Project plan - Ilmarinen
MS-E2177 - Seminar on Case Studies in Operations Research

Terho Mutikainen, Aaro Niinistö, Tuukka Pallonen, Panu Salonen

April 8, 2026

1 Background

Ilmarinen is a Finnish pension insurance company responsible for collecting pension insurance payments, investing those funds and paying out pensions for retirees. The managed investments are approximately 70 billion euros which are diversified into various types of securities - such as equity, real estate and bonds. Due to the nature of managed funds, the risk management of this portfolio is crucial. Thus, the allowed risks related to the investments are limited by legislation and is monitored by Finnish Financial Supervisory Authority (Finanssivalvonta).

Furthermore, Finnish Solvency Law (315/2015) provides a framework for setting a solvency limit for pension insurance companies to ensure the ability to pay out pensions even in a volatile or unfavorable market conditions. However, the solvency framework is more suitable for traditional investments such as equity or real estate. On the other hand, determining the solvency limit for a investment portfolio containing more complex instruments such as derivatives is not well-defined.

One example of these instruments are collateralized debt obligations (CDOs). CDOs are instruments which consist of a varying but usually large amount of different underlying debts - such as credit card loans, mortgages or car loans. However, the investors are not equal when it comes to receiving cash flows from these underlying assets: Investor can choose whether to invest in the senior, mezzanine or equity tranch. If some underlying loans default, the senior tranch will receive their cash flows first, then mezzanine and finally equity. Thus, senior tranch provides low risk but lower profit while equity tranch faces additional risk but higher profits.

2 Objectives

All in all, our objective for this project is to provide a comprehensive model to determine the risks associated with CDO investments. Concretely the client has requested that the model should output i) 97 % 1-year VaR of tranche loss and ii) expected return for a given CDO tranche. The model uses attachment/detachment, underlying quality (rating or rating mix), maturity/duration and possibly the number of loans as inputs. Developing such a risk model is necessary since CDOs face risk in various different risk classes (default risk, interest rate risk etc.) covered by solvency legislation. Thus, a different approach is necessary compared to the "traditional" investments more applicable to risk assessment purely based on solvency legislation. The model is going to be developed based on relevant literature.

The core risk engine will be based on Vasicek's asymptotic single risk factor (ASRF) / one-factor Gaussian copula model, with Monte Carlo simulation (and/or large-pool approximations) used to compute the tranche loss distribution and the requested outputs. Still, the key requirement from client is that the model must be simplified after development. In other words, end users conducting risk assessment wish to use a spreadsheet table or similar interface to communicate model outputs internally or to the Finnish Financial Supervisory Authority.

3 Tasks

3.a Scoping the problem and literature review

First, the aim is to familiarize ourself with the problem. We will get to know the relevant concepts like CDOs and study relevant legislation like the Finnish solvency law. This will be done in group meetings and individually by each group member. This will improve our understanding of the aims of the project. After the basic understanding, the aim is to dig deeper into following subjects: 1. Research the different credit rating systems and understanding the risk they define. 2. Familiarize ourselves with different ways of modeling combination of risky assets like Copula and American Monte-Carlo.

3.a.1 Research the different credit rating systems and understanding the risk they define

To model CDOs risk distribution we need to have information on the marginal risk distributions of the underlying assets and the covariance of those risks. The former links very strongly with credit ratings and we need to research how the different ratings can be transformed into two parameters: probability of default per year and then loss given default. This is needed to start developing the risk models in the next phase.

3.a.2 Familiarize ourselves with Vasicek/ASRF one-factor Gaussian copula modelling of portfolio credit loss and tranche loss mapping

There is plenty of literature on risk modeling of CDO and similar derivatives. Thus, it is important to first study different articles and research on the project topic to understand the different ways of approaching the problem and what are their requirements and benefits. Building a comprehensive understanding of the subject before making consequential decisions on the approach mitigates the risk of going in to wrong direction early on.

3.a.3 Researching mark-to-market risk, credit spread duration, and market volatility

To accurately capture the 1-year 97 % VaR required by the client, the risk model must account for mark-to-market losses rather than exclusively focusing on realized defaults over a single-year horizon. Consequently, we will also research the impact of credit spread duration and market volatility on tranche pricing, as suggested by our client. While CLO tranches typically possess near-zero interest rate duration due to their floating-rate nature, their market prices are highly sensitive to widening credit spreads. We will study how spread duration varies across the capital structure. Specifically, how a tranche's position in the sequential cash flow waterfall dictates its effective maturity and corresponding spread duration. Because a CLO is redeemed sequentially, higher-rated senior tranches are paid off faster and have a lower spread duration, whereas lower-rated mezzanine tranches stay in the structure longer and exhibit a higher spread duration. Incorporating these duration metrics and historical spread volatility, the model will be equipped to translate simulated macroeconomic stress into tranche-specific price drops.

3.a.4 Researching stochastic recovery rates and downturn LGD

A limitation of basic credit risk models is the assumption of a static, deterministic recovery rate, artificially reducing the variance of the portfolio loss distribution and underestimates extreme tail risks. To address this, we will research the implementation of stochastic recovery rates using the Beta distribution, a standard industry tool that restricts outcomes between 0 % and 100 % and can capture the bimodal or skewed nature of corporate loan recoveries. We will explore parameterizing this distribution to target an expected 50 % mean recovery rate by setting its shape parameters, α and β, equal to each other. We will also investigate the concept of "downturn LGD" (Loss Given Default) by modeling the empirically observed negative correlation between recovery rates and default rates during macroeconomic stress events. By linking recovery severity to the same systemic factors that drive default clustering, the model will properly generate the thicker tails necessary to compute the 1-year 97 % VaR required by solvency frameworks.

3.b Formulating the problem and developing models

After we have thorough understanding of the key concepts and regulation, we develop the risk model, design and write the code for the project. Also, we meet with the client to ensure we are proceeding in a way suitable for them.

3.c Model evaluation

We will evaluate the model through verification, stability diagnostics, sensitivity analysis, and benchmarking. For verification, we implement simple unit tests that confirm the correctness of the tranche loss mapping from portfolio loss to tranche loss using attachment and detachment points, and we will validate the implementation against special-case scenarios. Starting with the case $\rho = 0$ of independent defaults we ensure the loss allocation behaves as expected.

For Monte Carlo stability, we will assess simulation error by running convergence checks and reporting uncertainty measures for both the mean loss and the 97 % one-year VaR . For sensitivity analysis, we will quantify how results change under plausible ranges for key assumptions, including PD and rating-to-PD mappings, LGD/recovery distribution parameters, the asset correlation parameter ρ, the strength of downturn LGD.

Finally, for benchmarking, we will check whether the outputs satisfy economic intuition. Specifically, higher attachment should reduce tranche loss risk and VaR, thinner tranches should exhibit higher sensitivity (implicit leverage), and higher PD/LGD/ρ should increase tail loss. These evaluation steps will be used to identify model weaknesses and document the main drivers of tranche risk for Ilmarinen's intended use case.

3.d Model simplification

In the last phase, we convert the model into the table format requested by the client. At this point we have our third meeting with the client in other to make sure that the final product is what they were asking for. If some changes need to be made, earlier phases may be revisited.

3.e Deliverables

In addition to the steps, we work on the three deliverables made during the course: project plan, interim report and the final report. Also, we prepare presentations for the client and the course staff on these three deliverables.

4 Schedule

The project timeline (in Figure 1) is structured around three major course deliverables: the project plan, the interim report, and the final report. Before presenting the interim and final reports to the course staff, the team will meet with the client, Ilmarinen, to clarify the project direction and discuss the specifics of the model.

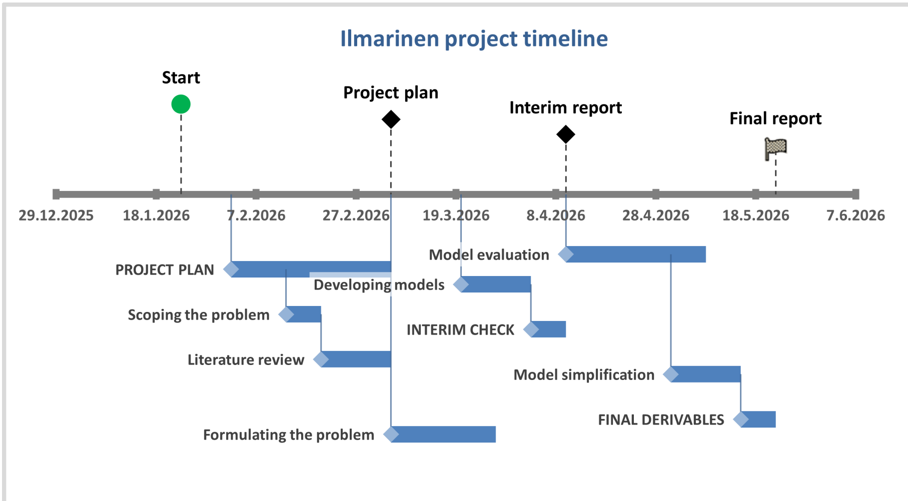

Figure 1: Project timeline

The project begins in mid-January 2026. The initial weeks are dedicated to scoping the problem within the context of Ilmarinen's extensive investment portfolio and familiarizing the team with the current Finnish solvency law requirements. Concurrently, the team conducts a thorough literature review focusing on Collateralized Debt Obligations, researching how these instruments bundle loans and slice them into various risk-based tranches, specifically the Senior, Mezzanine, and Equity tranches. This foundational research phase ends with the submission of the project plan on March 3, 2026.

Following the project plan approval, the team formulates the problem and begins developing the specific risk model. Throughout March, the focus is on ensuring the model successfully incorporates key variables such as underlying debt quality, duration, and tranche level. An internal interim check is conducted to evaluate the model's confidence levels against other asset-specific stresses in the existing solvency law. During this phase, the team also begins comparing initial results against established regulatory frameworks like Solvency II. This development phase wraps up with the interim report milestone on April 8, 2026.

The final phase kicks off in early April with rigorous model evaluation, testing the developed framework to ensure accuracy and reliability. A portion of May is dedicated to the model convertion task. Because the current risk presentation framework is quite simple, the team must approximate the complex model's results into a user-friendly format, such as a table or rule set. This step is vital to ensure the new methodology matches the simplicity of the other methods currently dictated by the solvency law. The project concludes with the packaging of all final deliverables and the submission of the final report on May 18, 2026.

5 Resources

Our project team consists of four M.Sc. students at Aalto University majoring in Mathematics and Operations Research. The team has a strong foundation in mathematical modeling, stochastic processes, and even financial engineering, supplemented by the programming skills essential for simulations. We are primarily supported by our client contact Joonas Lanne, Development Manager (Investment Risk Management) at Ilmarinen, who provides industry insights and ensures the model's practical relevance. Our team is also supported by our academic supervisor, Professor Ahti Salo, from the Department of Mathematics and Systems Analysis, who provides methodological guidance and academic oversight.

The project relies on public regulatory frameworks and possibly proprietary client data from Ilmarinen. We will design the model to work under two data regimes specified by the client: i) loan-level rating/PD inputs when available, and ii) tranche-level + portfolio-summary inputs when look-through is limited. We might also receive some proprietary data for testing our final model. The technical parameters for the model are derived from the Finnish Solvency Law (315/2015), Government Decree 447/2015.

There is much of academic research on modeling of CDO risks. Key academic sources could include David X. Li's (2000) seminal work on copula functions and papers related to Vasicek's Gaussian copula model utilization in credit risk context.

All computational modeling and documentation are performed using software tools. The risk engine is developed in Python and the scikit-learn library can be used to implement the regression-based proxy models. GitHub is used for version control and collaborative coding, while Overleaf is utilized for LaTeX-based documentation.

6 Risks

Table 1: Identified risks and corresponding mitigation actions.

Risk	Likelihood	Impact	Effect	Mitigation
Poor internal communication	Medium	Medium	Inefficient task allocation and possible misalignment of technical models.	Regular internal meetings and mandatory status updates via instant messaging.
Poor communication and collaboration with Ilmarinen	Low	High	Final risk model does not meet the specific regulatory or practical needs of the client.	Enough meetings with client contacts to validate assumptions and demonstrate progress early.
Model unsuitable for modelling risks of a CDO	Medium	High	For example, chosen copula model fails to capture tail risks effectively during stress testing.	Need to perform a comparative analysis between copulas to identify the most robust approach.
Complex model cannot be simplified	Medium	Medium	The resulting approximating model (table or rule set) is too inaccurate for daily calculations.	We can utilize the least squares monte carlo framework to fit polynomial functions with high-order validation.
Poor documentation and black box nature of the final model	Low	High	Ilmarinen cannot audit, maintain, or update the model parameters after project completion.	Maintain a clear, commented codebase and provide a detailed technical report and guide.
Inadequate personal task management	Medium	Medium	Project milestones and final deadlines are missed due to competing academic workloads.	Establish clear internal milestones and a clear task dependency map to monitor weekly progress.