rerun with the desired input assumptions rather than relying only on Excel interpolation.

9 Validation and sensitivity analysis

This section reports the validation results generated by the current Python model. The validation runs use the final VaR, fair-spread calculation and final VaR source diagnostics described above. Unless otherwise stated, values in the tables are percentages of tranche notional.

9.1 Purpose of validation

The model should produce sensible results under controlled scenarios. Since CDO risk measures are sensitive to model assumptions, rating inputs and tranche structure, the model should satisfy basic monotonicity properties. For example, increasing default probabilities should increase expected losses, while lower attachment points should produce larger tranche losses because a larger part of the portfolio loss distribution is allocated to the tranche through the waterfall [4]. In addition, erosion of junior subordination should reduce the mark-to-market value of mezzanine and senior tranches, since it increases expected future tranche losses and changes the value of the remaining premium and protection legs [10, 3].

9.2 Sensitivity parameters

In an ideal case, the model should be tested under alternative assumptions for the parameters that are most uncertain or most important for tranche risk. The main sensitivity parameters are: default correlation ρ, mean recovery rate, beta recovery dispersion, number of loans in the pool, contractual spread, spread volatility, maturity,attachment and detachment points.

The sensitivity analysis ideally identifies which assumptions are most important for Ilmarinen's use case.

9.3 Validation setup

The baseline validation grid uses 500 simulation paths, default correlation $\rho = 0.20$, mean recovery 50%, Beta concentration 20, five recovery states in the repricing layer and a loss grid size of 256. The grid consists of four rating buckets, three spread-duration buckets and three tranche buckets:

$$4 \times 3 \times 3 = 36$$

rating–duration–tranche scenarios. Additional sensitivity runs vary correlation, recovery, repricing and the number of simulation paths.

9.4 Final VaR source diagnostics

The model exports both `final_var_97_1y` and fields explaining where the final VaR comes from. Table 7 summarizes the source diagnostics in the validation output.

Table 7: Final VaR source counts in the validation output.

Final VaR source	Fallback reason	Number of rows
`return`	–	65
`tranche_loss`	`thin_premium_annuity`	3

Most validation rows use the simulated return distribution as the final VaR source. The three fallback rows correspond to the C/Unrated 0–10% tranche across the three duration buckets. In these cases the premium annuity is very small and the model uses tranche-loss VaR as the final VaR source, with the fallback reason `thin_premium_annuity`.

9.5 Fair-spread validation

Because the model solves a scenario-specific fair spread, validation checks both the fair-spread value and the risky premium annuity. Table 8 reports the 5.5-year validation cases. Very small or approximately zero fair spreads for remote senior tranches are not problematic as they indicate that the model-implied protection leg is essentially zero for that tranche under the scenario. By contrast, the C/Unrated 0–10% tranche has an extremely large fair spread and a very small premium annuity, which is a diagnostic of a distressed or poorly quoteable running-spread state for our model.

Table 8: Fair-spread diagnostics, 5.5-year validation grid. Final VaR is shown as percentage of tranche notional.

Rating	Tranche	Fair spread (bps)	Premium annuity	Final VaR	Source	Fallback reason
AA	0-10	10.03	5.48	7.75	`return`	–
AA	14-18	≈ 0	5.50	7.90	`return`	–
AA	60-100	0.00	5.50	7.90	`return`	–
BBB	0-10	71.39	5.38	13.22	`return`	–
BBB	14-18	≈ 0	5.50	11.31	`return`	–
BBB	60-100	≈ 0	5.50	11.31	`return`	–
B	0-10	2,483.62	3.01	79.31	`return`	–
B	14-18	0.03	5.50	13.58	`return`	–
B	60-100	≈ 0	5.50	12.07	`return`	–
C/Unrated	0-10	1.85e+06	0.01	100.00	`tranche_loss`	`thin_premium_annuity`
C/Unrated	14-18	11,141.42	0.90	74.80	`return`	–
C/Unrated	60-100	≈ 0	5.50	52.14	`return`	–

9.6 Rating monotonicity

The rating monotonicity test checks whether weaker ratings lead to higher modelled risk. Table 9 reports the 5.5-year spread-duration validation results.

Table 9: Rating monotonicity validation, 500 simulations, 5.5-year spread duration. Values are percentages of tranche notional except expected portfolio loss, which is a percentage of portfolio notional.

Rating	Tranche	Exp. port. loss	Exp. tr. loss	Loss VaR	Raw return VaR	Final VaR	Source
AA	0-10	0.009	0.09	1.33	7.75	7.75	`return`
AA	14-18	0.009	0.00	0.00	7.90	7.90	`return`
AA	60-100	0.009	0.00	0.00	7.90	7.90	`return`
BBB	0-10	0.062	0.62	5.61	13.22	13.22	`return`
BBB	14-18	0.062	0.00	0.00	11.31	11.31	`return`
BBB	60-100	0.062	0.00	0.00	11.31	11.31	`return`
B	0-10	1.49	14.78	70.14	79.31	79.31	`return`
B	14-18	1.49	0.00	0.00	13.58	13.58	`return`
B	60-100	1.49	0.00	0.00	12.07	12.07	`return`
C/Unrated	0-10	12.91	83.82	100.00	0.00	100.00	`tranche_loss`
C/Unrated	14-18	12.91	29.39	100.00	74.80	74.80	`return`
C/Unrated	60-100	12.91	0.00	0.00	52.14	52.14	`return`

The expected portfolio loss increases from approximately 0.01% for AA to 12.91% for C/Unrated. The junior 0–10% tranche shows the strongest rating sensitivity: final VaR increases from 7.75% for AA to 100% for C/Unrated. For mezzanine and senior tranches, final VaR is driven mostly by return-based mark-to-market effects unless the tranche-loss fallback is activated. The C/Unrated 14–18% and 60–100% tranches have high final VaR values, but their final VaR source remains `return`.

9.7 Tranche seniority and waterfall behaviour

The tranche seniority test checks whether losses are allocated through the waterfall in the correct order. Table 10 reports the B-rated 5.5-year validation result.

Table 10: Tranche seniority validation, rating B, 5.5-year spread duration, 500 simulations. Values are percentages of tranche notional except fair spread.

Tranche	Exp. tr. loss	Loss VaR	Raw return VaR	Final VaR	Fair spread (bps)	Exp. return	Endog. MTM
0-10	14.78	70.14	79.31	79.31	2,483.62	5.55	-2.50
14-18	0.00	0.00	13.58	13.58	0.03	-0.70	-0.53
60-100	0.00	0.00	12.07	12.07	≈ 0	-0.17	0.00

The waterfall behaviour is economically sensible. The junior 0–10% tranche has the highest direct loss exposure, with expected tranche loss of 14.78%, loss VaR of 70.14% and final VaR of 79.31%. The 14–18% and 60–100% tranches have zero loss VaR in this validation run, but positive final VaR due to the return-based mark-to-market channels.

9.8 Maturity and spread-duration sensitivity

The duration sensitivity test checks whether longer spread duration increases return-based risk. Table 11 reports the B-rated 14–18% tranche across the three duration buckets.

Table 11: Duration sensitivity validation, rating B, 14–18% tranche, 500 simulations. Values are percentages of tranche notional except fair spread and premium annuity.

Duration	Loss VaR	Raw return VaR	Final VaR	Fair spread (bps)	Premium annuity	Exp. return	Endog. MTM
2.50	0.00	5.49	5.49	≈ 0	2.50	-0.08	-0.00
5.50	0.00	13.58	13.58	0.03	5.50	-0.70	-0.53
9.00	0.00	64.81	64.81	37.26	8.96	-6.83	-6.92

The result confirms the expected duration effect. Realized loss VaR remains zero for this tranche across all three duration buckets, while final VaR rises from 5.49% at 2.5 years to 64.81% at 9.0 years. The increase is mainly driven by return-based mark-to-market risk and endogenous repricing of the remaining cash flows.

9.9 Conditional repricing on/off

The repricing switch test isolates the effect of the endogenous conditional repricing layer. Table 12 compares the B-rated 14–18% and 60–100% tranches with conditional repricing enabled and disabled.

Table 12: Conditional repricing validation, rating B, 5.5-year spread duration, 500 simulations. Values are percentages of tranche notional.

Tranche	Repricing	Loss VaR	Raw return VaR	Final VaR	Exp. return	Endog. MTM
14-18	No	0.00	12.07	12.07	-0.17	0.00
14-18	Yes	0.00	13.58	13.58	-0.70	-0.53
60-100	No	0.00	12.07	12.07	-0.17	0.00
60-100	Yes	0.00	12.07	12.07	-0.17	0.00

For the 14–18% tranche, enabling conditional repricing increases final VaR from 12.07% to 13.58% and decreases expected return from -0.17% to -0.70%. The effect is smaller for the 60–100% tranche in this validation run, where the senior tranche remains far from direct loss and the endogenous repricing component is essentially zero.

9.10 Correlation sensitivity

The correlation sensitivity test varies the common-factor correlation parameter:

$$\rho = 0.05, \quad 0.20, \quad 0.40.$$

Table 13 reports the B-rated 5.5-year cases.

Table 13: Correlation sensitivity validation, rating B, 5.5-year spread duration, 500 simulations. Values are percentages of tranche notional.

ρ	Tranche	Exp. tr. loss	Loss VaR	Raw return VaR	Final VaR	Exp. return
0.05	0-10	14.72	39.91	50.02	50.02	2.25
0.05	14-18	0.00	0.00	12.07	12.07	-0.18
0.05	60-100	0.00	0.00	12.07	12.07	-0.17
0.20	0-10	14.78	70.14	79.31	79.31	5.55
0.20	14-18	0.00	0.00	13.58	13.58	-0.70
0.20	60-100	0.00	0.00	12.07	12.07	-0.17
0.40	0-10	13.33	99.44	86.43	86.43	10.99
0.40	14-18	1.16	0.00	37.06	37.06	-3.12
0.40	60-100	0.00	0.00	12.07	12.07	-0.17

Higher correlation increases tail clustering. For the junior 0–10% tranche, loss VaR rises from 39.91% at $\rho = 0.05$ to 99.44% at $\rho = 0.40$, while final VaR rises from 50.02% to 86.43%. The 14–18% tranche is also sensitive at high correlation, with final VaR increasing to 37.06% at $\rho = 0.40$. The 60–100% tranche remains almost unchanged in this validation run because it is not directly reached by one-year realized losses and its risk is dominated by spread and return effects.

The effect of the default-correlation parameter is not expected to be strictly monotonic for every tranche-level metric. Elizalde [6] notes that while equity and senior tranche premiums have clearer monotonic relationships with default correlation, the effect on mezzanine tranches is not clear-cut. This is because increasing correlation changes the shape of the portfolio loss distribution. Increasing correlation may reduce the probability of moderate losses while increasing the probability of extreme losses. Therefore, the relevant validation point is that correlation changes the tail behaviour in an economically plausible way, especially for junior and mezzanine tranches.

9.11 Recovery sensitivity

The recovery mean sensitivity test varies the mean recovery assumption:

$$\mu_R = 0.40, \quad 0.50, \quad 0.60.$$

Table 14 reports the B-rated 5.5-year results.

Table 14: Recovery mean sensitivity validation, rating B, 5.5-year spread duration, 500 simulations. Values are percentages of tranche notional except expected portfolio loss and fair spread.

Mean recovery	Tranche	Exp. port. loss	Exp. tr. loss	Loss VaR	Raw return VaR	Final VaR	Fair spread (bps)
0.40	0-10	1.79	17.45	83.59	77.94	77.94	3,425.30
0.40	14-18	1.79	0.21	0.00	40.45	40.45	1.20
0.50	0-10	1.49	14.78	70.14	79.31	79.31	2,483.62
0.50	14-18	1.49	0.00	0.00	13.58	13.58	0.03
0.60	0-10	1.18	11.74	55.64	73.00	73.00	1,718.11
0.60	14-18	1.18	0.00	0.00	12.07	12.07	≈ 0

Lower mean recovery increases expected portfolio loss, expected tranche loss and loss VaR as expected. For the 0–10% tranche, the fair spread rises from 1718.11 bps at 60% mean recovery to 3425.30 bps at 40% mean recovery. For the 14–18% tranche, direct loss VaR remains zero, but final VaR is materially higher under the 40% recovery case because lower recoveries worsen conditional repricing and return tails.

The recovery-dispersion check varies the Beta concentration parameter:

$$\nu_R = 10, \quad 20, \quad 50.$$

The results for the B-rated 14–18% tranche are shown in Table 15.

Table 15: Recovery concentration sensitivity validation, rating B, 14–18% tranche, 5.5-year spread duration, 500 simulations. Values are percentages of tranche notional.

Beta concentration	Exp. tr. loss	Loss VaR	Raw return VaR	Final VaR	Exp. return
10	0.05	0.00	15.63	15.63	-0.78
20	0.00	0.00	13.58	13.58	-0.70
50	0.01	0.00	14.47	14.47	-0.72

The relationship is not monotone, which is reasonable for a lightweight 500-simulation tail-risk validation run. The important conclusion is that recovery dispersion affects the return tail and should remain part of sensitivity analysis.

9.12 Simulation convergence

The simulation convergence test compares 500, 1000, 2000 and 5000 simulation paths for the B-rated 14–18% tranche with 5.5-year spread duration. The results are shown in Table 16.

Table 16: Simulation convergence validation, rating B, 14–18% tranche, 5.5-year spread duration. Values are percentages of tranche notional.

Simulations	Exp. tr. loss	Loss VaR	Raw return VaR	Final VaR	Exp. return	Endog. MTM
500	0.00	0.00	13.58	13.58	-0.70	-0.53
1000	0.10	0.00	13.92	13.92	-0.78	-0.51
2000	0.08	0.00	13.40	13.40	-0.62	-0.52
5000	0.07	0.00	12.78	12.78	-0.54	-0.49

The final VaR estimate remains in a relatively narrow range across the convergence runs: 13.58%, 13.92%, 13.40% and 12.78% for 500, 1000, 2000 and 5000 simulations, respectively. Expected tranche loss remains small and loss VaR remains zero. This supports the use of the simulation engine for proxy table, while also showing that tail estimates for thin mezzanine tranches remain noisy.

9.13 Summary of validation findings

The validation results support the main economic logic of the model. Weaker ratings increase portfolio loss and junior tranche risk. The tranche waterfall behaves correctly, with junior tranches absorbing losses before mezzanine and senior tranches. Longer spread duration increases return-based risk when direct realized loss VaR is not binding. Conditional repricing has a measurable effect on mezzanine tranches. Correlation and recovery assumptions affect tail risk and should therefore remain part of the sensitivity analysis.

The fair-spread diagnostics provide an additional valuation check. Regular scenarios produce interpretable fair running spreads, while distressed or nearly exhausted tranches can produce very large spreads because the risky premium annuity is small. These cases are identified through the premium-annuity and final VaR source diagnostics rather than hidden by capping the raw spread.

Some non-monotonic behaviour remains in validation runs as the simulation count is small for example. This is expected because CDO tranche risk is nonlinear and because tail estimates are sensitive to simulation noise. Such deviations should be interpreted as simulation error, tranche position, spread duration, conditional repricing and final VaR source diagnostics rather than automatically model errors.

10 Numerical results

This section presents the numerical results generated by the Python model and the Excel proxy workbook. The purpose is to demonstrate how the model produces the core outputs requested by the client: a one-year downside risk measure and an expected return estimate.

The main simplified reporting metric is

`final_var_97_1y`.

This value is selected in the Python model and then displayed in Excel. All values in the simplified table are reported as percentages of tranche notional.

10.1 Baseline setup

The baseline numerical run uses the parameter choices described in Section 7. The main setup is summarized in Table 17. The spread-duration buckets are interpreted as short, medium and long spread-duration cases. The three tranche buckets represent equity-style, mezzanine-style and senior-style tranche positions.

Table 17: Baseline numerical setup for the simulation run.

Input	Baseline value
Number of simulations	5,000
Number of loans	200
Default correlation ρ	0.20
Mean recovery	50%
Recovery model	Beta-distributed stochastic recovery
Recovery states in repricing	5
Loss grid size	256
Batch size	1000
Observation horizon	1 year
Rating buckets	AA, BBB, B, C/Unrated
Spread-duration buckets	2.5, 5.5 and 9.0 years
Tranche buckets	0–10%, 14–18% and 60–100%
Headline simplified metric	`final_var_97_1y`

10.2 Simplified final VaR table

Table 18 presents the simplified 4-by-9 final VaR table from the 5,000-simulation final output grid. Each cell reports

`final_var_97_1y`

as a percentage of tranche notional.

Table 18: Simplified 4-by-9 final VaR table from the 5,000-simulation final output grid. Values are percentages of tranche notional.

Rating	Bucket	0–10%			14–18%			60–100%		
		2.5y	5.5y	9.0y	2.5y	5.5y	9.0y	2.5y	5.5y	9.0y
AA	[AAA, AA-]	3.57	7.48	12.09	3.30	7.27	11.89	3.30	7.27	11.89
BBB	[A+, BBB-]	6.99	12.19	18.28	4.73	10.40	17.01	4.73	10.40	17.01
B	[BB+, CCC]	67.23	75.70	71.85	5.11	12.88	60.56	5.05	11.10	18.17
C/Unrated	Unrated	100.00	100.00	100.00	57.60	75.73	98.77	21.79	47.95	78.46

The table is a one-year downside risk matrix for the protection seller. In regular scenarios, the final VaR is driven by the simulated holding-period return distribution. In distressed or weakly quoteable for our model scenarios, the source and fallback-reason fields identify whether the reported value is based on a fallback rule. In the final output grid, the fallback is activated for the C/Unrated 0–10% tranche, where the premium annuity is very small and the tranche-loss source is used.

The values show that tranche seniority, rating quality and spread duration are the main drivers of the reported risk. The 0–10% tranche is dominated by direct credit risk in the weaker rating buckets. The 14–18% tranche is especially sensitive to the combination of lower rating and longer maturity: in the B-rated row, final VaR increases from 5.11% at 2.5 years to 60.56% at 9.0 years. The 60–100% tranche has lower direct loss exposure but still reports positive final VaR because the return distribution includes spread-driven mark-to-market risk and conditional repricing.

10.3 Comparison across tranche seniority

The simplified table confirms that tranche seniority is one of the dominant drivers of reported risk. The 0–10% tranche is the most junior tranche in the simplified grid and is therefore the first tranche to absorb portfolio losses. This explains why the weaker rating buckets show very high final VaR values for this tranche. In the B-rated row, the 0–10% final VaR is between 67.23% and 75.70% across the three duration buckets. In the C/Unrated row, the 0–10% tranche reaches 100% in all three duration buckets.

The 14–18% tranche behaves differently. For the AA and BBB rows, the tranche is not materially affected by direct one-year realized tranche losses in the final grid. Instead, the final VaR is mainly driven by the return-based mark-to-market channels. This is visible from the duration pattern: for example, the BBB 14–18% tranche increases from 4.73% at 2.5 years to 17.01% at 9.0 years. In the B row, the same tranche increases from 5.11% at 2.5 years to 60.56% at 9.0 years. This shows that mezzanine-style tranches can become highly sensitive when lower credit quality is combined with a long remaining horizon.

The 60–100% tranche is the most senior simplified tranche. It has lower direct one-year loss exposure than the junior tranche, but its final VaR is still positive because the model includes mark-to-market risk. For AA and BBB, the senior-tranche values are close to the corresponding mezzanine values because both are mainly driven by spread-duration effects rather than direct realized loss. In the C/Unrated row, the 60–100% tranche increases from 21.79% at 2.5 years to 78.46% at 9.0 years. This indicates that even senior-style exposure can have material one-year downside risk when the underlying credit state is distressed and the remaining cash flows are long-dated.

Overall, the table shows that the waterfall behaves as expected. Junior tranches are dominated by direct credit loss, while mezzanine and senior tranches are more sensitive to mark-to-market effects, spread volatility, spread duration and conditional repricing.

10.4 Comparison across ratings

Rating quality affects the model primarily through the default-probability input. Moving from AA to BBB, B and C/Unrated increases the expected loss of the underlying pool and therefore increases the probability that portfolio losses reach the lower tranches. This effect is clearest for the 0–10% tranche. At 5.5 years, final VaR increases from 7.48% for AA to 12.19% for BBB, 75.70% for B and 100.00% for C/Unrated.

For mezzanine and senior tranches, the rating effect is less mechanical because final VaR combines realized loss, fair-spread carry, endogenous repricing and exogenous spread-driven mark-to-market effects. The 14–18% tranche still shows a clear deterioration pattern. For example for maturity of 5.5 years, final VaR is 7.27% for AA, 10.40% for BBB, 12.88% for B and 75.73% for C/Unrated. The senior 60–100% tranche also increases heavily in the distressed bucket, from 11.10% for B to 47.95% for C/Unrated at 5.5 years.

The table is not interpreted as requiring strict monotonicity in every single cell. The final VaR is a tail statistic of a simulated holding-period return distribution, and it contains several interacting components. As a result, local non-monotonicity can occur. For example, the B-rated 0–10% tranche is 75.70% at 5.5 years and 71.85% at 9.0 years. This does not imply that longer maturity is safer in general. This just reflects the nonlinear interaction between the realized-loss tail, premium carry, spread shocks and conditional repricing in a finite Monte Carlo sample. The important pattern is that weaker rating buckets generally increase portfolio loss, junior-tranche loss exposure and distressed-state final VaR.

10.5 Comparison across maturities and spread durations

The model uses spread duration to scale the exogenous spread-driven mark-to-market component. However, the final VaR is not driven by spread-duration scaling alone. It is the 97% downside measure of the full one-year return distribution,

$$R_j = C_j + \Delta V_{j,\text{spread}} + \Delta V_{j,\text{endo}} - L_{j,u}.$$

Thus, longer maturity can increase VaR both because spread shocks are applied over a longer spread-duration exposure and because the conditional repricing layer revalues more years of remaining premium and protection-leg cash flows.

The duration effect is visible in the AA and BBB rows. For example, the BBB 60–100% tranche increases from 4.73% at 2.5 years to 17.01% at 9.0 years, even though this tranche is not directly hit by one-year portfolio losses in the final output grid. The effect is especially strong for mezzanine-style tranches under weaker credit quality. For the B-rated 14–18% tranche, final VaR increases from 5.11% at 2.5 years to 12.88% at 5.5 years and 60.56% at 9.0 years.

The B-rated 14–18% 9-year cell has an economic explanation. Its final VaR is not a statement that there is a 3% probability of losing 60.56% of principal directly within one year. In the final output, the one-year realized tranche-loss VaR for this cell is 0.00%. The 60.56% value is instead the empirical 3% lower-tail return quantile,

$$\texttt{return\_p03\_1y} = -0.605571487997273, \qquad \texttt{return\_var\_97\_1y} = 0.605571487997273.$$

With 5,000 simulation paths, this corresponds approximately to the 150th worst simulated return path.
For that 3% quantile path, the return decomposition is shown in Table 19.

Table 19: Return decomposition for the B-rated 14–18% tranche, 9-year maturity, at the empirical 3% return quantile.

Component	Decimal return	Percentage of tranche notional
Premium carry	0.0037258	0.37%
Spread-driven MTM	-0.1481439	-14.81%
Endogenous tranche MTM	-0.4611533	-46.12%
Realized one-year tranche loss	0.0000000	0.00%
Total return	-0.6055715	-60.56%

The large component is the endogenous mark-to-market loss. On this path, the portfolio has 5.11% realized one-year portfolio loss and 20 defaults. This is not enough to hit the 14–18% tranche directly, but it is enough to leave the remaining eight-year tranche cash flows much riskier. The model therefore revalues the remaining premium and protection legs under the bad observed credit state. In this path, the conditional premium annuity is 7.0619, the fair spread is 37.26 bps, and the conditional premium leg is approximately

$$0.003726 \times 7.0619 \approx 0.0263.$$

The conditional protection leg is 0.4875, so the conditional tranche value is approximately

$$0.0263 - 0.4875 = -0.4612.$$

This equals the endogenous mark-to-market loss in the quantile-path decomposition. The 60.56% final VaR is therefore mostly caused by a bad observed one-year credit state, which makes the remaining eight-year expected tranche loss much larger.
The expected total tranche loss by future payment date in this conditional path is shown in Table 20.

Table 20: Conditional expected total tranche loss by future payment date for the B-rated 14–18% tranche, 9-year quantile path.

Payment date	Conditional expected total tranche loss
2y	0.0000000
3y	0.00000035
4y	0.0001395
5y	0.0042156
6y	0.0337663
7y	0.1250109
8y	0.2875337
9y	0.4874644

10.6 Effect of conditional repricing

The baseline results include conditional repricing. This means that after the one-year observation horizon, the remaining tranche value is updated conditional on the realized portfolio loss, realized tranche loss, surviving pool and realized common factor.

Conditional repricing is most relevant for mezzanine-style tranches. A mezzanine tranche can have zero one-year loss VaR but still have a large final VaR if the observed one-year credit state materially worsens the distribution of future losses. The B-rated 14–18% 9-year case illustrates this mechanism. The tranche is not directly hit at the one-year 97% loss quantile, but the bad quantile path has 20 defaults and a poor macro factor. The conditional repricing layer then projects a substantially higher expected future tranche loss over the remaining eight years, producing a large negative endogenous mark-to-market component.

The senior 60–100% tranche is less exposed to direct loss, but it still has positive final VaR because the return distribution includes exogenous spread-driven mark-to-market risk and conditional repricing. This is why senior tranches can have low realized loss probability over one year while still having material market-value risk.

10.7 Regulatory-style interpretation of model outputs

The model is not intended to calculate Ilmarinen's full regulatory capital requirement directly. Instead, it produces CDO/CLO-specific risk and return inputs that can be interpreted in a solvency-style framework. The simplified table provides a model-implied one-year downside risk measure for each rating, tranche and duration bucket:

$$S^{\mathrm{CDO}}_{j,\mathrm{proxy}} = \texttt{final\_var\_97\_1y}.$$

This proxy differs from a pure realized-loss assumption because it may include mark-to-market losses as well as realized tranche losses. This is appropriate for an economic one-year risk measure, since the value of a CDO tranche can decline before the tranche suffers principal write-down. In the final table, the C/Unrated 0–10% cells reach 100%, while the largest mezzanine value is the C/Unrated 14–18% 9-year cell at 98.77%.

The expected one-year holding-period return can similarly be interpreted as a model-implied expected return input. The final mapping into Ilmarinen's full solvency framework requires portfolio-level exposure information, client-specific calibration and a decision on how market-value effects are translated into the regulatory loss-assumption format.

10.8 Interpretation of current results

The final results support four main conclusions. First, tranche seniority is a dominant driver of risk. Junior tranches are exposed to the first portfolio losses and therefore show large final VaR values in the weaker rating buckets. This is clearest in the B and C/Unrated 0–10% tranche results.

Second, spread duration and remaining maturity matter most when the tranche is not already dominated by direct realized loss. Investment-grade mezzanine and senior tranches show increasing final VaR as duration increases, because the value of their remaining cash flows is more sensitive to spread shocks and conditional repricing. This is visible in the B-rated 14–18% tranche, where final VaR rises sharply at the 9-year duration point.

Third, the final VaR metric is more informative than a pure loss VaR for this project scope. Several mezzanine and senior tranches have zero one-year loss VaR but still have positive final VaR. This reflects the economic fact that a tranche can lose market value because of spread widening or subordination erosion even before suffering realized principal loss.

Fourth, the distressed C/Unrated cases require careful interpretation. The 0–10% tranche reaches 100% final VaR across all duration buckets because the final VaR source is tranche loss under the thin premium-annuity fallback. The C/Unrated 14–18% and 60–100% tranches remain return-source cases, with high final VaR values driven by the simulated holding-period return distribution.

The results are model-based proxy outputs rather than market-calibrated tranche prices. The model uses simplified assumptions for default correlation, recovery, risk-neutral valuation and spread dynamics. For material use, the scenario matrix can be rerun with client-specific calibration inputs, repeated random seeds and explicit confidence intervals for the reported tail estimates.

11 Limitations and future work

The objective of the project was to build a defensible model that captures the main risk drivers of CDO/CLO tranches and can be implemented in a practical Excel proxy. This required a balance between realism, interpretability and computational feasibility. The resulting framework should therefore be interpreted as a structured-credit risk proxy.

The first limitation is the one-factor Gaussian copula dependence structure. The model links obligor defaults through a single common systematic factor, which makes the model tractable and easy to explain. However, real CDO and CLO portfolios may be affected by several sectoral, regional and macroeconomic factors. A single factor cannot fully capture these sources of dependence. In addition, the default-correlation parameter ρ is fixed in the baseline model and tested only through sensitivity cases. Since tranche risk is highly sensitive to correlation assumptions, especially for mezzanine and senior tranches, this remains an important source of model risk.

The second limitation concerns calibration. The model uses public rating-level default probabilities, a stylized stochastic-recovery specification and rating-based spread-volatility inputs. These assumptions are suitable for a transparent project-level proxy, but they are not equivalent to instrument-specific market calibration. In particular, the model does not calibrate to observed tranche prices, market-implied base correlations, liquidity premia or Ilmarinen-specific portfolio marks. Therefore, the results should be interpreted as model-based risk estimates rather than market prices.

The fair-spread calculation improves internal consistency, because the running spread used for premium carry is solved from the same risk-neutral loss distribution as the premium and protection legs. However, the resulting fair spread is still model-implied rather than market-observed. Very large fair spreads may occur when the risky premium annuity is small. Such cases are useful diagnostics of distressed or weakly quoteable running-spread states, but they should not be interpreted as tradable coupon estimates.

The third limitation is the treatment of spread risk. The model does not assign the contractual running spread directly from a broad rating-spread table. Instead, the running spread is computed as a scenario-specific fair spread. Rating information is still used in the exogenous spread-risk overlay through the rating-based one-year spread-volatility parameter $\sigma_s(r)$. This captures a simplified form of market spread uncertainty, but it is not a complete dynamic model of tranche spreads, liquidity premia or market-implied risk premia. The exogenous spread shock should therefore be understood as a spread-risk sensitivity layer.

The fourth limitation is the treatment of recovery and conditional repricing. Stochastic recovery is included through a Beta distribution, which improves realism compared with a fixed recovery assumption. However, the recovery parameters are not estimated from loan-level recovery data. The conditional repricing layer captures subordination erosion and updated expected future losses, but it remains simplified. It does not fully model dynamic base-correlation changes, liquidity effects, market-implied risk premia or investor-specific funding and accounting effects.

The fifth limitation is data granularity. The implementation can support finite and heterogeneous portfolios, but the final proxy is built on representative rating, tranche and spread-duration buckets. Without full look-through data, the model cannot capture name-level concentration risk, sector concentrations, manager effects, collateral-quality tests or idiosyncratic exposure differences. Two portfolios with the same average rating may therefore have different true risk profiles if their underlying loan compositions differ.

The sixth limitation is the Excel proxy. The workbook is designed to make the Python model usable in a simple client-facing format. Exact grid points correspond to Python-generated scenario outputs, but off-grid results rely on lookup or interpolation. Consequently, the Excel tool should be used primarily within the generated scenario grid. Material decisions or unusual tranche structures should be analysed by rerunning the Python model with the relevant assumptions rather than relying only on Excel interpolation.

The seventh limitation concerns simulation precision. The simplified 4-by-9 final VaR matrix is based on the current validation grid. The convergence test shows that increasing the number of simulations stabilizes the B-rated 14–18% tranche result, but tail estimates for thin tranches remain noisy. For production use, the full scenario matrix should be rerun with a larger simulation count and, preferably, with repeated seeds or confidence intervals for the reported VaR estimates.

Future work should therefore focus on calibration, data granularity and numerical robustness. The most useful extensions would be:

- calibrating PD, recovery and spread-volatility assumptions to Ilmarinen's internal data or preferred market sources;
- testing the model on real or representative CDO/CLO portfolios;
- rerunning the full simplified matrix with a larger simulation count;
- comparing Excel-interpolated outputs against direct Python model runs;
- introducing sector, region or asset-class factors if look-through data are available;
- testing alternative copula, correlation, recovery and spread-risk assumptions;
- validating fair-spread diagnostics against observed tranche quotes where available;

- validating final VaR outputs against historical tranche performance or market stress scenarios.

Overall, the main limitation is not the absence of more complex modelling techniques, but the need to keep the final tool transparent, auditable and practical. The current framework should therefore be viewed as an extendable baseline for structured-credit risk assessment.

12 Conclusion

This project developed a simplified but scientifically defensible model for assessing the risk and return of CDO/CLO tranches for Ilmarinen. The model is based on a one-factor Gaussian copula framework, where rating-based default probabilities are mapped into correlated default scenarios. Portfolio losses are then allocated to tranches through attachment and detachment points, which captures the nonlinear waterfall structure of CDO risk.

The model extends a pure static loss engine in four important ways. First, it allows for a finite pool representation rather than relying only on an infinitely diversified homogeneous portfolio. Second, it includes stochastic recovery, so default severity can vary across scenarios. Third, it includes conditional repricing after the one-year observation horizon. This repricing layer captures the effect that defaults in junior parts of the structure can reduce the market value of mezzanine or senior tranches even before those tranches suffer realized principal losses. Fourth, the model computes a scenario-specific fair running spread from the premium-leg and protection-leg valuation rather than assigning the running spread directly from a broad rating-spread table.

The final implementation consists of a Python model engine and an Excel proxy workbook. Python performs the simulation, loss allocation, fair-spread calculation, conditional repricing and final risk metric selection. Excel provides a practical input-output interface where the user can select rating, spread duration, attachment and detachment. The main simplified reporting metric is the Python-selected one-year final VaR,

`final_var_97_1y`.

The model also exports

`final_var_source` and `final_var_fallback_reason`,

so that the reported final VaR can be traced either to the simulated return distribution or to an explicitly identified fallback rule.

The current results show that tranche seniority, rating quality and spread duration are the main drivers of the simplified final VaR table. Junior tranches are dominated by direct credit losses, while mezzanine and senior tranches can also be materially affected by spread risk and conditional repricing. The validation results support the consistency of the model: weaker ratings increase portfolio loss, lower recovery increases loss severity, higher correlation changes tail dependence and conditional repricing increases return-based risk for mezzanine tranches.

The fair-spread diagnostics provide an additional valuation check. Regular scenarios produce interpretable model-implied running spreads, while distressed or nearly exhausted tranches may produce very large fair spreads because the risky premium annuity is small. These cases are not hidden by capping the raw fair spread. Instead, they are identified through the premium-annuity, fair-spread and final-VaR-source diagnostics.

The model is not intended to calculate Ilmarinen's full regulatory capital requirement or to replace market-standard pricing systems. Instead, it provides a transparent structured- credit risk framework that produces CDO/CLO-specific risk and return estimates in a format that can be connected to solvency-style analysis. With additional client data and calibration inputs, the framework can be refined further while preserving its main advantage, which is a clear link between model assumptions, tranche structure, fair spread, conditional repricing and reported final VaR.

13 Reflection

All in all, we feel that the project met the goals set at the start of course. First and foremost, each group member participated actively and attended close to every single team meeting, client catch-up and excursion. Thus, the exchange of information and new ideas was efficient and we were able to iterate initial solutions quickly. We also feel that the team members were always willing to do a bit more than expected to ensure the team's success instead of offering a minimum effort required. The initial project plan and timetable also held, and we were able to complete the project in most parts couple weeks before the final deadline, leaving space for final touches and details.

On the other hand, we found that efficient teamwork with the technical implementation would have benefited from a version control system instead of building the risk engine and running simulations on one team members computer. Thus, the workload regarding the implementation part clustered around one member but this issue was recognized and transparently discussed - this arrangement was found still the most efficient one and the individual also emphasized that this was an acceptable way of working. If we could start the implementation from scratch, the optimal workflow would have been setting a version control at the beginning, iterating the code together and every team member would have been able to run their own simulations.

This issue only concerned the technical implementation, and the workload was equally divided when it comes to contributing in meetings, presenting results at excursions and writing the deliverables (project plan, interim report) during the course. At the final reflections after our last internal meeting, whole team felt that other members were easy to work with, the learning goals for the course were reached and the project topic sparked genuine interest to solve a real issue presented by Ilmarinen.

References

[1] Leif Andersen and Jakob Sidenius. Extensions to the gaussian copula: Random recovery and random factor loadings. *Journal of Credit Risk*, 1(1):29–70, 2004/2005.

[2] Basel Committee on Banking Supervision. International convergence of capital measurement and capital standards: a revised framework. https://www.bis.org/publ/bcbs128.pdf. Accessed: 2026-03-15.

[3] Petra Buzková and Petr Teplý. Collateralized debt obligations' valuation using the one factor gaussian copula model. *Prague Economic Papers*, (1), 2012.

[4] Areski Cousin and Jean-Paul Laurent. An overview of factor models for pricing cdo tranches, January 2008. Working paper, 14 January 2008.

[5] Eduskunta. Laki eläkelaitoksen vakavaraisuusrajan laskemisesta ja vastuuvelan kattamisesta 315/2015. https://www.finlex.fi/fi/lainsaadanto/2015/315, 2015. Accessed: 2026-02-23.

[6] Abel Elizalde. Credit risk models iv: Understanding and pricing cdos. Working Paper 0608, CEMFI, April 2006.

[7] Ingo Fender and Janet Mitchell. Cdo rating methodology: Some thoughts on model risk and its implications. Technical report, Bank for International Settlements, 2005.

[8] John Hull and Alan White. Valuation of a cdo and an nth to default cds without monte carlo simulation. *Journal of Derivatives*, 12(2), 2004. Winter 2004.

[9] John Hull and Alan White. Valuing credit derivatives using an implied copula approach. *Journal of Derivatives*, 2006.

[10] John Hull and Alan White. Dynamic models of portfolio credit risk: A simplified approach. *Journal of Derivatives*, 15(4):9–28, 2008.

[11] Sunju Hwang and Jinsub Kim. Value at risk of portfolios using copulas. *Journal of Information and Communication Convergence Engineering*, 2021.

[12] Jean-Paul Laurent and Jon Gregory. Basket default swaps, cdo's and factor copulas, September 2003. Working paper.

[13] David X. Li. On default correlation: A copula function approach. Working Paper 99-07, The RiskMetrics Group, 44 Wall St., New York, NY 10005, April 2000. This draft: April 2000; first draft: September 1999.

[14] Sosiaali- ja terveysministeriö. Sosiaali- ja terveysministeriön asetus eläkelaitoksen vakavaraisuusrajan laskemisesta 447/2015. https://www.finlex.fi/fi/lainsaadanto/2015/447, 2015. Accessed: 2026-02-23.

[15] S&P Global Ratings. Default, transition, and recovery: 2025 annual global corporate default and rating transition study, 2026. Accessed: 2026-03-26.

[16] Nikola A Tarashev and Haibin Zhu. Credit fundamentals, ratings and value-at-risk: Cdos versus corporate exposures. Technical report, Bank for International Settlements, 2008.

[17] Yang Yi. Collateralized debt obligation (cdo) modelling and its model comparisons. BMI paper, Vrije Universiteit.