DELFT INSTITUTE OF APPLIED MATHEMATICS

Quantifying the Impact of Physical Climate Risk on ABN AMRO's Mortgage Portfolio

Yfke de Heus
4977505

Supervisor
Prof.dr.ir. C. Vuik

September 19
2025

Quantifying the Impact of Physical Climate Risk on ABN AMRO's Mortgage Portfolio

by

Yfke de Heus

Master Thesis

MSc Applied Mathematics - Financial Engineering

Professor:	Prof.dr.ir. C. Vuik, Delft Institute of Applied Mathematics, TU Delft
Supervision:	Riccardo Peli, ABN AMRO Bank
Supervision:	Bart Bakker Schut, Accenture
Date:	October 17, 2025
Faculty:	Applied Mathematics, TU Delft

Abstract

Climate risks are becoming an increasingly important element of credit risk modelling, both because physical hazards such as floods or earthquakes can severely damage residential properties and the probability of such events occurring is rising globally. Additionally, recent regulatory guidelines require banks to assess and manage climate and environmental (CER) risks, as they can affect both mortgage borrowers and lenders. This research examines the current climate risk framework of the Dutch bank ABN AMRO and proposes a methodology to address one of its limitations, which assumes that climate damages are reflected only in collateral values and do not affect the creditworthiness of the associated mortgages, thereby placing the entire impact on the Loss Given Default (LGD). After outlining the theoretical background in both credit and climate risk, and analysing the bank's three-step climate risk quantification process, a behavioural model extension is introduced that captures how borrowers may respond to climate-related physical damage. The approach distinguishes four possible borrower actions: no repair of climate-related damage, self-funded repair, repair financed through an additional ABN AMRO loan, or repair financed through an external loan. These pathways link climate damage to credit risk parameters through collateral value, disposable income and loan exposure, with behavioural probabilities—calibrated using expert input—that ultimately affect not only the Loss Given Default (LGD), but also the Probability of Default (PD) and Exposure at Default (EAD). Results show that including behavioural choices could increase estimated climate impacts on Risk-Weighted Assets (RWA) and Expected Loss (EL) by nearly a factor of two. At mortgage portfolio level, the impacts remain well below 1% and do not exceed the bank's materiality thresholds, which is consistent with the low probability of extreme climate events in the Netherlands. Overall, the findings highlight that extending the current framework with behavioural choices can substantially alter outcomes, underscoring the importance of underlying assumptions. The proposed framework contributes to an improved impact estimation of physical climate risk on the mortgage portfolio, offering a foundation for further research and potential application in regions where climate risks are more severe.

Preface

This master thesis was conducted under the supervision of *Delft University of Technology*, in collaboration with the companies *Accenture* and *ABN AMRO*. Each of these parties has contributed in their own way to this final product, and their involvement has been essential to the completion of this work.

I started my graduation internship at Accenture in December 2024 with a broad interest in risk management, but without a predefined thesis topic. I was eager to explore the mathematical foundations of financial risk models and to gain insight into the challenges encountered in practice. During this initial stage, which I often refer to as the 'swimming phase', Accenture provided me with great support. Numerous meetings were arranged with different departments, gradually shaping my focus toward the integration of climate risk into credit risk modelling. I am especially grateful to Bart Bakker Schut for never dismissing any question as trivial, for always making the effort to connect me with the right people, and for allowing me the time and space to develop my research in a way that suited me.

By February 2025, after a first round of literature review, it became clear that my research would focus on the integration of physical climate risks into credit risk models, specifically in the context of residential mortgages. I soon realized that working with real-world data would make the research far more relevant and impactful, and I was therefore very pleased that ABN AMRO welcomed me to carry out my work within the bank. The engaged and supportive team provided a pleasant environment for conducting this research. My special thanks go to Riccardo Peli, whose continuous involvement, critical feedback, and collaborative spirit made this project both stronger and more enjoyable. His support in brainstorming ideas, reviewing drafts and evaluating each step has been truly invaluable.

I would also like to express my sincere gratitude to my TU Delft supervisor, Prof. Kees Vuik, who has guided me throughout the entire process. Having also supervised my bachelor thesis, he once again provided clarity and direction whenever needed. He was always open to my new ideas, and whenever I was stuck he went out of his way to connect me with the right people so I could move forward. In our weekly meetings, I greatly appreciated his advice on both theoretical and practical issues, as well as his ability to help me refocus on the bigger picture, which always left me reassured and with a clear way forward.

Finally, I wish to thank my family, partner and friends for their constant encouragement, for reading sections of my work and for being there throughout this journey.

I hope that this thesis proves to be of value and relevance, and that it may serve as a foundation for further research both within and beyond the bank. I wish you a pleasant read.

Yfke de Heus

Delft, September 2025

List of Symbols

Symbol	Definition
K	Capital requirement
L	Loss variable
VaR	Value at Risk
α	Confidence level
$C\text{-}VaR$	Credit Value at Risk
Φ	CDF of standard Gaussian distribution
Φ^{-1}	Quantile function of standard Gaussian distribution
ρ	Systematic risk correlation parameter
$VaR_{\alpha}(PD)$	Value-at-Risk of the default probability at confidence level α
PD_{α}	Critical value of PD at confidence level α (Worst-Case Default Rate (WCDR) or downturn PD)
A_t	Market value of a firm's assets
D	Debt level (default threshold)
A_0	Current asset value
σ_A	Asset volatility
T	Maturity
ε	Idiosyncratic factor
Δr	Difference in return between receiving payments as originally scheduled and receiving them with delay
P_i^{flood}	Probability of a flood event affecting building i
$D_{i,\text{flood}}^{\text{physical}}$	Estimated damage factor (physical)
$M_{i,\text{flood}}$	Mitigation or insurance factor
$D_{i,\text{flood}}^{\text{perception}}$	Damage function (flood perception risk)
$P_i^{\text{foundation}}$	Annual probability of foundation damage
$D_{i,\text{foundation}}$	Damage function (foundation risk)
$P_i^{\text{earthquake}}$	Annual probability of a damaging natural earthquake in the relevant region
$D_{i,\text{earthquake}}$	Damage function (earthquake risk)
C	Creditworthiness variable
f	Function mapping the creditworthiness proxy to predicted PD
$\widehat{\text{PD}}_{i,\text{non-CER}}$	Predicted default probability for borrower i (based on non-CER-adjusted creditworthiness)
$\widehat{\text{PD}}_{i,\text{CER}}$	Predicted default probability for borrower i (based on CER-adjusted creditworthiness)
$V_{i,\text{CER}}^{\text{no repair}}$	CER-adjusted collateral value for borrower i (no-repair path)
$LTV_{i,\text{CER}}^{\text{no repair}}$	CER-adjusted Loan-to-Value ratio for borrower i (no-repair path)
g	Function mapping the LTV to LGD (representing the LGD model)
$\text{LGD}_{i,\text{CER}}^{\text{no repair}}$	CER-adjusted LGD for borrower i (no repair path)
$\text{LGD}_{i,E(\text{CER})}^{\text{no repair}}$	Expected CER-adjusted LGD for borrower i (no repair path)
$C_{i,\text{CER}}^{\text{self}}$	CER-adjusted creditworthiness proxy for borrower i (self-funded path)
$LTC_{i,\text{CER}}^{\text{self}}$	CER-adjusted Loan-to-creditworthiness ratio for borrower i (self-funded path)
$\lambda_{i,\text{CER}}^{\text{self}}$	Relative CER impact factor for borrower i (self-funded path)

Symbol	Definition
$\mathrm{PD}_{i,\mathrm{CER}}^{\mathrm{self}}$	CER-adjusted PD for borrower i (self-funded path)
$\mathrm{PD}_{i,E(\mathrm{CER})}^{\mathrm{self}}$	Expected CER-adjusted PD for borrower i (self-funded path)
$C_{i,\mathrm{CER}}^{\mathrm{loan\ AAB}}$	CER-adjusted creditworthiness proxy for borrower i (AAB loan path)
$LTC_{i,\mathrm{CER}}^{\mathrm{loan\ AAB}}$	CER-adjusted Loan-to-Creditworthiness (LTC) ratio for borrower i (AAB loan path)
$\mathrm{PD}_{i,\mathrm{CER}}^{\mathrm{loan\ AAB}}$	CER-adjusted PD for borrower i (AAB loan path)
$\lambda_{i,\mathrm{CER}}^{\mathrm{loan\ AAB}}$	Relative CER impact factor for borrower i (AAB loan path)
$\mathrm{PD}_{i,E(\mathrm{CER})}^{\mathrm{loan\ AAB}}$	Expected CER-adjusted PD for borrower i (AAB loan path)
$\mathrm{EAD}_{i,\mathrm{CER}}^{\mathrm{loan\ AAB}}$	CER-adjusted EAD for borrower i (AAB loan path)
$\mathrm{EAD}_{i,E(\mathrm{CER})}^{\mathrm{loan\ AAB}}$	Expected CER-adjusted EAD for borrower i (AAB loan path)
$V_{i,\mathrm{CER}}^{\mathrm{repair}}$	CER-adjusted collateral value for borrower i (repair path)
$LTV_{i,\mathrm{CER}}^{\mathrm{repair}}$	CER-adjusted Loan-to-Value ratio for borrower i (repair path)
$\mathrm{LGD}_{i,\mathrm{CER}}^{\mathrm{loan\ AAB}}$	CER-adjusted LGD for borrower i (AAB loan path)
$\mathrm{LGD}_{i,E(\mathrm{CER})}^{\mathrm{loan\ AAB}}$	Expected CER-adjusted LGD for borrower i (AAB loan path)
$k^{\mathrm{loan\ ext}}$	Scaling factor for the annual creditworthiness impact of an external repair loan
$C_{i,\mathrm{CER}}^{\mathrm{loan\ ext}}$	CER-adjusted creditworthiness proxy for borrower i (external loan path)
$LTC_{i,\mathrm{CER}}^{\mathrm{loan\ ext}}$	ER-adjusted Loan-to-Creditworthiness (LTC) ratio for borrower i (external loan path)
$\mathrm{PD}_{i,\mathrm{CER}}^{\mathrm{loan\ ext}}$	CER-adjusted PD for borrower i (external loan path)
$\lambda_{i,\mathrm{CER}}^{\mathrm{loan\ ext}}$	Relative CER impact factor for borrower i (external loan path)
$\mathrm{PD}_{i,E(\mathrm{CER})}^{\mathrm{loan\ ext}}$	Expected CER-adjusted PD for borrower i (external loan path)
$M_{i,\mathrm{CER}}$	Borrower-level CER-adjustment factor
$\overline{M}$	Portfolio-level CER-adjustment factor
$P_{\mathrm{no\ repair}}$	Probability that a borrower chooses not to repair CER-related damage
P_{repair}	Probability that a borrower chooses to repair CER-related damage
$P_{\mathrm{self}\vert\mathrm{repair}}$	Conditional probability that a borrower self-funds CER-related repair costs
$P_{\mathrm{loan\ AAB}\vert\mathrm{repair}}$	Conditional probability that a borrower finances CER-related repair costs through an additional loan from ABN AMRO
$P_{\mathrm{loan\ ext}\vert\mathrm{repair}}$	Conditional probability that a borrower finances CER-related repair costs through an external loan
$M_{i,\mathrm{PD},CER}$	Borrower-level CER-adjustment factor for PD
$M_{i,\mathrm{LGD},CER}$	Borrower-level CER-adjustment factor for LGD
$M_{i,\mathrm{EAD},CER}$	Borrower-level CER-adjustment factor for EAD
$\overline{M}_{\mathrm{PD}}$	Portfolio-level CER-adjustment factor for PD
$\overline{M}_{\mathrm{LGD}}$	Portfolio-level CER-adjustment factor for LGD
$\overline{M}_{\mathrm{EAD}}$	Portfolio-level CER-adjustment factor for EAD
$\mathrm{PD}_{\mathrm{CER}}$	CER-adjusted portfolio-level PD
$\mathrm{LGD}_{\mathrm{CER}}$	CER-adjusted portfolio-level LGD
$\mathrm{EAD}_{\mathrm{CER}}$	CER-adjusted portfolio-level EAD
I_0	Income at loan application
r_i	Annual income adjustment factor
π	Inflation correction factor
χ_e	Age boundary (early-career segment)

Symbol	Definition
r_e	Annual income adjustment factor (early-career segment)
χ_m	Age boundary (mid-career segment)
r_m	Annual income adjustment factor (mid-career segment)
χ_l	Age boundary (late-career segment)
r_l	Yearly rate (late-career segment)
χ_p	Age boundary (post-retirement segment)
r_p	Yearly factor (post-retirement segment)
s	Snapshot date: the reporting point when loan and borrower data are recorded
a_s	Age at snapshot date
a_0	Age at loan origination
I_s	Estimated income at the snapshot date
y_e, y_m, y_l, y_p	Number of years spent in the early-, mid-, late-, and post-career segments respectively
L_s	Outstanding loan amount on snapshot date
$\widehat{\text{LTI}}_s$	estimated LTI ratio on snapshot date
$\widehat{\text{LTI}}$	Estimated Loan-to-Income (LTI) ratio
y_{logit}	Logit-transformed probability of default
$\widehat{\text{PD}}$	PD estimate
X_1	Estimated Loan-to-Income (LTI) input variable
β_1	Estimated Loan-to-Income (LTI) coefficient
β_0	Estimated Loan-to-Income (LTI) intercept
$y_{\text{logit, base}}$	Logit-transformed probability of default, base model
X_2	Terms in arrears variable
X_3	BKR variable
β_2	Terms in arrears coefficient
$\beta_{3,i}$	BKR coefficient
$\mathbb{1}\{X_3 = i\}$	Binary indicator BKR variables per category
τ	Kendall's τ
γ	Goodman–Kruskal's γ
n_c	Number of concordant pairs
n_d	Number of discordant pairs
$\widehat{\text{PD}}_i$	Predicted PD for borrower i
PD_i	True PD for borrower i
$\lambda_{i,\text{CER}}^{\text{behaviour}}$	Relative CER impact factor for borrower i (general behaviour)
$max\ damage_t$	Maximum structural damage per m²
ϕ_p^S	Fraction of a property's collateral value lost due to flood damage under scenario S for property p
$property\ value_p$	Collateral value for property p (current market value)
$damage_p^S$	Total flood-induced damage to property p in scenario S
$\theta(h)_t^S$	Damage function for flood depth h
h	Inundation depth.
S	Flood scenario
t	Property type
A_p	Floor area (m^2) of property p
ι	Inflation correction factor
p	Property (index)

Symbol	Definition
LTV_i^S	LTV of a loan i under flood scenario S
LTV_i^0	Starting-point LTV (for loan i)
LGL_i	Loss Given Loss for loan i
LGL_i^S	LGL under flood scenario S (for loan i)
$sales\ ratio_p^S$	Liquidation value / market value ratio
$sales\ ratio_p^0$	Sales ratio before flooding
LGD_i^S	LGD of loan i in flood scenario S
m_{LGD}^S	Scenario-specific LGD multiplier
w_b	Weight of bank b in system exposure
w_i	Weight of loan i in bank exposure
$y_{i,b,t}$	dependent variable, representing the default status of borrower i at bank b at time t
c_b	Baseline default probability constant
$\mathbf{Z}_{i,b,t}$	Vector of primary independent variables (e.g. LTV, interest rate, GDP growth)
β'	Vector of coefficients (for Z)
$\mathbf{X}_{i,b,t}$	Vector of additional variables (control variables)
$u_{i,b,t}$	Error term
m_{PD}^S	System-wide PD multiplier for scenario S
RWA^S	Scenario-specific RWA
m_{RW}^S	Scenario-specific RWA multiplier
$RWA_{i,b}^S$	RWA for loan i of bank b under flood scenario S
$RWA_{i,b}$	Initial RWA (for loan i of bank b)
$K_{i,b}^S$	Scenario-specific capital requirement factor (for loan i of bank b)
$K_{i,b}$	Initial capital requirement factor (for loan i of bank b)
K^S	Scenario-specific capital requirement factor
$\Delta CET1\ ratio^S$	CET1 ratio change under scenario S
ΔEL^S	Difference between the expected loss under the flood scenario and the starting-point expected loss
ΔRWA^S	Scenario-specific change in risk-weighted assets
$CET1\ ratio^0$	Initial CET1 ratio
$CET1ratio^S$	CET1 ratio under flood scenario S
$CET1^S$	CET1 capital under flood scenario S
r_{DV}	Damage-to-value ratio
r_{DI}	Damage-to-income ratio
κ	Cut-off parameter controlling the repair probability threshold
ξ	Curvature parameter shaping the repair probability function
ζ	Slope parameter controlling the growth of loan-based repair probability
$q_{\text{self}\mid\text{repair}}$	Conditional probability of self-funded repair
$q_{\text{loan}\mid\text{repair}}$	Conditional probability of loan-funded repair
$g(r_{DI})$	Monotone function $1 - e^{-\zeta r_{DI}}$ mapping r_{DI} to loan-funded repair probability
π_{AAB}	Constant share of ABN AMRO within loan-based repair solutions

Table 1: List of Symbols.

List of Abbreviations

Abbreviation	Definition
AA	Appropriate Adjustments
AAB	ABN AMRO Bank
A-IRB	Advanced IRB
AUC	Area under the Curve
AUC-ROC	Area Under the Receiver Operating Characteristic Curve
BAG	Basisregistratie Adressen en Gebouwen
BCBS	Basel Committee on Banking Supervision
BIS	Bank for International Settlements
BKR	Bureau Krediet Registratie
CAR	Capital Adequacy ratio
CCF	Credit Conversion Factor
CEBS	Committee of European Banking Supervisors
CER	Climate-related and environmental risks
CET1	Common Equity Tier 1
COREP	Common Reporting
CRD	Capital Requirements Directive
CRR	Capital Requirements Regulation
CSRD	Corporate Sustainability Reporting Directive
DNB	De Nederlandsche Bank
DSTI	Debt Service-to-Income (ratio)
DTI	Debt-to-Income ratio
E&S	Environmental and Social (risks)
EAD	Exposure at Default: The amount of money that is at risk at the moment the borrower defaults
EBA	European Banking Authority
ECB	European Central Bank
EL	Expected Loss: The amount of money the bank expects to lose on a loan
EPCs	Energy Performance Certificates
ESG	Environmental, Social en Governance
F-IRB	Foundation IRB
FINREP	Financial Reporting
FN	Number of false negatives
FP	Number of false positives
FPR	False Positive rate
GEM	Global Earthquake Model
GFDRR	Global Facility for Disaster Reduction and Recovery
GHG	Greenhouse Gases
Gini	Gini coefficient
IRB	Internal Ratings-Based (approach)
LGD	Loss Given Default: The percentage of the loan that is lost if the borrower defaults, after taking into account any money the bank recovers, such as by selling collateral (e.g. a house)
LGL	Loss Given Loss

Abbreviation	Definition
LIWO	Landelijk Informatiesysteem Water en Overstromingen
LRA DR	Long-Run Average Default Rate
LTV	Loan-to-Value ratio
MSE	Mean Squared Error
MoC	Margin of Conservatism
NHG	National Mortgage Guarantee
NIBUD	Nationaal Instituut voor Budgetvoorlichting
PD	Probability of Default: The likelihood that a borrower will fail to repay their loan within a specific period, usually one year
RIi	Council for the Environment and Infrastructure
RRE	Residential Real Estate
RWA	Risk-Weighted Assets
SA	Standardised Approach
SSM	Single Supervisory Mechanism
SSM2017	Schade- en Slachtoffer Module 2017
TN	Number of true negatives
TP	Number of true positives
TPR	True Positive rate
UL	Unexpected Loss: The part of the loss that goes beyond what the bank expected
USGS	US Geological Survey
WOZ	Waardering Onroerende Zaken

Table 2: List of Abbreviations.

Contents

1

Introduction

Risk management supports the stability and integrity of the banking industry. Within this field, credit risk measures the potential losses incurred when borrowers fail to meet their financial obligations, playing a pivotal role in the process. Effective credit risk management is crucial, influencing not only the health of financial institutions but also the broader economic ecosystem.

Traditionally, credit risk in the banking sector has been viewed through a regulatory or financial lens. In recent years, climate and environmental risk has emerged as a relevant external driver that affects traditional risk categories, including credit risk. Climate risk is defined as the potential for problems for societies or ecosystems from the impacts of climate change. The assessment of climate risk is based on formal analysis of the consequences, likelihoods and responses to these impacts. This broad category can be divided into two types of risks. Physical risks arise from the direct impacts of climate-related events, such as increased flooding, storms and wildfires. Transition risks are associated with the economic adjustments required as societies shift towards low-carbon technologies and policies.

The increasing prevalence of climate risks is undeniable, translating into both physical impacts [1–5] and transitional consequences [6, 7]. There are significant implications across many sectors, including credit risk management within the Dutch banking industry. Recognizing this, regulatory frameworks have begun to adapt. The Basel Accords are the official regulatory standards issued by the Basel Committee on Banking Supervision, designed to ensure that financial institutions maintain adequate capital and follow sound risk management practices. The latest iteration of these regulations, known as Basel IV, was finalized at the beginning of 2023. These globally applicable regulations aim to establish a general framework for robust risk modelling, within which the analysis and integration of climate risks is becoming an increasingly important topic. The integration of climate and environmental risk into credit risk models is primarily being shaped at the European level. The European Banking Authority (EBA), the EU agency responsible for ensuring effective and consistent regulation and supervision across the European banking sector, plays a key role in this development. In recent years, several regulatory publications have specifically addressed how climate-related and broader ESG risks should be incorporated into the banking sector. The most recent of these is the *EBA Guidelines on ESG Risk Management (2025)* [8], which provide practical guidance on how banks are expected to measure, manage and integrate ESG risks into their internal processes. These guidelines will become binding from January 2026 onwards.

Despite these regulatory guidelines, the development and adoption of climate risk within bank's internal credit risk models are still in their early stages. Implementation varies widely among financial institutions, some are already testing and validating integration strategies, while others remain in denial about the need to adapt [9]. This difference in responses highlights that the integration of climate risk in the banking sector is still in its preliminary phase.

A particularly interesting aspect within this context is the exposure class of residential mortgages. The consequences for banks must be considered when collateral—often regarded as the bank's main security when granting a mortgage loan—is affected by physical climate-related disasters or when new climate policies significantly impact borrowers' repayment capacity. While both physical and transitional risks may affect this exposure class, physical climate risks are especially relevant due to their direct potential to damage the underlying property. Although residential mortgages represent a critical exposure class, the development of robust models that integrate climate and environmental risks remains in an early stage. This is not only due to limited data availability, but also because methodological approaches are still evolving. The scope of this thesis therefore concerns the integration of physical climate risk into credit risk models for residential mortgages, specifically focusing on potential improvements in existing quantification methods.

For several years, banks have been exploring the integration of climate-related risks into credit risk management. One of the major Dutch banks involved in this field is ABN AMRO (ABN AMRO Bank N.V.). Like any modelling journey, this process involves challenges, one of which is quantifying the impact of these risks. This thesis will support ABN AMRO in refining the assessment of climate risk within credit risk modelling for residential mortgages.

Research Questions and Scope

This thesis investigates how physical climate risks can be quantified and meaningfully integrated into internal credit risk models for retail mortgages. The project was carried out in collaboration with ABN AMRO and focuses on evaluating and improving the current modelling approach. The main research question is formulated as follows

> *Are there assumptions within ABN AMRO's current framework for physical climate risk quantification that represent a limitation, to what extent could such a limitation affect its integration into retail mortgage credit risk models, and how can this be evaluated and improved?*

To answer this central question, the following sub-questions are defined. These are structured to reflect the logical sequence of the research process, from understanding the existing setup to identifying and implementing a model improvement. Importantly, while the research was tailored to ABN AMRO's framework, the questions are framed in general terms to ensure broader relevance.

1. *What is the structure of ABN AMRO's current framework for physical climate risk quantification in retail mortgage credit risk models?*
    - How is physical climate risk damage estimated and translated into collateral value impact?
    - How are these damage estimates subsequently integrated into credit risk parameters (PD, LGD, EAD)?
2. *Which limitations and assumptions become apparent in the current framework?*
    - Which assumptions play a central role in the estimation process?
    - To what extent do these assumptions diverge from methodological or practical realism?
3. *How can a model refinement or extension be designed and implemented?*
    - Which design principles can be applied to improve the framework?
    - What data and methods are required to carry this out?
4. *How can such a refinement be evaluated and validated?*
    - Which criteria can be used to assess the performance of the model?
    - How do the outcomes affect the interpretation of climate-related credit risk impacts?

This report is structured to first provide the necessary theoretical background, presented in the form of a literature study (Chapter 2). This literature section is organised to review the foundational aspects of credit risk and climate risk within the banking sector. It begins with an introduction to credit risk modelling, including key concepts, theoretical foundations, and essential metrics such as Probability of Default (PD) and Loss Given Default (LGD). The specific exposure class of residential mortgages is also briefly discussed. This is followed by an overview of relevant regulatory frameworks—starting with the general risk modelling guidelines established in the Basel Accords [10–13], and subsequently covering more recent developments in climate and environmental risk (CER) regulation within the banking sector [8, 14, 15]. Once this theoretical basis has been established, the methodology section follows (Chapter 3). This part outlines the structure and progression of the modelling approach and describes each phase in detail, along with the corresponding theoretical and modelling considerations. The methodology is closely aligned with the sequence of sub-research questions that guided this study. While this description may still appear abstract at this stage, it is important to note that ABN AMRO's existing climate risk modelling framework served as the primary starting point, and that the methodological structure of this thesis evolved in direct response to the insights gained during the study. The sub-research questions reflect this iterative process and provide a more concrete structure to the modelling phases that followed. After all necessary background and methodology has been introduced, the results section presents the final empirical findings (Chapter 4). These results are subsequently critically reviewed in the discussion (Chapter 5), where implications are explored and directions for future research are proposed, and then drawn together in the conclusion (Chapter 6).

2

Literature study

Before the research questions can be addressed, first a solid knowledge foundation must be established across several domains: credit risk and climate risk. This includes a more in-depth understanding of how residential real estate exposures are modelled within credit risk frameworks, as well as a focused introduction to flood risk as a key component of physical climate risk. Finally, it is essential to explore existing research at the intersection of these two domains. Establishing this theoretical foundation is the primary objective of the initial phase of this research—of which this literature report is the main deliverable.

2.1 Credit risk modelling

Credit risk modeling is a fundamental aspect of financial risk management, providing methodologies to quantify and mitigate risks associated with lending activities. These models aim to estimate the financial losses that may occur if borrowers are unable to meet their repayment obligations. Given its central role in ensuring financial stability, credit risk modeling has evolved into a structured field, incorporating both theoretical frameworks and empirical applications.

This section provides an overview of the key components of credit risk modeling. It begins with a set of fundamental concepts and terminology that are essential for understanding the structure and function of credit risk models (§ 2.1.1). Next, the theoretical foundation is discussed, including the statistical and mathematical principles underlying capital requirements and loss estimation (§ 2.1.2). This is followed by a high-level overview of common modeling techniques for Probability of Default (PD) and Loss Given Default (LGD), with a focus on practices observed in the Dutch banking sector (§ 2.1.3). Finally, the section explores the specific characteristics of residential mortgage exposures, identifying key risk drivers that influence credit risk modeling within this exposure class (§ 2.1.4).

To establish a foundational understanding, it is essential to first revisit the concept of financial risk in general. Risk is formally defined as follows:

Definition 1. *Risk is any event or action that may adversely affect an organization's ability to achieve its objectives and execute its strategies.*

This definition can be categorized into distinct types of risk [16]:

- **Market Risk**: The risk of changes in the value of a financial position or portfolio due to movements in underlying components such as stock prices, bond prices, exchange rates, or commodity prices .

- **Credit Risk**: The risk of not receiving promised repayments on outstanding investments due to the default of a borrower. This risk is the focus of this research.
- **Operational Risk**: The risk of losses resulting from inadequate or failed internal processes, people, systems, or external events.
- **Liquidity Risk**: The risk associated with the lack of marketability of an investment, where it cannot be sold quickly enough to prevent a loss.
- **Model Risk**: The risk of using a misspecified or inappropriate model to measure financial risk.

This research specifically addresses credit risk, which represents the likelihood of a borrower failing to meet contractual obligations, resulting in financial losses for the lender. Credit risk is crucial in the context of financial institutions, particularly in lending activities such as mortgages, loans and credit facilities.

2.1.1 Fundamentals

The field of credit risk modeling covers a wide variety of models, exposure classes and methodologies that often differ significantly between institutions. Despite this diversity, establishing a clear and consistent framework of foundational terminology is essential to ensure effective communication and understanding. This section introduces a set of key terms that frequently appear in credit risk modeling, including throughout this report. A simple illustrative example is also included to clarify how these metrics might look in practice.

- **Exposure Class**: A regulatory category that groups similar types of loans based on their characteristics and risk profiles. Common exposure classes include retail mortgages, corporate loans and exposures to governments or financial institutions. Each class follows different regulatory rules and modeling requirements. For example, residential mortgages form a separate exposure class with their own capital requirements and risk drivers.
- **Probability of Default (PD)**: The likelihood that a borrower will fail to repay their loan within a specific period, usually one year. It reflects how risky the borrower is. PD is one of the key metrics in credit risk modeling. The way in which this value is determined depends on the regulatory approach a bank is required to follow. Under more advanced regulatory approaches, banks are allowed to develop their own PD models using internal historical data. These models are typically tailored to specific exposure classes and reflect the characteristics of the bank's own portfolio.
- **Exposure at Default (EAD)**: The amount of money that is at risk at the moment the borrower defaults. For a simple loan, this is just the outstanding balance. For credit cards or other flexible credit lines, it also includes the unused part that the borrower might still withdraw before defaulting. In that case, the amount is adjusted using a so-called Credit Conversion Factor (CCF). So in general the EAD is given by

  $$\text{EAD} = \text{Drawn Balance} + (\text{Undrawn Balance} \cdot \text{CCF})$$

  Estimating EAD accurately is important, especially for products where the borrower still has room to borrow more before defaulting. For the exposure class of residential mortgages specifically, EAD is typically set equal to the outstanding loan balance, as no additional credit can be drawn.
- **Loss Given Default (LGD)**: The percentage of the loan that is lost if the borrower defaults, after taking into account any money the bank recovers, such as by selling collateral (e.g. a house). It is calculated as

  $$\text{LGD} = 1 - \frac{\text{Recovery Amount}}{\text{Exposure at Default}}.$$

  The recovery amount depends on many factors, like the value of the collateral and how easy it is to sell. A higher LGD means that the bank loses more money if things go wrong. Just like PD, LGD can be estimated using internal models, based on historical recovery data, or can be set by regulation.

- **Expected Loss (EL)**: The amount of money the bank expects to lose on a loan. It combines three components: how likely a borrower is to default (PD), how much is lost if that happens (LGD) and how much money is at risk (EAD), so it is given by the following

$$\mathrm{EL} = \mathrm{PD} \cdot \mathrm{LGD} \cdot \mathrm{EAD} \tag{2.1}$$

  This expected loss is treated as a regular cost of doing business. Banks usually include it in the interest rate they charge or account for it by setting money aside (called provisions).

- **Unexpected Loss (UL)**: The part of the loss that goes beyond what the bank expected. Even if the expected loss is low, actual losses can sometimes be much higher, especially during a crisis. UL is used to calculate how much extra capital a bank needs to keep as a safety buffer. The idea is: if things go worse than expected, the bank should still be able to absorb the shock without getting into trouble.

- **Stress Testing**: A way to check what would happen to the bank's loans if the economy suddenly takes a turn for the worse. For example, what if house prices fall sharply, or unemployment rises? In a stress test, banks simulate such scenarios and estimate how much the PD, LGD, EAD, and the corresponding losses would increase. This helps them prepare for extreme but plausible situations and is also required by regulators.

To make these definitions more concrete, the following example illustrates how the different components (PD, LGD, EAD, and EL) interact in a simplified mortgage case. It shows how a bank estimates the potential loss on a loan by combining these key metrics.

Example 2.1.1. *Consider the following simplified mortgage case: a borrower receives a mortgage of $400,000 to buy a house worth $500,000. This means the initial Loan-to-Value (LTV) ratio is 80%. Over time, the borrower repays part of the loan, and at some specific time the outstanding balance is reduced to $360,000. This remaining amount is the Exposure at Default (EAD) — the amount that would still be at risk if the borrower defaults.*

Now assume the borrower defaults. The bank repossesses the property and sells the house. Let's say it manages to sell the house for $342,000, which is lower than the original house value, typically due to the urgency and price pressure associated with a forced sale, or adverse economic cycle conditions. This means the bank recovers part of the exposure, but not all of it. There is a loss of:

$$Loss = \$360{,}000 - \$342{,}000 = \$18{,}000$$

Based on this, the Loss Given Default (LGD) can be calculated as the percentage of the EAD that is lost:

$$LGD = 1 - \frac{Recovery\ Amount}{Exposure\ at\ Default} = 1 - \frac{342{,}000}{360{,}000} = 0.05$$

Next, assume the Probability of Default (PD) is estimated to be 25%, i.e. 0.25. This means there is a 1 in 4 chance that a similar borrower will default within a year.

Now that the PD, LGD and EAD are obtained, it concludes with the calculation of the Expected Loss (EL):

$$EL = PD \cdot LGD \cdot EAD = 0.25 \cdot 0.05 \cdot 360{,}000 = \$4{,}500$$

This means that on average, the bank expects to lose $4,500 on this mortgage. This expected loss is usually covered by the interest charged on the loan or through provisions set aside for credit losses.

With a clearer understanding of the fundamental components of credit risk modeling, the focus now shifts to a deeper exploration of the theoretical framework behind these concepts. The following section discusses the mathematical and regulatory principles that form the foundation of credit risk models.

2.1.2 Theoretical foundation

A key objective of credit risk modeling is to determine how much economic capital a bank should hold internally to remain solvent during severe financial stress. This internally determined capital serves as a cushion against unexpected losses. In parallel, banks are also required to hold regulatory capital — a minimum amount of capital mandated by international regulatory standards. This regulatory capital underpins the capital ratio that banks must meet to safeguard financial stability.

The required amount of regulatory capital is expressed in terms of Risk-Weighted Assets (RWA), which adjust a bank's exposures for their associated risk. A fixed percentage of these RWA must be covered by regulatory capital; currently this requirement is set at 8% [17]. In other words, the RWA figure appears in the denominator of the capital ratio and plays a key role in determining whether a bank meets its regulatory obligations.

The calculation of RWA relies on the Capital Requirement K, which represents the risk-sensitive capital buffer required per unit of exposure. Informally, K reflects the proportion of an exposure deemed sufficiently risky to require additional capital. There are different levels of flexibility in how banks may determine these values, which correspond to different approaches. The so-called Standardized Approach (SA) applies fixed regulatory risk weights, whereas the Internal Ratings-Based (IRB) approach allows banks to use internal models to estimate risk components such as PD, LGD and EAD. These approaches will be discussed in more detail in regulatory section 2.2.

Given that these concepts may be difficult to grasp without prior exposure to credit risk modeling, this section aims to provide the theoretical and mathematical foundations necessary for their understanding. The key concepts presented here will also serve as the basis for Section 2.2, which provides a broader overview of the regulatory framework and its historical development. Since the capital formulas used in IRB models consist of several components, this section starts from the most fundamental element: the potential losses a financial institution might face. This naturally leads to the question of how such losses are distributed and what underlying assumptions shape their behaviour.

Loss distribution function

The estimation of economic capital relies on the probability distribution function of credit losses, commonly referred to as the loss distribution of a credit portfolio or credit loss function [18, 19]. This distribution represents the range of possible losses a bank may incur due to credit risk. A visual representation of this distribution is given in Figure 2.1.

The loss distribution is characterized by its skewed nature, where small losses are significantly more probable than large ones. The x-axis represents the potential credit losses, while the y-axis denotes the probability density of these losses occurring. Several key risk measures are directly linked to this distribution. Expected Loss (EL) corresponds to the average credit loss a bank expects over a given time horizon. It is considered a normal cost of doing business and is typically covered by pricing strategies and provisions. In contrast, Unexpected Loss (UL) represents the deviation from the expected loss due to uncertainty in credit outcomes. Banks must hold capital to absorb UL, to be able to withstand crisis credit events.

A critical measure associated with stress scenarios is the Value at Risk (VaR), which is defined as the maximum potential loss a bank is expected to sustain with a given confidence level over a specific time horizon. Mathematically, VaR at confidence level α is given by

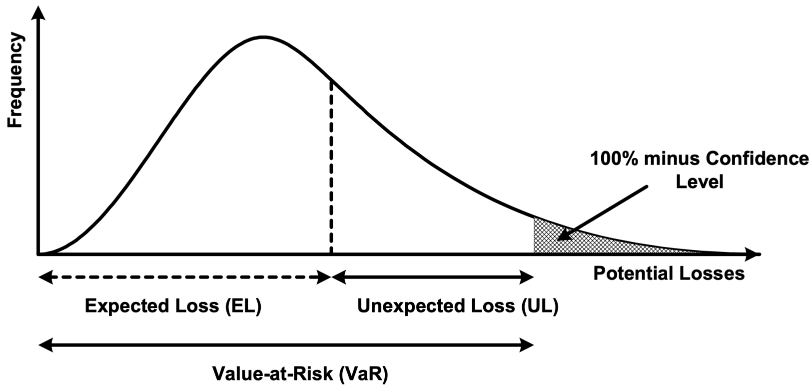

Figure 2.1: Credit Loss Function [19].

$$P(L > \mathrm{VaR}_\alpha) = 1 - \alpha, \tag{2.2}$$

where L represents the loss variable under consideration. In credit risk management, VaR is often referred to as C-VaR, the Credit Value at Risk, a convention that will also be followed here. By regulation, a confidence level of $\alpha = 0.999$ is generally applied in credit risk assessments to ensure adequate capital buffers against unexpected losses.

The difference between the EL and the C-VaR defines the stress loss, representing the probability that actual losses exceed both EL and UL. This corresponds to the right tail of the loss distribution. The relationship between UL, EL and C-VaR naturally links to the regulatory capital requirements. The next section explains how these requirements are calculated under the Basel framework.

Capital requirement and risk weighted assets

Under the Basel framework, capital requirements are directly linked to unexpected losses (UL). How this relationship is established—along with a simplified outline of the derivation of regulatory capital—will be explained throughout this section. It is important to note that under the Standardized Approach (SA), prescribed regulatory risk weights are applied to determine capital requirements. In contrast, under the Internal Ratings-Based (IRB) approach, capital requirements are calculated based on risk parameters such as PD and LGD. For the IRB approach, the regulatory formula defined in CRE31, paragraph 31.14 [17], for retail residential mortgage exposures that are not in default, is given by the following equation:

$$\textit{Capital requirement} = K = LGD \cdot \Phi\left(\frac{\Phi^{-1}(PD)}{\sqrt{1-\rho}} + \sqrt{\frac{\rho}{1-\rho}} \cdot \Phi^{-1}(0.999)\right) - PD \cdot LGD, \tag{2.3}$$

where Φ denotes the cumulative distribution function (CDF) of the standard Gaussian distribution and Φ^{-1} denotes its inverse, the quantile function of the standard Gaussian, which maps a probability p to the corresponding quantile value. PD denotes the probability of default, LGD represents the loss given default and ρ is referred to as the correlation, a parameter that captures systematic risk and serves as a proxy for exposure to the general economy. Under the currently applicable regulatory framework, a fixed correlation value of $\rho = 0.15$ is prescribed for retail residential mortgage exposures. To illustrate, for corporate loans this correlation is generally set to a higher value by regulators, as corporate loans are assumed to be more sensitive to the general economy than residential mortgages [1].

[1] Note that: ρ cannot be zero, as this would imply $K = 0$. To prevent banks from underestimating systematic risk, Basel II established minimum values for ρ across different exposure classes.

The capital requirement equation (2.3) may seem to appear out of nowhere, but its derivation is based on a quite extensive mathematical framework. A detailed exploration is not necessary for the scope of this study, a simplified derivation can be provided to offer some intuitive understanding, using the following definitions [20]:

$$K = VaR_{\alpha}(PD) \cdot LGD, \tag{2.4}$$

$$VaR_{\alpha}(PD) = PD_{\alpha} - PD, \tag{2.5}$$

$$PD_{\alpha} = \Phi\left(\frac{\sqrt{\rho}\Phi^{-1}(\alpha) + \Phi^{-1}(PD)}{\sqrt{1-\rho}}\right). \tag{2.6}$$

where $VaR_{\alpha}(PD)$ denotes the Value-at-Risk of the default probability at confidence level α, and PD_{α} is the critical value of PD at confidence level α, commonly referred to as the Worst-Case Default Rate (WCDR) or downturn PD.

Using these expressions, the capital requirement as specified by regulation follows directly:

$$\begin{aligned}
K &= VaR_{\alpha}(PD) \cdot LGD, \\
&= (PD_{\alpha} - PD) \cdot LGD, \\
&= \left(\phi\left(\sqrt{\frac{1}{1-\rho}}\Phi^{-1}(PD) + \sqrt{\frac{\rho}{1-\rho}}\Phi^{-1}(\alpha)\right) - PD\right) \cdot LGD, \\
&= LGD \cdot \phi\left(\sqrt{\frac{1}{1-\rho}}\Phi^{-1}(PD) + \sqrt{\frac{\rho}{1-\rho}}\Phi^{-1}(\alpha)\right) - PD \cdot LGD.
\end{aligned}$$

Finally, the connection between the capital requirement (2.3) and the loss distribution presented in Figure 2.1 can be further clarified. The right-hand side of Equation (2.3) corresponds to the expression for the Expected Loss, i.e. $EL = PD \cdot LGD$. It is important to note that, for explanatory purposes, EL is now expressed as a percentage of the EAD, whereas the actual EL can be represented as an absolute value (2.1). This absolute value is obtained by multiplying the percentage by the EAD. The left-hand side, $LGD \cdot \Phi\left(\frac{\Phi^{-1}(PD)}{\sqrt{1-\rho}} + \sqrt{\frac{\rho}{1-\rho}} \cdot \Phi^{-1}(0.999)\right)$, corresponds to the so-called conditional expected loss, which is equivalent to the previously introduced Credit Value at Risk (C-VaR). This can also be expressed as the product of the conditional probability of default and loss given default, i.e. $C\text{-}VaR = PD_{\alpha} \cdot LGD$, with $\alpha = 0.999$. Taking this into account, referring back to Figure 2.1, as previously mentioned, the unexpected loss is obtained by subtracting the expected loss from the C-VaR, i.e. $UL = C\text{-}VaR - EL$. This completes the framework, as this exactly corresponds to the subtraction observed in the capital requirement (2.3). It follows that the capital requirement corresponds to the unexpected loss component, confirming the earlier statement that these metrics are directly linked.

After determining the capital requirement (2.3) as a percentage of the exposure, the next step is to calculate the Risk-Weighted Assets (RWA). This is done by multiplying K by the Exposure at Default (EAD) and a factor of 12.5. This factor represents the reciprocal of the minimum capital ratio of 8% [2], ensuring consistency with regulatory requirements. The formula is given by:

$$RWA = 12.5 \cdot K \cdot EAD. \tag{2.7}$$

Another key metric is the Common Equity Tier 1 (CET1) ratio [14], it is used to assess a bank's financial strength. It reflects the proportion of a bank's capital, consisting of its highest quality assets, relative to its risk-weighted assets. The formula for calculating the CET1 ratio is:

[2] The factor of 12.5 ensures consistency with the regulatory capital framework set by the ECB and is based on the reciprocal of the minimum capital ratio of 8%. This is discussed in more detail in Section 2.2. See also [17].

$$CET1 \text{ ratio} = \left(\frac{CET1 \text{ capital}}{RWA}\right) \cdot 100\% \tag{2.8}$$

where the *CET1 capital* represents the bank's core equity capital, such as it's common shares and share surplus, retained earnings and capital from subsidiaries.

Regulatory requirements require that banks maintain a minimum CET1 ratio to ensure they have sufficient capital to absorb losses and continue operating during periods of financial stress. The minimum CET1 ratio requirement is $> 4.5\%$.

To make the capital requirement formula more tangible, the following example illustrates its application using a hypothetical retail mortgage exposure and standard Basel parameters.

Example 2.1.2. *For retail mortgage exposures, banks are allowed to develop their own internal models for estimating the Probability of Default (PD), Loss Given Default (LGD), and Exposure at Default (EAD). This is part of the so-called Advanced Internal Ratings-Based (A-IRB) approach. More details on this are provided in the regulatory section 2.2. For this example, the input values are assumed to be given.*

To illustrate the calculation of regulatory capital and Risk-Weighted Assets (RWA) under the IRB approach, a simplified case is considered in which the bank has only one retail mortgage exposure. In reality, the capital requirement is computed per individual loan, after which the results are aggregated across all exposures to determine the total RWA. However, for illustrative purposes, the focus here is on a single exposure.

Regarding the confidence level, a value of $\alpha = 0.999$ *is commonly used in IRB-based credit risk capital calculations to ensure sufficient capital is held for extreme but plausible credit events.*

Let's assume a bank has a retail mortgage exposure with the following characteristics

- *Probability of Default (PD):* 1% *or* 0.01
- *Loss Given Default (LGD):* 40% *or* 0.40
- *Exposure at Default (EAD): €100,000*
- *Correlation* $\rho = 0.15$ *(Basel-prescribed for retail mortgages)*
- *Confidence level* $\alpha = 0.999$,

The capital requirement K *is given by the IRB formula*

$$K = LGD \cdot \Phi\left(\frac{\Phi^{-1}(PD)}{\sqrt{1-\rho}} + \frac{\sqrt{\rho}}{\sqrt{1-\rho}} \cdot \Phi^{-1}(0.999)\right) - PD \cdot LGD,$$

recall equation (2.3). This leads to the following calculations

$$\Phi^{-1}(PD) = \Phi^{-1}(0.01) \approx -2.33$$

$$\Phi^{-1}(0.999) = 3.09$$

$$\sqrt{1-\rho} = \sqrt{0.85} \approx 0.922, \quad \sqrt{\rho} \approx 0.387$$

$$\frac{-2.33}{0.922} \approx -2.527, \quad \frac{0.387}{0.922} \cdot 3.09 \approx 1.297$$

$$\Phi(-2.527 + 1.297) = \Phi(-1.23) \approx 0.109$$

$$K = 0.40 \cdot 0.109 - 0.01 \cdot 0.40 = 0.0436 - 0.004 = 0.0396$$

So the capital requirement is approximately 3.96% *of the EAD. Next, the Risk-Weighted Assets (RWA) are computed:*

$$RWA = 12.5 \cdot K \cdot EAD = 12.5 \cdot 0.0396 \cdot €100,000 = €49,500$$

Finally, if the bank has €8,000 in CET1 capital for this exposure:

$$\textit{CET1 ratio} = \left(\frac{€8,000}{€49,500}\right) \cdot 100\% \approx 16.16\%$$

Since this is above the minimum required CET1 ratio of 4.5%, *the bank satisfies the regulatory requirement.*

Banks have some degree of flexibility in modeling PD, LGD and EAD, leading to potential variation in capital requirements across institutions. As evident from the capital requirement (2.3), accurate estimation of PD and LGD is crucial for banks to balance capital efficiency and risk coverage, ensuring that they hold sufficient capital without excessive reserves that could limit financial performance.

2.1.3 Probability of default (PD) and loss given default (LGD) modelling

The development of Probability of Default (PD) and Loss Given Default (LGD) models is tailored to each individual bank. This section provides a general overview of examples of typical structures used by Dutch banks. While implementation details may vary, the core framework and regulatory context share many common elements across institutions operating within the Netherlands. The two key metrics, PD and LGD, will both be discussed separately to provide a general understanding of their typical modeling approaches.

Probability of Default (PD) modeling

The Probability of Default (PD) is a key metric in credit risk modelling, used not only in the calculation of regulatory capital under the Internal Ratings-Based (IRB) approach, as discussed in Section 2.1.2, but also in areas such as credit acceptance, economic capital modelling, loan pricing, provisioning and portfolio monitoring. Although the way PD is modelled is bank-specific, there is a common structure in terms of general steps that are typically applied. At the start of the PD modelling process, a selection of relevant risk drivers is made. Based on this selected set of drivers, a statistical model is used to produce a certain score. This score is used to define different risk profiles—also referred to as 'pools'. These pools are then calibrated, followed by final adjustments, ultimately resulting in the final PD used in each of the aforementioned applications. The PD estimation process is therefore generally structured in two phases: risk differentiation and risk quantification, both of which will be briefly explained in this section.

Phase 1: **Risk differentiation (ranking)**
In the first phase, statistical models such as logistic regression are employed to rank borrowers according to their credit risk. These models take loan- or borrower-specific risk drivers $x_1, x_2, \dots, x_k$ as input and produce a continuous score that reflects the relative likelihood of default. For example, in a logistic regression setting, the output score is given by

$$\text{score}(x) = \frac{1}{1 + \exp(-(\beta_0 + \beta_1 x_1 + \beta_2 x_2 + \cdots + \beta_k x_k))}$$

Although this score lies between 0 and 1 and resembles a probability, it does not yet represent the final PD. Rather, it provides a risk ranking that allows the institution to assign exposures to so-called discrete rating grades or pools, each bounded by lower and upper score thresholds.

Phase 2: **Risk quantification (calibration)**
In the second phase, each rating grade is assigned a PD estimate that reflects the average default experience observed historically for exposures in that grade. This is typically based on the *Long-Run Average Default Rate* (LRA DR), computed as the average of annual default rates over a sufficiently long and representative period of time. The use of long-run averages ensures that the PDs reflect both benign and adverse credit conditions.

After assigning the LRA DR, institutions may apply *appropriate adjustments (AA)* to account for structural breaks, data inconsistencies or observed jumps in default rates. Finally, a *Margin of Conservatism (MoC)* is added to reflect estimation uncertainty, particularly in cases of data limitations or model uncertainty. The final output—referred to as the *regulatory PD*—is then used in capital requirement calculations.

To make the process more intuitive, a hypothetical example is provided below.

Example 2.1.3. *Consider a bank estimating the probability of default (PD) for residential real estate (RRE) exposures under the IRB approach. The institution collects historical data on 100,000 mortgage loans, which may include borrower- and loan-specific variables such as loan-to-value ratio (LTV), borrower age, debt-to-income ratio (DTI) and delinquency history.*

Step 1: Risk differentiation: A logistic regression model is trained to estimate a credit risk score for each borrower. For a particular borrower, the model outputs a score of 0.014. Although this value lies between 0 and 1, it does not yet represent a final PD—it serves to rank borrowers by relative credit risk.

Step 2: Score segmentation: Based on the score distribution, the bank defines ten rating grades (pools), with pools A through I equally spaced between 0.0000 and 0.0400, and pool J capturing all scores above 0.0400. The segmentation is shown below:

Rating grade	**Score range**
A	*[0.0000 – 0.0040)*
B	*[0.0040 – 0.0080)*
C	*[0.0080 – 0.0120)*
D	*[0.0120 – 0.0160)*
⋮	⋮
H	*[0.0280 – 0.0320)*
I	*[0.0320 – 0.0400)*
J	*[0.0400 – 1]*

Our example borrower (score = 0.014) is assigned to rating grade D.

Step 3: Calibration using the LRA DR: The bank calculates the long-run average default rate (LRA DR) for each pool using historical default data over a 10-year period. For pool D, the LRA DR is estimated at 0.15%.

Step 4: PD best estimate: Based on current performance data and model calibration, the bank estimates the best estimate PD for pool D to be slightly lower than the LRA DR, at 0.14%.

Step 5: Margin of Conservatism (MoC): To account for data limitations and estimation uncertainty, a conservatism margin of 0.02 percentage points is added.

The final regulatory PD for rating grade D is therefore:

$$PD_{final} = 0.14\% + 0.02\% = 0.16\%$$

This PD is subsequently used as an input in the IRB capital requirement calculation.

Each exposure class relies on a distinct set of risk drivers for PD estimation. These variables serve as inputs to the risk differentiation model, which produces a credit risk score used to assign exposures to rating pools. Historical performance data is then used to calibrate PDs for each grade, with appropriate adjustments and conservative uplifts applied as needed, in line with ECB expectations and Basel requirements.

Loss Given Default (LGD) modeling

Loss Given Default (LGD) quantifies the percentage loss a bank incurs in the event of a borrower's default. While the general definition is consistent across institutions, the modelling of LGD can differ significantly, depending on factors such as data availability, portfolio characteristics and internal methodologies. This section illustrates an example of a commonly applied modelling technique, in which LGD is decomposed into three components: the *probability of cure*, the *Loss Given No Cure* (LGNC), and the *Loss Given Cure* (LGC).

The rationale for this decomposition lies in the fact that, when a borrower defaults, it does not necessarily result in a full loss of the outstanding exposure. Several mitigating factors can reduce the realised loss, including the possibility that the borrower cures (i.e. returns to a performing status), enters into a restructured agreement, benefits from third-party guarantees, or provides collateral that can be liquidated to recover part of the exposure. As a result, an LGD model must distinguish between different default outcomes and their respective loss implications.

The first component, the *probability of cure*, estimates the probability that a borrower who has defaulted will return to a non-default status within a defined period. This probability, known as the *cure rate*, reflects the fraction of defaulted exposures that do not result in a write-off. It is a crucial input to determine the weighting between LGNC and LGC in the overall LGD formula. The cure rate is often estimated using logistic regression models, similar to Probability of Default (PD) models. These models require careful risk driver selection, including borrower characteristics, loan terms and macroeconomic indicators.

The second component, the *LGNC*, estimates the expected percentage loss for loans that do not cure. Rather than focusing directly on losses, this component can be framed in terms of the earnings the lender expects to recover from the exposure at default. These earnings may include proceeds from the sale of the collateral, remaining savings or deposits, recoveries from insurance policies, or any other financial inflows following default. LGNC is then expressed as the complement of the expected recovery ratio

$$\text{LGNC} = \frac{\text{Total expected loss}}{\text{EAD}} = 1 - \frac{\text{Total expected earnings}}{\text{EAD}}$$

This ratio expresses the share of the exposure at default that remains unrecovered in non-cured cases.

The third component, the *LGC*, captures the small but mostly non-zero loss that may still occur in cured cases. Even though the borrower repays the full principal, a loss may arise due to delayed cashflows. The lender receives the expected payments later than planned, resulting in an economic loss driven by the time value of money. This can be represented by:

$$\text{LGC} = \Delta r$$

Here, Δr reflects the difference in return between receiving payments as originally scheduled and receiving them with delay. Since no principal is lost and other costs are limited, LGC tends to be low and relatively stable.

The three components are then finally combined into a single LGD estimate using a weighted average structure

$$\text{LGD} = (1 - \text{Cure Rate}) \cdot \text{LGNC} + \text{Cure Rate} \cdot \text{LGC}$$

This structure ensures that the overall LGD reflects the probability-weighted losses across both cured and non-cured paths. To make this more intuitive, an illustrative example is presented below.

Example 2.1.4. *Consider a portfolio of residential mortgage loans, where borrowers are classified as in default after missing three consecutive monthly payments. Let the outstanding amount at the moment of default, i.e. the Exposure at Default (EAD), to be €200,000.*

Historical data indicates that 20% of defaulted loans cure within the defined observation period, resulting in a cure rate of 0.20. Accordingly, 80% of the loans are classified as non-cured.

For the non-cured loans, the lender expects to recover €130,000 through collateral liquidation, offsetting deposits, insurance payouts or other sources. The resulting loss is €70,000, leading to:

$$LGNC = \frac{200{,}000 - 130{,}000}{200{,}000} = \frac{70{,}000}{200{,}000} = 0.35$$

For the cured loans, the full principal is eventually recovered, but a small loss is still incurred due to delayed payments. Suppose this timing-related loss is estimated at €6,000 per loan. Then:

$$LGC = \frac{6{,}000}{200{,}000} = 0.03$$

The final LGD is obtained by combining the two components according to the weighted average structure:

$$LGD = (1 - 0.20) \cdot 0.35 + 0.20 \cdot 0.03 = 0.28 + 0.006 = 0.286$$

$$\Rightarrow LGD = 28.6\%$$

While banks often rely on internally developed LGD models tailored to their specific portfolios, the underlying metric is generally structured around some form of combination of the three components discussed: the probability of cure, LGNC and LGC. This structure ensures that the resulting LGD reflects both the larger loss in non-cured cases and the smaller, timing-related loss in cured cases. Moreover, the use of EAD as a denominator ensures comparability across loans of varying sizes.

2.1.4 Residential real estate (RRE) exposure

Since the focus of this research is on the exposure class of residential mortgages—i.e. residential real estate (RRE)—it is useful to first provide some background on factors that could influence credit risk models within this class. These factors are generally referred to as **risk drivers**, as they may affect key credit risk metrics.

Each bank develops its own set of significant risk drivers, with internal modelling choices and portfolio characteristics guiding their selection and weighting. Nevertheless, it is possible to outline a generic set of drivers that are often considered relevant in the context of RRE portfolios [17, 21, 22].

These risk drivers can be grouped into four broad categories: borrower characteristics, loan and transaction characteristics, property characteristics, and external factors. Each category captures a different dimension of the risk profile and could potentially influence the overall assessment of residential real estate exposures.

In the following, generalized examples of such drivers are provided and briefly explained for each category. It should be noted that banks typically define and apply their own, more specific sets of risk drivers, tailored to portfolio characteristics and data availability. The overview presented here is therefore not exhaustive; its aim is to illustrate the types of factors that could plausibly affect credit risk in residential mortgage portfolios.

I. Borrower characteristics

- *Creditworthiness of the borrower*: Includes credit history, existing debt and behavioural scoring models. For RRE, the borrower's creditworthiness could potentially influence the probability of default (PD) in internal rating models.
- *Income and employment status*: A stable income may be important for mortgage repayment. A commonly used metric is the *Loan-to-Income (LTI)* ratio, defined as

$$\text{LTI} = \frac{\text{Loan Amount}}{\text{Gross Annual Income}}. \tag{2.9}$$

  Higher LTI values could indicate increased repayment pressure and thus potentially higher credit risk.
- *Type of borrower*: Whether the borrower is a salaried employee, self-employed or retired could affect income stability and thereby the exposure to credit risk in RRE lending.

II. Loan and transaction characteristics

- *Loan amount*: A larger mortgage relative to income or property value could increase the risk of default and the potential loss given default (LGD) in case of foreclosure.
- *Loan-to-Value (LTV) ratio*: A frequently applied metric in RRE modelling, defined as

$$\text{LTV} = \frac{\text{Loan Amount}}{\text{Collateral Value}}. \tag{2.10}$$

  Higher LTVs reduce the collateral buffer and could therefore lead to higher LGD if property values decline.
- *Loan maturity*: Longer maturities may increase exposure to macroeconomic volatility, potentially affecting both default timing and prepayment behaviour.
- *Interest rate type*: Variable-rate mortgages (VRMs) expose borrowers to interest rate risk, which could increase the likelihood of payment shocks and defaults.
- *Presence of guarantees*: Guarantees (e.g. from a government scheme) may reduce the effective credit risk to the lender in case of borrower default.

III. Property characteristics

- *Property value*: The market and stressed value of the property are central in estimating LGD and are a key input for LTV calculations.
- *Property quality and maintenance*: Poor maintenance could lower property value and thereby reduce recoverability in case of foreclosure.
- *Energy performance*: Higher energy efficiency may increase property value and reduce household expenses, which could improve affordability and lower default risk. In the EU, Energy Performance Certificates (EPCs) are commonly used to capture this effect.

IV. External factors

- *Economic conditions*: Variables such as GDP growth, interest rates and unemployment could influence both default probabilities and borrower affordability in RRE portfolios.
- *Housing market developments*: Supply-demand imbalances, price trends and market liquidity may affect collateral values and potential loss severity.
- *Regulations*: Loan caps, minimum down payments and macroprudential measures could indirectly shape the risk profile of RRE portfolios.

- *Social risks*: Demographic shifts (e.g. ageing population, urbanisation) or behavioural changes (e.g. preference for renting) may influence demand and thus property price dynamics.

Understanding these potential risk drivers provides a broad overview of factors that could affect credit risk in residential mortgage portfolios. This overview serves as a conceptual basis for later analysis of how climate-related variables may interact with such drivers.

2.2 Regulations

To safeguard the financial system from excessive risk-taking and instability, regulatory frameworks have evolved to ensure that banks maintain sufficient capital buffers, manage risks effectively and operate transparently. This section outlines key regulatory structures relevant to this research, providing a clear foundation for the regulatory terminology used throughout.

2.2.1 Regulatory framework

One of the key institutions in global banking regulation is the *Bank for International Settlements (BIS)*, founded in 1930 and headquartered in Basel, Switzerland. The BIS acts as a hub for central banks, promoting international financial cooperation and monetary stability [23]. Within the BIS, the *Basel Committee on Banking Supervision (BCBS)* was established to develop consistent regulatory standards for banks worldwide. This committee has introduced several regulatory frameworks, known as the *Basel Accords*, which define capital adequacy and risk management requirements.

The Basel Accords provide a series of internationally recognized guidelines designed to enhance financial stability by standardizing capital and risk management requirements across the banking sector [23]. These frameworks have evolved over time in response to financial crises, addressing emerging risks and regulatory shortcomings. However, since the BCBS is not a regulatory authority, its frameworks must be implemented by national and regional regulators.

In the European Union (EU), banking regulation builds upon the Basel Accords and is further shaped by several key institutions. The *European Central Bank (ECB)* and the *European Banking Authority (EBA)* play central roles in ensuring financial stability and regulatory compliance. The ECB, apart from its role in monetary policy, is responsible for banking supervision through the *Single Supervisory Mechanism (SSM)*. The SSM acts as a framework that standardizes banking oversight across the eurozone. It ensures that all participating countries apply banking regulations consistently and that large, cross-border banks are monitored centrally [24].

Within this framework, the ECB directly supervises large and systemically important banks — that is, banks whose failure could threaten the stability of the European financial system. Smaller banks remain under the supervision of national regulators, who operate within the SSM but retain responsibility at the national level. In the Netherlands, this role is fulfilled by *De Nederlandsche Bank (DNB)*, which oversees smaller Dutch banks and assists the ECB in supervising larger institutions such as ING, ABN AMRO, and Rabobank [25]. By working together within the SSM, the ECB and national regulators like the DNB aim to maintain a stable and resilient European banking sector.

The *European Banking Authority (EBA)*[3], established in 2011, plays a crucial role in harmonizing banking supervision across EU member states. It helps ensure that regulatory frameworks, such as the Basel Accords, are applied consistently within the EU. The legal rules themselves—such as the *Capital Requirements Directive (CRD)* and the *Capital Requirements Regulation (CRR)*—are drafted by the European Commission and adopted by the European Parliament and Council. These texts are based on the Basel Accords but

[3] Previously coordinated by the *Committee of European Banking Supervisors (CEBS)*, but after the 2008 financial crisis reshaped into the EBA for stronger regulatory enforcement.

are adapted to fit the European legal and institutional framework. Based on the CRD and CRR, the EBA develops *EBA Guidelines* to support uniform interpretation and application across member states. In addition, the ECB sometimes publishes its own non-binding guides to clarify how it interprets these rules in its supervision. Unlike the ECB, the EBA does not directly supervise banks, but instead provides a common set of supervisory guidelines that both the ECB and national regulators, such as DNB, are expected to follow.

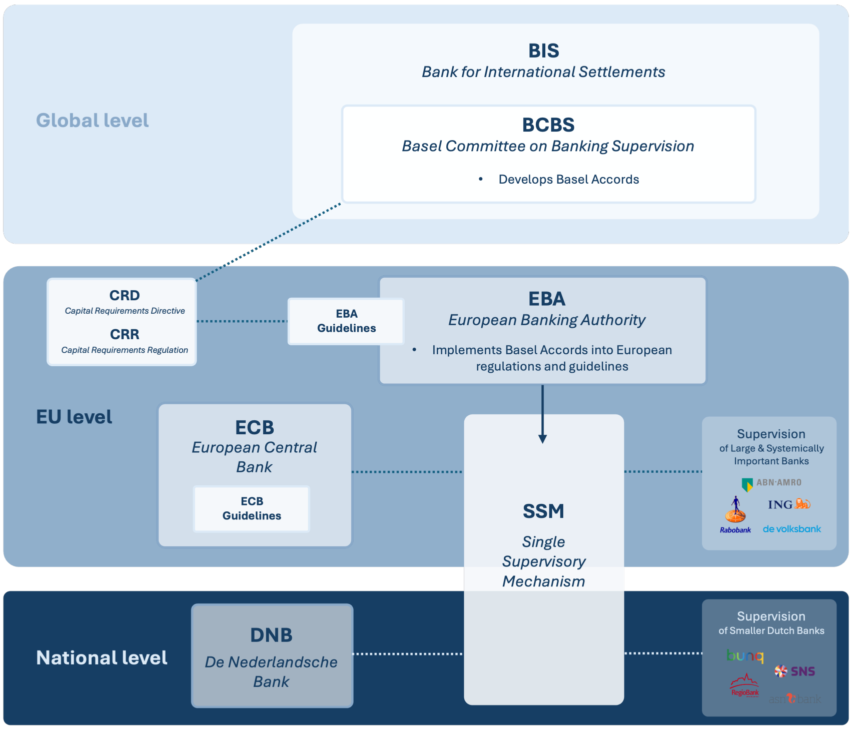

Figure 2.2: Overview of the Banking Regulatory and Supervisory Framework in the Netherlands

2.2.2 Basel accords

The Basel Accords are the international regulatory frameworks developed by the Basel Committee on Banking Supervision (BCBS). This section provides a high-level overview of key developments, focusing on aspects relevant to this research, aiming to establish a basic understanding of the currently permitted credit risk models.

Basel I

Basel I, introduced in 1988, marked the first global effort to establish minimum capital requirements for banks. It primarily focused on credit risk, categorizing assets into broad risk classes with fixed risk weights. Under this framework, banks were required to maintain a minimum capital adequacy ratio (CAR) of 8%, meaning they had to hold at least 8 cents of capital for every euro of risk-weighted assets (RWA) [10]. For example, residential mortgages were assigned a fixed risk weight of 50%, meaning that only half of the mortgage exposure counted toward the bank's total RWA, regardless of the borrower's creditworthiness [26]. While this approach provided a simple and standardized method for capital regulation, it lacked risk sensitivity, as it did not differentiate between loans with different default probabilities or loss severities.

Basel II

To address these limitations, Basel II was introduced in 2004 [11], offering a more risk-sensitive approach. The assumption that all loans within the same category should receive a fixed risk weight was abandoned. Instead, banks were given two options to determine capital requirements for credit risk:

- *Standardised Approach (SA)*: This method assigns fixed risk weights to exposures based on external credit ratings, without considering the bank's own risk assessment.
- *Internal Ratings-Based (IRB) approach*: Banks were allowed to use internal models to estimate their own risk weights, making capital requirements more sensitive to actual credit risk.

The IRB approach introduced two variants [22]:

- *Foundation IRB (F-IRB)*: Banks were permitted to estimate the probability of default (PD) for their exposures, but were required to use standard supervisory values for other parameters, such as the loss given default (LGD) and exposure at default (EAD), as defined by regulators.
- *Advanced IRB (A-IRB)*: Banks were allowed to fully estimate their own risk parameters, including the probability of default (PD), loss given default (LGD), and exposure at default (EAD), based on their internal data and models.

Basel II additionally introduced three pillars: 1 - *Minimum Capital Requirements*, 2 - *Supervisory Review*, and 3 - *Market Discipline*. In short, the first two pillars ensure that, even within the IRB approaches, minimum capital requirements per exposure class remain in place and that models are regularly assessed to ensure their reliability, even under stressed conditions. In addition to establishing a more risk-sensitive framework, Basel II also formalized the capital requirement (K), risk-weighted assets (RWA) and Common Equity Tier 1 ($CET1$) calculations, as discussed previously in equations (2.3), (2.7) and (2.8) for retail mortgage exposures specifically. These formulas, particularly relevant under the IRB approach, remain fundamental in regulatory capital assessments today, linking capital requirements directly to estimated credit risk parameters PD, LGD, and EAD. The third pillar establishes disclosure requirements, which for example includes mandating that banks publicly report whether they use the SA or IRB approach, as well as details on their capital structure and sectoral exposures. This transparency allows market participants and regulators to better assess a bank's financial stability.

Basel III

The 2008 financial crisis exposed weaknesses in Basel II, particularly its reliance on internal risk models, which underestimated systemic risk and led to insufficient capital buffers. In response, Basel III was introduced in 2010, maintaining the three-pillar structure of Basel II but introducing improvements to strengthen financial stability [12]. In short, the improvements under the first pillar included higher capital requirements, the introduction of the *capital conservation buffer* and *countercyclical buffer*, and the implementation of a *leverage ratio* to prevent excessive risk-taking. Under the second pillar, regulatory oversight was strengthened by introducing mandatory stress testing and incorporating liquidity risk assessments. For the third pillar, transparency requirements were expanded, obligating banks to disclose their leverage ratios, capital composition and liquidity coverage to ensure greater market discipline.

Basel IV

The refinements introduced under Basel IV, often referred to as the finalisation of Basel III [13], were finalised at the end of 2017 and came into effect at the beginning of 2023. It again builds upon the existing three-pillar structure. Below is an overview of the key improvements introduced under each pillar.

Pillar 1 – Minimum Capital Requirements: The already existing Standardized Approaches (SA) have been refined under Basel IV. In short, the SA refinements include the creation of more specific risk categories and the incorporation of additional parameters in the calculation of risk weights within these categories. For example, for certain exposure classes, the creditworthiness of the borrower now plays a greater role in determining risk weights. Another refinement includes tighter conditions for the use of external credit ratings (i.e. ratings from agencies such as Moody's or S&P).

Additionally, Basel IV introduces an *output floor*, requiring that internally modeled risk-weighted assets (RWAs) cannot be lower than *72.5%* of the standardized approach calculations. This measure ensures that banks relying on internal models do not significantly underestimate their risks. For retail mortgage exposures specifically, under the Standardized Approach, the risk weights depend on the loan-to-value (LTV) ratio of the asset loan (2.10), with the corresponding values shown in Figure 2.3. These exposures refer to residential mortgages where repayment is primarily based on the borrower's income rather than rental or investment cash flows, which aligns with the exposure class that this research focuses on. Consequently, due to the output floor requirement, the IRB risk weights must be at least 72.5% of the risk weights presented in the table.

Whole loan approach risk weights for regulatory residential real estate exposures that are not materially dependent on cash flows generated by the property						
	LTV ≤ 50%	50% < LTV ≤ 60%	60% < LTV ≤ 80%	80% < LTV ≤ 90%	90% < LTV ≤ 100%	LTV > 100%
Risk weight	20%	25%	30%	40%	50%	70%

Figure 2.3: Risk Weights for Retail Mortgage Exposures under the Standardized Approach, based on the Loan-to-Value (LTV) Ratio of the Underlying Asset [27].

Furthermore, Basel IV restricts the use of the *Internal Ratings-Based (IRB) approach* for certain asset classes with low default rates, such as *large corporates with revenues exceeding €500 million* and financial institutions. These exposures must now be assessed using the standardized approach, ensuring more consistent and reliable risk calculations across banks.

Finally, stricter *data requirements* have been introduced for internal models, particularly regarding *historical data coverage* and *input floors* for Probability of Default (PD) and Loss Given Default (LGD) estimates. This includes a minimum *PD input floor* of 0.05% for most exposure classes [28].

Pillar 2 – Supervisory Review Process: Supervisors are provided with detailed guidelines to assess banks' internal models and risk management practices. Basel IV introduces new *validation requirements* for internal models, including stricter *back-testing* and *benchmarking against external data*. Banks must now conduct more frequent reviews of their Probability of Default (PD) and Loss Given Default (LGD) estimates, ensuring that assumptions remain conservative and aligned with real-world default data [28].

Pillar 3 – Market Discipline: Basel IV expands the scope of disclosures related to credit risk. Banks are now required to provide more detailed transparency on their risk-weighted asset calculations, including a comparison between their internally modeled RWAs and those calculated under the standardized approach. Disclosures must also cover the *impact of the output floor* on capital requirements and the assumptions used in internal models [13].

2.2.3 Climate-related and environmental risks (CER) regulation

Understanding the regulatory landscape surrounding Climate and Environmental Risks (CER) is essential for this research, as it defines the expectations for how internal credit risk models should incorporate these risks.

Three key regulatory documents have been selected based on their significance and timeline. The ECB Guide on Climate-Related and Environmental Risks (2020) [14] provides the earliest complete published regulatory perspective, outlining broad expectations for how banks should manage CER. Rather than prescribing methodologies, it sets foundational principles that later regulations build upon. The EBA Report on the Role of Environmental and Social Risks in the Prudential Framework (2023) [15] evaluates the extent to which ESG risks have been integrated into the prudential framework and presents recommendations for further improvement. The most recent document, the EBA Guidelines on ESG Risks Management (2025) [8], introduces practical methodologies for banks to measure, manage, and integrate ESG risks into their credit risk models.

As all three regulatory documents are quite extensive and broadly applicable across various areas of the financial sector, it is important to clarify that this section does not aim to provide a comprehensive regulatory analysis. Instead, a focused selection has been made and summarized to highlight the aspects most relevant to this research, particularly those concerning the integration of CER into internal credit risk models for residential mortgages. Each of the three regulatory documents will be discussed separately in this section, after which a schematic summary can be found in Figure 2.4.

ECB Guide on climate-related and environmental risks (2020)

In November 2020 [14], the European Central Bank (ECB) published its first comprehensive final document outlining supervisory expectations regarding climate-related and environmental risks (CER) within the financial sector. This guide emphasizes that climate risks should not be viewed as standalone risks but rather as drivers of existing risk categories, such as credit risk, operational risk, market risk and liquidity risk. The document consists of 13 main expectations, several of which specifically address credit risk. The aspects derived from the report, particularly relevant to the context of this research, can be summarized as follows:

1. *Consideration of CER at all relevant stages of the credit risk modeling process*, [14] - *Expectation 8*.
   - Climate risks may lead to cash outflows or a reduction in liquidity buffers, requiring adjustments in liquidity risk management.
   - PD and LGD may increase in sectors or regions exposed to physical risks, such as real estate in flood-prone areas, due to declining collateral values.
   - Stricter energy efficiency regulations can result in higher adaptation costs and lower corporate profits, increasing PD and reducing collateral values.
2. *Integration of CER into stress testing and scenario analysis*, [14] - *Expectation 11*.
   - Stress testing and scenario analysis should account for:
     - the impact of physical and transition risks on financial exposures.
     - how climate risks may evolve under different future scenarios, acknowledging that historical data may not fully capture them.
     - the potential manifestation of climate risks over short, medium and long-term horizons.
3. *Data collection and disclosure of CER-related information*, [14] - *Expectation 13*.
   - Institutions must define how they assess material climate risks, including disclosure frequency and methods.
   - Credit risk exposures and collateral values should be disclosed by geography, highlighting areas with high physical risk.

EBA Report on Environmental & Social Risks in Prudential Framework (2023)

In October 2023 [15], the European Banking Authority (EBA) published a report assessing the role of environmental and social (E&S) risks within the prudential framework. The report aims to evaluate the extent to which these risks have been integrated into existing risk management practices and capital requirements, while also identifying areas where further regulatory guidance may be necessary.

The report states that at the time of publication, there were no banks which already explicitly incorporated E&S risks into their internal ratings-based (IRB) credit risk models. Where such risks were considered, they were primarily integrated through qualitative adjustments, such as expert judgment overrides, rather than through structural modifications to probability of default (PD) or loss estimation models. Many institutions assumed that these risks would be indirectly captured through existing model inputs, such as collateral valuation or financial indicators, rather than requiring dedicated E&S risk factors. The absence of a standardized approach meant that implementation varied significantly across institutions, and no clear industry consensus had been reached regarding how to systematically integrate these risks into IRB models.

While some banks had begun exploring the inclusion of environmental and social (E&S) risk drivers in credit models, the process remained in its early stages. A key challenge is the lack of historical data linking environmental risk factors to credit performance, making it difficult to quantify their impact in a statistically robust manner. Moreover, climate and environmental risks are inherently forward-looking, meaning that even if historical data were available, they would not necessarily provide reliable insights into future credit outcomes. Some institutions had developed climate-informed shadow PDs to supplement existing models, while others relied on manual overrides to adjust risk assessments where necessary. Additionally, banks were considering ways to incorporate forward-looking assessments into their risk frameworks but faced difficulties in determining the appropriate assumptions and methodologies. Despite these efforts, practical implementation remained limited, as banks struggled with both data availability and regulatory uncertainty.

To address these gaps, the EBA recommended that banks take a cautious approach when dealing with missing or unreliable E&S risk data. The report suggested that financial institutions should consider applying adjustments where uncertainty exists, ensuring that climate and environmental factors are not underestimated in risk assessments. Furthermore, the EBA proposed that expert-based qualitative variables could be used in rating systems where quantitative metrics were not yet fully developed. In the longer term, the EBA intends to explore the potential for formally incorporating specific E&S risk drivers into existing regulatory guidelines for credit risk modeling. However, further empirical research and industry collaboration will be necessary to establish a more standardized approach to integrating these risks into prudential frameworks.

EBA Guidelines on the management of ESG risks - Final Report (2025)

The 2023 EBA report highlighted the need for a more formal and conceptual approach to ESG risk management [8]. In response, in January 2025, the EBA published the *Guidelines on the Management of ESG Risks* [8] final report, introducing explicit reference methodologies to standardize the identification, measurement, management and monitoring of ESG risks. This report is based on Article 87a (5) of the Capital Requirements Directive (CRD VI), which mandates the EBA to issue guidelines on these matters. For large institutions, the guidelines will generally apply from 11 January 2026, while small and non-complex institutions will have an extended transition period until 11 January 2027.

A structured way to summarize the content of this report is by distinguishing between modeling requirements and data requirements. The aspects most relevant to this research can be summarized as follows [29]:

Modeling Requirements

- *ESG factors must be explicitly embedded into credit underwriting, risk classification, and portfolio management.* Banks can no longer treat ESG risks separately but must integrate them into existing credit risk models.
- *Materiality assessments should analyze the financial impact of ESG risk drivers on counterparties, sectors, regions, and loan (sub-)portfolios.* Banks must assess how environmental risks impact different regions and loan portfolios. For example, flood risks may affect collateral values, increasing credit risk in certain areas.
- *The guidelines stress the importance of quantifying environmental risks, including physical and transition risks.* Banks must move from qualitative assessments to measurable financial impacts, integrating

climate risk scores and adjusting PD/LGD estimates accordingly.

- *Clarity on using proxies and scenario analysis, allowing flexibility as data improves.* Since ESG data is evolving, banks may use proxies and scenario-based analysis to estimate risks. Overlays can be applied temporarily, similar to how banks handle novel risk factors.

Data Requirements

- *Data collection is now based on the ESG risk materiality assessment, allowing flexibility in granularity.* Instead of collecting all ESG data, banks may prioritize data based on a materiality assessment, ensuring focus on key risk areas.
- *The alignment with CSRD disclosures is emphasized, meaning banks should use publicly available data, especially on emissions and climate plans.* Banks should rely on CSRD-reported ESG data, which simplifies data collection but also requires frequent model updates as new data becomes available.
- *The use of proxies when data is unavailable is allowed, but a reduction over time is expected.* Banks may use proxies for missing ESG data, but these should be gradually phased out as direct ESG reporting improves.
- *Specific data points are required, such as GHG emissions, energy consumption, social standards and governance issues.*[4] Banks must collect GHG emissions, energy usage, and governance-related factors, ensuring borrower engagement where necessary. CSRD data should be prioritized for compliance.

A schematic summary of the three discussed documents is presented in Figure 2.4.

Figure 2.4: Overview and Timeline of Climate-Related and Environmental Risk (CER) Regulations.

Overall, the actual adoption of Climate and Environmental Risk (CER) integration within the Dutch banking industry remains limited. However, with the upcoming enforcement of stricter regulatory requirements, the urgency for further research on this topic within the Dutch banking sector becomes increasingly important. This necessity sets the stage for the next section, which examines the current state of progress within a specific area of climate and environmental risk.

[4] GHG emissions refer to the amount of greenhouse gases released by an entity's operations. Social standards include factors like employee rights, working conditions and community impact. Governance issues relate to internal controls, board structure and transparency. The CSRD (Corporate Sustainability Reporting Directive) is an EU regulation that mandates standardized sustainability disclosures from companies.

2.3 Climate and environmental risks

This section introduces the concept of climate and environmental risk (CER) and outlines its relevance for financial stability. It defines the main categories of CER, distinguishes between physical and transition risks, and introduces how these risks interact with traditional financial risk types. This provides the conceptual foundation for the subsequent analysis of their role within credit risk modelling.

2.3.1 Risk taxonomy

To establish a foundational understanding, it is essential to define climate and environmental risk (CER) in a structured way. Climate-related and environmental risks are generally understood to be defined as follows

Definition 2. *Climate and Environmental Risk (CER) is the potential financial impact resulting from climate change and environmental degradation, influencing financial stability through direct and indirect channels, affecting asset valuations, business operations and overall market dynamics. [14]*

Climate-related and environmental risks can be broadly categorized into two categories: *physical risks* and *transition risks*, each with distinct implications for financial stability and economic activity. Both categories are briefly introduced below.

Physical risks result from the direct impacts of climate change and environmental degradation on economic assets and infrastructure. These risks can be divided into **acute risks**, which result from extreme weather events such as hurricanes, floods and drought, and **chronic risks**, which result from long-term shifts in climate patterns, including sea level rise, temperature rise and biodiversity loss. The financial implications of physical risks include direct asset damage, business disruptions and supply chain vulnerabilities, which can ultimately lead to broader macroeconomic instability.

Transition risks arise from the process of shifting to a lower-carbon and more sustainable economy. These risks are caused primarily by regulatory changes, technological advances, shifts in market sentiment and changing consumer preferences. Institutions may face financial losses due to changes in asset valuations, increased operating costs or reduced market demand for carbon-intensive industries. If the transition occurs in a disordered manner - for example, through abrupt policy implementations or rapid shifts in investor behavior - the financial system may face increased volatility and systemic risk.

Both physical and transition risks interact with existing financial risk categories, such as credit risk, market risk, liquidity risk and operational risk. The extent of their impact depends on institutions' exposure to climate-sensitive sectors and the effectiveness of risk mitigation strategies. To better illustrate the impact of climate and environmental risks on financial institutions, Figure 2.5 presents a table of the key risk drivers categorized into physical and transition risks, published by the ECB [14]. In particular, credit risk is influenced by these risk factors, as changes in climate conditions or policy transitions can alter default probabilities (PD), loss given default (LGD) and collateral valuations. In this research, the focus is specifically on physical risks, meaning climate-related hazards that directly affect residential properties through physical damage—such as flooding, soil subsidence or earthquakes. Since the components impacted by these physical climate events are often risk drivers within credit risk models (e.g. collateral value), this creates an indirect link to the corresponding credit risk metrics (e.g. LGD) that use these risk drivers as input.

For one specific physical risk—flood risk—additional independent research was conducted. Since the overall focus of this study is on a broader set of physical climate risks, this independent analysis was not included in the main literature review. Instead, it has been added as an appendix (Appendix A), serving as a supplementary element rather than a central component. The aim was to conduct a standalone assessment of at least one of the climate risk drivers used later in the modelling phase, in order to better evaluate the underlying assumptions made by the bank.

Examples of climate-related and environmental risk drivers

Risks affected	Physical		Transition	
	Climate-related	Environmental	Climate-related	Environmental
	• Extreme weather events • Chronic weather patterns	• Water stress • Resource scarcity • Biodiversity loss • Pollution • Other	• Policy and regulation • Technology • Market sentiment	• Policy and regulation • Technology • Market sentiment
Credit	The probabilities of default (PD) and loss given default (LGD) of exposures within sectors or geographies vulnerable to physical risk may be impacted, for example, through lower collateral valuations in real estate portfolios as a result of increased flood risk.		Energy efficiency standards may trigger substantial adaptation costs and lower corporate profitability, which may lead to a higher PD as well as lower collateral values.	
Market	Severe physical events may lead to shifts in market expectations and could result in sudden repricing, higher volatility and losses in asset values on some markets.		Transition risk drivers may generate an abrupt repricing of securities and derivatives, for example for products associated with industries affected by asset stranding.	
Operational	The bank's operations may be disrupted due to physical damage to its property, branches and data centres as a result of extreme weather events.		Changing consumer sentiment regarding climate issues can lead to reputation and liability risks for the bank as a result of scandals caused by the financing of environmentally controversial activities.	
Other risk types (liquidity, business model)	Liquidity risk may be affected in the event of clients withdrawing money from their accounts in order to finance damage repairs.		Transition risk drivers may affect the viability of some business lines and lead to strategic risk for specific business models if the necessary adaptation or diversification is not implemented. An abrupt repricing of securities, for instance due to asset stranding, may reduce the value of banks' high quality liquid assets, thereby affecting liquidity buffers.	

Source: ECB.

Figure 2.5: Examples of climate-related and environmental risk drivers, categorized into physical and transition risks [14].

2.3.2 Synthesis and transition

This literature study has laid the theoretical foundation for the next phases of the thesis project. It has explored key components of credit risk modelling, reviewed relevant regulatory developments and introduced the definition of climate risk. By bringing together insights from these domains, it provides the necessary background to begin evaluating the assumptions underlying ABN AMRO's current physical climate risk impact estimation.

While this report was written independently from ABN AMRO's internal data and modelling practices, it provides a structured overview of state-of-the-art, publicly available modeling approaches relevant to this research. The methods and insights discussed—particularly some available in the Appendix A—serve as a conceptual and methodological reference point. While the approach discussed is insightful, it also has limitations, such as the omission of mitigating factors and the focus on scenario-based impacts without considering flood probabilities. These very limitations also contribute to the ability to objectively analyze how ABN AMRO has incorporated climate risks into its models. In other words, they provide a useful basis for properly answering the research sub-questions: examining the structure of ABN AMRO's current climate risk quantification framework (SQ1), identifying its limitations and underlying assumptions (SQ2), and subsequently designing, implementing, and evaluating a potential refinement (SQ3–SQ4).

The next phase of this research will focus on identifying and mapping the key assumptions underlying ABN AMRO's internal climate risk quantification, and on assessing how these assumptions influence the translation of physical climate impacts into credit risk parameters.

3
Methodology

In this research, a specifically predefined modelling objective was not present at the outset of the collaboration with ABN AMRO. Rather than following a predefined modelling template, the research was shaped through iterative analysis and refinement. Given the open-ended nature of the assignment, it was essential to follow a well-structured and adaptive process. In simple terms, the overarching aim was to analyse ABN AMRO's current Climate and Environmental Risk (CER) impact estimation methodology, critically evaluate the assumptions involved and determine whether—and how—this framework could be improved.

To maintain structure throughout the project, the methodology has been divided into four key phases. These phases are summarised below and visualised in Figure 3.1. Each phase corresponds to a distinct part of the modelling journey and collectively forms the backbone of this research.

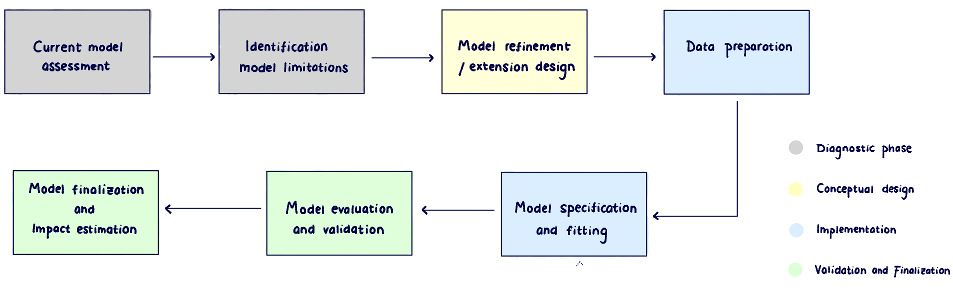

Figure 3.1: Methodological Framework

The first phase focused on fully understanding and analysing ABN AMRO's existing CER estimation approach. What assumptions were made? Which of them are simply debatable nuances, and which could be considered intuitively or practically questionable? At this stage, several directions were still possible—from a broad sensitivity analysis to a targeted model refinement or a more comprehensive extension. This phase involved extensive exploration, brainstorming and critical reflection, ultimately guiding the decision to focus on a limitation considered particularly impactful within the scope of this research.

Once this modelling limitation was identified, the second phase focused on translating the limitation into a set of possible improvements. The central question in this stage became: how can this limitation be addressed in a meaningful and theoretically consistent way? As reflected in the final research question and sub-questions, it was at this point that the scope of the study took shape into a model extension. The challenge then shifted to defining this extension in a structured and coherent design.

The third phase marked the transition from conceptual design to implementation. Based on the theoretically defined extension, it was necessary to explore available data sources, prepare the relevant inputs, define any required new variables and ensure integration into ABN AMRO's existing model infrastructure. The goal was to arrive at a tangible and operational version of the extended model, suitable for comparison with the current setup.

With tangible results from the third phase, the final stage of the process involved evaluating the modelling choices and validating the outcomes. This included reflecting on key assumptions, selecting appropriate input parameters, and—where relevant—adjusting the design. Additionally, expert knowledge played an important role in this stage to assess the realism and consistency of the proposed modelling approach.

Together, these four phases define the full methodological trajectory of this research. The different steps also directly align with the research sub-questions: the current model assessment provides the foundation for examining the structure of ABN AMRO's framework (SQ1); the identification of model limitations addresses the assumptions and constraints within this framework (SQ2); the refinement and implementation phase corresponds to the design of a potential model extension (SQ3); and the evaluation phase links to the validation of this refinement and the interpretation of its outcomes (SQ4). As the project progressed, the scope and structure of the final model naturally took shape through the completion of each phase. The remainder of this chapter is structured along the key components of the methodological process, moving from the current model to its proposed extension, followed by implementation and evaluation. Altogether, this provides a detailed breakdown of the methodological decisions and technical implementation steps taken throughout the study.

3.1 Current model assessment

The first step in the methodological process involved analysing the current setup of ABN AMRO's climate risk impact estimation framework, specifically the part relating to physical climate risks affecting the residential real estate portfolio.

The starting point of the current approach is a relevance assessment: which physical risk drivers are considered relevant within the Dutch banking context? This *risk identification* step involved compiling a list of all climate- and nature-related risks that could potentially affect the mortgage portfolio. The selection was based on evidence from literature and consultations with subject-matter experts, with some climate risks being flagged as immaterial already at this stage. In line with the definition of Climate and Environmental Risk (Definition 2) introduced earlier, the term *channel* will be used throughout the remainder of this report to refer to a distinct transmission pathway through which a specific climate or environmental risk type can affect financial stability. Based on the findings of this preliminary risk identification step, the risks considered relevant for further impact assessment were earthquake risk, foundation risk and flood risk (further subdivided into collateral flood risk and perception flood risk). These risks are discussed each individually later in this section to provide context on their definitions at a high level.

Before introducing the individual climate risk channels, it makes sense to first outline the overarching framework that determines how the physical climate risks are translated into credit risk impact. This is the part that addresses the general question: *which credit risk metrics are selected as the basis for quantifying the impact of CER?* The answer to this closely aligns with the regulatory framework discussed earlier (see Section 2.1.2), as the final outputs of the climate risk estimation process are translated into three key metrics: Risk-Weighted Assets (RWA) and Expected Loss (EL), which will both briefly be revisited in this section for completeness. These regulatory metrics ultimately depend on the three core risk parameters: Probability of Default (PD), Loss Given Default (LGD) and Exposure at Default (EAD). It is important to note that banks develop internal models for PD, LGD and EAD separately. Due to confidentiality, these internal model specifications will not be discussed in detail. However, for the purpose of this research, a general understanding of credit risk modelling concepts is sufficient (see Section 2.1.3).

The remainder of this section outlines how ABN AMRO's current climate risk impact framework processes a general physical risk driver and quantifies its downstream effect on these key credit risk metrics. First, the general logic applicable across all climate risk drivers is explained, after which each relevant climate risk is briefly discussed separately.

3.1.1 Climate and environmental risk impact quantification

This section describes ABN AMRO's current methodology for assessing the impact of physical climate risk drivers on credit risk metrics. The quantification approach consists of three main steps, covering both standalone estimation outputs and intermediate calculations that serve as building blocks within the broader modelling structure.

The first step involves estimating the total yearly costs, calculated separately for each CER channel. These represent the costs incurred by clients for each respective transmission channel. In the second step, part of the methodology from the first step is reused to translate these costs into their impact on the Loss Given Default (LGD). Finally, the output from the second step is used to derive the downstream impact of CER-adjusted LGD on the higher-level credit risk metrics: Expected Loss (EL), Risk-Weighted Assets (RWA), and Economic Capital (EC).

As mentioned in the second step, the climate risk impact is translated solely into an adjustment of the LGD. This is the result of a modelling assumption: when a homeowner experiences damage to the property due to a flood, earthquake or foundation issue, it is assumed that the damage will not be repaired. Under this no-repair assumption, the impact of CER is placed exclusively on the LGD.

The underlying flow in the *no-repair* behavioural path can be summarised as follows. When a climate-related event (CER) causes physical damage to the collateral and the borrower chooses not to repair it, the market value of the property decreases. In case a borrower subsequently defaults, the unrepaired property yields a lower recovery value for the bank. As a result, the share of the exposure that remains unrecovered—captured by the Loss Given No Cure (LGNC) component—increases. This leads directly to a higher Loss Given Default (LGD), as further explained in the general LGD modelling Section 2.1.3. Ultimately, this higher LGD feeds into the expected loss (EL) and capital requirements (RWA) through the standard regulatory formulas.

The rationale behind this no-repair assumption was to maintain a simplified version of the model, in line with the modelling scope previously defined by ABN AMRO. Capturing more complex behavioural responses—such as variations in borrower decisions following property damage—was considered outside the bank's intended scope for the original framework.

Together, the three steps form the complete structure of ABN AMRO's current CER impact quantification process. Each step combines internal, portfolio-specific mortgage data with external, publicly available CER data. This approach is consistent with the expectations set out in the EBA Guidelines on the management of ESG risks (2025) [8]. A visual overview of the current modelling approach is provided in Figure 3.2, where Step I is illustrated using flood risk as an example. The abbreviations and definitions shown in this step—here specified for flood risk—are explained in more detail in Section 3.1.2, where each physical climate risk channel is described. The purpose of this figure is to provide an initial schematic sketch of the current structure. The annotations in Step I should not be overinterpreted; at this stage they are included only to illustrate that this step already involves multiple modelling elements, which differ for each physical climate risk. The current model will subsequently be described in full mathematical detail, while this figure primarily serves as a schematic overview to which the reader can refer back. It additionally illustrates how the defined steps connect to the horizontal behavioural flow at the bottom, which will be explained in more detail in the remainder of this section.

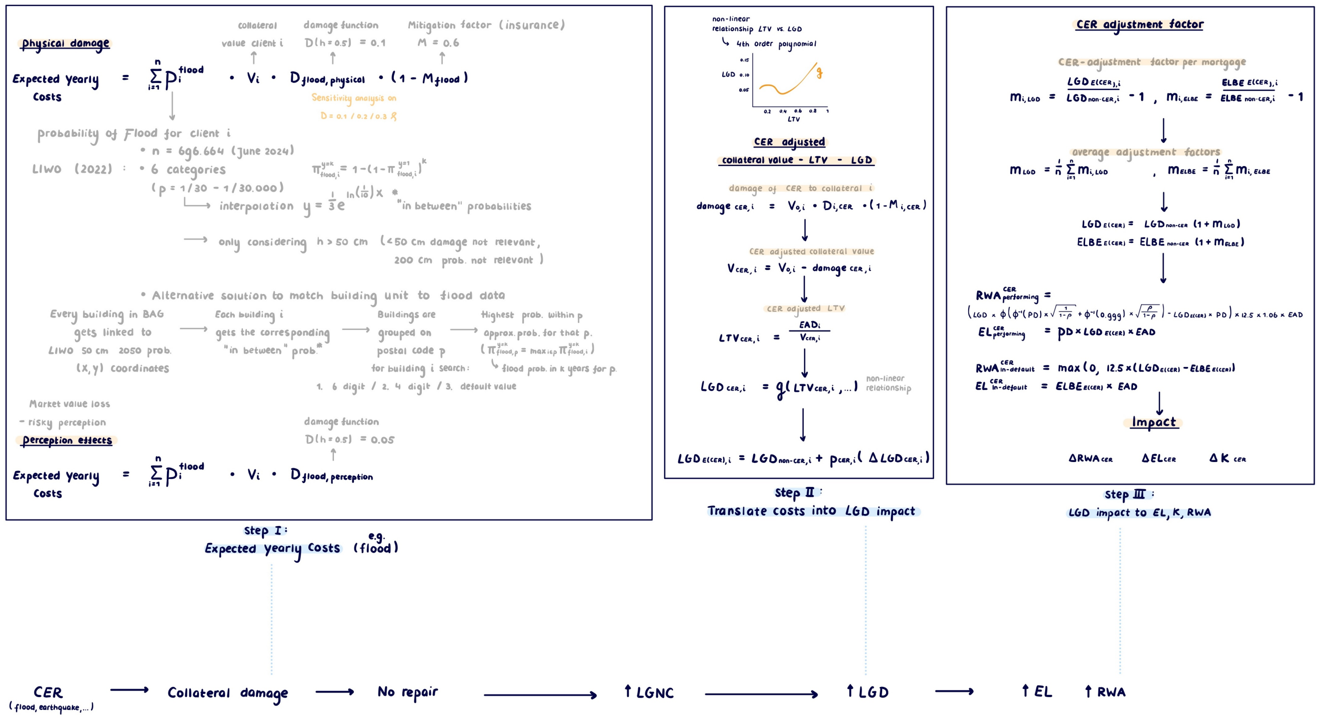

Figure 3.2: ABN AMRO's current CER impact quantification

The remainder of this section will go through these three modelling steps in more detail, beginning with the estimation of total yearly costs.

I: Yearly climate risk costs estimation

The first step is channel-specific and focuses on calculating the average total yearly costs associated with each physical climate risk considered separately in the model. While the exact formulation differs depending on the nature of the risk, the general structure of the calculation remains consistent and is based on the following form

$$\text{Expected yearly costs} = \sum_i \left(p_{i,\text{CER}} \times \text{damage}_{i,\text{CER}}\right).$$

where $p_{i,\text{CER}}$ denotes the annual probability of a CER event affecting mortgage i, and $\text{damage}_{i,\text{CER}}$ refers to the estimated collateral damage in such a case.

Following the bank's terminology, this step is referred to as *channel-specific*, with each *channel* representing a distinct physical climate risk type—such as flooding, foundation degradation, or earthquakes. Throughout the remainder of the report, the term *channel* will therefore be used interchangeably with *climate risk type*.

Accordingly, this step provides a channel-specific view of the overall impact across the mortgage portfolio. Further elaboration on the details of this calculation will be given in the subsections in 3.1.2 dedicated to each specific risk channel.

II: Translation of costs into LGD

As explained earlier, the current CER modelling approach includes assumptions that restrict the climate risk impact solely to the LGD. This is based on the idea that, following a CER event such as a flood, earthquake, or foundation issue, the value of the underlying collateral is affected. Since collateral value is one of the components in the LGD calculation, this damage is reflected in the LGD outcome.

The first part of this step therefore involves calculating the collateral damage conditional on the climate event. This is defined as the original collateral value V_i, reduced by the estimated damage caused by the climate event

$$V_{i,\text{CER}} = V_i - \text{damage}_{i,\text{CER}}. \tag{3.1}$$

Having defined the CER-adjusted collateral value $V_{i,\text{CER}}$, the next step is to assess its effect on a variable directly linked to the LGD calculation—namely, the Loan-to-Value (LTV) ratio (2.10). After incorporating a CER-related damage, this ratio reflects the extent to which the collateral value is affected, expressed relative to the outstanding loan. In this sense, it can also be interpreted as a normalised proxy for the borrower's residual exposure after climate-related property damage. Formally, it can be defined as

$$LTV_{i,\text{CER}} = \frac{EAD_i}{V_{i,\text{CER}}}. \tag{3.2}$$

where EAD_i is the Exposure at Default, which for residential mortgages corresponds to the outstanding loan amount.

Given the CER-adjusted LTV, it is now possible to translate this into a corresponding LGD impact. Mainly due to confidentiality reasons—and also because understanding the precise formulation is not required for the purposes of this study—this relationship is not further elaborated. For notation purposes, let $g(\cdot)$ denote the internally defined non-linear function that determines the LGD, which among other factors depends on the LTV. This can formally expressed as follows

$$LGD_{i,\text{CER}} = g(LTV_{i,\text{CER}}, \dots). \tag{3.3}$$

The value above represents the LGD in the hypothetical case where the CER event occurs with probability one. In practice, such events are rare, and the associated probability $p_{\text{CER},i}$ must be taken into account. More details on how these probabilities are determined can be found in the upcoming channel-specific sections in 3.1.2. This leads to the definition of the expected CER-adjusted LGD, calculated as follows

$$\begin{aligned} LGD_{i,E(\text{CER})} &= LGD_{i,\text{non-CER}} + p_{i,\text{CER}} \cdot (LGD_{i,\text{CER}} - LGD_{i,\text{non-CER}}) && (3.4)\\ &= LGD_{i,\text{non-CER}} + p_{i,\text{CER}} \cdot \Delta LGD_{i,\text{CER}}. && (3.5) \end{aligned}$$

where:

- $LGD_{i,\text{CER}}$: the LGD for mortgage i after accounting for collateral damage due to CER;
- $LGD_{i,\text{non-CER}}$: the original LGD before any climate impact;
- $p_{i,\text{CER}}$: the probability per year of a CER event occurring for mortgage i;
- $\Delta LGD_{i,\text{CER}}$: the incremental change in LGD due to the CER event.

This completes the second step in the current estimation framework, in which physical damage to the collateral is translated into an expected credit risk impact on the LGD.

III: Integration into regulatory risk metrics (EL, RWA)
Now that the expected CER-adjusted LGD has been derived, the next step is to translate this into the corresponding impact on regulatory key credit risk metrics.

It is important to note that a fundamental distinction is made between performing loans—where borrowers are up to date with their payments—and defaulted loans, where payment obligations have not been met.

This classification is essential, as it influences how each exposure is treated in the risk calculations. Without going into full technical detail, this distinction has consequences specifically for the LGD-related calculations. As a result, the computations are executed separately: once on data for performing clients, and once on clients who are already classified as in-default. The key difference lies in the role of the Expected Loss Best Estimate (ELBE): for performing exposures, only the LGD is used in the RWA calculations, whereas for defaulted exposures both the LGD and ELBE are required. While the LGD model computes ELBE values for all exposures, it is only relevant in this context for those already classified as in-default.

So far, all calculations have been carried out at the mortgage level—that is, for each individual client i. This has yielded, per client, an expected LGD and ELBE, denoted respectively as $LGD_{i,E(\text{CER})}$ and $ELBE_{i,E(\text{CER})}$. To obtain a portfolio-level impact, these individual CER-affected values need to be aggregated. This is done by first constructing a so-called "CER adjustment factor per mortgage" and subsequently computing the portfolio average of these individual adjustment factors. This results in a single, generalised expected CER-adjusted LGD and ELBE estimate. Formally, this is defined as

$$M_{i,LGD,\text{CER}} = \frac{LGD_{i,E(\text{CER})}}{LGD_{i,\text{non-CER}}} - 1 \quad \rightarrow \quad \overline{M}_{LGD,\text{CER}} = \frac{1}{n}\sum_{i} M_{i,LGD,\text{CER}}, \tag{3.6}$$

$$LGD_{E(CER)} = LGD_{\text{non-CER}}(1 + \overline{M}_{LGD,\text{CER}}). \tag{3.7}$$

$$M_{i,ELBE,\text{CER}} = \frac{ELBE_{i,E(\text{CER})}}{ELBE_{i,\text{non-CER}}} - 1 \quad \rightarrow \quad \overline{M}_{ELBE,\text{CER}} = \frac{1}{n}\sum_{i} M_{i,ELBE,\text{CER}}, \tag{3.8}$$

$$ELBE_{E(CER)} = ELBE_{\text{non-CER}}(1 + \overline{M}_{ELBE,\text{CER}}). \tag{3.9}$$

In the current approach, these expected values are then translated into the key regulatory metrics: Expected Loss (EL) and Risk-Weighted Assets (RWA). For further background on the regulatory setup of these metrics in the context of performing exposures, see Section 2.1.2. The CER-adjusted versions of these metrics are fully aligned with the regulatory framework for residential mortgage portfolios, with the only change being that the CER-adjusted LGD and/or ELBE are used in place of the standard parameters. For the performing exposures, these are defined as follows

$$RWA^{\text{performing}}_{\text{CER}} = 12.5 \cdot K_{E(CER)} \cdot EAD \tag{3.10}$$

$$EL^{\text{performing}}_{\text{CER}} = PD \cdot LGD_{E(CER)} \cdot EAD \tag{3.11}$$

with

$$K_{E(CER)} = LGD_{E(CER)} \cdot \Phi\left(\frac{\Phi^{-1}(PD)}{\sqrt{1-\rho}} + \sqrt{\frac{\rho}{1-\rho}} \cdot \Phi^{-1}(0.999)\right) - PD \cdot LGD_{E(CER)}, \tag{3.12}$$

for definitions of all variables, see (2.3). In the case of in-default exposures, the corresponding metrics are defined as

$$RWA^{\text{in-default}}_{\text{CER}} = \max\left\{0,\ 12.5 \cdot \left(LGD_{E(\text{CER})} - ELBE_{E(\text{CER})}\right) \cdot EAD\right\}, \tag{3.13}$$

$$EL^{\text{in-default}}_{\text{CER}} = ELBE_{E(CER)} \cdot EAD. \tag{3.14}$$

To provide additional clarification on how the RWA formulations arise in the cases of performing and defaulted portfolios, a brief explanation is included. Recall from Section 2.1.2 the regulatory definition of capital requirements under the IRB approach: banks are required to hold at least 8% of their Risk-Weighted Assets (RWA) as capital. This gives the general identity RWA $= \frac{1}{0.08} \cdot$ Capital Requirement $= 12.5 \cdot \text{UL} \cdot \text{EAD}$, where UL denotes the unexpected loss. This expression forms the basis for both the performing and defaulted exposure RWA formulations. For performing exposures, the capital requirement is driven by the uncertainty about whether a loss will occur. The unexpected loss component is derived using the Basel IRB risk weight function K, which— as shown in Section 2.1.2 — precisely reflects the UL in expectation, conditional on default. For defaulted exposures, the probability of default is considered to be 1, and the remaining uncertainty lies in the magnitude of the actual loss. In this case, the expected loss is given by the best estimate ELBE, while the downturn-adjusted LGD serves as a conservative upper bound. The unexpected loss is thus defined as the difference between LGD and ELBE. In both cases, the formulation ensures consistency with the same capital logic: *RWA always reflects the capital required to cover unexpected losses*, either incorporating the probability of default or conditioned on the fact that default has already occurred.

The final CER impact is then computed as the difference between the CER-adjusted metric and its corresponding non-CER baseline, which follows the same formulaic structure as shown above, but uses the original LGD and/or ELBE values without incorporating any CER-related effects. In full, this yields the total absolute CER impact on each of the three regulatory measures

$$EL_{\text{total CER impact}} = \Delta EL_{\text{CER}} = EL_{E(CER)} - EL_{\text{non-CER}}, \tag{3.15}$$

$$RWA_{\text{total CER impact}} = \Delta RWA_{\text{CER}} = RWA_{E(CER)} - RWA_{\text{non-CER}}, \tag{3.16}$$

This completes the final step in the current CER impact estimation process used for ABN AMRO's residential mortgage portfolio.

3.1.2 Physical climate risk channels

Up to this point, each physical climate risk has been referred to under the general label "CER". As previously explained, this label covers a set of physical risk types classified as relevant in the Dutch banking context—flood risk, earthquake risk and foundation risk.

To provide additional context, this subsection briefly outlines, for each of these climate risk types, how the corresponding damage estimate $\text{damage}_{i,\text{CER}}$ has been constructed. Where applicable, further information is also provided on how the associated probabilities of the CER events occurring have been derived.

Flood risk

The first climate risk channel to be discussed is flood risk. Within this category, a distinction is made between physical flood risk and perception-based flood risk. The former refers to the direct financial loss resulting from physical damage to the property due to a flood. In contrast, the perception component refers to the possible decrease in property values in nearby areas affected by a flood—driven by increased perceived risk among homeowners and potential buyers, even if no actual damage occurred to a specific property [30–33].

For both cases, a brief explanation is provided below on how the damage estimations and associated event probabilities have been constructed.

For each building i, the expected yearly physical damage is calculated as

$$\sum_i P_i^{\text{flood}} \cdot V_i \cdot D_{i,\text{flood}}^{\text{physical}} \cdot (1 - M_{i,\text{flood}}),$$

where

- P_i^{flood}: the probability of a flood event affecting building i,
- V_i: the collateral value of the building,
- $D_{i,\text{flood}}^{\text{physical}}$: the estimated damage factor (chosen: 0.1 for 10% damage),
- $M_{i,\text{flood}}$: the mitigation or insurance factor (chosen: 0.6 reflecting 60% covered).

The flood probabilities are derived from the LIWO database (Landelijk Informatiesysteem Water en Overstromingen) [34], with interpolation applied across six defined flood categories. Only floods above a certain threshold (50 cm) are considered relevant, as smaller events do not cause significant damage and larger events have insignificant probability of happening. The chosen damage factor of 10% is based on estimates of average flood losses relative to collateral value, as reported in [35]. This value is intended to represent the order of magnitude of major flood events in the Netherlands and is consistent with previous scenario analyses. The mitigation factor of 60% originates from literature on the role of insurance coverage in reducing effective losses [36]. It reflects the assumption that, on average, a significant share of building damages is compensated by insurance schemes.

In addition to physical damage, perception effects are modelled as

$$\sum_i P_i^{\text{flood}} \cdot V_i \cdot D_{i,\text{flood}}^{\text{perception}} \tag{3.17}$$

where $D_{i,\text{flood}}^{\text{perception}}$ represents a smaller devaluation factor (chosen: 0.05). This perception damage factor of 5% is based on literature reporting average decreases in property values following flood events [37–39]. In this perception channel, no mitigation factor is applied, as insurance does not cover value losses driven by market sentiment.

Foundation risk

Foundation risk is the second relevant physical climate risk identified. It refers to the progressive degradation of building foundations over time, primarily due to subsidence or biological deterioration. In the Netherlands, a large number of homes are built on wooden pile foundations, which are particularly vulnerable when groundwater levels drop. Prolonged periods of drought or long-term shifts in soil moisture caused by climate change can expose these wooden piles to oxygen, leading to bacterial rot and structural weakening. For buildings without pile foundations, changing soil conditions—such as the shrinkage of clay or peat soils—can cause uneven settlement and cracks in walls or floors. These processes can reduce the structural integrity and value of a property.

The estimated financial impact of foundation risk is primarily based on recent empirical work by Hommes et al. (2023) [40]. Their study investigates how reported foundation damage affects housing prices in the Dutch market. Using transaction data and natural language processing on property listings, they find that when foundation problems are explicitly mentioned, the sale price drops by an average of 12% compared to model estimates. When compared to properties where the foundation has already been repaired (which show a +2% uplift), the price difference widens to over 14%. This relative decrease is used as a proxy for the financial damage caused by foundation degradation, as it reflects market expectations about future repair costs and associated risks.

The annual probability of foundation damage is estimated at 0.4%, based on the national outlook provided by the Council for the Environment and Infrastructure (Rli) [41]. Their 2024 report estimates that approximately 425,000 buildings, or 6% of the national building stock, will be affected by foundation problems between now and 2035. Assuming a uniform distribution over this 15-year period, this translates to an annual incidence rate of 0.4%. While this is a simplified assumption, it provides a consistent basis for estimating expected damage across large mortgage portfolios.

Using these externally based sources, the implicit assumption is made that the AAB mortgage portfolio is representative of general residential properties in the Netherlands. Following the same general structure as in the impact estimation for flood risk, the total expected yearly physical damage for each building i is calculated as

$$\sum_{i} P_i^{\text{foundation}} \cdot V_i \cdot D_{i,\text{foundation}} \tag{3.18}$$

where, as explained, $P_i^{\text{foundation}}$ is set to the general value of 0.004 [41], V_i is the collateral value of client i, and the damage function $D_{i,\text{foundation}}$ is set to 0.14 [40].

It is important to note that, unlike flood risk, foundation risk is considered a *chronic* physical risk. As highlighted by Deltares [42], damage to foundations does not result from a single extreme event, but rather accumulates slowly over time due to recurring droughts, subsidence and long-term changes in groundwater levels. For the Netherlands specifically, foundation risk currently happens to have the highest incidence rate among the considered climate risks. However, the financial impact is not solely driven by this relatively high likelihood of occurrence, but also by this persistent and structural nature of the risk. As will be shown later, this risk leads to by far the highest estimated damages. This chronic characteristic helps explain why foundation risk accounts for a larger share of the overall expected physical climate risk.

Earthquake risk

The last relevant physical climate risk identified for the Netherlands is *earthquake risk*, which here refers specifically to *natural earthquakes*—thus excluding earthquakes induced by gas extraction. The only significant natural earthquake in recent Dutch history occurred in Roermond in 1992 [43]. For the purpose of impact estimation, this event is used as the sole reference point, under the assumption that a similar earthquake in the same region would result in comparable damage.

Based on estimates, the Roermond earthquake caused a total of approximately €165 million in damage [44](adjusted for inflation [45]). Using data from the Dutch BAG register [46], Roermond had around 16,561 residential buildings at the time. Dividing the total damage by this number gives an average estimated damage of €9,963 per residential property. This is a conservative estimate, as the damage figure also includes non-residential buildings and infrastructure. However, due to the limited availability of relevant data on natural earthquakes in the Netherlands, this approach provides a pragmatic benchmark.

Moving to the probability of such an event occurring, regional earthquake risk assessments show that only the provinces of Noord-Brabant and Limburg face a notable risk of natural earthquakes. For the rest of the Netherlands, this probability is therefore assumed to be zero. According to the Global Facility for Disaster Reduction and Recovery (GFDRR) [47], there is a 10% chance that a potentially damaging earthquake will occur in one of these southern provinces over the next 50 years. This estimate is based on global earthquake hazard models, such as those developed by the Global Earthquake Model (GEM) [48] and the US Geological Survey (USGS) [49], which combine historical earthquake data, tectonic fault lines and local ground conditions. Assuming this probability is evenly spread over time, the annual likelihood of such an earthquake is estimated at approximately 0.21%.

Following the same approach as for the other relevant risks, the expected annual physical damage due to earthquake risk is estimated by combining the assumed per-property damage with the number of affected mortgages and the annual probability of occurrence. Based on the earlier described assumptions, the expected yearly earthquake damage is defined as

$$\sum_{i \in \text{LB}} P_i^{\text{earthquake}} \cdot D_{i,\text{earthquake}} \tag{3.19}$$

Here, $P_i^{\text{earthquake}}$ denotes the annual probability of a damaging natural earthquake in the relevant region, set to 0.0021 [47]; and $D_{i,\text{earthquake}}$ is the assumed damage per residential property, set to €9,963 [44, 45]. The summation is taken over all mortgage exposures in the provinces of Limburg and Noord-Brabant.

3.1.3 Identified limitations

Once a full understanding of the current CER impact estimation framework was established, the next step in the methodological process involved identifying potential limitations in the existing setup. Despite the already careful and well-structured approach—characterised by a high level of detail—certain assumptions and modelling choices could still benefit from further review or refinement through additional research.

At this point in the process, multiple possible directions were available for further exploration. Several simplifications and assumptions have been made throughout the modelling process—from parameter-specific assumptions within each CER channel to higher-level structural modelling choices. While various limitations were reviewed and discussed, this section focuses exclusively on the one that was ultimately selected to be addressed. It was decided to investigate a limitation rising from one of the assumptions made early in the modelling framework—specifically, the assumption of *no repair*.

This assumption implies that whenever climate-related damage occurs, the property is not repaired. As a result, several behavioural responses are excluded from the modelling rationale, and the impact of CER is restricted solely to the LGD. However, this assumption may not fully reflect how homeowners typically respond to physical damage in practice. For instance, in the case of minor property damage, it is reasonable to expect that many homeowners would explore repair options rather than immediately abandoning the property. While the no-repair assumption was likely made to maintain a more simplified modelling structure, it inherently limits the range of possible behavioural responses. For that reason, it was identified as a key modelling simplification worth further exploration in this research.

Figure 3.3: Behavioural flow no-repair assumption

To provide further clarity on the current setup, Figure 3.3 recalls the existing behavioural flow—from the moment of damage to the resulting impact on the regulatory metrics. In the no-repair scenario, unrepaired damage leads to a lower collateral value, increasing the LGNC component and ultimately resulting in higher LGD, expected loss (EL) and risk-weighted assets (RWA).

When aiming to expand the current flow to incorporate multiple behavioural pathways, a central question arises: *how would the behavioural flow change if the option to repair the property were explicitly incorporated into the model?* This question provided the foundation for a more in-depth conceptual investigation, which ultimately guided the formulation of the extended model design outlined in the next section.

3.2 Model extension: framework

This section presents the extended model framework. The core idea of the extension is to incorporate the option of 'repair' into the current setup, along with the behavioural choices that follow from this addition. The section will first focus on the structure of the extended behavioural flow, and provide a high-level overview of the types of credit risk impact it introduces. This will then be followed by a more detailed explanation of how the behavioural structure translates into the extended mathematical formulation of the CER impact estimation.

3.2.1 Behavioural dependence

Once the decision was made to extend the model by introducing behavioural variation resulting from the possibility to repair, the first step was to define what behavioural options should be considered. This behavioural structure forms the foundation for the remainder of the model extension, making it essential to clearly understand the possible pathways following climate-related damage to a property.

The most intuitive way to approach this is to consider the funding source: if a property is damaged, financial resources are needed to cover the repair. Therefore, the behavioural options can be defined by the source from which this funding is obtained. In this study, three categories were distinguished: (1) self-funded repairs, (2) repairs funded through an additional loan from ABN AMRO Bank (AAB), and (3) repairs financed externally—either through a loan from another financial institution or from private sources. These three categories are considered to comprehensively capture all realistic repair funding pathways.

The key question is: how do these repair strategies affect the core credit risk parameters? The remainder of this subsection provides a high-level overview of the expected impact of each behavioural path on the credit risk measures PD, LGD and EAD. In all repair scenarios, the collateral value remains unchanged, as it is assumed that any loss in value caused by the damage is offset by the repair, restoring the property's value to its original level.

Self-funded In this case, the homeowner covers the repair cost entirely using personal financial resources, without taking on additional debt. Whether this is done through available income or through previously accumulated savings is not yet specified at this point, but in either case, the homeowner's available financial capacity is directly reduced. This reduction in available resources can be expected to negatively affect the borrower's creditworthiness. While the exact features used in PD models are not discussed here, it can be broadly understood that creditworthiness intuitively directly relates to default probability: when financial capacity decreases, the ability to meet mortgage obligations becomes more strained. Therefore, this behavioural path is assumed to have a direct and increasing effect on the Probability of Default (PD), with no impact on LGD or EAD.

Bank-funded (loan AAB) Under this scenario, the borrower takes out an additional loan from ABN AMRO to finance the repair. This introduces several effects that influence different parts of the risk framework. First, taking on a new loan increases the total exposure to the bank. In the context of residential mortgages, this directly translates into a higher Exposure at Default (EAD). In addition, because the LTV (Loan-to-Value) ratio is a key driver within the LGD model, an increase in exposure also increases the LTV, which subsequently leads to a higher Loss Given Default (LGD). Moreover, the additional loan also impacts the borrower's financial position. Since the total outstanding debt increases, the borrower's capacity to service debt relative to income declines. As a result, creditworthiness decreases, which—similar to the self-funded case—translates into a higher PD.

Externally funded In the third behavioural path, the borrower seeks funding outside of ABN AMRO—for example, via another bank or personal contacts. In this case, the additional loan does not increase ABN AMRO's recorded exposure, and thus does not directly affect EAD or LGD. However, it is still assumed that the borrower's creditworthiness is affected. Even when borrowing externally, there are typically repayment obligations (interest and/or principal), which reduce the borrower's future disposable income. Therefore, this path also implies a creditworthiness decline, and hence an increase in PD.

In all three behavioural scenarios, the strength of the impact—whether on creditworthiness or exposure—depends on the nature of the financing method. The purpose of this section was to outline the behavioural differentiation framework on a conceptual level. The mathematical implementation and exact modelling of these effects will be addressed in the next subsection. A summary overview of the behavioural pathways and their qualitative impact is visualised in Figure 3.4.

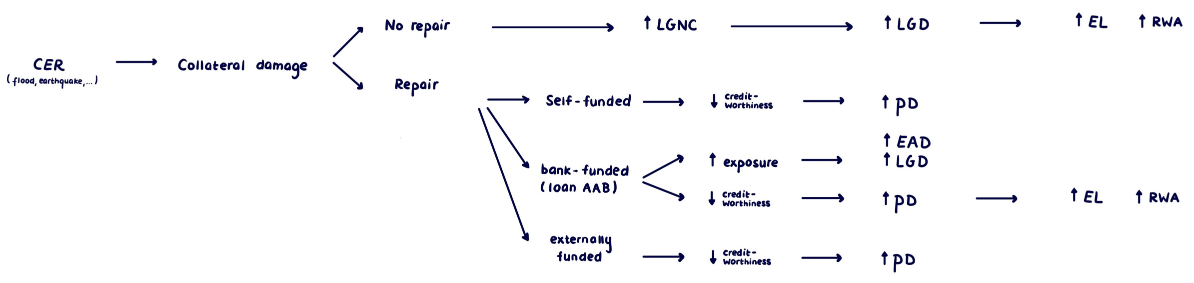

Figure 3.4: Behavioural flow including repair assumption: four distinct behavioural pathways (no-repair, self-funded repair, AAB-funded repair, externally-funded repair)

3.2.2 Mathematical formulation of the extended framework

Now that the behavioural flow has been clearly defined, the next step is to formalise each behavioural pathway within the structure of the existing CER impact estimation framework. Since the current setup already offers a comprehensive and well-organised structure, the modelling extension was designed to build upon this foundation rather than replacing it entirely.

In practice, this means that the original three-step structure—ranging from the estimation of yearly climate-related costs, to the translation of those costs into credit risk parameters, and finally to the integration into overall risk metrics such as EL, RWA, and EC—is maintained. The objective of this section is to incorporate the newly introduced behavioural dependencies into these stages.

It is important to note that the behavioural differentiation only becomes relevant from Step II onward, specifically in the translation of estimated damage per risk channel into credit risk parameters. In the original setup, this step only affected the LGD; in the extended framework, the impact is now distributed across LGD, PD and EAD. Step I of the model remains unchanged under the proposed extension, as it concerns the estimation of channel specific damage costs—an aspect not affected by the introduced behavioural path dependencies. For this reason, it is not repeated here and the formulation begins directly from Step II.

A visual representation of the complete extended framework using the added behavioural paths, is provided in Figure 3.5.

II (Extension): Path-dependent translation of costs into PD, LGD and EAD

In this step, the behavioural nature of the model extension becomes most apparent, as each financing strategy leads to a slightly different integration of CER impact. A climate-related effect must be quantified across all three core credit risk metrics: Probability of Default (PD), Loss Given Default (LGD) and Exposure at Default (EAD).

For LGD, the existing input feature—Loan-to-Value (LTV)—can be used directly, as it is known to have a clearly defined relationship with LGD in the underlying model as described in Section 3.1.1. This is consistent with how LGD sensitivity is also captured in the initial setup, and the same relation is therefore retained within the extended framework. For EAD, the required adjustment is likewise relatively straightforward and follows directly from the structure of the behavioural setup.

The only metric for which this is less straightforward is PD. The defined driver influencing PD within this behavioural framework is the borrower's creditworthiness. Since a change in creditworthiness occurs under multiple behavioural paths, this concept must first be introduced before proceeding further.

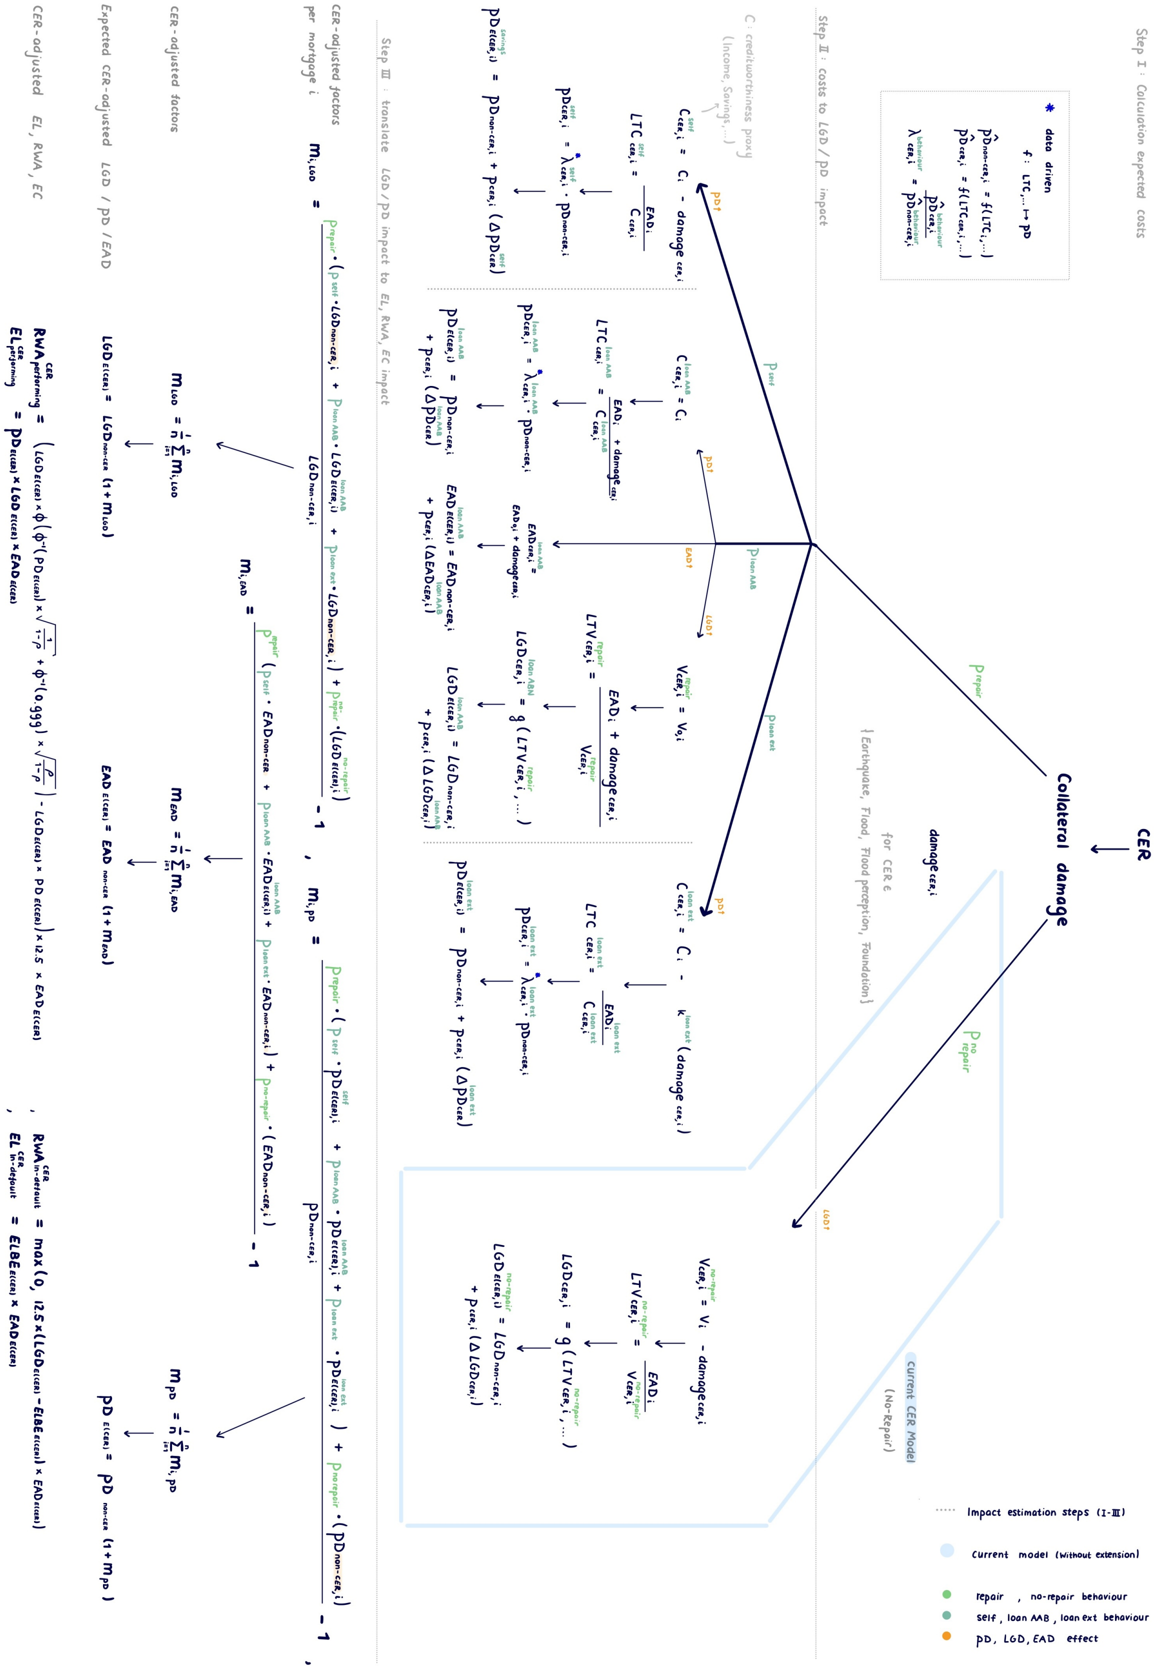

Figure 3.5: Model extension design

Firstly, the model requires a variable that can act as a proxy for creditworthiness, which must then be scaled relative to the size of the outstanding loan—much like the LTV does for property value. At the time of designing the model extension, it was not yet possible to finalise this creditworthiness variable, as this heavily depended on the availability and structure of the data—something that was planned to be determined during the subsequent data exploration phase. Therefore, in this model extension design section, creditworthiness is denoted by the placeholder variable C, and its scaled version as LTC ("Loan-to-Creditworthiness"). These serve as structural placeholders that will later be replaced with concrete specifications once the data analysis has been completed.

Secondly, an underlying assumption as discussed is that a borrower's probability of default is intuitively linked to their creditworthiness. While the exact form and estimation of this relationship will be elaborated upon in the implementation section, the remainder of this subsection introduces the notation used to formalise this dependency in the extension framework.

It is assumed that a function f exists, mapping the creditworthiness proxy—represented by LTC—and potentially other explanatory variables, to a predicted probability of default

$$f : (\mathrm{LTC}, \dots) \mapsto \mathrm{PD}. \tag{3.20}$$

If this function were applied to both the original (non-CER-adjusted) creditworthiness and the CER-adjusted creditworthiness, the resulting predicted default probabilities—based on the relationship defined by f—would be expressed as follows

$$\hat{\mathrm{PD}}_{i,\text{non-CER}} = f(\mathrm{LTC}_i, \dots), \tag{3.21}$$

$$\hat{\mathrm{PD}}_{i,\mathrm{CER}} = f(\mathrm{LTC}_{i,\mathrm{CER}}, \dots). \tag{3.22}$$

This functional form captures the behavioural dependence in PD through shifts in creditworthiness, which will later be further specified based on observed data.

To maintain clarity, the four behavioural paths are described separately in the paragraphs that follow, with key modelling decisions justified where relevant.

No-Repair The first behavioural path to consider is the already existing one: the *no-repair* option. In terms of structure and underlying assumptions, this path is fully consistent with the current CER impact estimation introduced earlier.

As defined earlier in the baseline setup, this path assumes that the physical damage caused by a CER event remains unrepaired. Although previously introduced, the formulation is briefly restated here under the behavioural label *no-repair* to ensure consistency across all paths in the extended framework. The damage reduces the collateral value, which in turn affects the Loan-to-Value (LTV) ratio. This revised LTV feeds into the LGD model, ultimately resulting in a CER-adjusted LGD and its expected counterpart. These steps are formalised below.

$$V_{i,\mathrm{CER}}^{\text{no repair}} = V_i - \mathrm{damage}_{i,\mathrm{CER}} \tag{3.23}$$

$$LTV_{i,\mathrm{CER}}^{\text{no repair}} = \frac{\mathrm{EAD}_i}{V_{i,\mathrm{CER}}^{\text{no repair}}} \tag{3.24}$$

$$\mathrm{LGD}_{i,\mathrm{CER}}^{\text{no repair}} = g(LTV_{i,\mathrm{CER}}^{\text{no repair}}, \dots) \tag{3.25}$$

$$\mathrm{LGD}_{i,E(\mathrm{CER})}^{\text{no repair}} = \mathrm{LGD}_{i,\text{non-CER}} + p_{i,\mathrm{CER}} \cdot (\Delta \mathrm{LGD}_{i,\mathrm{CER}}^{\text{no repair}}) \tag{3.26}$$

From this point onward, the behavioural pathways that form the core of the model extension are introduced. To clarify which components constitute the actual extension to the existing framework, Figure 3.6 presents a zoomed-in view of Step II—focusing specifically on the loan behaviour pathways within the repair branch of the model.

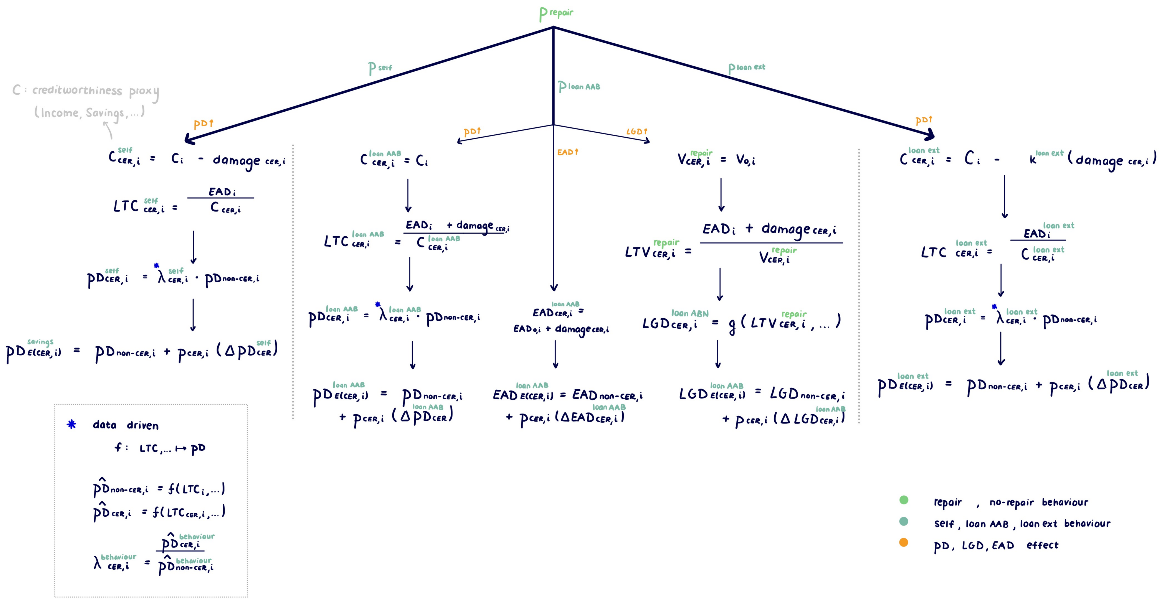

Figure 3.6: Zoomed-in view of Step II, highlighting the loan behaviour pathways within the repair segment of the model extension

Repair — Self-funded The first behavioural path under the repair scenario considers the case in which the borrower decides to fully cover the repair costs out of pocket. As previously discussed, this choice affects only the probability of default (PD), through a direct impact on the borrower's creditworthiness.

At this stage of the model extension, it is assumed that the cost of repair is subtracted directly from the absolute creditworthiness proxy C. The magnitude of this deduction equals the estimated damage amount. As such, creditworthiness under the CER scenario is expressed as

$$C^{\text{self}}_{i,\text{CER}} = C_i - \text{damage}_{i,\text{CER}}. \tag{3.27}$$

As in the current setup, this absolute measure is then scaled by the outstanding loan amount, resulting in the loan-to-creditworthiness ratio (LTC). In the case of residential mortgages, the outstanding loan is equal to the exposure at default (EAD), yielding

$$LTC^{\text{self}}_{i,\text{CER}} = \frac{\text{EAD}_i}{C^{\text{self}}_{i,\text{CER}}}. \tag{3.28}$$

The adjusted LTC is subsequently mapped to the PD using the predictive function f, introduced earlier in equations (3.21–3.22). From this, a relative CER impact factor λ can be defined, representing the proportional change in PD under the self-funded repair scenario

$$\lambda^{\text{self}}_{i,\text{CER}} = \frac{\widehat{\text{PD}}^{\text{self}}_{i,\text{CER}}}{\widehat{\text{PD}}_{i,\text{non-CER}}}. \tag{3.29}$$

This formulation serves as the general structure for calculating the relative impact factor λ, which is similarly applied to the other behavioural scenarios (e.g. AAB- and externally-funded repair), with only the superscript adjusting to reflect the behavioural path. This factor is applied to the actual observed non-CER PD to obtain the CER-adjusted PD for borrower i

$$\text{PD}^{\text{self}}_{i,\text{CER}} = \lambda^{\text{self}}_{i,\text{CER}} \cdot \text{PD}_{i,\text{non-CER}}. \tag{3.30}$$

Finally, the expected PD is computed by incorporating the probability of the CER event, consistent with the methodology applied in the original setup

$$\text{PD}^{\text{self}}_{i,E(\text{CER})} = \text{PD}_{i,\text{non-CER}} + p_{i,\text{CER}} \cdot \left(\Delta\text{PD}^{\text{self}}_{i,\text{CER}}\right). \tag{3.31}$$

Repair — Bank-funded (loan AAB) The second behavioural category within the repair scenario involves covering the repair costs by taking an additional loan from ABN AMRO. As discussed in the behavioural dependence section, this path results in a threefold impact: it affects the probability of default (PD), the loss given default (LGD) and the exposure at default (EAD). For each of these components, the relevant driver of impact will be introduced and its modification under this behavioural scenario explained.

- **PD**

  In the case where the borrower chooses to finance the repair through an additional loan from the bank, it is implicitly assumed that this does not come at the expense of personal equity. Although in reality such a loan would typically involve interest payments and future repayments, it was decided to keep the absolute creditworthiness proxy constant at its initial value. This simplification is justified by the wide variety of mortgage types and repayment structures (e.g. interest-only, annuity, linear, savings-based, investment-linked mortgages), which would otherwise introduce unnecessary complexity into the model.

  More importantly, the dominant financial impact under this repair path occurs through the increase in the borrower's Loan-to-Creditworthiness (LTC) ratio. Specifically, the numerator of the LTC increases due to the newly issued loan, which is assumed to cover the full amount of the climate-related damage. As a result, the exposure at default (EAD) increases by exactly the amount of the damage, while the absolute creditworthiness proxy C remains unchanged

$$C^{\text{loan AAB}}_{i,\text{CER}} = C_i, \tag{3.32}$$

$$LTC^{\text{loan AAB}}_{i,\text{CER}} = \frac{\text{EAD}_i + \text{damage}_{i,\text{CER}}}{C^{\text{loan AAB}}_{i,\text{CER}}}. \tag{3.33}$$

  The corresponding adjustment in the PD follows the same structure as described in the previous repair scenario. A to be defined relationship function f again drives the borrower-specific relative change factor λ, resulting in

$$\text{PD}^{\text{loan AAB}}_{i,\text{CER}} = \lambda^{\text{loan AAB}}_{i,\text{CER}} \cdot \text{PD}_{i,\text{non-CER}} \tag{3.34}$$

$$\text{PD}^{\text{loan AAB}}_{i,E(\text{CER})} = \text{PD}_{i,\text{non-CER}} + p_{i,\text{CER}} \cdot (\Delta\text{PD}^{\text{loan AAB}}_{i,\text{CER}}) \tag{3.35}$$

- **EAD**

  The change in the exposure at default (EAD) is more straightforward. As mentioned previously, in the case of residential mortgages, the outstanding loan amount is typically equal to the EAD. Therefore, if the borrower takes an additional loan equal to the amount of the climate-related damage, the EAD increases directly by that same amount.

Unlike the effects on LGD and PD, no relationship function is needed to define this impact — the effect on EAD can be formulated directly as follows

$$\text{EAD}_{i,\text{CER}}^{\text{loan AAB}} = \text{EAD}_i + \text{damage}_{i,\text{CER}}, \tag{3.36}$$

$$\text{EAD}_{i,E(\text{CER})}^{\text{loan AAB}} = \text{EAD}_{i,\text{non-CER}} + p_{i,\text{CER}} \cdot (\Delta\text{EAD}_{i,\text{CER}}^{\text{loan AAB}}). \tag{3.37}$$

The expected CER impact is again derived in a similar manner as before, incorporating the probabilities of the respective CER events occurring.

- **LGD**
  Lastly, the impact on LGD is also considered in this path. However, it is not defined in the same way as in the no-repair case, where a reduction in the collateral value is assumed. In the repair scenario, the underlying assumption is that the property is fully restored, and thus its value remains unchanged—i.e. the denominator of the LTV remains equal to the original, non-CER-adjusted collateral value. What does change in this context is the exposure—the numerator of the LTV—as the borrower is assumed to take out a new loan from ABN AMRO to cover the repair costs. This leads to the following set of formulations

$$V_{i,\text{CER}}^{\text{repair}} = V_i \tag{3.38}$$

$$LTV_{i,\text{CER}}^{\text{repair}} = \frac{\text{EAD}_i + \text{damage}_{i,\text{CER}}}{V_{i,\text{CER}}^{\text{repair}}} \tag{3.39}$$

$$\text{LGD}_{i,\text{CER}}^{\text{loan AAB}} = g(LTV_{i,\text{CER}}^{\text{repair}}, \dots) \tag{3.40}$$

$$\text{LGD}_{i,E(\text{CER})}^{\text{loan AAB}} = \text{LGD}_{i,\text{non-CER}} + p_{i,\text{CER}} \cdot (\Delta\text{LGD}_{i,\text{CER}}^{\text{loan AAB}}) \tag{3.41}$$

Repair — Externally funded The final behavioural path considered is the case in which the borrower proceeds with repair, but the required funds are obtained externally. This may include financing through another bank, or borrowing from personal connections such as friends or family.

From a modelling perspective, this path presents a high degree of uncertainty. No reliable information or supporting data is available to quantify behavioural or financial impacts associated with this scenario. However, one aspect can be stated with certainty: since the repair is funded externally, there is no impact on the exposure held at ABN AMRO. As a result, this path leads to no change in either EAD or LGD.

Nevertheless, a change in creditworthiness must still be accounted for. Given the uncertainty and lack of data surrounding external financing structures, a simplified assumption is made: the creditworthiness is negatively affected by a constant-weighted share $k^{\text{loan ext}}$ of the damage amount. This assumption reflects the typical repayment obligations that would follow from taking on a new loan—whether from another bank or informal sources. Although considerable uncertainty remains around repayment terms, the damage amounts involved are generally modest; therefore, it is assumed that an external loan would carry a relatively short maturity of less than five years. The annual impact on creditworthiness is conservatively assumed to be 25% of the damage amount, which corresponds to applying a scaling factor of $k^{\text{loan ext}} = 0.25$ to quantify the additional debt burden on the borrower's credit capacity.

$$C_{i,\text{CER}}^{\text{loan ext}} = C_i - k^{\text{loan ext}} \cdot \text{damage}_{i,\text{CER}}, \tag{3.42}$$

$$LTC_{i,\text{CER}}^{\text{loan ext}} = \frac{\text{EAD}_i}{C_{i,\text{CER}}^{\text{loan ext}}}. \tag{3.43}$$

The resulting impact on PD is again modelled using a relative change ratio $R^{\text{loan ext}}$, followed by the same approach to incorporating the probabilities of the respective CER events to arrive at the expected CER-adjusted PD

$$\mathrm{PD}_{i,\mathrm{CER}}^{\text{loan ext}} = \lambda_{i,\mathrm{CER}}^{\text{loan ext}} \cdot \mathrm{PD}_{i,\text{non-CER}}, \tag{3.44}$$

$$\mathrm{PD}_{i,E(\mathrm{CER})}^{\text{loan ext}} = \mathrm{PD}_{i,\text{non-CER}} + p_{i,\mathrm{CER}} \cdot (\Delta\mathrm{PD}_{i,\mathrm{CER}}^{\text{loan ext}}). \tag{3.45}$$

This concludes the second step, in which the impact of climate-related events on all core credit risk metrics has been defined and quantified for each behavioural path. In the next step, these individual CER-adjusted metrics will be integrated into a single framework — incorporating the probabilities associated with each behavioural response.

III (Extension): Path-dependent integration into regulatory risk metrics (EL, RWA)

In this final step of the total CER impact estimation process, the same structural setup as in the current model is maintained. More specifically, the idea of computing a borrower-specific CER-adjustment factor $M_{i,\mathrm{CER}}$, and aggregating these to obtain a general adjustment factor $\overline{M}$, remains central to the extended framework. However, in contrast to the current setup, the updated approach incorporates both behavioural variation and the fact that CER may affect multiple core risk parameters (PD, LGD, and EAD).

This was made possible by introducing behavioural path probabilities, representing the likelihood that a borrower in ABN AMRO's residential mortgage portfolio follows a particular behavioural response to climate-related damage. The following behavioural probabilities are introduced

- $P_{\text{no repair}}$: probability that the borrower chooses not to repair the CER-related damage;
- P_{repair}: probability that the borrower opts for repair;
- $P_{\text{self}|\text{repair}}$: conditional probability that, given repair, the borrower self-funds the cost;
- $P_{\text{loan AAB}|\text{repair}}$: conditional probability that the borrower finances the repair through an additional loan from ABN AMRO;
- $P_{\text{loan ext}|\text{repair}}$: conditional probability that the borrower finances the repair externally (e.g. another financial institution or informal sources).

Using these probabilities, the borrower-level CER-adjustment factors can be expressed as a weighted average over all behavioural scenarios, with weights corresponding to their respective probabilities. This allows for the integration of each scenario's expected impact on the three core risk parameters. The resulting adjustment factors are defined as follows:

$$M_{i,\mathrm{PD},CER} = \frac{P_{\text{repair}} \cdot \left(P_{\text{self}|\text{repair}} \cdot \mathrm{PD}_{i,\mathrm{E(CER)}}^{\text{self}} + P_{\text{loan AAB}|\text{repair}} \cdot \mathrm{PD}_{i,\mathrm{E(CER)}}^{\text{loan AAB}} + P_{\text{loan ext}|\text{repair}} \cdot \mathrm{PD}_{i,\mathrm{E(CER)}}^{\text{loan ext}}\right) + P_{\text{no repair}} \cdot \mathrm{PD}_{i,\text{non-CER}}}{\mathrm{PD}_{i,\text{non-CER}}} - 1 \tag{3.46}$$

$$M_{i,\mathrm{LGD},CER} = \frac{P_{\text{repair}} \cdot \left(P_{\text{self}|\text{repair}} \cdot \mathrm{LGD}_{i,\text{non-CER}} + P_{\text{loan AAB}|\text{repair}} \cdot \mathrm{LGD}_{i,\mathrm{E(CER)}} + P_{\text{loan ext}|\text{repair}} \cdot \mathrm{LGD}_{i,\text{non-CER}}\right) + P_{\text{no repair}} \cdot \mathrm{LGD}_{i,\mathrm{E(CER)}}}{\mathrm{LGD}_{i,\text{non-CER}}} - 1 \tag{3.47}$$

$$M_{i,\mathrm{EAD},CER} = \frac{P_{\text{repair}} \cdot \left(P_{\text{self}|\text{repair}} \cdot \mathrm{EAD}_{i,\text{non-CER}} + P_{\text{loan AAB}|\text{repair}} \cdot \mathrm{EAD}_{i,\mathrm{E(CER)}} + P_{\text{loan ext}|\text{repair}} \cdot \mathrm{EAD}_{i,\text{non-CER}}\right) + P_{\text{no repair}} \cdot \mathrm{EAD}_{i,\text{non-CER}}}{\mathrm{EAD}_{i,\text{non-CER}}} - 1 \tag{3.48}$$

In cases where a given behavioural path is assumed to have no effect on a particular risk parameter, the non-CER value is used in the weighted average. For instance, in the self-funded repair scenario, only the PD is affected. Therefore, the non-CER values for LGD and EAD are retained when computing their respective adjustment factors.

Following the same approach as in the current setup, the mortgage-level adjustment factors are aggregated into portfolio-level averages for each of the core credit risk parameters. These average adjustment factors are computed as follows

$$\overline{M}_{\text{PD}} = \frac{1}{n}\sum_{i=1}^{n} M_{i,\text{PD}}, \quad \overline{M}_{\text{LGD}} = \frac{1}{n}\sum_{i=1}^{n} M_{i,\text{LGD}}, \quad \overline{M}_{\text{EAD}} = \frac{1}{n}\sum_{i=1}^{n} M_{i,\text{EAD}}. \tag{3.49}$$

These averages are then used to derive the final CER-adjusted portfolio-level values for each of the core credit risk metrics. This is done by applying the corresponding average adjustment factor as a multiplicative increase to the respective non-CER values

$$\text{PD}_{\text{CER}} = \text{PD}_{\text{non-CER}} \cdot (1 + \overline{M}_{\text{PD}}), \tag{3.50}$$

$$\text{LGD}_{\text{CER}} = \text{LGD}_{\text{non-CER}} \cdot (1 + \overline{M}_{\text{LGD}}), \tag{3.51}$$

$$\text{EAD}_{\text{CER}} = \text{EAD}_{\text{non-CER}} \cdot (1 + \overline{M}_{\text{EAD}}). \tag{3.52}$$

Finally, the translation of these CER-adjusted inputs into their regulatory impact—specifically on Expected Loss (EL) and Risk-Weighted Assets (RWA)—follows exactly the same procedure as outlined in the current model, as defined in equations (3.10)–(3.14). The only modification lies in the substitution of the original input parameters with the CER-adjusted values for PD, LGD, and EAD, as derived from the behavioural extension framework.

This concludes the full model extension design. The goal throughout has been to build on the existing structure wherever possible, while carefully adding the behavioural dependencies in a clear and structured way. The design aims to find a balance between keeping the model practical and making it more realistic, by including behavioural assumptions in a way that is thoughtful but not overly complex.

3.2.3 Challenges

Now that the full structure of the model extension has been defined, it is helpful to briefly reflect on a specific element within the framework: the earlier mentioned creditworthiness proxy. As noted previously, this component was left intentionally open in the initial model design, with the idea that it would be specified during the data exploration phase of the project.

Even before examining any data, certain challenges were already expected when working with a proxy for creditworthiness. The core issue is that a bank does not have complete information about a customer's full financial position. While banks typically hold extensive data on outstanding loans and repayment behaviour, they do not necessarily have insight into a customer's total assets or liquidity across institutions.

Consider the following: a customer might have a mortgage at ABN AMRO, but their income or savings could be managed through other financial providers. In the ideal scenario, a customer holds their mortgage, checking account and savings account all at ABN AMRO. In that case, the bank could form a relatively complete picture of their financial health. However, this situation is often not the norm. Many individuals spread their financial products across multiple banks or institutions [50].

As such, even without analysing the specific data available for this study, it was already clear that this creditworthiness element would present a modelling challenge. In fact, this limitation was one of the reasons why behavioural paths were excluded in the initial modelling approach. When certain assumptions require data that is either unavailable or difficult to approximate reliably, model simplicity is often favoured.

Addressing this challenge became one of the key motivations for the next step of this research: an in-depth exploration of the available data, with the goal of identifying suitable features that could serve as a reasonable proxy for borrower creditworthiness.

3.3 Model extension: data-driven implementation and evaluation

With the extended modelling framework formally defined and its associated challenges outlined, the next step involves transitioning to the data-driven implementation phase. This includes exploring the available data, selecting and preprocessing relevant variables, performing model fitting and finally integrating the results into the existing CER impact estimation structure. This section covers all key elements of this phase and concludes with a dedicated evaluation subsection, where appropriate performance metrics are introduced and discussed.

3.3.1 Data description

The first step of this next phase involved a detailed exploration of the available data. The aim was to identify which information could serve as a suitable proxy for borrower creditworthiness.

Generally, banks collect extensive and granular datasets for the estimation of their internal Probability of Default (PD), Loss Given Default (LGD) and Exposure at Default (EAD) models. These datasets contain detailed information on borrower characteristics, loan conditions and collateral features. For the purposes of this project, full access to the datasets underlying these three models was granted, which also includes all possible data points that could possibly support the construction of a creditworthiness proxy.

Due to confidentiality constraints, this section only provides a high-level summary of the data sources used in the model extension. While many variables were explored during the data analysis phase, the focus here is restricted to the final variable selection and the corresponding rationale. Table 3.1 presents the relevant variables considered in the extension, structured by the core credit risk metric they belong to.

It is important to mention that each dataset contains numerous time-stamped observations for a large number of mortgage contracts, some of which date back multiple decades while others are relatively recent. The table shows only a subset of the variables used in this project, focusing on those most relevant to the context. Specifically within the PD category, these variables form the foundation for defining the creditworthiness proxy, which will be elaborated in the next section.

3.3.2 Creditworthiness proxy

As discussed earlier, one of the core challenges in this study is to establish a relationship between the probability of default (PD) and a variable that reliably reflects borrower creditworthiness. However, identifying such a variable is not straightforward. After evaluating several data-driven alternatives—which will not all be mentioned due to confidentiality reasons—the final choice was to base the proxy on borrower income.

Income estimation

Among the many data points available, one stood out for its reliability: the income at the time of loan application. This is a required component in any mortgage application, as it forms a basis for assessing whether a borrower qualifies for a mortgage and, if so, how much can be borrowed. It is therefore reasonable to assume that the income reported at application is accurate and well-recorded in ABN AMRO's internal

Metric	Variable	Description
LGD	snapshot_date_i	The date of the recorded measurement for borrower i
	outstanding_i	Open remaining loan amount at the time of measurement for i
	$\text{collateral\_value}_i$	Estimated collateral value for borrower i
	LGD_i	LGD score for borrower i, without any CER effect
PD	snapshot_date_i	The date of the recorded measurement for borrower i
	$\text{application\_date}_i$	Application date of the mortgage loan for borrower i
	outstanding_i	Open remaining loan amount at the time of measurement for i
	age_i	Borrower i's age at time of snapshot
	$\text{income\_at\_application}_i$	Declared income at the moment of application
	PD_i	Estimated PD score for borrower i, without any CER effect
	default_t1_i	Binary indicator of whether borrower i defaulted within one year following the snapshot date
EAD	snapshot_date_i	The date of the recorded measurement for borrower i
	outstanding_i	Open remaining loan amount at the time of measurement for i
	EAD_i	EAD value for borrower i, without any CER effect

Table 3.1: Overview of relevant variables available per credit risk model

systems. Given the importance of data quality for proxy construction, this was a strong argument in favour of using this variable.

Once the decision was made to use income at loan application (I_0), the next challenge arose: how to estimate the borrower's income at the snapshot date, which is the target variable needed to construct the creditworthiness proxy. Since the PD model reflects current credit risk, the proxy must also be based on a current (i.e. snapshot) income estimate.

To address this, a simplified income forecasting model was designed. It is based on the assumption that income evolves over a person's lifecycle in a predictable manner, increasing in early years, stabilising at mid-career, and eventually declining around and after retirement. Four age-based segments were defined, each associated with a corresponding annual adjustment factor. Publicly available salary data from Statista was used to empirically derive four lifecycle segments, each with a corresponding annual income adjustment factor r_i, as well as an inflation correction factor $\pi = 1.02$. The underlying derivations and parameter values are detailed in Appendix B.

- **Early-career segment:** For individuals below $\chi_e = 35$, income grows by a factor of $r_e = 1.06$ per year.
- **Mid-career segment:** Between ages $\chi_e = 35$ and $\chi_m = 55$, income continues to grow more modestly, with $r_m = 1.01$.
- **Late-career segment:** Between ages $\chi_m = 55$ and $\chi_l = 65$, income begins to decline gradually at a yearly rate of $r_l = 0.99$.
- **Post-retirement segment:** For individuals aged above $\chi_p = 65$, income declines further with a yearly factor $r_p = 0.97$. This smooth annual decline approximates the typical drop in income after retirement, which in reality is often more abrupt.

Each borrower's estimated income at the snapshot date, corresponding to age a_s, is derived by applying the appropriate growth or decay factors for the number of years spent in each age segment since the loan origination age a_0. Additionally, a uniform inflation correction is applied using a fixed inflation index $\pi = 1.02$, compounded over the same time span. This results in the following formula for the estimated income:

$$I_s = I_0 \cdot (r_e)^{y_e} \cdot (r_m)^{y_m} \cdot (r_l)^{y_l} \cdot (r_p)^{y_p} \cdot \pi^{(a_s - a_0)}, \tag{3.53}$$

where:

- I_0 is the reported income at loan origination,
- I_s is the estimated income at the snapshot moment,
- y_e, y_m, y_l, y_p are the number of years spent in the early-, mid-, late-, and post-career segments respectively.

where the number of years spent in each segment is defined as

$$y_e = \max(0, \min(\chi_e, a_s) - a_0), \tag{3.54}$$

$$y_m = \max(0, \min(\chi_m, a_s) - \max(\chi_e, a_0)), \tag{3.55}$$

$$y_l = \max(0, \min(\chi_l, a_s) - \max(\chi_m, a_0)), \tag{3.56}$$

$$y_p = \max(0, a_s - \max(\chi_l, a_0)). \tag{3.57}$$

This approach provides a deliberately simplified, yet structured method to estimate a borrower's income on the snapshot date, which serves as the basis for defining the creditworthiness proxy. The base-case values used for the segment boundaries and growth factors were not data-driven but selected for initial model testing. These values will later be revisited and validated through expert judgement in subsequent stages of the research.

Estimated Loan-to-Income Ratio (LTI)

Finally, the only remaining step was to scale the estimated income (on the snapshot date) to the outstanding loan. Let I_s again denote the estimated income as calculated in the previous step, and let L_s represent the outstanding loan amount on the snapshot date. The corresponding estimated LTI ratio is then defined as:

$$\widehat{\mathrm{LTI}}_s = \frac{L_s}{I_s} \tag{3.58}$$

This completes the estimation of the Loan-to-Income ratio, which will serve as the proxy for creditworthiness in the remainder of the modelling framework. Since certain implementation choices were made during this estimation step—particularly related to data structure and availability—these will be specified in the following subsection.

Data prepossessing

An important consideration during the implementation of the LTI estimation, was the discovery of certain data limitations. One of the required input variables for the estimation is the borrower's age at the time of loan application. During implementation, it became clear that this variable contains a substantial number of missing values. While missing entries can occur for a variety of reasons, in this case, a specific explanation applies to a large portion of the dataset: income at application was only systematically recorded from a certain year onward. As a result, data entries for loans originated before that point do not include income information at all.

Figure 3.7 provides an overview of the missing values, binned by year of loan application, highlighting the period during which these entries are absent.

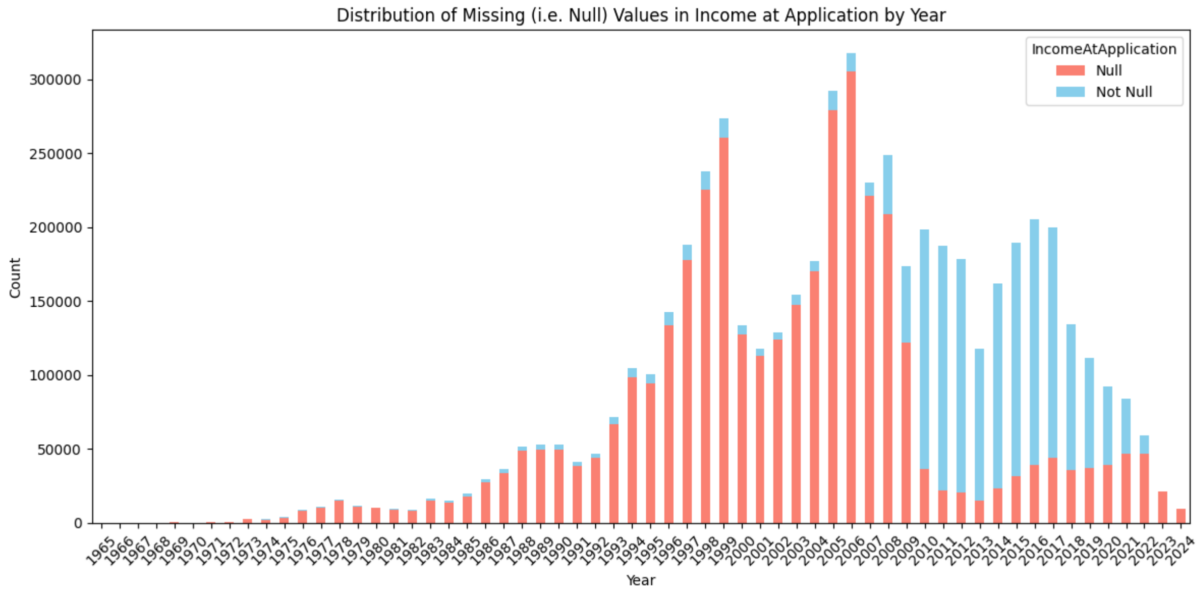

Figure 3.7: Missing data entries *Income at Application* data

To guide decisions in situations involving incomplete data, a commonly applied rule of thumb is used: if more than 75% of entries for a key variable are missing, the variable should be excluded from the analysis. This threshold is also consistent with ABN AMRO's internal data quality requirements. Fortunately, the share of missing values in this case remained below that threshold, with 62% of the entries lacking income data. Despite this limitation, as the filtered dataset still retains 1,774,212 observations, it can be considered a robust basis for model estimation.

Furthermore, it is worth reflecting on the nature of the remaining sample. While the retained observations are biased toward more recent years, this is not necessarily a disadvantage. This consideration becomes particularly relevant in the final stage of the impact estimation process—specifically in step III. Recall that in this step, an average adjustment factor is constructed (see Equation 3.49) and subsequently applied to calculate the final CER-adjusted PD (see Equation 3.50). Whereas all earlier stages rely on datasets with the maximum available historical coverage, the goal in this final step is to reflect the impact on RWA and EL as accurately as possible for the current portfolio composition. To that end, the most recent full year of data—2022—was used for the final impact estimation. Therefore, the fact that the adjustment factor was derived from relatively recent historical data is consistent with the aim of producing a timely and representative impact assessment.

Based on these considerations, the decision was made to proceed with the income-at-application variable despite its higher level of incompleteness. As a consequence, the LTI estimates—and later the derived relationship between LTI and PD—are constructed from a sample consisting primarily of more recent mortgage years (2010–2022). Given the context and final use case, this was not considered a critical limitation.

This concludes the first implementation step, which lays a foundational component for the remainder of the model extension. The next subsection will define a functional relationship between the estimated Loan-to-Income ratio $\widehat{\mathrm{LTI}}$ and the probability of default (PD).

3.3.3 Model fitting

This section outlines the approach taken to define the functional relationship between the estimated LTI and the probability of default (PD). The goal here is not to present the empirical outcomes—these are discussed in the results section—but rather to establish the theoretical modelling foundation on which those results are based.

As previously explained, the PD is a probability, and therefore by definition constrained to the interval $[0, 1]$. As discussed in Section 2.1.3, PDs are generally derived from internal credit scoring models, which produce discrete rating grades that are then mapped to PD values. As a result, the available PD data is discretised—each distinct PD level corresponding to a rating bucket—see also Example 2.1.3. This discreteness presents an additional challenge when attempting to fit a continuous model to the data.

While several different model specifications were tested across various samples, the final approach was deliberately kept as simple as possible to preserve model interpretability. It was therefore decided to fit one model on the full available dataset, using two different sets of input features. The remainder of this section first introduces the functional form of the selected model, followed by a brief explanation of the two input specifications considered.

The final model chosen for further implementation is a linear regression applied to the logit-transformed probability of default (PD). This choice is motivated by the fact that modelling PD directly using linear regression is problematic due to the bounded nature of the target variable, which by definition must lie within the interval $[0, 1]$. A standard linear model does not naturally respect these bounds and may produce predicted values outside the admissible probability range. To address this, a common approach is to apply a logit transformation, which maps the probability onto the real number line. The fitted target variable is therefore defined as

$$y_{\text{logit}} = \text{logit}(\text{PD}) = \ln\left(\frac{\text{PD}}{1 - \text{PD}}\right). \tag{3.59}$$

Accordingly, the PD estimate resulting from the model is given by

$$\widehat{\text{PD}} = \frac{1}{1 + e^{-y_{\text{logit}}}}. \tag{3.60}$$

By modelling the logit of PD as a linear function of the input variables, the resulting estimates can be transformed back to valid probability values through the inverse logit function. This approach maintains mathematical consistency while also improving interpretability, as the effect of each input variable is captured on the log-odds scale—a common metric in credit risk modelling frameworks.

The structure of y_{logit} depends on the selected input variables. Two input feature sets were evaluated. The first specification models a univariate dependence on the estimated $\widehat{\text{LTI}}$ (3.58), while the second introduces a multivariate dependence that incorporates two additional variables used in PD modelling. Both feature sets are described in more detail in the following subsections.

Single dependence model: LTI

The target of interest in the single dependence specification is the estimated Loan-to-Income ($\widehat{\text{LTI}}$) ratio (3.58). The objective is to investigate and formalise its relationship with the probability of default. Within this setup, the regression model is fitted using the estimated $\widehat{\text{LTI}}$ as the sole explanatory variable. While the results of this fit will be discussed in the results section 4, it is essential to first present a clear description of the underlying regression formulation and an initial examination of the data to be fitted. This section outlines both the mathematical specification of the regression and the relevant data characteristics of the LTI distribution used for model fitting.

Consistent with the general formulation introduced in (3.59), the predicted probability of default in this single-dependence model is defined as

$$\widehat{\text{PD}} = \frac{1}{1 + e^{-(\beta_0 + \beta_1 X_1)}}, \tag{3.61}$$

where X_1 denotes the estimated loan-to-income (LTI) value, β_1 is the associated model coefficient, and β_0 represents the intercept. Accordingly, the target variable in the regression is the logit-transformed observed PD

$$y_{\text{logit, base}} = \text{logit}(\text{PD}) = \ln\left(\frac{\text{PD}}{1-\text{PD}}\right) = \beta_0 + \beta_1 X_1. \tag{3.62}$$

To provide additional intuition about the structure of the fitted data, in Figure 3.8 three diagnostic plots are presented: a scatter plot, a hexbin plot and a violin plot. Together, these visualisations aim to offer insight into the relationship between estimated LTI and observed PD values across the dataset.

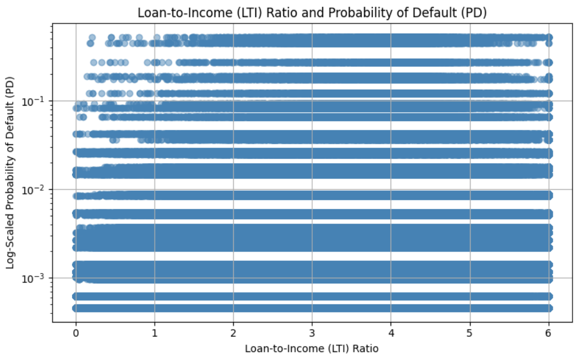

(a) Scatter plot between LTI and log-scaled PD

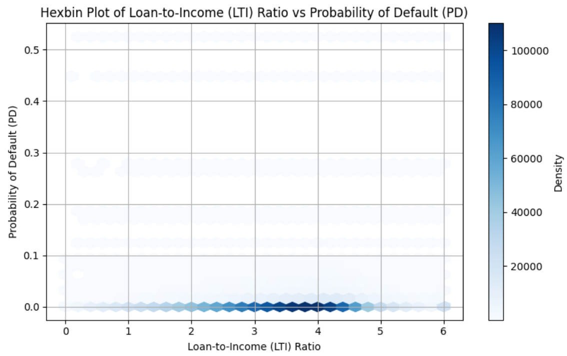

(b) Hexbin plot of LTI and PD

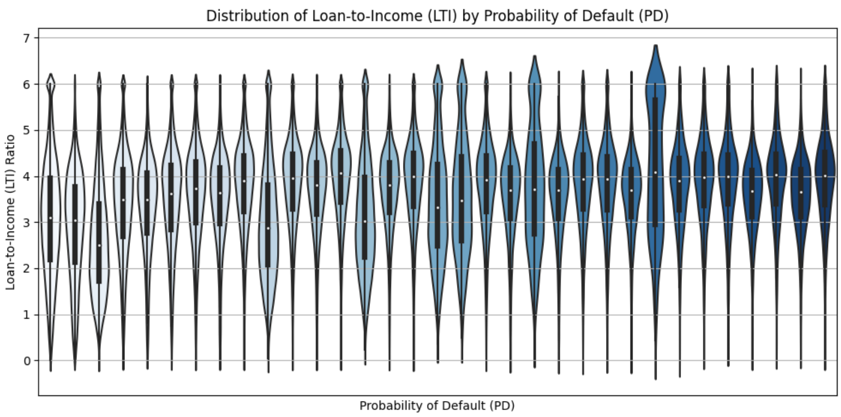

(c) Distribution of LTI across discretized PD values (violin plot). The PD range spans [0, 0.5], with most density concentrated below 0.005. Specific bin values are omitted for confidentiality; darker violins indicate higher discretized PD values.

Figure 3.8: Visualizations of the relationship between the loan-to-income ratio LTI (3.58) and the probability of default (PD).

In the first plot (Figure 3.8a), the direct relationship between PD and the estimated LTI is visualised through a scatter plot. Due to the large volume of data points (approximately 1.8 million) and the discretised nature of PD values—recall the binned structure of PD models producing a fixed set of PD buckets—it is difficult to observe the density distribution across individual PD levels. However, one visible pattern is a relatively lower density in the top-left corner of the plot, corresponding to combinations of low LTI and high PD values. This observation is in line with intuition: lower LTI ratios (i.e. higher creditworthiness) are less frequently associated with higher probabilities of default. In simpler terms, this potentially supports the expected behavioural relationship that a decrease in creditworthiness corresponds to an increase in PD.

Given the data density and high degree of skewness, a hexbin plot was additionally created (Figure 3.8b). A *hexbin* plot is a two-dimensional histogram that groups observations into hexagonal bins instead of traditional rectangular ones. It is particularly useful for visualising the density of large datasets in scatter-like plots, especially when data points overlap heavily. In this case, it reveals how frequently specific combinations of LTI and PD values occur, with darker regions representing higher densities. This plot illustrates the distribution more clearly, revealing a substantial skew toward very low PD values. More specifically, approximately *75%* of the dataset falls below a PD of 0.005. While this is reassuring—as it indicates that most borrowers in the residential mortgage portfolio have a low likelihood of default—it also presents a modelling challenge. The imbalance may lead to limited influence from observations in the upper range of PD values during regression fitting. This implies that while PD values above 0.005 may carry important behavioural meaning, their scarcity in the dataset could cause them to be underemphasised in the resulting regression fit.

Lastly, a violin plot was included (Figure 3.8c) to provide a complementary view of the distribution. A *violin* plot combines a box plot and a kernel density plot, displaying both the distribution and probability density of the data. In this context, it shows the full distribution of LTI values within discrete PD intervals. Wider sections indicate higher concentrations of LTI values, while the shape of the violin reveals potential multimodal patterns within each PD bucket. In simpler terms, this plot shows the spread of the LTI values within each discretised PD bucket. Because the PD buckets are bank-specific and therefore confidential, their exact numerical values are omitted from the x-axis. To provide some general orientation: the leftmost violin corresponds to the lowest PD bucket (close to zero), while the rightmost violin reflects the highest PD bucket, with most PD values falling within the interval $[0, 0.5]$. Additionally, darker coloured violins indicate higher discretized PD values within the approximate range just described. As also suggested by the scatter plot, the violin plot shows that lower LTI values are less prevalent in higher PD buckets. In addition, the violin plot reveals a complementary pattern: in the lowest PD buckets, there is a visibly higher density of lower LTI values. Together, these observations align with intuitive expectations and provide preliminary evidence of a potential moderate relationship between LTI and PD that is worth modelling more formally.

Multiple dependence model: LTI, BKR score, Terms in Arrears

To gain a more complete understanding of the statistical significance of LTI—and to test whether it remains relevant in the presence of other variables known to influence PD—a second model was constructed. This model incorporates two additional features alongside the LTI estimate: the number of *terms in arrears* and the *BKR score*. Both are features that are known to have material relevance within internal PD modelling practices. First the mathematical formulation of this extended model will be defined and after this the construction and characteristics of each included feature will be explained.

The predicted PD in the case of this multiple-dependence linear regression (applied on the logit of PD) is defined as

$$\widehat{\mathrm{PD}} = \frac{1}{1 + e^{-\left(\beta_0 + \beta_1 X_1 + \beta_2 X_2 + \sum_{i=1}^{g} \beta_{3,i} \cdot \mathbb{1}\{X_3 = i\}\right)}}, \tag{3.63}$$

where X_1 denotes the estimated LTI value, X_2 represents the number of terms in arrears, and X_3 defines the categorical BKR score. The corresponding target regression variable is formulated as

$$y_{\text{logit, base}} = \mathrm{logit}(\mathrm{PD}) = \ln\left(\frac{\mathrm{PD}}{1 - \mathrm{PD}}\right) = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \sum_{i=1}^{g} \beta_{3,i} \cdot \mathbb{1}\{X_3 = i\}. \tag{3.64}$$

As shown, the BKR score X_3 is treated differently from the other input variables due to its categorical nature. BKR, short for *Bureau Krediet Registratie*, is the Dutch credit registration authority that maintains records of borrowers' credit histories. Each registration corresponds to a discrete letter-based category. Letters early in the alphabet (e.g. `A`, `B`) typically indicate a clean or near-clean payment history, while those further down (e.g. `K`, `L`) reflect more severe credit issues.

In addition to the letter categories `A` through `L`, the dataset also included seven additional high-risk indicators, `P1` through `P7`, each representing increasingly severe forms of payment problems. These were grouped into a single consolidated category P to limit the dimensionality of the resulting regression design matrix.

Moreover, in this case missing BKR values carry specific behavioural meaning: a missing registration may indicate that the borrower has no recorded credit history. Historical modelling experience shows that such cases are often associated with low-risk borrowers. For this reason, missing values were not removed but treated as a separate category (the `Null` category).

Because BKR classifications represent distinct qualitative categories rather than ordinal or continuous data, they were one-hot encoded. This is a common encoding technique used in regression models to represent categorical variables. It creates separate binary indicator variables for each category, allowing the model to assign a separate coefficient to each BKR type. More specifically, this encoding introduces binary variables of the form $\mathbb{1}\{X_3 = i\}$ for which $i \in \{\mathtt{A}, \mathtt{B}, \ldots, \mathtt{L}, \mathtt{P}, \mathtt{Null}\}$, each associated with a coefficient $\beta_{3,i}$, allowing the model to independently capture the marginal contribution of each BKR category to the log-odds of default.

To avoid multicollinearity in this setup, one BKR category must be excluded and treated as the reference. In this implementation, category B was selected as the reference. This choice does not affect the model's predictive performance but provides a natural baseline for comparison, as B is one of the more frequently occurring categories in the dataset.

Now that the treatment of the BKR variable has been clarified, it is time to briefly examine the corresponding data used. In line with the earlier visualizations for LTI, Figure 3.9 presents a scatter plot, hexbin plot and violin plot showing the relationship between PD and the BKR categories.

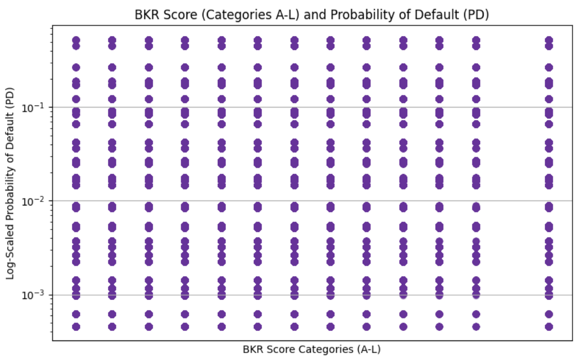

(a) Scatter plot between BKR and log-scaled PD

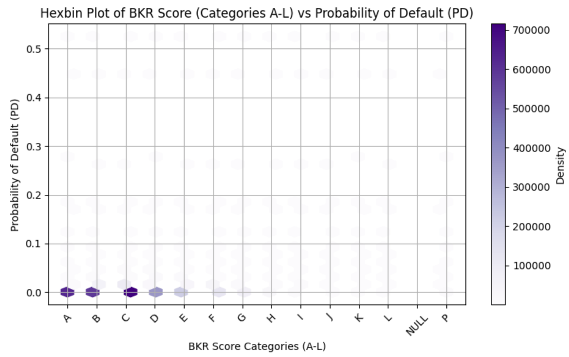

(b) Hexbin plot of BKR and PD

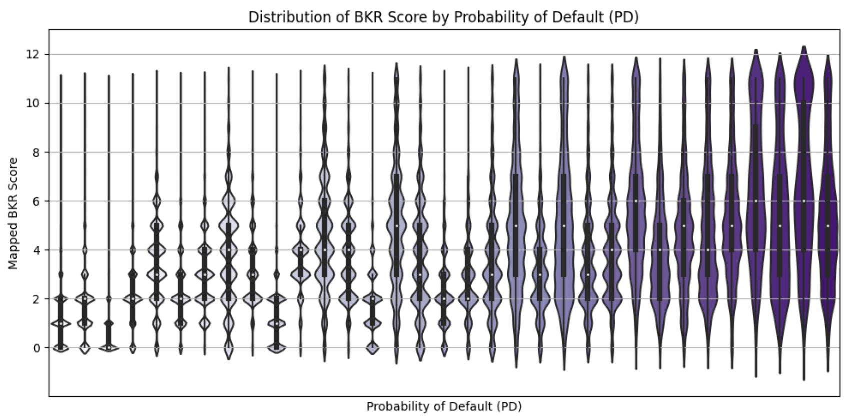

(c) Distribution of BKR across discretized PD values (violin plot). The PD range spans [0, 0.5], with most density concentrated below 0.005. Specific bin values are omitted due to confidentiality; darker violins indicate higher discretized PD values.

Figure 3.9: Visualizations of the relationship between BKR and the probability of default (PD).

The scatter plot (Figure 3.9a) shows that data points exist for each BKR category across the entire range of discretised PD values. However, it provides limited insight into how the distribution of PD varies across BKR scores. The hexbin plot (Figure 3.9b) confirms the familiar pattern of high data density for low PD values and additionally reveals that this density is concentrated primarily among the lower-risk BKR categories—namely, A, B and C. This indicates that the majority of borrowers in the portfolio hold relatively healthy credit records.

The violin plot (Figure 3.9c) offers more detailed distributional insight: it clearly shows that lower PD values are more prevalent among borrowers with lower-risk BKR classifications, while higher PD values tend to occur more frequently within higher-risk BKR categories. Taken together, these visualisations suggest that the BKR variable is likely to exhibit a statistical significance in the regression model—a result that would be consistent with the fact that BKR classification is already a known input factor in PD modelling.

Now turning to the final additional input feature considered in the multiple dependence model: Terms in Arrears, specifically measured over the 12 months preceding the observation date. This variable captures the number of missed or overdue payments within a one-year period and serves as a well-established, highly significant predictor in PD modelling frameworks. In this context, it reflects how often a borrower has failed to meet their payment obligations in the year prior to the snapshot. Figure 3.10 presents the associated data plots.

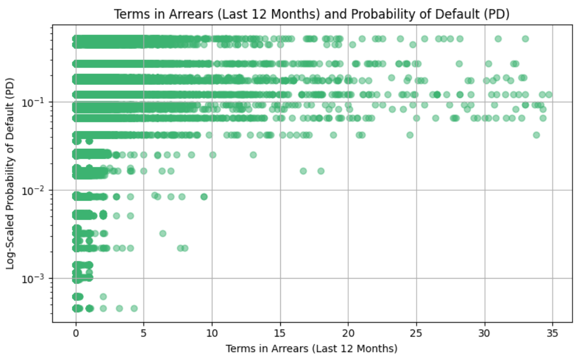

(a) Scatter plot between Arrears and log-scaled PD

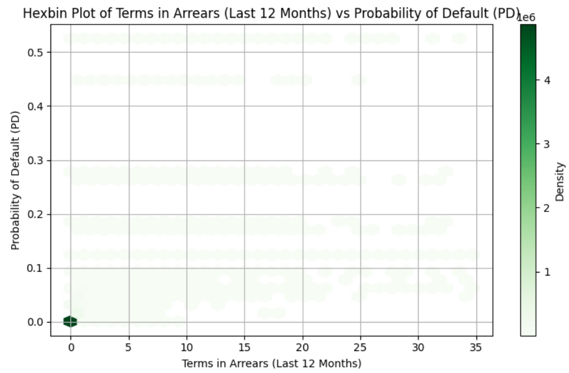

(b) Hexbin plot of Arrears and PD

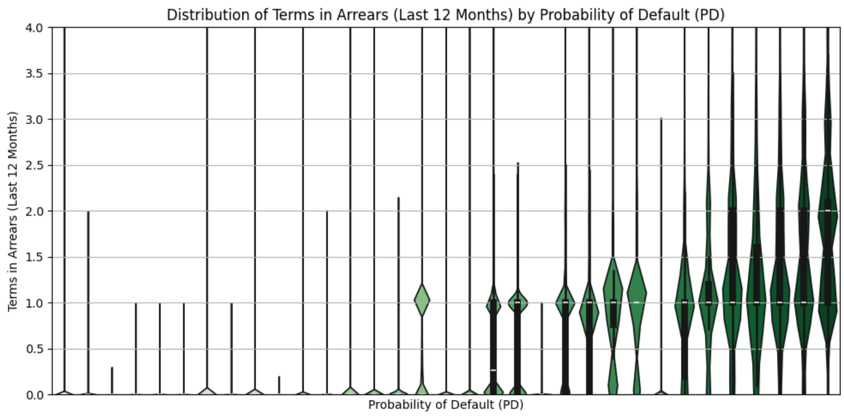

(c) Distribution of Arrears across discretized PD values (violin plot). The PD range spans [0, 0.5], with most density concentrated below 0.005. Specific bin values are omitted due to confidentiality; darker violins indicate higher discretized PD values.

Figure 3.10: Visualizations of the relationship between Terms in Arrears (last 12 Months) and the probability of default (PD).

The scatter plot (Figure 3.10a) reveals a clear dependency between PD and Terms in Arrears. Specifically, there are virtually almost no observations with more than two missed payments among borrowers with low PD values. This alone suggests a strong correlation between delinquency and default risk. The hexbin plot (Figure 3.10b) confirms that the distribution is heavily left-skewed, with the overwhelming majority of borrowers having zero missed payments. This reflects a generally healthy portfolio in terms of recent repayment behaviour.

The violin plot (Figure 3.10c) reinforces this insight. It shows a sharp increase in the number of arrears as PD values rise, supporting the expectation of a strong positive relationship between Terms in Arrears and default risk. This aligns well with the inclusion of this variable in ABN AMRO's internal PD modelling practices.

Imputation rules

As a final step before fitting the model, the imputation rules applied to each of the considered input variables must be outlined. These are discussed below per input feature.

Loan-to-Income (LTI) Since the LTI is already a constructed estimate, additional care was taken to define realistic bounds based on observed distributions and regulatory guidance. As can be seen in the data visualisation (Figure 3.8c), most values fall between 1 and 4, with a long right tail due to occasional outliers. To cap extreme values, a maximum value of 6 was imposed.

This decision aligns with Dutch Loan-to-Income (LTI) limits recommended by the *Nationaal Instituut voor Budgetvoorlichting (NIBUD)*, as documented by *De Nederlandsche Bank (DNB)* [51]. In their 2014 example table, LTI limits are reported to range from 2.6 to 5.7 depending on household income and interest rate scenarios. The adopted cap of 6 used in this study lies slightly above these thresholds, providing a conservative upper bound that accommodates possible yearly adjustments or exceptional borrower circumstances.

The following imputation rules were implemented:

- All LTI values within the interval [0, 6] were retained.
- Negative values, which are not meaningful in this context, were excluded.
- Values above 6 were set equal to 6.

BKR score As discussed earlier, the BKR classification comprises 14 categories: `A`, `B`, . . . , `L`, `P`, `Null`. The `P` category aggregates seven distinct high-risk payment problem types (P1 through P7) into one generalised class. Additionally, a separate `Null` category is included to potentially represent borrowers without any credit registration, as a missing BKR value can indicate the absence of a credit history—often associated with low-risk borrowers.

Terms in Arrears (12-month) This variable reflects the number of missed payments in the 12 months preceding the snapshot date. The bank's maximum for this parameter was set at 35. Following their approach, the following imputations were applied:

- Values within the range [0, 35] were retained.
- Values outside this range were set to 0, under the assumption that such entries are the result of data errors. In these cases, the conservative assumption is made that no arrears occurred.

Overall, based on the discussed data visualisations (Figures 3.8-3.10) and the more pronounced relationships observed between PD and the additional input variables—particularly the *BKR score* and *Terms in Arrears*—it is reasonable to expect that the multiple dependence model will outperform the single-variable specification in terms of its ability to approximate observed PD values. However, before turning to the results in Section 4, the next subsection first outlines the appropriate model performance metrics, as these are not necessarily straightforward in the context of the highly left-skewed and discretised probability data.

3.3.4 Model performance metrics

Following the implementation phase, the next step in the methodology involves validating the model outputs. A key component of this validation is assessing the performance of the different PD estimation models. Two model specifications were developed and tested—one baseline model and one extended model incorporating multiple explanatory variables. The aim is to determine which of the two offers stronger predictive performance, and whether either can be considered sufficiently accurate to proceed within the subsequent final impact estimation phase.

To enable this comparison, appropriate performance metrics have to be selected. Given the discretised and highly left-skewed nature of the PD data, traditional accuracy metrics such as R^2 and Mean Squared Error (MSE) are not well-suited for evaluating model performance. These measures are designed for continuous targets and rely heavily on the squared residuals between predicted and actual values. In the presence of strong skewness, this leads to the drawback that outliers can disproportionately dominate the score. In such cases, R^2 may largely reflect a model's performance on a small number of extreme observations, rather than providing an accurate assessment of overall predictive quality—particularly in the bulk of the distribution where the majority of data points lie.

To obtain a meaningful assessment of model performance under these distributional characteristics, two alternative evaluation metrics are used: the *Gini coefficient* and concordance-based measures, specifically *Kendall's tau* and *Goodman–Kruskal's gamma*. The remainder of this section defines these performance metrics formally and discusses their suitability within the context of default probability modelling.

Gini coefficient

The Gini coefficient is a widely used performance metric for binary classification models. It quantifies how well a model distinguishes between two outcome classes by assessing its ability to rank observations. In the context of credit risk, it is commonly applied to assess the discriminatory power of PD models, i.e. how effectively the model assigns higher predicted default probabilities to borrowers who actually defaulted.

The Gini coefficient is derived from the Area Under the Receiver Operating Characteristic Curve (AUC-ROC), which aggregates model performance across all possible classification thresholds. The ROC curve is constructed by plotting the true positive rate (TPR) against the false positive rate (FPR), where both are defined as functions of a given threshold t:

- True Positive Rate (TPR):

$$\mathrm{TPR}(t) = \frac{\text{Number of true positives at threshold } t}{\text{Total number of actual positives}} = \frac{TP(t)}{TP(t) + FN(t)}, \tag{3.65}$$

- False Positive Rate (FPR):

$$\mathrm{FPR}(t) = \frac{\text{Number of false positives at threshold } t}{\text{Total number of actual negatives}} = \frac{FP(t)}{FP(t) + TN(t)}. \tag{3.66}$$

For completeness, note that TP, FP, TN, and FN represent the counts of true positives, false positives, true negatives, and false negatives, respectively. By varying the threshold t from 0 to 1 and plotting $\mathrm{TPR}(t)$ against $\mathrm{FPR}(t)$, the ROC curve is obtained. The area under this curve — the AUC — summarizes the model's overall ability to distinguish between the positive and negative classes. It can be expressed as

$$\mathrm{AUC} = \int_0^1 \mathrm{TPR}(\mathrm{FPR})\, d(\mathrm{FPR}). \tag{3.67}$$

Alternatively, the AUC can be interpreted probabilistically as

$$\mathrm{AUC} = P\left(\hat{p}_i > \hat{p}_j \mid y_i = 1, y_j = 0\right), \tag{3.68}$$

where in this context $\hat{p}_i$ is the predicted PD score for borrower i, and y_i is the corresponding true binary PD label.

The Gini coefficient rescales the AUC from the $[0.5, 1]$ interval to the $[0, 1]$ range via

$$\text{Gini} = 2 \cdot \text{AUC} - 1. \tag{3.69}$$

A Gini value of 0 implies no discriminatory power, while 1 indicates perfect rank ordering. In practice, values above 0.6 are typically considered acceptable and values above 0.7 as strong [52, 53].

Ordinal rank-based performance: Kendall's τ and Goodman–Kruskal's γ

As a second set of performance metrics, two rank-based concordance measures are considered: Kendall's τ and Goodman–Kruskal's γ. These statistics assess the degree to which the predicted PD values preserve the ordinal relationship of the observed outcomes. In contrast to error-based metrics, they do not rely on the numerical distance between values, but instead capture whether the model correctly ranks higher-risk borrowers above lower-risk ones. This makes them especially well-suited for evaluating model quality in the context of discretised and heavily skewed probability distributions.

Let (x_i, y_i) and (x_j, y_j) denote two predicted-observed pairs from the dataset. The following definitions apply

- A pair is concordant if the ranking agrees: $(x_i - x_j)(y_i - y_j) > 0$,
- A pair is discordant if the ranking disagrees: $(x_i - x_j)(y_i - y_j) < 0$,
- A pair is tied if $x_i = x_j$ or $y_i = y_j$.

Based on this classification, the two metrics are defined as follows

$$\tau = \frac{n_c - n_d}{\binom{n}{2}}, \quad \gamma = \frac{n_c - n_d}{n_c + n_d}, \tag{3.70}$$

where n_c and n_d are the number of concordant and discordant pairs, respectively, and $\binom{n}{2}$ is the total number of distinct observation pairs.

Kendall's τ considers all possible observation pairs and adjusts for ties by reducing the total possible association score. In contrast, Goodman–Kruskal's γ excludes tied pairs from the calculation entirely. As a result, γ focuses only on pairs where a clear ranking difference exists, making it more sensitive to directional agreement in datasets with many tied values—such as PD data with a limited number of rating grades.

Both metrics range from -1 to $+1$, and their values can be interpreted as follows

- A score of $+1$: perfect agreement in ranking (strong positive association),
- A score of 0: no association between predicted and actual values,
- A score of -1: perfect inverse ranking (strong negative association).

In the context of this study, each pair represents two mortgage observations, each with a predicted PD and an actual observed PD value. A high value of τ or γ indicates that when borrower i is assigned a higher predicted PD than borrower j, this aligns with their true PDs. For example, if $\widehat{\text{PD}}_i > \widehat{\text{PD}}_j$ and also $\text{PD}_i > \text{PD}_j$, the model has made a concordant prediction. These rank-based measures therefore capture how well the model preserves the risk ordering of borrowers—arguably more meaningful than numeric precision in a discretised probability setting.

Unlike the Gini coefficient, there are no universally accepted threshold values for Kendall's τ or Goodman–Kruskal's γ that define what constitutes a "strong" or "acceptable" level of association. Interpretations of these metrics are typically context-dependent. Within the context of this study, the primary purpose of these performance metrics is to serve as an additional comparison tool across model specifications. The key takeaway is that the model that yields the highest values for these metrics can be considered the better performer in this setting, as it demonstrates stronger ordinal alignment between predicted and observed default behaviour.

By applying both performance metrics to the single-variable and multiple-variable model configurations, the resulting scores enable an interpretable comparison of which specification better captures the underlying PD structure. In addition, they provide insight into whether either model performs sufficiently well to justify proceeding to the impact estimation phase. If a model yields both acceptable and relatively higher scores in terms of Gini and rank concordance, this provides a solid foundation for proceeding with that model in the remainder of the study. The results obtained from applying the discussed model performance metrics are presented in the following Results section.

4

Results

This section presents the main empirical results of the extended credit risk model. It is structured into two parts. Section 4.1 covers the Probability of Default (PD) estimation models, including both the single-dependence base model and the extended multivariate specification. For each model, relevant outputs are reported, including parameter estimates, calibration plots, CER-induced relative impact factors and model performance metrics. The latter include the Gini coefficient, Kendall's τ, and Goodman–Kruskal's γ, which are used to evaluate the discriminatory power and rank-order alignment of each model. Based on these metrics, the extended model is selected for further use. Section 4.2 then presents the results of the CER impact estimation, using this selected model to analyse a range of behavioural scenarios—including full and partial repair, multiple financing methods, and sensitivity analyses. Together, these results provide the basis for the model validation and final impact conclusions presented later in this report.

4.1 PD estimation models

This section presents the results of the Probability of Default (PD) estimation for both model specifications introduced in Section 3.3.3. For each model, the following outputs are reported: the estimated intercept (β_0) and regression coefficients, a scatter plot comparing predicted versus observed PD values, a calibration plot assessing alignment between predicted and actual outcomes, the associated relative influence scores (λ) and a set of performance metrics. These metrics include the Gini coefficient and two rank-based measures—Kendall's τ and Goodman–Kruskal's γ—to facilitate the comparison of model performance.

4.1.1 Base model results – single dependence (LTI)

The outcome of the linear regression on logit(PD), using only the estimated LTI as the explanatory variable X_1, results in the following parameter estimates

- Intercept $\beta_0 = -6.7587$,
- LTI Coefficient $\beta_1 = 0.1739$.

Table 4.1 presents all regression results of the base model, with LTI as the sole explanatory variable.

Table 4.1: Base model regression results (log-odds scale)

Variable	Estimate	Std. Error	t-statistic	p-value	CI lower	CI upper
Intercept	-6.7586	0.0024	-2795.11	<0.001	-6.7633	-6.7539
LTI	0.1739	0.0008	220.59	<0.001	0.1723	0.1754

It follows from these results that the LTI coefficient is positive and statistically significant; its 95% confidence interval lies strictly above zero and the reported p-value is below 0.001. The small standard error indicates that the estimate is precise. From this it can be concluded that, based on the used dataset, there is indeed a positive relationship between the estimated LTI and the probability of default.

Using these parameters, this yields the following PD estimation function

$$\widehat{\text{PD}} = \frac{1}{1 + e^{-(-6.7587+0.1739 \cdot X_1)}}. \tag{4.1}$$

Figure 4.1 presents the visual findings for this model fit. The left graph shows the scatter plot comparing actual versus estimated PD values on a logarithmic scale. Due to this scaling, the logistic curve appears linear in shape. From this graph, it is also evident that the relationship between PD and the estimated LTI is positive. This was an important initial result, as a positive correlation was almost a prerequisite for continuing with the defined LTI as a creditworthiness proxy. A negative relationship—where a higher LTI would correspond to a lower PD—would be counterintuitive, as in reality, lower creditworthiness should be associated with a higher probability of default. This finding therefore provided an important first confirmation to proceed with the chosen definition.

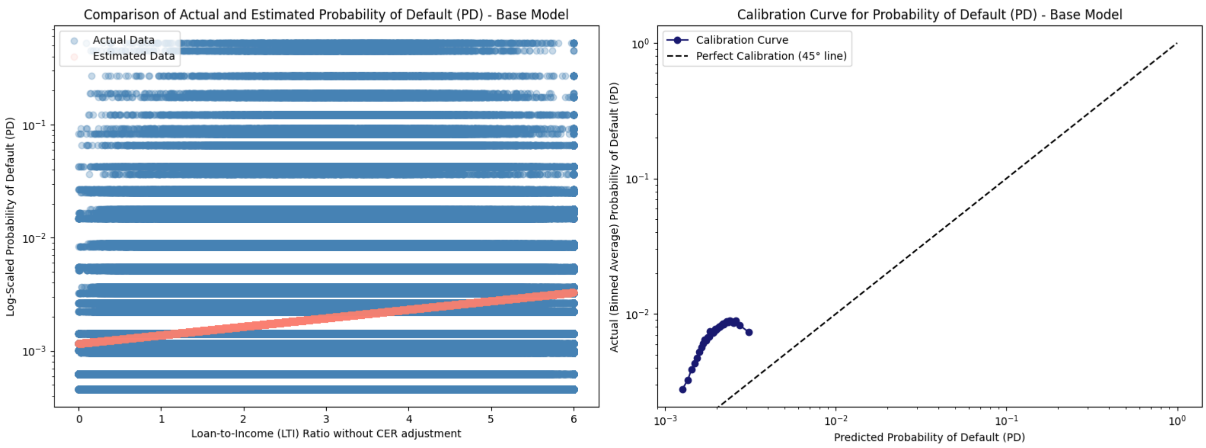

Figure 4.1: Comparison between actual and estimated PD under the base model. Left: log-scaled scatter plot of observed vs. predicted PDs. Right: calibration curve showing the model's ability to match average observed PDs.

The right graph shows the calibration curve, which reflects how closely the predicted PD aligns with the observed average PD values. For the observed data, the mean PD was computed within predefined bins of predicted PD values to provide a smoothed empirical estimate of default likelihood across the range. It shows that the estimated PDs underestimate the true PD values across the prediction range. Additionally, the curve ends abruptly at a predicted PD of approximately 1.5×10^{-2}. This corresponds to the saturation of the fitted function in the scatter plot to the left: due to the predefined maximum LTI value of 6, the linear fit on the logit scale does not produce estimated PD values beyond this point. As a result, no higher predicted PD values are available to populate the upper bins of the calibration curve, leading to its early truncation.

Moreover, the curve exhibits a local decline—i.e. an interval of non-monotonicity. Since the calibration curve is expected to be monotonically increasing—higher predicted PDs should correspond to higher actual average PDs—this behaviour suggests a misalignment between the model and the true ranking structure in the data. This likely reflects the limitation of relying on a single input variable (LTI) in a highly skewed dataset. These findings collectively signal that the model is an oversimplification and suggest that extending it with additional drivers of default would likely improve its predictive performance.

Figure 4.2 includes the associated relative CER-induced changes in PD, denoted by λ, recall that for a general behavioural pathway, this was defined as follows

$$\lambda_{i,\text{CER}}^{\text{behaviour}} = \frac{\widehat{\text{PD}}_{i,\text{CER}}^{\text{behaviour}}}{\widehat{\text{PD}}_{i,\text{non-CER}}^{\text{behaviour}}}. \tag{4.2}$$

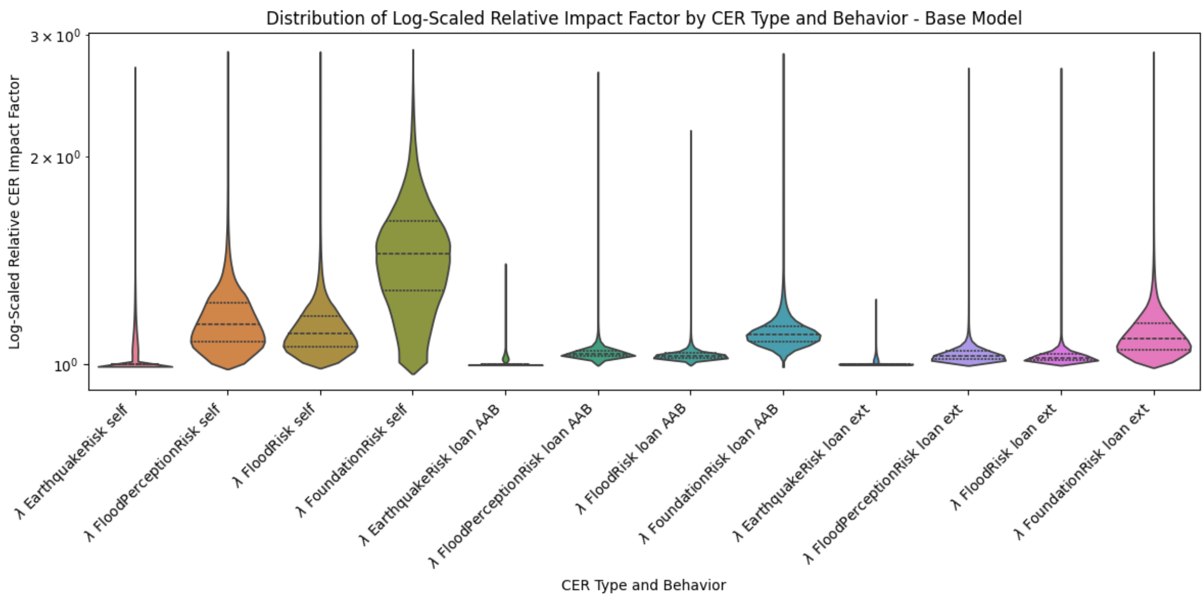

Figure 4.2: Distribution of log-scaled relative impact factor λ for each CER type and behavioural path, based on the base PD model. Higher values indicate stronger increases in risk metrics due to CER-driven behavioural responses.

This factor captures the relative change in predicted PD values under the CER-adjusted model, compared to the baseline (non-CER) model, for each physical risk type (earthquake, flood, flood perception and foundation) and each loan behaviour (self-funding, AAB-funded and externally-funded). It is important to note that these are conditional relative change factors—i.e. they represent the effect of a CER event assuming it occurs with certainty, independent of its probability. As such, some values may appear large, but this reflects the unweighted scenario-specific impact.

The most striking observation is that the self-funding scenario results in the highest relative impact across all risk types. This is consistent with expectations: as shown in the model extension flowchart (recall Figure 3.6), the self-funding path directly impacts PD via a full deduction of the repair cost from disposable income. This direct burden naturally results in the largest upward adjustment in predicted default probability.

Additionally, the risk-specific severity is clearly reflected. Foundation risk consistently leads to the highest relative impact across all behaviours, which aligns with its comparatively high damage costs. Earthquake risk, by contrast, shows the smallest effect, again consistent with the underlying damage calculations.

While none of these outcomes as presented above are surprising—and they align with theoretical expectations—they do highlight this model's limitation in nuance. These relative factors represent a useful first approximation, but a more refined modelling approach is needed to capture the complex relationship between creditworthiness and PD. This motivates the next section, where the inclusion of additional explanatory variables aims to improve this behavioural-PD mapping.

4.1.2 Extended model results — multiple dependence (LTI, Terms in Arrears, BKR)

With the inclusion of two additional input features—implemented as outlined in Section 3.3.3—the following regression outcomes were obtained for the multivariate linear model on the logit-transformed PD

$$\widehat{\mathrm{PD}} = \frac{1}{1 + e^{-\left(-7.2216 + 0.0912 \cdot X_1 + 1.7344 \cdot X_2 + \sum_{i=1}^{g} \beta_{3,i} \cdot \mathbb{1}\{\mathrm{BKR}_i = i\}\right)}}, \tag{4.3}$$

where

- $X_1 = \mathrm{LTI}_i$: the estimated Loan-to-Income ratio,
- $X_2 = \mathrm{Arrears}_i$: the number of missed payments in the past 12 months,
- $X_3 = \mathrm{BKR}_i$: the credit registration category (dummy encoded, with category B as the reference),
- $\beta_0 = -7.2216$: intercept term,
- $\beta_1 = 0.0912$, $\beta_2 = 1.7344$: coefficients for LTI and arrears respectively,
- $\beta_{3,i}$: coefficients for each BKR category $i \in \{\mathrm{A, C, D, ..., L, P, Null}\}$.

The estimated coefficients for the BKR categories (relative to reference category B) are as follows

$$\begin{aligned}
\beta_{3,\mathrm{A}} &= -0.8401 \\
\beta_{3,\mathrm{C}} &= 0.2806 \\
\beta_{3,\mathrm{D}} &= 0.6610 \\
\beta_{3,\mathrm{E}} &= 0.8909 \\
\beta_{3,\mathrm{F}} &= 1.1184 \\
\beta_{3,\mathrm{G}} &= 1.3274 \\
\beta_{3,\mathrm{H}} &= 1.7624 \\
\beta_{3,\mathrm{I}} &= 1.9534 \\
\beta_{3,\mathrm{J}} &= 2.0956 \\
\beta_{3,\mathrm{K}} &= 2.4102 \\
\beta_{3,\mathrm{L}} &= 2.7575 \\
\beta_{3,\mathrm{P}} &= 2.8544 \\
\beta_{3,\mathrm{Null}} &= 0.7440
\end{aligned}$$

For completeness, Table 4.2 reports the regression results of the extended model, which includes LTI, arrears and BKR categories as explanatory variables.

It follows from these results that the LTI coefficient remains positive and statistically significant, with its 95% confidence interval again lying strictly above zero. Compared to the base model, the estimated effect size of LTI is reduced, suggesting that part of its explanatory content overlaps with arrears and BKR categories.

The arrears variable exhibits a large and statistically significant positive coefficient. While direct comparison of coefficient magnitudes is not meaningful due to different variable scales, the positive sign and the size of the coefficient for *Terms in Arrears* indicate that higher arrears are associated with higher log-odds of default (and thus higher predicted PD values). This is consistent with the earlier visual analysis in Figure 4.1, where a clear relationship with observed PD values was visible.

Regarding the BKR indicators, the only category with a negative coefficient is `A`, implying lower log-odds of default relative to the reference category `B`. This aligns with the BKR classification, where an `A` rating reflects a healthy credit history. Most other categories have positive and statistically significant coefficients, meaning that—relative to `B`—they are associated with higher predicted default probabilities. Coefficients generally increase for higher BKR categories.

Table 4.2: Extended model regression results (log-odds scale)

Variable	Estimate	Std. Error	t-statistic	p-value	CI lower	CI upper
Intercept	-7.2216	0.0028	-2542.28	<0.001	-7.2271	-7.2160
LTI	0.0912	0.0007	124.32	<0.001	0.0897	0.0926
Arrears (last 12M)	1.7344	0.0540	32.09	<0.001	1.6285	1.8403
BKR A	-0.0409	0.0032	-12.70	<0.001	-0.0472	-0.0346
BKR C	0.2806	0.0029	98.09	<0.001	0.2750	0.2862
BKR D	0.6011	0.0035	171.14	<0.001	0.5943	0.6080
BKR E	0.8098	0.0040	201.14	<0.001	0.8019	0.8177
BKR F	1.1187	0.0048	233.27	<0.001	1.1093	1.1281
BKR G	1.3238	0.0063	211.41	<0.001	1.3115	1.3361
BKR H	1.7063	0.0087	195.92	<0.001	1.6893	1.7234
BKR I	1.9534	0.0118	165.06	<0.001	1.9302	1.9766
BKR J	2.0953	0.0155	135.21	<0.001	2.0649	2.1256
BKR K	2.4103	0.0180	133.88	<0.001	2.3750	2.4456
BKR L	2.7575	0.0222	124.21	<0.001	2.7140	2.8011
BKR P (combined)	2.8544	0.0378	75.45	<0.001	2.7802	2.9285
BKR Null	0.7440	0.0026	290.69	<0.001	0.7390	0.7491

Taken together, these results suggest that, even after adding variables known to be related to PD, the estimated LTI retains some explanatory power in this multivariate specification.

Figure 4.3 shows the scatter plot and calibration curve for the results of the multiple dependence model. A key observation in the scatter plot (left) is the appearance of multiple linear bands in the estimated PD values. This is expected: due to the inclusion of discrete variables—most notably the BKR dummies and Terms in Arrears—combinations of input values lead to distinct shifts in the predicted logit scale, resulting in these banded patterns across the LTI range.

These structural changes in the prediction function are also reflected in the calibration curve. Both issues observed under the single dependence model—namely, local non-monotonicity and truncation at the upper end—are largely resolved. The predicted PDs now cover almost the full range of observed PD values, and the curve exhibits a consistently monotonically increasing trend. However, some degree of underestimation is still visible across most of the prediction range.

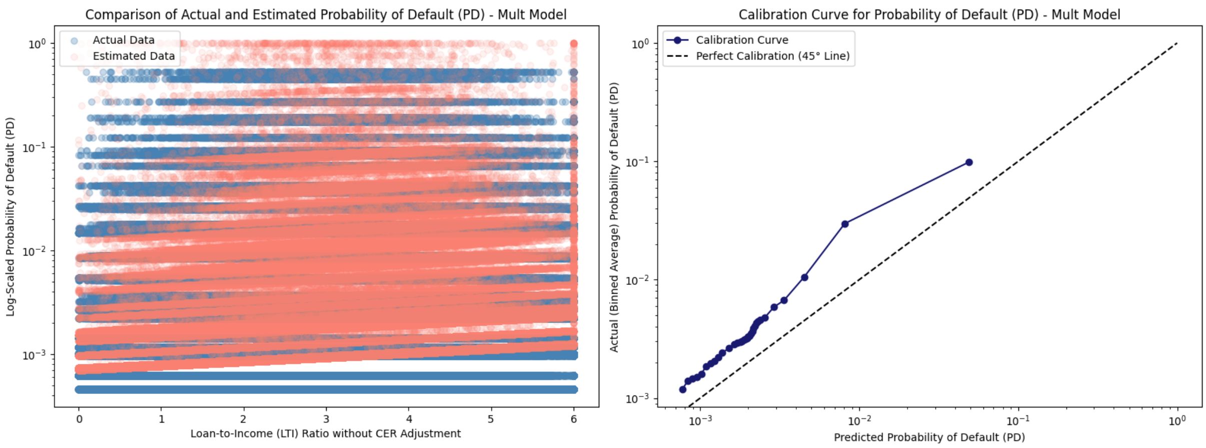

Figure 4.3: Comparison between actual and estimated PD under the multiple-dependence model. Left: log-scaled scatter plot of observed vs. predicted PDs. Right: calibration curve showing the model's performance across PD buckets.

Lastly, the relative CER impact factors under the multiple dependence model are shown in Figure 4.4. Many of the patterns observed under the single-dependence specification are still present—for instance, the highest impact is again found for foundation risk and self-funded repair behaviour. However, the overall relative impact factors appear slightly lower compared to the single-dependence results. This outcome is expected, as the addition of other known significant input features helps better distribute the explanatory power across the model. These lower relative factors likely offer a more realistic approximation of the true relative change in PD conditional on a CER event occurring.

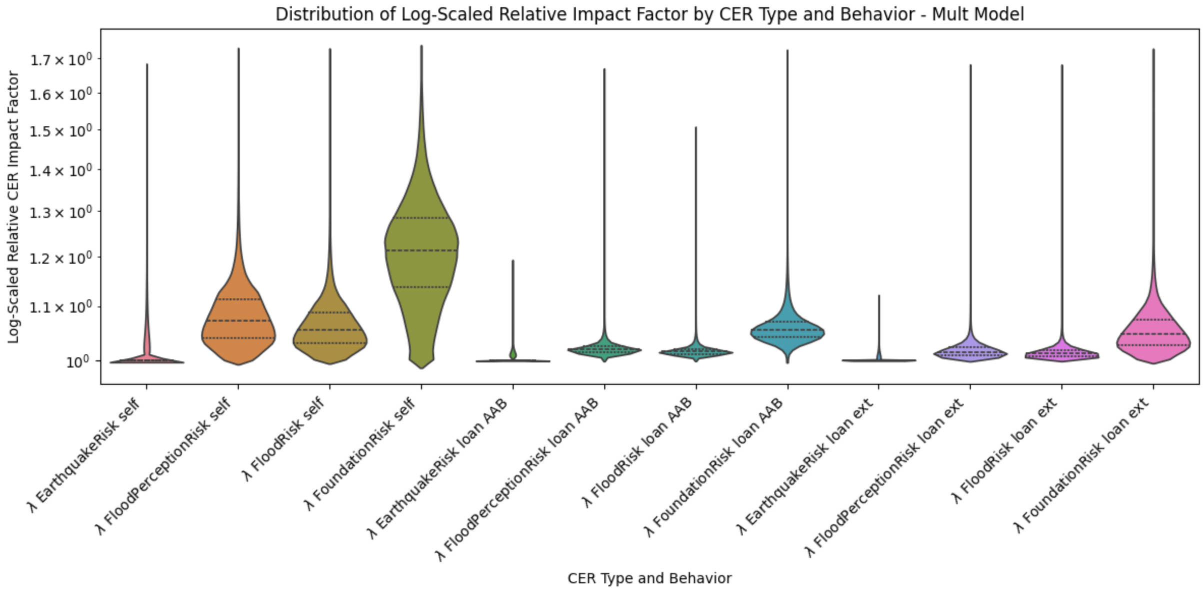

Figure 4.4: Distribution of log-scaled relative impact factor λ for each CER type and behavioural path, based on the extended multiple-dependence PD model. This model captures additional behavioural risk drivers.

Overall, the results demonstrate improved alignment with the true data distribution compared to the single-variable model. Combined with the performance metrics discussed in the next section, these findings support the case for adopting the extended model in the remainder of the impact estimation framework.

4.1.3 Model performance evaluation

To complement the earlier visual analysis, both rank-based and discrimination-focused performance metrics were used to evaluate the predictive quality of the different PD models. Specifically, Kendall's τ and Goodman–Kruskal's γ are used to assess the ordinal association between predicted and observed PD rankings, while the Gini coefficient evaluates the model's ability to correctly distinguish between defaulted and non-defaulted borrowers. These metrics are theoretically defined in Sections 3.3.4 and 3.3.4. Higher values across all metrics indicate stronger agreement between predicted and observed credit risk levels, and therefore better model performance.

Table 4.3 presents the results for both the base model and the extended multiple-variable model.

Metric	**Base model**	**Multiple-variable model**
Kendall's τ	0.1467	0.3570
Goodman–Kruskal's γ	0.1528	0.3719
Gini coefficient	0.19	0.68

Table 4.3: Comparison of model performance metrics for the base and extended PD model specifications

The results clearly show that the multiple-variable model outperforms the base model across all three metrics. Both Kendall's τ and Goodman–Kruskal's γ increase to more than twice their base values for the multiple-dependence model, indicating stronger preservation of rank order. The improvement in γ, which is more sensitive to directional agreement in the presence of ties, is particularly relevant given the discretised nature of the predicted PD outputs.

In addition, the Gini coefficient increases from 0.19 to 0.68, suggesting that the extended model exhibits significantly stronger discriminatory power. Additionally, the Gini value of 0.68 also falls within the "acceptable" performance range defined by the European Banking Authority (EBA), where Gini scores above 0.6 are considered acceptable and those approaching 0.7 are regarded as strong indicators of discrimination quality [53]. These results indicate that the multiple-variable model is better at ranking borrowers by risk and can be reliably used for further impact estimation.

Taken together, these results provide sufficient evidence that incorporating additional features—such as arrears history and BKR registration—improves the model's ability to differentiate risk across the borrower population. In line with this, and consistent with earlier findings from the calibration analysis, the multiple-variable model—defined in Equation (4.3)—is selected for use in the subsequent CER impact estimation phase. All final results presented in the remainder of this study are based on the PD estimates produced by this extended model specification.

4.2 Impact estimation

This section presents the initial iteration of the CER impact estimation results, generated according to the extended model framework defined in the methodology. The focus is placed on a selected set of behavioural scenarios that isolate individual behavioural pathways, along with one benchmark scenario reflecting the current best estimate of behavioural probabilities. The results cover the following cases:

- **No-repair scenario:** The first scenario sets $P_{\text{norepair}} = 1$, representing the situation where no borrowers choose to repair the damage. This effectively reflects the original model setup prior to extension and serves as a validation check, ensuring consistency with previously established results.
- **Repair scenarios with isolated funding behaviours:** In the next set of scenarios, repair is assumed to be universally chosen (i.e. $P_{\text{repair}} = 1$), while one specific funding behaviour is isolated at a time:
  - *Self-funded repair only:* $P_{\text{self|repair}} = 1$, with the other conditional probabilities set to zero.
  - *Bank-funded repair (ABN AMRO loan):* $P_{\text{AAB|repair}} = 1$, isolating the case where all borrowers finance repair by taking a new loan from ABN AMRO.
  - *Externally funded repair:* $P_{\text{ext|repair}} = 1$, representing the scenario where borrowers obtain funding from external sources.
- **Repair with mixed funding assumptions:** One additional repair-only scenario includes a mixed set of funding probabilities, reflecting a plausible distribution of borrower behaviour based on expert input obtained in a consultation within the bank: $P_{\text{self|repair}} = 0.2$, $P_{\text{AAB|repair}} = 0.7$, $P_{\text{ext|repair}} = 0.1$. The process of the expert consultation and the rationale behind these values are briefly outlined below.
- **Best estimate scenario:** Finally, a current best estimate scenario is included in which $P_{\text{norepair}} = 0.02$ and $P_{\text{repair}} = 0.98$, derived from the same expert consultation. The conditional financing probabilities within the repair segment are identical to those used in the preceding mixed scenario. This scenario is intended to produce a single set of results that, within the scope of this study, is considered the most realistic estimate of the CER impact under the best-judged behavioural distribution. Again, a brief explanation of the expert's reasoning behind these specific probability values is provided below.

Since the model is applied to a portfolio consisting exclusively of ABN AMRO clients, the most suitable approach for refining these behavioural assumptions is to gather insights from within the bank itself. Therefore, the best estimate probability set was derived through discussions with experts from *ABN AMRO Hypotheken*

Groep, a subsidiary of ABN AMRO responsible for the bank's mortgage activities and one of the largest mortgage providers in the Netherlands. The consultation focused on the experts' expectations regarding the likely behaviour of their mortgage holders in the event of climate-related damage—specifically, whether repairs would be undertaken, and if so, the most probable method of financing those repairs. A first key insight from the discussion was that, for approximately 98% of ABN AMRO mortgage customers, it would generally be possible to obtain an additional loan from ABN AMRO to cover the costs of climate-related repairs. Based on this, the assumption was made that if funds are available—either through a loan or self-financing—homeowners will always prefer to proceed with repairs rather than leave the damage unaddressed. This reasoning underpins the best estimate values of $P_{\text{norepair}} = 0.02$ and $P_{\text{repair}} = 0.98$. When zooming in on the choice of financing method, the experts noted that the decision to self-finance or to take a loan would, in reality, strongly depend on the size of the damage relative to factors such as the household's income. During the final stage of this thesis a side-analysis was initiated to explore such relationships, and preliminary results were generated; these are presented in Appendix C. As indicated there, this work is still at an early stage and will require further parameter tuning and refinement. To maintain model simplicity in the main study, the behavioural probabilities were kept static and not linked to borrower-specific data. Accordingly, the main results reported in this section are based on these static probabilities. Based on the expert advice for these static probabilities, it was concluded that obtaining an additional loan from ABN AMRO is feasible in the vast majority of cases. Repairs were therefore considered most likely to be financed through such a loan, leading to the advised probability of $P_{\text{AAB|repair}} = 0.7$. The experts further assessed that, because this internal financing option is so widely available, the likelihood of seeking external financing would be considerably lower, and recommended a value of $P_{\text{ext|repair}} = 0.1$. The remaining share was assigned to self-financing, resulting in $P_{\text{self|repair}} = 0.2$. As these values directly reflect the expert advice, they were adopted without modification as the best estimate probability set used in this study.

To present a complete and structured overview, the results have been grouped into three main categories. First, a full isolation of the *no-repair* path is presented, serving both as a validation check and as a benchmark for comparison. Second, a range of scenarios are shown in which only the *repair* path is active, further broken down into its three financing types—self-funded, ABN AMRO-funded, and externally funded—as well as a composite case based on a best-estimate parameter combination. Finally, the third group presents an overall best-estimate scenario, in which both repair and no-repair behaviours are included proportionally. This categorisation reflects the structure of the extension framework and determines the layout of the subsections that follow.

As explained in the methodology section, the estimated CER impact is first expressed in terms of changes to the core credit risk metrics: LGD, PD, EAD and ELBE. These components are then used to compute the final impact on Risk Weighted Assets (RWA) and Expected Loss (EL). All reported results are shown as percentage changes.

4.2.1 Current scenario: no-repair only

The first scenario reflects the original assumption embedded in the current model setup, in which all borrower behaviour follows the no-repair pathway. This is represented by setting $P_{\text{norepair}} = 1$, with all other behavioural probabilities equal to zero. The results are shown in Table 4.4.

As expected, the only impacted core credit risk components under this scenario are LGD and ELBE, which is consistent with the definitions used in the model extension. These outcomes are effectively identical to those of the current model and will primarily serve as a point of reference when evaluating the results of more behaviourally scenarios in the remainder of this section.

	Flood Physical	Flood Perception	Foundation	Earthquake
LGD % change (performing)	0.00407	0.00524	0.07638	0.00153
LGD % change (in-default)	0.00142	0.00189	0.04719	0.00041
PD % change	0.00000	0.00000	0.00000	0.00000
EAD % change	0.00000	0.00000	0.00000	0.00000
ELBE % change (in-default)	0.00106	0.00138	0.02898	0.00026
RWA % change	0.00390	0.00504	0.07628	0.00125
EL % change	0.00247	0.00319	0.05119	0.00075

Table 4.4: **No-repair scenario:** CER impact estimation results on LGD, PD, EAD, ELBE, RWA and EL under defined behavioural scenario parameters: $P_{\text{norepair}} = 1$, $P_{\text{repair}} = 0$, $P_{\text{self|repair}} = 0$, $P_{\text{AAB|repair}} = 0$, $P_{\text{ext|repair}} = 0$; income estimation parameters: $\chi_e = 35$, $r_e = 1.06$, $\chi_m = 55$, $r_m = 1.01$, $\chi_l = 65$, $r_o = 0.99$, $r_p = 0.97$; loan external scaling: $k^{\text{loan ext}} = 0.25$.

4.2.2 Repair-only scenario: comparison across financing behaviours

Having considered the baseline scenario in which all borrowers follow the no-repair assumption, the focus now shifts to the first scenario that isolates the alternative: full repair behaviour. Specifically, this first case assumes that all affected borrowers choose to repair the damage and fund the associated costs entirely out-of-pocket—i.e. without borrowing. The corresponding results are presented in Table 4.5.

Compared to the fully no-repair case shown in Table 4.4, a clear difference is observed: for all CER types, the impact (in RWA and EL) more than doubles. This amplification can be attributed solely to changes in PD, as this is the only core credit risk metric affected under this behavioural path. As expected, the increase in PD translates into a higher RWA, reflecting the functional dependency between the two in the IRB capital requirement formula 2.3. While this relationship is not linear, it results in a clear upward shift in RWA consistent with the observed PD changes. Although the absolute percentage changes may seem small, the relative increase—especially compared to the no-repair case—represents a meaningful shift in the risk profile.

The second scenario explores the isolated case in which all borrowers choose to finance repair costs by taking out an additional loan from ABN AMRO. This is referred to as the *AAB-funded repair* path. Intuitively, this pathway is expected to yield the highest overall impact on EL and RWA, as it is the only behavioural branch that affects all three core credit risk metrics simultaneously—namely LGD, PD and EAD. Both RWA and EL are strongly dependent on EAD, making this the only path that influences them through this channel. The corresponding results are shown in Table 4.6.

As anticipated, the magnitude of the impact is substantially higher than in the self-funded repair scenario—almost an order of magnitude greater across considered credit risk types. This outcome aligns with expectations: if all borrowers within the affected portfolio were to cover their damages through additional loans provided by the bank itself, the resulting increase in credit exposure and probability of default would represent a significantly greater risk to the bank, compared to a situation where borrowers fund the repairs independently.

The next scenario considers the isolated case in which all borrowers choose to finance the repair costs through external sources—such as another financial institution or personal contacts—rather than through ABN AMRO. This behavioural pathway is referred to as the *externally funded repair* case. The results for this scenario are presented in Table 4.7.

Among all repair-path isolations, this scenario shows the smallest impact—again in line with expectations.

Like the self-funded repair case, this pathway only affects the PD component. However, the creditworthiness adjustment in this context is derived from a constant-weighted reduction in disposable income, rather than a full subtraction of the total damage amount (recall Equation (3.42)). As a result, the downward adjustment to income—and thus the impact on PD—is more moderate than in the self-funded case.

Finally, within the fully repair-based scenarios, Table 4.8 presents the results for a scenario in which all borrowers choose to repair, while the associated funding behaviour is distributed across a mixed set of probabilities. This setup aims to provide a preliminary approximation of a potentially realistic behavioural mix, and serves as a useful comparison point against the fully no-repair benchmark presented earlier in Table 4.4.

Compared to current ABN AMRO model setup, the impact on RWA and EL magnitudes under this full-repair scenario is, on average, approximately seven times higher. This outcome aligns with theoretical expectations: in the repair branch, multiple core credit risk components—PD, LGD and EAD—are simultaneously affected, leading to a broader transmission of risk.

From an intuitive standpoint, the reasoning is slightly more nuanced. One might initially expect that repairing damage would reduce risk exposure. However, in this model extension, repairs are assumed to restore the collateral value to its original level, so the net collateral value remains unchanged under the repair behaviour. Instead, repair activity introduces new financial obligations—such as additional loans or depletion of savings—that negatively affect borrower creditworthiness and elevate repayment burdens. These behavioural effects ultimately increase the modelled probability of default and thus the overall risk profile. In terms of comparison with the other isolated loan behaviours, the overall impact falls somewhere in between: lower than in the high-risk loan-from-ABN-AMRO scenario, but higher than in the lowest-risk case of externally funded repairs. This is in line with expectations, as the different funding types affect the bank's exposure to varying degrees.

	Flood Physical	**Flood Perception**	**Foundation**	**Earthquake**
LGD % change (performing)	0.00000	0.00000	0.00000	0.00000
LGD % change (in-default)	0.00000	0.00000	0.00000	0.00000
PD % change	0.00858	0.00838	0.09846	0.00134
EAD % change	0.00000	0.00000	0.00000	0.00000
ELBE % change (in-default)	0.00000	0.00000	0.00000	0.00000
RWA % change	0.00387	0.00495	0.05810	0.00113
EL % change	0.00308	0.00393	0.04615	0.00010

Table 4.5: **Self-funded repair scenario:** CER impact estimation results on LGD, PD, EAD, ELBE, RWA and EL under defined behavioural scenario parameters: $P_{\text{norepair}} = 0$, $P_{\text{repair}} = 1$, $P_{\text{self|repair}} = 1$, $P_{\text{AAB|repair}} = 0$, $P_{\text{ext|repair}} = 0$; income estimation parameters: $\chi_e = 35$, $r_e = 1.06$, $\chi_m = 55$, $r_m = 1.01$, $\chi_l = 65$, $r_o = 0.99$, $r_p = 0.97$; loan external scaling: $k^{\text{loan ext}} = 0.25$.

	Flood Physical	Flood Perception	Foundation	Earthquake
LGD % change (performing)	0.00598	0.00753	0.10827	0.00215
LGD % change (in-default)	0.00214	0.00283	0.06504	0.00061
PD % change	0.00181	0.00227	0.02941	0.00040
EAD % change	0.03138	0.03923	0.68773	0.00800
ELBE % change (in-default)	0.00106	0.00138	0.02898	0.00026
RWA % change	0.03815	0.04791	0.81525	0.00317
EL % change	0.03554	0.04458	0.76822	0.00245

Table 4.6: **Loan ABN AMRO repair scenario:** CER impact estimation results on LGD, PD, EAD, ELBE, RWA and EL under defined behavioural scenario parameters: $P_{\text{norepair}} = 0$, $P_{\text{repair}} = 1$, $P_{\text{self|repair}} = 0$, $P_{\text{AAB|repair}} = 1$, $P_{\text{ext|repair}} = 0$; income estimation parameters: $\chi_e = 35$, $r_e = 1.06$, $\chi_m = 55$, $r_m = 1.01$, $\chi_l = 65$, $r_o = 0.99$, $r_p = 0.97$; loan external scaling: $k^{\text{loan ext}} = 0.25$.

	Flood Physical	Flood Perception	Foundation	Earthquake
LGD % change (performing)	0.00000	0.00000	0.00000	0.00000
LGD % change (in-default)	0.00000	0.00000	0.00000	0.00000
PD % change	0.00144	0.00182	0.02335	0.00029
EAD % change	0.00000	0.00000	0.00000	0.00000
ELBE % change (in-default)	0.00000	0.00000	0.00000	0.00000
RWA % change	0.00085	0.00108	0.01555	0.00013
EL % change	0.00067	0.00085	0.01235	0.00002

Table 4.7: **Loan externally repair scenario:** CER impact estimation results on LGD, PD, EAD, ELBE, RWA and EL under defined behavioural scenario parameters: $P_{\text{norepair}} = 0$, $P_{\text{repair}} = 1$, $P_{\text{self|repair}} = 0$, $P_{\text{AAB|repair}} = 0$, $P_{\text{ext|repair}} = 1$; income estimation parameters: $\chi_e = 35$, $r_e = 1.06$, $\chi_m = 55$, $r_m = 1.01$, $\chi_l = 65$, $r_o = 0.99$, $r_p = 0.97$; loan external scaling: $k^{\text{loan ext}} = 0.25$.

	Flood Physical	Flood Perception	Foundation	Earthquake
LGD % change (performing)	0.00410	0.00527	0.07579	0.00151
LGD % change (in-default)	0.00150	0.00198	0.04553	0.00043
PD % change	0.00273	0.00345	0.04292	0.00057
EAD % change	0.02197	0.02746	0.48141	0.00560
ELBE % change (in-default)	0.00074	0.00096	0.02029	0.00018
RWA % change	0.02757	0.03463	0.58374	0.00225
EL % change	0.02556	0.03206	0.54816	0.00174

Table 4.8: **Mixed repair scenario:** CER impact estimation results on LGD, PD, EAD, ELBE, RWA and EL under defined behavioural scenario parameters: $P_{\text{norepair}} = 0$, $P_{\text{repair}} = 1$, $P_{\text{self|repair}} = 0.2$, $P_{\text{AAB|repair}} = 0.7$, $P_{\text{ext|repair}} = 0.1$; income estimation parameters: $\chi_e = 35$, $r_e = 1.06$, $\chi_m = 55$, $r_m = 1.01$, $\chi_l = 65$, $r_o = 0.99$, $r_p = 0.97$; loan external scaling: $k^{\text{loan ext}} = 0.25$.

4.2.3 Mixed scenario: best-estimate behavioural composition

Finally, the results of the best estimate scenario are presented in Table 4.9. As mentioned earlier, this probability set was derived from expert consultations with the *ABN AMRO Hypotheken Groep*, reflecting their expectations of mortgage holders' behaviour in the event of climate-related damage. The resulting best estimate composition assumes that repairs are almost always undertaken, with the majority financed through an additional ABN AMRO loan, and smaller shares through self-financing or external funding. This section presents the model outcomes under this probability set, which is considered the most realistic representation within the scope of this study.

The results can first be compared to the baseline case: the current setup under the full no-repair assumption. Although the magnitude of effects varies by risk type, the best estimate scenario yields, on average, more than a sevenfold increase in EL and RWA relative to the existing model. This is based on the assumption that 98% of the affected portfolio opts for repair and 2% does not. Among those who do repair, 70% are assumed to finance the costs through an ABN AMRO loan, 20% through self-funding, and the remaining 10% through external funding. Given that the loan-from-ABN-AMRO path was shown to produce the highest impact, it is not surprising that this mixed scenario still results in a notable increase in overall impact.

	Flood Physical	**Flood Perception**	**Foundation**	**Earthquake**
LGD % change (performing)	0.00410	0.00527	0.07580	0.00151
LGD % change (in-default)	0.00150	0.00198	0.04556	0.00043
PD % change	0.00267	0.00338	0.04206	0.00056
EAD % change	0.02153	0.02691	0.47178	0.00549
ELBE % change (in-default)	0.00075	0.00097	0.02046	0.00018
RWA % change	0.02709	0.03404	0.57358	0.00223
EL % change	0.02510	0.03148	0.53822	0.00172

Table 4.9: **Best estimate mixed scenario:** CER impact estimation results on LGD, PD, EAD, RWA and EL under defined behavioural scenario parameters: $P_{\text{norepair}} = 0.02$, $P_{\text{repair}} = 0.98$, $P_{\text{self|repair}} = 0.2$, $P_{\text{AAB|repair}} = 0.7$, $P_{\text{ext|repair}} = 0.1$; income estimation parameters: $\chi_e = 35$, $r_e = 1.06$, $\chi_m = 55$, $r_m = 1.01$, $\chi_l = 65$, $r_o = 0.99$, $r_p = 0.97$; loan external scaling: $k^{\text{loan ext}} = 0.25$.

An average increase of this magnitude is by no means negligible and therefore requires careful interpretation. The key question is: what precisely drives this large change in RWA and EL? A closer look at the results shows that a substantial share of the RWA increase stems from the rise in EAD. This raises the question: what does it actually mean when the primary cause of the impact is higher exposure? Upon reflection, an increase in EAD represents more an expansion of the ABN AMRO mortgage portfolio than the existing portfolio itself becoming riskier. In other words, a higher RWA is not always equivalent to a higher level of risk, and the bank would not necessarily view this as negative. When an increase in RWA is driven by higher PD and/or LGD, it directly signals a deterioration in portfolio quality, as the risk per euro of exposure rises. By contrast, when the rise in RWA is largely explained by higher EAD, the interpretation is more nuanced: greater exposure means the bank has more loans outstanding, which raises the absolute amount at risk but does not necessarily indicate higher risk *per euro lent*. Simply put, from a business perspective, a higher EAD is generally perceived as positive, as it means more lending and potentially more revenue.

Following this reasoning, it was considered essential to also re-examine the best estimate scenario from another angle—namely, by isolating the impact of changes in PD and LGD while holding EAD constant at its non-CER value. This allows for a more direct assessment of the true increase in the portfolio's risk profile. The results of this adjusted analysis are presented in Table 4.10. Note that when only percentage changes are evaluated while keeping EAD constant, the reported change in RWA is identical to the change in the portfolio-average risk weight, earlier denoted by K (see Equation 2.3) and also commonly abbreviated as RW (Risk-Weights) in credit-risk literature. For consistency, the table lists this as RWA (% change), but know that under these conditions it conveys the same information as the RW. No equivalent established term exists for EL excluding EAD, so this is shown simply as EL (% change).

	Flood Physical	Flood Perception	Foundation	Earthquake
RWA (% change)	0.00704	0.00900	0.11988	0.00135
EL (% change)	0.00474	0.00605	0.08086	0.00077

Table 4.10: **Best estimate mixed scenario (constant EAD)**: CER impact estimation results with EAD kept constant for RWA (i.e. RW) and EL under defined behavioural scenario parameters: $P_{\text{norepair}} = 0.02$, $P_{\text{repair}} = 0.98$, $P_{\text{self|repair}} = 0.2$, $P_{\text{AAB|repair}} = 0.7$, $P_{\text{ext|repair}} = 0.1$; income estimation parameters: $\chi_e = 35$, $r_e = 1.06$, $\chi_m = 55$, $r_m = 1.01$, $\chi_l = 65$, $r_o = 0.99$, $r_p = 0.97$; loan external scaling: $k^{\text{loan ext}} = 0.25$.

	Flood Physical	Flood Perception	Foundation	Earthquake
RWA (% change)	0.00390	0.00504	0.07628	0.00125
EL (% change)	0.00247	0.00319	0.05119	0.00075

Table 4.11: **No-repair scenario**: CER impact estimation results with EAD kept constant for RWA (i.e. RW) and EL under defined behavioural scenario parameters: $P_{\text{norepair}} = 1$, $P_{\text{repair}} = 0$, $P_{\text{self|repair}} = 0$, $P_{\text{AAB|repair}} = 0$, $P_{\text{ext|repair}} = 0$; income estimation parameters: $\chi_e = 35$, $r_e = 1.06$, $\chi_m = 55$, $r_m = 1.01$, $\chi_l = 65$, $r_o = 0.99$, $r_p = 0.97$; loan external scaling: $k^{\text{loan ext}} = 0.25$.

For ease of comparison, Table 4.11 again presents the RWA and EL results of the no-repair scenario, i.e. the baseline against which all other scenarios were also evaluated. As can be seen—and as expected—both RWA and EL still increase across all risk types when moving from the no-repair baseline to the fully mixed scenario. While the changes in impact are clearly lower than in the case where CER-related EAD changes were incorporated, there is still a noticeable increase even when keeping exposure constant. When placing the results of both tables side by side, the relative increases become evident: for Flood Physical and Flood Perception, RWA rises by approximately 80% and EL by roughly 90%. Foundation also shows increases, with both RWA and EL growing by nearly 60%. By contrast, the Earthquake category exhibits only minor differences, with RWA rising by about 8% and EL by less than 3%. This comparison highlights that the overall impact of the best-estimate behavioural composition still is consistently higher than the no-repair baseline. At the same time, excluding EAD from the equation leads to substantially different results for RWA and EL. From a risk perspective, these adjusted outcomes are arguably more appropriate for drawing conclusions, as they reflect the risk profile of the unchanged portfolio rather than an expansion of exposures.

This concludes the empirical results from the CER impact estimation model. To better understand and assess the sensitivity of the model outcomes to the behavioural probability assumptions, an additional sensitivity analysis was performed. The results of this analysis are presented in the following part.

4.2.4 Sensitivity analysis

To better understand the effect of varying behavioural probability assumptions on the impact estimation results, a sensitivity analysis was performed in addition to just presented results. The analysis systematically evaluated the effect of different combinations of the four key behavioural probabilities: P_{norepair}, $P_{\text{self|repair}}$, $P_{\text{AAB|repair}}$ and $P_{\text{ext|repair}}$. Each of these probabilities was assigned a value from the set $\{0, 0.25, 0.5, 0.75, 1\}$, resulting in a total of 35 valid combinations that sum to one. For each combination, the model was run to observe the resulting variation in key impact metrics.

The outcomes are visualised using 3D tetrahedral plots, where each corner of the tetrahedron represents one of the four behavioural pathways being assigned the full probability mass (i.e. a probability of 1, with all others equal to 0). All other internal points represent valid combinations of the four probabilities, constrained to the discrete values defined earlier. The colour of each point reflects the resulting impact estimate for a given metric under that specific behavioural assumption set, with the corresponding value indicated by the colour bar legend on the right side of the plot.

The sensitivity analysis covered five core impact metrics—probability of default (PD), exposure at default (EAD), loss given default (LGD) for both performing and in-default loans, expected loss (EL), and risk-weighted assets (RWA). Although the procedure was repeated for all four climate-risk drivers (flood risk, flood perception, foundation risk, and earthquake risk), only the results for foundation risk are presented. Foundation risk showed the largest impacts and thus represents the most extreme case; the other drivers were analysed and displayed similar or weaker sensitivities, so their plots are omitted here. The results of the foundation risk sensitivity analysis are presented in Figure 4.5. For each of the credit risk metrics, the corresponding sensitivity visualisation will be briefly discussed.

As can be seen, the impact on Loss Given Default (LGD) is primarily concentrated between the *No Repair* and *Repair AAB* corners of the tetrahedron, as shown in Subfigures 4.5a and 4.5b. This is consistent with the model assumptions: these are the only two behavioural pathways that directly affect LGD. In the case of *No Repair*, the damage is not mitigated, leading to a larger effect. Under the *Repair AAB* scenario, the LGD increases due to the additional bank loan, which raises the borrower's loan-to-value ratio. Across all other regions of the plot—especially closer to the *Repair Ext* and *Repair Self* corners—the impact on LGD remains lower, confirming that the model behaves as intended.

Subfigure 4.5c illustrates that the Probability of Default (PD) is most sensitive to the *Repair Self* scenario. In this behavioural pathway, the full damage amount is subtracted from the borrower's disposable income, which leads to a sharp increase in default probability. This is reflected in the red colour gradient concentrated around that corner. As the assumed probability mass shifts toward the *Repair Self* behaviour, the model consistently shows higher PD values, confirming the expected directional sensitivity. Other behavioural mixes result in existing but significantly lower PD impacts, which is consistent with expectations based on the model design. While the *Repair AAB* and *Repair Ext* pathways also affect PD, their impact is by definition less severe than in the *Repair Self* case.

The Exposure at Default (EAD) impact is most concentrated around the *Repair AAB* corner, as shown in Subfigure 4.5d. This aligns with the model design, as this is the only behavioural pathway in which a new loan is added to the outstanding exposure. All other scenarios in which the *Repair AAB* probability is set to zero—i.e. the plane defined by *No Repair*, *Repair Self*, and *Repair Ext*—leave EAD unchanged, as these behavioural choices are explicitly defined to have no effect on this metric.

Ultimately, the most critical metrics from the bank's perspective are the Expected Loss (EL) and Risk-Weighted Assets (RWA), as many capital-related decisions are directly based on these values. Put simply, EL reflects the amount of capital that must be held for expected credit events, while RWA serves as a buffer for unexpected losses. Understanding how sensitive these two metrics are to the inclusion of climate risks—and specifically to this model extension—is therefore an important outcome of the analysis.

The sensitivity plots for both the EL and RWA shown in Subfigures 4.5e and 4.5f, reveal similar patterns in how the impact estimates respond to changes in behavioural probabilities. In both cases, the highest impact concentrations are found in the direction of the *Repair AAB* vertex. This observation is consistent with how the model is defined, as this behavioural path is the only one that affects all three underlying credit risk components simultaneously.

The overall range of impact values further indicates that the model outcomes are meaningfully sensitive to the assumed distribution of behavioural responses. For specifically foundation risk, it can be seen that moving

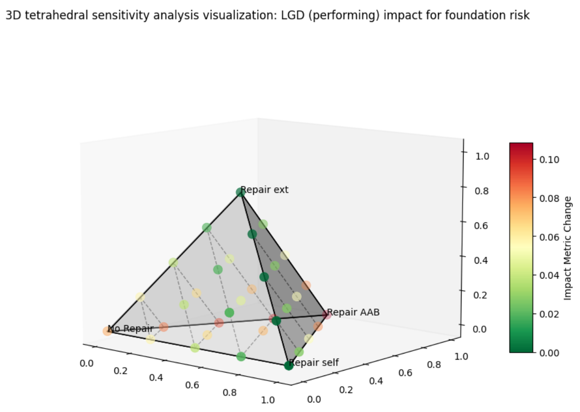

(a) LGD (performing) sensitivity

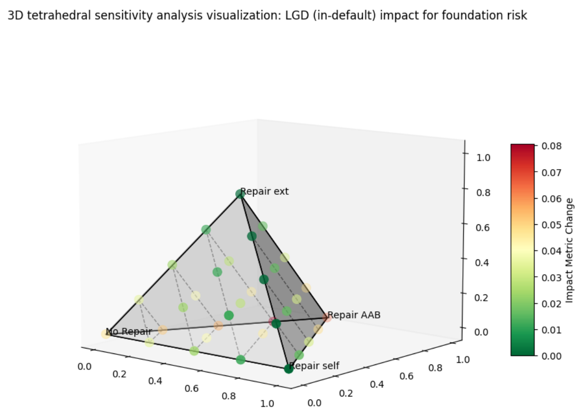

(b) LGD (in-default) sensitivity

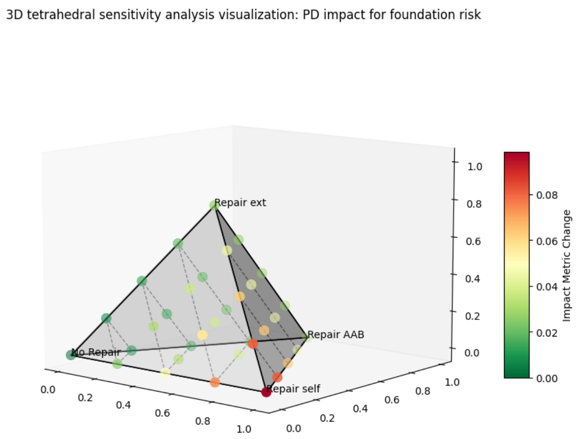

(c) PD sensitivity

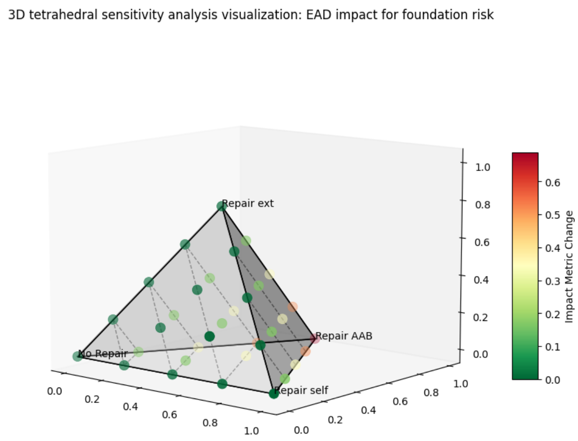

(d) EAD sensitivity

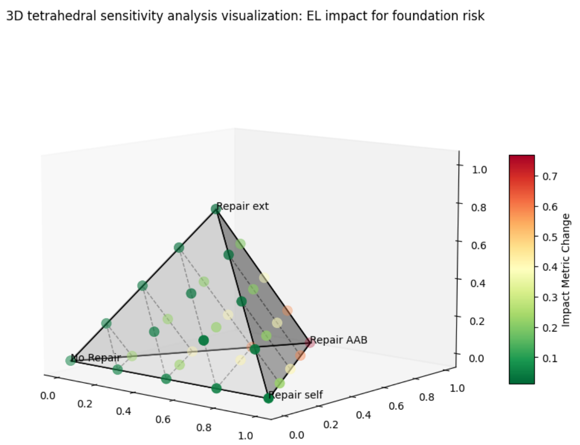

(e) Expected Loss (EL) sensitivity

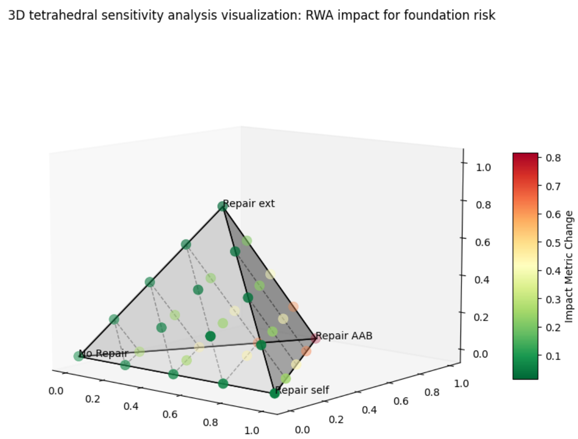

(f) Risk-Weighted Assets (RWA) sensitivity

Figure 4.5: Sensitivity analysis results for the foundation risk scenario. Each subfigure shows how changes in behavioural repair probabilities affect a specific credit risk or regulatory impact metric (in % change).

from a scenario that fully excludes the *Repair AAB* option to one that assigns it a high probability, results in an EL increase of approximately tenfold. In the case of RWA, this corresponds to an increase of almost eightfold compared to the outcome under the scenario excluding the *Repair AAB* pathway. This contrast highlights that the choice of behavioural assumption set does influence the change in impact. As such, the analysis underscores the importance of either carefully justifying the selected behavioural probabilities or complementing the model with expert input.

4.2.5 Portfolio-level materiality assessment

Beyond the technical estimation of climate impacts, an important question is what these results actually imply in practice. Within banks, materiality thresholds are applied to determine how risks are managed and if adjustments to internal models are required. Here, *materiality* refers to whether an effect is large enough to be considered relevant for decision-making or regulatory purposes. If in this context such thresholds are exceeded, climate-related impacts may need to be integrated into internal models for PD, LGD and EAD, which would constitute a substantial change with implications not only for the bank itself but also for its borrowers. Although the primary objective of this research was never to determine whether the climate impact is material, but rather to propose a more robust framework, it is nevertheless valuable to also analyse the results from this broader perspective, which will therefore be addressed in this final results section.

While absolute EL and RWA values of the mortgage portfolio cannot be disclosed for confidentiality reasons, it is important to note that they amount to several billions of euros. This means that even seemingly small relative changes, such as those reported in Sections 4.2.1–4.2.3, can correspond to sizeable monetary amounts when expressed in absolute terms. In this final results section, the estimated absolute impacts are expressed as percentages of the total portfolio, based on the scenario that produces the largest effects in absolute terms. This 'worst-case scenario' corresponds to all behavioural probability weight being placed on the repair option financed through ABN AMRO loans. The outcomes of this analysis are presented in Table 4.12.

Category	EL (% of total)	RWA (% of total)
Flood Physical	0.008	0.010
Flood Perception	0.011	0.013
Foundation	0.164	0.190
Earthquake	0.001	0.001

Table 4.12: Relative portfolio-level impact of climate risks (as % of total portfolio) of worst-case scenario.

The results show that, even in the most adverse scenario considered, the increase in EL and RWA remains well below 1% of the total portfolio. In other words, while the behavioural modelling approach leads to measurable differences compared to the previous framework, the absolute impacts remain immaterial at portfolio level under ABN AMRO's internal materiality thresholds.

This means that, based on the results of the behavioural model extension as well, the conclusion remains that it is not yet appropriate to integrate climate impacts into ABN AMRO's PD, LGD and/or EAD models. This finding is consistent with the initial no-repair scenario, in which the materiality thresholds were also not exceeded. The impacts observed under the behavioural framework are certainly higher than under the current approach, but still not sufficiently large for the bank to justify a internal model revision.

4.2.6 Regulatory perspective

To further interpret the finding of immateriality from the last section, one possible explanation is that the probability of extreme climate events occurring in the Netherlands is relatively low compared to many other European countries [54]. This observation also provides a moment to reflect on the broader objective of the European regulators in mandating research into climate-related risks, even in countries where the likelihood of such events is minimal. While one might question the necessity of conducting such assessments in the

Dutch context, the ultimate aim extends beyond national borders. The ambition of the European financial regulators is to ensure that the European financial system as a whole becomes more resilient to climate risks, while also enabling the exchange of best practices across member states [55].

In the Dutch case, the probability of severe events is limited, which helps explain why the impacts observed here remain below materiality thresholds. However, in other European countries the likelihood and potential severity of climate-related events is significantly higher. In those contexts, internal assessments are expected to lead to more substantial adjustments, thereby reinforcing their relevance. By requiring all banks to perform such analyses, the European regulators aim to advance a harmonised and robust framework for integrating climate risks into credit risk modelling, ultimately safeguarding both banks and their borrowers across Europe.

Thus, although both ABN AMRO's initial assessments and the results of this research indicate that climate risks in the Netherlands are not currently material enough to warrant adjusting internal models, the bank's contribution still forms part of the collective effort to strengthen the resilience of the European banking sector in the face of increasing climate-related risks.

5

Discussion

This section provides a chronological summary of the main components of the research. Along the way, key limitations are discussed at each stage, and where relevant, suggestions for future research are proposed to address them.

The first step of this research was to analyse ABN AMRO's existing framework, which already quantifies four types of physical climate risk: flood risk, flood risk perception, foundation risk and earthquake risk. Estimated damages, derived from public data, are linked to these risks at the individual mortgage level. While this step involves simplifications, the focus of this thesis was on a later stage of the pipeline: the translation of damages into credit risk metrics. In the current setup, it is assumed that homeowners do not repair damage, resulting in an effect only on LGD through reduced collateral value. This 'no-repair' assumption was selected as the focus for model refinement. On one hand, decreasing the collateral value to account for damages caused by an environmental event can be considered a valid approach. On the other hand, it is also plausible that households respond to such damage by repairing the property. Therefore, the research shifted towards designing a behavioural model extension, with each component becoming a result on its own and naturally bringing along certain new limitations. The remainder of this section recaps the full model extension structure, along with the key outcomes that will be discussed and certain limitations at each modelling step.

Model extension design and creditworthiness proxy

A central feature of the proposed extension is the introduction of behavioural dependence into the model, replacing the current deterministic assumption with a set of alternative repair behaviours. Specifically, three financing options were defined: (1) self-funded repair using income or savings, (2) repair financed through an additional loan from ABN AMRO Bank (AAB), and (3) repair financed externally, such as through another financial institution or family support. These alternatives broaden the model's scope, as not only LGD but also PD and EAD are affected, each in distinct ways depending on the chosen repair behaviour. Central to this behavioural framework is the concept of *creditworthiness*, reflecting the idea that repairing damage may reduce financial capacity and thus increase the probability of default. Because present-day income or savings data were not available for this research, a proxy was constructed based on the borrower's income at origination of the mortgage. This was projected forward using age-based income growth parameters derived from public CBS data, producing personalised income forecasts per borrower. *Limitation 1:* While this method allowed for a data-driven income estimate, it still relies on population averages and does not account for individual income dynamics. Borrowers with atypical career or income paths may therefore be misrepresented. Additionally, income growth was assumed to be constant and unaffected by macroeconomic shifts or household-specific shocks. *Further research* could include calibrating the forecasting model using bank-internal borrower data or incorporating external macroeconomic scenarios to better reflect individual income volatility.

PD estimation models

The income estimates from the previous step were transformed into a Loan-to-Income (LTI) ratio, serving as a proxy for creditworthiness. Because LTI is not an existing input in the current PD model, a new relationship between LTI and PD had to be established. Several model options were tested, and a logistic regression on the logit of PD was ultimately selected for its simplicity and bounded output. Two variants were compared: a base model using only LTI, and an extended model including additional information available to the bank — specifically, BKR categories and terms in arrears. Performance was evaluated using the Gini coefficient, Kendall's τ, and Gamma concordance, with the extended model consistently outperforming the base version. Given that the Gini increased from 0.19 to 0.68 — surpassing the EBA's benchmark threshold of 0.6 — the extended model was deemed sufficiently accurate for continued use in the broader modelling framework. *Limitation 2:* In the extended PD model, creditworthiness is allowed to vary following climate damage — via adjusted income or debt and the resulting LTI. However, other model inputs such as BKR status and terms in arrears are assumed to remain fixed. In practice, these variables may be jointly affected by financial stress, creating an internal inconsistency in the model logic. *Further research* could involve dynamic estimation of these related inputs, exploring whether and how indicators like arrears and credit registry scores co-move with income or debt shocks, instead of treating them as independent variables.

Additionally, a broader limitation, resulting from the timeframe in which this research had to be conducted, is that developing a comprehensive integrated creditworthiness-based PD model was not feasible. The creation of such an alternative PD model, including a creditworthiness variable and following the bank's actual modelling procedures, was therefore beyond the scope of this thesis. For this reason, the analysis was limited to two simplified alternatives. *Further research* within the bank could involve a full investigation of risk driver selection, including more robust performance testing.

Damage integration into credit risk metrics

With a proxy for creditworthiness now established through the estimated LTI, the next step was to integrate the impact of climate-related damage into this proxy. For the no-repair scenario, this followed the original approach: damage reduced the collateral value, which increased the Loan-to-Value (LTV) ratio and thereby affected LGD through a CER-adjusted LGD calculation. For the three repair-related behavioural paths, the mechanism differed by behaviour. In the self-funding scenario, the damage amount was subtracted from estimated disposable income (i.e., the denominator of LTI), leading to a lower creditworthiness proxy and hence a higher CER-adjusted PD. In the internal loan case, the total loan amount increased, raising both the LTI numerator and LTV, while also increasing the bank's exposure — affecting EAD. The external loan path assumed unchanged exposure, but still reduced disposable income due to repayments, thus impacting PD via LTI. *Limitation 3:* The behavioural impact of climate damage on PD was modelled solely via changes in the Loan-to-Income (LTI) ratio. This implicitly assumes that all financial consequences of damage — whether from income loss or increased liabilities — are fully captured through this single proxy. In reality, the link between financial strain and default risk may also largely depend on factors such as available savings or a household's willingness to take financial risks. One alternative — also considered during this research — would be to assume that borrowers hold a savings buffer equal to a fixed percentage of their annual gross income. In that case, the damage could first be deducted from savings, and only if this is insufficient, disposable income (and thus LTI) would be impacted. This could potentially provide a more accurate reflection of how households typically manage repair costs, but it also would introduce another uncertainty about the actual size and use of savings. For this reason, the more conservative approach of directly subtracting the damage from disposable income was used. *Further research* could compare such alternative approaches and validate them by obtaining household-level data on financial decisions following climate-related damage, for example through surveys or existing statistical sources.

Behavioural probabilities and final impact estimation

Together, the previous steps detail how different behavioural responses to climate damage can affect LGD, PD, and in some cases EAD, through distinct mechanisms. To compute the expected values of these CER-adjusted metrics, the model incorporated both the probability of each climate risk occurring — based on ABN AMRO's existing assumptions — and the probability that a homeowner repairs the damage and how. The combination of four behavioural scenarios — no repair, self-funded repair, internal loan and external loan — enabled a flexible framework in which different credit risk metrics are affected in different ways. To aggregate the outcomes, behavioural probabilities were applied at mortgage level, resulting in CER-adjusted values for LGD, PD, and EAD. These were then translated into regulatory risk metrics — Risk-Weighted Assets (RWA) and Expected Loss (EL) — using the IRB framework.

To assess the sensitivity of the model to these behavioural inputs, a grid-based sensitivity analysis was performed across all four behavioural probabilities, using values in $\{0, 0.25, 0.5, 0.75, 1\}$, resulting in 35 unique combinations. For each, the impact on LGD, PD, EAD, RWA, and EL was evaluated. Results indicated a notable degree of sensitivity. For example, in the case of foundation risk, the RWA impact ranged from a 0.08% change under the no-repair assumption to a 0.82% change when assuming full internal loan financing—a tenfold increase. Given this sensitivity, expert input was used to construct a *best estimate set* of behavioural probabilities, intended to reflect the most plausible distribution across the four paths. Even under this expert-informed configuration, the model estimated roughly seven times the impact compared to the no-repair baseline, highlighting that behavioural dynamics significantly influence final impact assessments. *Limitation 4:* The behavioural probabilities used in the final output — such as the likelihood of repair or the preferred financing method — were derived through qualitative discussions with experts from the bank's mortgage modelling team. While this ensured practical relevance and internal alignment, no structured expert judgment approach was used to guide uncertainty assessment, address differing views, or combine opinions quantitatively. This limits the robustness and transparency of the final estimates, especially considering their direct influence on regulatory metrics. *Further research* could apply structured expert elicitation techniques — such as the Cooke method or Bayesian belief networks — to derive a more defensible and transparent behavioural probability set. These methods allow for systematic uncertainty quantification and expert aggregation, thereby strengthening the credibility of climate-related impact assessments in regulatory contexts.

Another limitation emerged after the expert consultation, when the experts indicated that in their view, the best estimate probability set should be dependent on the amount of damage caused by the climate event. A suggestion for *further research* would therefore be to make the probability of repairing or not repairing dependent on the ratio between the damage amount and the original collateral value—since it is reasonable to assume that if the damage is close to the property's total value, the no-repair choice becomes more likely. Secondly, a suggestion would be to make the financing probabilities within the repair path dependent on the ratio between the damage amount and the homeowner's income. For example, a lower damage-to-income ratio could be associated with a higher likelihood of self-funding, whereas a higher ratio could shift the likelihood towards obtaining an additional loan from ABN AMRO.

This concludes the discussion section, in which the research has been critically reviewed and several suggestions for further research have been proposed. The aim has been to provide both a reflective assessment of the work conducted and to highlight opportunities for methodological refinement in future studies. The next section presents the main conclusions of this thesis, summarising the key takeaways and outlining the most important insights derived from the research.

6 Conclusion

With the increasing frequency and severity of climate-related disasters worldwide, and with regulatory requirements increasingly demanding integration of climate risk into credit risk modelling frameworks, the need for better methods to quantify and incorporate these risks is growing. This is equally relevant for bank *ABN AMRO*, where the integration of climate risk into credit risk modelling is becoming an increasingly important area of internal research and development. As part of the bank's emerging internal efforts to better quantify climate risk within its mortgage portfolio, the aim of this research was to contribute a meaningful extension to their current approach.

To structure this contribution, four sub-questions were defined, each guiding a distinct phase of the research process. The first sub-question focused on understanding the existing framework—specifically, how physical climate-risk damage is estimated and translated into collateral-value impact, and how these damage estimates are subsequently integrated into credit-risk parameters. This was addressed by unpacking ABN AMRO's current three-step method, from physical risk estimation to credit-impact calculation. With this foundation in place, the second sub-question examined the key assumptions and limitations inherent in the current approach. A broad set of assumptions was identified, and after careful consideration the no-repair assumption emerged as having significant potential for refinement because of its limited realism in practice, which motivated the development of a model to address this limitation. This, in turn, laid the groundwork for the third and fourth sub-questions: how a model refinement or extension could be designed, implemented, and evaluated. These were answered by designing and implementing a model extension that introduces behavioural dependence, allowing climate-related events to affect not only LGD but also PD and EAD. The extension was evaluated through sensitivity analysis and by comparing the results with those of the case without the extension, and finally by situating the findings within the broader risk framework to assess their portfolio-level materiality.

Returning to the main research question — *Are there assumptions within ABN AMRO's current framework for physical climate risk quantification that represent a limitation, to what extent could such a limitation affect its integration into retail mortgage credit risk models, and how can this be evaluated and improved?* — several conclusions can now be drawn.

First, regarding the impact of the current *no-repair* assumption: the results have shown that this assumption meaningfully affects the climate impact estimation. On average, incorporating behavioural repair decisions into the model leads to roughly a seven-fold increase in the percentage changes of the credit-risk metrics (PD, LGD, EAD, RWA, and EL) compared with the case in which the no-repair assumption is retained. The primary driver of this considerable increase is the rise in exposure (i.e. EAD), which reflects an expansion of the existing mortgage portfolio rather than an increase in the underlying riskiness of that portfolio. When the results were examined from another perspective—focusing purely on the risk profile while holding the exposure constant—there was still a clear deterioration of the portfolio's risk profile. On average, the no-repair baseline underestimates the RW by approximately 85% for flood risk (both physical and perception) and by around 60% for foundation risk on a risk-per-euro basis. In other words, in some cases the portfolio's risk profile nearly doubles relative to the current impact estimation, which reflects a pure increase in risk and is therefore taken as the principal result of this research. These findings indicate that the modelling assumption made in the current framework does not merely simplify the setup, but also potentially limits the accuracy of the current climate estimation.

Second, on the possibility of improving the current framework: this research has successfully demonstrated that such an improvement is both feasible and practically implementable. A new model extension was developed, introducing behavioural dependence in a way that reflects the reality of individual financial behaviour. This behavioural approach acknowledges that mortgage holders may respond differently when faced with climate-related damage, and that their creditworthiness — and therefore their credit risk — is shaped by these decisions.

While the extended framework leads to clearly different and overall higher climate-impact results, it is also important to consider the magnitude of these changes in absolute terms. Even under worst-case behavioural and climate scenarios, the additional RWA and EL resulting from climate-adjusted LGD, PD and EAD remain below materiality thresholds. As such, there is no immediate requirement to change existing internal models based on these outcomes. However, the purpose of this research has not been to demonstrate material climate risk impacts on the Dutch mortgage portfolio. Rather, the objective was to explore the development of a framework that offers a more robust approach than the current setup. Additionally, this lack of materiality should not be seen as a reason for inaction. In fact, a likely driver behind this non-materiality outcome, particularly for the Netherlands, is the relatively low probability of extreme climate events occurring here. In regions with higher exposure to physical climate risk — or if these probabilities increase in the future in the Netherlands — the importance of a robust modelling framework will become much more pronounced. For that reason, and in line with the expectations of European financial regulators, it remains important to invest in the development of such frameworks. This thesis aimed to contribute to that effort by proposing and demonstrating a first step towards behavioural-based climate risk modelling in the context of retail mortgage credit risk.

A

Flood risk background

Flood risk is commonly defined as a function of two key components: the probability of a flood event occurring and the impact it would generate if realized. While flood probability depends on geographic and climate factors, the impact is determined by the extent of economic and financial damages inflicted on affected areas. The interplay between these factors makes flood risk assessment an important component of disaster preparedness, urban planning and financial stability analysis.

Given the Netherlands' unique vulnerability to flooding—particularly due to its below-sea-level geography and the exclusion of flood damage from standard property insurance policies—the potential implications for financial institutions are considerable. Although the Dutch government actively invests in flood defense and mitigation measures [56], the risk of extreme flood events remains relevant. Existing studies have attempted to quantify how extreme flood events could impact bank capital positions, offering valuable insights into the intersection of climate risk and financial stability.

Understanding flood risk at a granular level is essential for assessing its potential consequences for property values, credit markets and financial institutions. The following section outlines a state-of-the-art methodology used to model flood exposure, estimate financial losses and evaluate their impact on key banking metrics in the Netherlands. A specific study will be used as a reference point, the discussion paper *Floods and Financial Stability: Scenario-based Evidence from Below Sea Level* by Francesco G. Caloia, Kees van Ginkel, and David-Jan Jansen (2023) [35], which examines the potential financial stability risks posed by floods in the Netherlands. This paper was conducted by a joint research group affiliated with Erasmus University Rotterdam, the University of Amsterdam (UvA), and Vrije Universiteit Amsterdam (VU). Additionally, the study was later published by De Nederlandsche Bank (DNB) as a working paper, aiming to provide insights relevant to financial supervision and policy development.

The study employs a scenario-based methodology to assess how flood-related property damages could impact bank capital through increased credit risk. It builds upon findings from major studies on flood scenario modeling, flood damage estimation and financial exposure assessment, integrating insights from both hydrological and economic research. The relevance of this study is underscored by its large-scale dataset, covering approximately EUR 650 billion in real estate exposures—referring to outstanding loan amounts—across more than three million residential properties.

Scenario-based Approach

As can be inferred from its title, this paper operates based on a set of stress scenarios. In the context of flooding in the Netherlands, it is natural to define these stress scenarios in terms of dike breaches, as failures in the flood defense system represent the primary mechanism through which large-scale inundations occur.

This approach is also adopted in this study, where a total of 38 flood scenarios are analyzed. Of these, 32 scenarios represent single-breach flood events, while the remaining 6 correspond to extreme multi-breach scenarios, as identified in the research by Dutch flood experts [57]. The single-breach scenarios are based on the *Landelijk Informatiesysteem Water en Overstromingen* (LIWO) [34], an open-source system that models thousands of potential flood events across the country. From this dataset, the study selects scenarios with the highest expected economic impact, specifically those where estimated property damages exceed EUR 500 million. In contrast, the extreme multi-breach scenarios originate from an expert study conducted in 2007, which examines the consequences of simultaneous dike failures across multiple regions.

This section provides a detailed overview of the methodology employed in the paper. First, the dataset used in the study is introduced, outlining the sources and characteristics of the financial and geographic data. Next, the study's flood damage estimation process is examined, explaining how flood scenarios were selected and how property damages were calculated. Finally, the methodology used to quantify the impact on credit risk modeling and bank capital adequacy is discussed, focusing on how flood-induced property devaluations influence credit risk metrics, such as loss-given-default (LGD) and probability of default (PD), ultimately impacting banks' capital positions.

This methodological overview is particularly relevant for this thesis, as it informs the questions concerning how flood risk is currently quantified and incorporated into credit risk models, while also illustrating a set of assumptions made in this process. By reviewing an established scenario-based approach, this section provides a publicly available reference point for identifying potential directions for model refinement.

Data

The methodology in this study is based on a combination of three key data sources, each contributing to different aspects of the flood risk assessment and its financial implications. By integrating granular loan-level data, administrative property data, and regulatory bank disclosures, the study provides a comprehensive foundation for analyzing the credit risk impact of flood-related property devaluations.

Loan-level data. A core component of the study is the use of loan-level data, which provides detailed information on mortgage and commercial real estate exposures of Dutch banks. These datasets, generally sourced from financial institutions, contain specific information on individual loan contracts, such as the outstanding loan amount, repayment structures, and borrower-specific characteristics (e.g. income and credit history). Additionally, this part of the dataset captures loan developments, allowing for the tracking of trends in loan-to-value (LTV) ratios and other credit risk parameters. This data is crucial for assessing how flood-induced property devaluations affect mortgage portfolios.

Property Microdata. To estimate the impact of flooding on property values, the study incorporates administrative microdata from *Statistics Netherlands (CBS)*. This dataset contains essential property-level information, including the official property valuation (WOZ value), which is determined by municipalities for taxation purposes. Additionally, it includes structural characteristics of properties, such as floor area (m^2), geographic location at the postal-code level, and classification into residential or commercial real estate. By integrating this data with flood scenario modeling, the study can estimate the extent of property damage in different locations and assess how such damages translate into changes in collateral values.

Supervisory Bank-specific data. To quantify the financial implications of flood risk, the study utilizes supervisory data reported by Dutch banks under the Common Reporting (COREP) and Financial Reporting (FINREP) frameworks. These regulatory filings provide critical insights into the financial health of banks, including balance sheet compositions, capital adequacy metrics, and asset quality assessments. Specifically, this dataset includes information on banks' total loan exposures, profitability measures related to mortgage lending and the level of capital held against potential credit losses. By linking this data with loan-level and property valuation data, the study evaluates how flood-related shocks could affect key banking stability indicators.

Data Category		
Loan-Level Data	**Property Microdata**	**Supervisory Bank-Specific Data**
Loan contract details (amount, outstanding debt)	WOZ property valuation (municipal assessment)	Bank balance sheets (total loan exposures)
Loan type (fixed vs. variable rate)	Structural characteristics (floor area, number of floors)	Profitability metrics (loan performance, earnings)
Insurance coverage information	Geographic location (postal code level)	Capital adequacy (CET1 ratios, buffers)
Borrower characteristics (income, credit history)	Property classification (residential vs. commercial)	Asset quality assessments (risk-weighted exposures)
Quarterly loan trend data		

Table A.1: Overview of data sources used in the study, categorized into three main types, with corresponding examples of data entries.

Table A.1 provides a summarized overview of the data sources used in the study, grouped into categories, with examples of relevant data entries included for each. By combining three complementary data sets, the study enables a detailed assessment of how extreme flood events could translate into financial vulnerabilities for the banking sector. While the paper does not specify the technical integration procedure in detail, it outlines how each source contributes to distinct steps in the modeling chain: property damage estimation, credit risk parameter calculation and capital adequacy analysis. The granularity of the data allows for scenario-based stress testing, providing a forward-looking analysis of the potential credit risk implications of climate-related flooding.

Flood Damage Methodology

As previously mentioned, the methodology by Caloia et al. [35] is based on stress-testing flood scenarios, considering floods originating from either the sea or major rivers, specifically in areas currently protected by flood defenses. The methodology relies on two primary scenario types, each differing in severity and source.

The first type consists of 32 single-breach flood scenarios, in which localized failures in flood defense systems result in the inundation of specific areas. These scenarios were obtained from the Landelijk Informatiepunt Water en Overstromingen (LIWO), an open-source national database containing over 5,000 flood scenarios. To select the most relevant scenarios, two criteria were applied. First, the study includes only areas vulnerable to breaches in primary flood defense systems, classified as 'type B' floods. In other words this refers to flooding in areas where the water system is classified as "main" (e.g. major rivers such as the Rhine and Meuse) and where flood protection is present. Despite the presence of flood defenses, these areas remain at risk due to extreme weather events or failures in the protection system. Notably, property damages resulting from this flood type are typically excluded from coverage by standard insurance policies. Second, within each region, the scenario leading to the highest estimated property damage—exceeding a threshold of EUR 500 million—was selected. This approach ensures a focus on tail risks, highlighting the worst-case impacts on financial stability rather than the average expected damages.

Additionally, as second type the study incorporates six extreme multiple-breach flood scenarios. These scenarios, developed by Dutch flood experts in 2007, depict instances where multiple dike breaches occur simultaneously, representing highly unlikely but still conceivable extreme flood events. The objective of including these cases is to assess the potential implications of severe flooding on the financial sector under worst-case conditions.

To provide a visual perspective, Figures A.1 and A.2 illustrate the extent of flooding in the case of different scenarios.

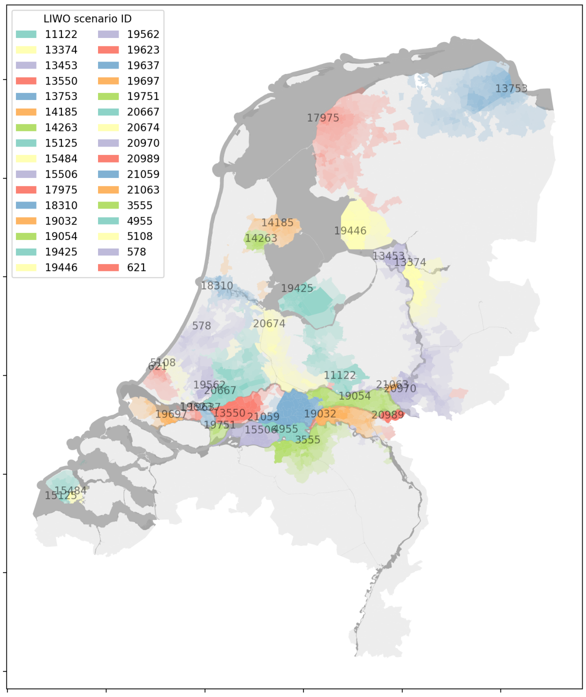

Figure A.1: Scenario set for single-breach floods.

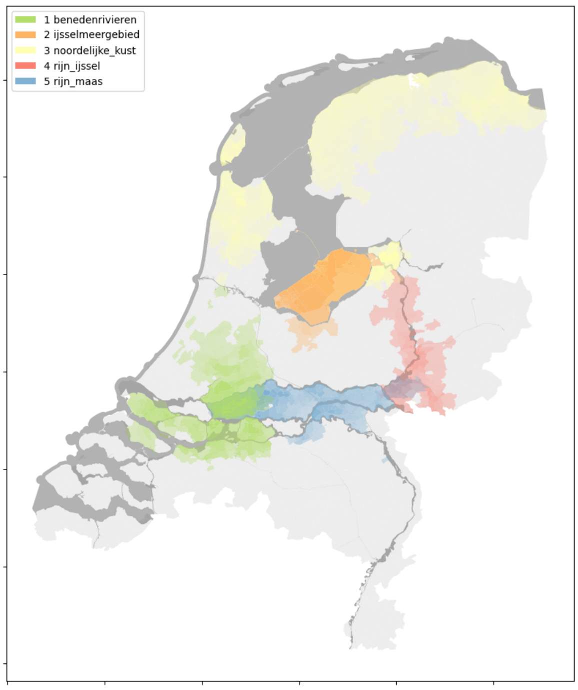

(a) Extreme scenarios 1 to 5

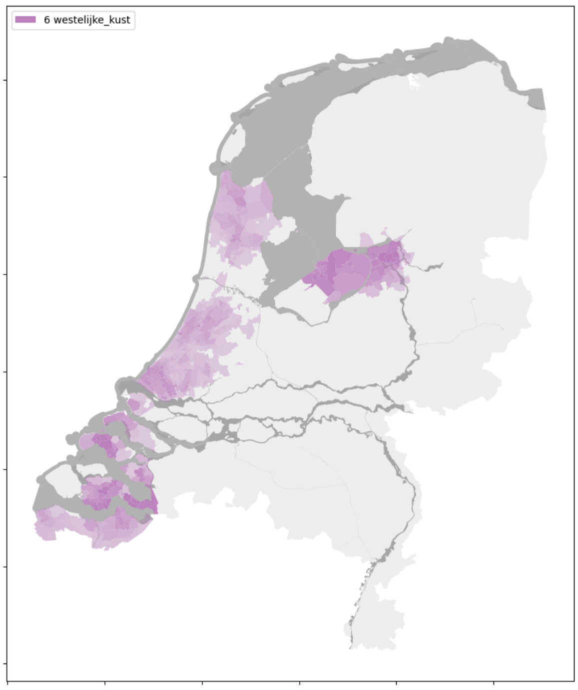

(b) Extreme scenario 6

Figure A.2: Scenario set for extreme multiple-breach floods. [35]

A crucial aspect of the impact assessment is the estimation of flood depths in the classified affected areas.

The LIWO system provides data at a high spatial resolution, offering flood depth estimates at a minimum scale of 100×100 meters. However, since the financial data used in this paper is only available at the level of four-digit postal codes, the flood depth data had to be aggregated accordingly, despite originally being available at a much finer resolution of 100×100 meters. Therefore, this study uses a mean water depth per postal code to align the flood impact data with financial exposure data. While this aggregation allows for compatibility between both datasets, it introduces a limitation.

To estimate the inundation depth at the postal-code level, the study employs a multi-step aggregation method. First, it isolates built-up areas within each postal-code zone, excluding land use types such as agricultural fields and infrastructure, which are not directly relevant for property damage estimation. The mean water depth is then computed over the remaining built-up area. Notably, locations within a postal-code region that remain dry during a flood event are assigned a depth of zero, ensuring that the computed average depth reflects the overall conditions affecting properties within the area. While this simplification is necessary for linking flood data with financial exposures, it introduces another limitation: localized variations in water depth within a postal-code area are not captured, potentially leading to an overestimation or underestimation of damages in certain sub-regions.

A key step in quantifying the financial impact of flood risk on retail exposures involves establishing a parameter that links flood risk to collateral value. Part of this parameter, is a flood damage parameter adopted from a previous national study, specifically the most recent *Standard Method for Calculating Flood Damage in The Netherlands* (SSM2017) (Slager and Wagenaar, 2017 [58]). This methodology estimates the maximum possible damage to a residential property, expressed per square meter, depending on the type of property. A distinction is made between structural damage to the building itself and damage to the household contents. Since this study focuses solely on the collateral value of properties, only the structural damage component to the building itself is considered. This parameter is denoted as '$max\ damage_t$' in the following formulas.

To quantify the flood-induced reduction in collateral value, this study defines a parameter ϕ_p^S, representing the fraction of a property's collateral value lost due to flood damage under scenario S for property p. This is given by

$$\phi_p^S = \min\left(\frac{damage_p^S}{property\ value_p}, 1\right), \tag{A.1}$$

where $property\ value_p$ represents the observed collateral value of the property in the loan-level dataset, i.e. its current market value. This upper bound of 1 makes sense, as the reduction in a property's collateral value due to flood damage cannot exceed its total value. Within the previous formula, the total flood-induced damage to property p in scenario S is computed as

$$damage_p^S = \theta(h)_t^S \cdot max\ damage_t \cdot A_p \cdot \iota, \tag{A.2}$$

where the parameters can be described as follows

- $\theta(h)_t^S$ is the *damage function*, a value between 0 and 1 that determines the fraction of the maximum possible damage as a function of the inundation depth h. Different property types have distinct damage curves, originally specified by Slager and Wagenaar (2017). The function $\theta(h)_t^S$ is scenario-specific and property-type-specific, meaning that for each flood scenario S and property type t, it assigns a damage fraction based on the floodwater depth at that location. Due to data limitations, this study applies a small adjustment specifically for apartments by using a weighted average of the damage functions for ground-floor and first-floor apartments, rather than directly adopting the original values from SSM2017. To illustrate the damage functions applied in this study, Figure A.3 presents the relationship between inundation depth and the fraction of maximum possible damage for residential real estate. The figure shows that the damage fraction increases non-linearly with water depth, with apartments exhibiting a steeper increase compared to single-family homes.

- *max damage*$_t$ represents the *maximum structural damage* (in euros per m^2, in 2011 prices) that a given property type t can sustain in the event of flooding. As previously mentioned, this value is derived from SSM2017, considering only the component related to direct damages to the building itself.
- A_p denotes the *floor area* (in m^2) of property p, sourced from administrative microdata provided by Statistics Netherlands (CBS).
- ι is an *inflation correction factor* that adjusts the estimated damage values from their original reference year (2011) to align with 2020 price levels, ensuring consistency with the loan-level data.

The computed parameter $\theta(h)_t^S$ is crucial as it directly links to the estimation of loan-to-value (LTV) ratios under flood scenarios, forming the basis for subsequent credit risk modeling. Notably, the values of ϕ_p^S are restricted to the interval $[0, 1)$ to prevent cases where flood-induced losses would exceed the entire collateral value of a property.

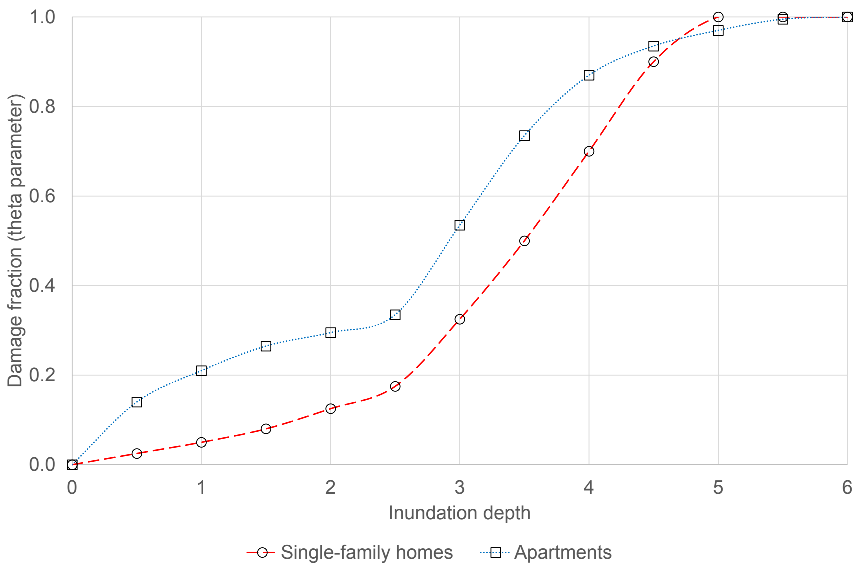

Figure A.3: Relationship Between Inundation Depth and Damage Fraction for Residential Properties. [35]

Example A.0.1. *To illustrate the application of the damage calculation methodology, consider a hypothetical flood scenario in Rotterdam. The damage estimation framework is applied to a specific property affected by the flood, using the relevant parameters and equations outlined in the methodology.*

- *Scenario (S): A single-breach flood scenario in Rotterdam, indexed as $S = 10$, corresponding to for example LIWO single-breach scenario ID 19637 (see Figure A.1).*
- *Property (p): Within this flooded region, let a specific property be selected as property number 42, located in a residential area of Rotterdam. The parameter $p = 42$ serves as the property index, uniquely identifying this property.*
- *Type of property (t): Property number 42 is classified as a single-family home. The parameter t represents the property type, in this case, "single-family home".*
- *Inundation depth (h): The inundation depth at property number 42 corresponds to the water level at that specific location. This data is derived from flood simulations (SSM2017) and presented in water*

depth maps. Assume that the water depth at property index 42 is $h = 3.25$ *meters.*

- *Damage factor ($\theta(h)_t^S$): The damage factor depends on the property type (t) and inundation depth (h). Suppose a reference table indicates that for a single-family home (t) at a water depth of 3.25 meters (h), the corresponding damage factor is* $\theta(h) = 0.4$.
- *Maximum damage (max $damage_t$): The maximum possible damage per square meter also depends on the property type (t). According to Slager and Wagenaar (2017), the maximum structural damage for a single-family home is €2,500 per m^2 (in 2011 price levels).*
- *Property area (A_p): The total floor area of property number 42, the single-family home, is* $A_p = 120$ m^2.
- *Price level correction factor (ι): To adjust for inflation and align damage estimates with 2020 price levels, a correction factor of* $\iota = 1.15$ *is applied.*

Using these parameters, the estimated flood damage for property number 42 ($p = 42$) is calculated as:

$$
\begin{aligned}
damage_{42}^{10} &= 0.4 \cdot 2500 \text{ €}/m^2 \cdot 120 \ m^2 \cdot 1.15 \\
&= 138,000\text{€}
\end{aligned}
$$

Thus, the estimated damage for this specific scenario ($S = 10$), property ($p = 42$), being a single-family home (t = single-family home), in this flood scenario equals €138,000.

Using this, the flood-induced decline in the collateral value can be calculated. In this case, the estimated damage is $damage_p^S = 138,000$. *Let the property value be assumed to be €600,000. Substituting these values gives* $\phi_p^S = \min\left(\frac{138,000}{600,000}, 1\right) = 0.23$. *This implies that the collateral value of property index 42 has decreased by 23% as a result of the flood event in scenario* $S = 10$.

The methodology outlined above aligns with prior flood damage modeling approaches but introduces key modifications tailored to the financial sector. Unlike standard SSM2017 applications, which estimate damages for all properties in an affected area, this method focuses exclusively on properties serving as collateral for bank loans. Additionally, the use of mean inundation depths per four-digit postal-code areas—rather than high-resolution grid-based flood data—reflects the need to align the geographical scale of flood data with that of bank exposure datasets. Although the paper does not provide full transparency on the technical integration between datasets, it implicitly raises important questions about the assumptions underlying this step. Understanding the nature and justification of these assumptions can be important for evaluating the robustness of the resulting risk estimates and is directly relevant to the broader research question of how flood risk is operationalized within credit risk models.

Credit Risk Impact Methodology

The methodology by Caloia et al. [35] further quantifies flood risks specifically in relation to banking metrics, utilizing key credit risk parameters. In the flood damage methodology, a connection is established between collateral value and flood risk, which in the context of credit risk translates into effects on the credit risk parameters Loss Given Default (LGD) and Probability of Default (PD), as discussed in Section 2.1.3. Additionally, a link is made to another key measure in financial supervision, the CET1 ratio, as described in (2.8). The analysis is conducted using end-2020 data as the starting point, with a one-year horizon.

1. LGD Impact

A key component in assessing the credit risk impact of flood-induced damages is the Loss Given Default (LGD). Flood-related property devaluations increase the Loan-to-Value (LTV) ratio, which in turn raises

LGD estimates. The study defined a new parameter LTV_i^S, the *LTV of a loan i under flood scenario S* as follows

$$LTV_i^S = LTV_i^0 \cdot \frac{1}{1 - \phi_p^S} \tag{A.3}$$

where LTV_i^0 is the starting-point LTV, and ϕ_p^S denotes the flood-induced decline in collateral value, as previously defined (A.1).

To determine whether a bank can recover its outstanding loan exposure through collateral liquidation, the metric *Loss Given Loss (LGL)* is introduced. The LGL measures the fraction of the exposure that remains uncovered after selling the property in a distressed sale scenario. It is defined as

$$LGL_i = \max\left(0, \frac{exposure - liquidation\ value}{exposure}\right) \tag{A.4}$$

where the $liquidation\ value$ is the estimated value at which the property can be sold post-damage. The LGL under flood scenario S is further refined as

$$LGL_i^S = \max\left(0, \frac{LTV_i^S - sales\ ratio_p^S}{LTV_i^S}\right) \tag{A.5}$$

Here, the sales ratio denotes the ratio between the liquidation value under flood-scenario S and the current market value of the property. The market value of a property is the price it would fetch under normal conditions. The liquidation value, on the other hand, is the price actually obtained in a forced sale, which is often lower than the market value due to the urgency of the sale. In the paper, it is assumed that the calculation of liquidation value additionally takes into account the costs required to prepare the property for sale due to flood damage. In other words, it can be calculated as follows

$$sales\ ratio_p^S = sales\ ratio_p^0 \cdot (1 - \phi_p^S) = \frac{\text{pre-flood liquidation value}}{\text{current value of the property}} \cdot (1 - \phi_p^S). \tag{A.6}$$

Using the computed LGL, the LGD of loan i in flood scenario S is given by

$$LGD_i^S = (1 - probability\ of\ cure) \cdot LGL_i^S + costs \tag{A.7}$$

where:

- $1 - probability\ of\ cure$ denotes the fraction of loans that remain in default after restructuring and management of arrears.
- $Costs$ refer to the administrative expenses incurred by the bank when selling the collateral property, which represent a fraction of the current exposure.

To ultimately link the scenario-specific LGD to the actual LGD, the final step introduces the *scenario-specific LGD multiplier* m_{LGD}^S. This multiplier reflects how the overall LGD changes under a given flood scenario S, accounting for differences in exposure across banks and loan types. The multiplier is computed as

$$m_{LGD}^S = \sum_b \sum_i w_b w_i \frac{LGD_{i,b}^S}{LGD_{i,b}}, \tag{A.8}$$

where $\frac{LGD_{i,b}^S}{LGD_{i,b}}$ represents the relative change in LGD for loan i at bank b under flood scenario S. The weights w_b and w_i play an important role in this calculation, as the objective of the paper is to estimate the overall impact of floods on Dutch bank capital. The *bank weight* w_b represents how much exposure a particular bank has relative to the total exposure across all banks in the system. Larger banks with more outstanding loans will naturally have a greater weight, as they hold a larger share of the total credit risk. Because this paper focuses on systemic risk, it uses a sample of eight major Dutch banks, covering the vast majority of the Dutch mortgage market. The *loan weight* w_i accounts for the relative size of an individual loan within a bank's total exposure. This means that larger loans, or loans with a higher outstanding balance, contribute more to the final LGD calculation than smaller ones.

By combining these weights, the LGD multiplier captures how each bank is affected by the flood scenario in proportion to its exposure. A bank with significant lending in flood-prone areas will experience a stronger impact than one with limited exposure to these regions. To provide further intuition, the following presents a hypothetical example illustrating the impact of LGD, using the collateral value parameter $\phi_p^S = 0.23$ from the previous example.

Suppose that the initial loan characteristics and flood-induced impact are defined as follows

- Loan-to-Value ratio (LTV_i^0): The initial Loan-to-Value ratio is 60%, i.e. 0.6.
- Flood-induced decline in collateral value (ϕ_p^S): The reduction in property value due to flooding is assumed to be 0.23.
- Exposure: The outstanding loan balance for the given property is €250,000.
- Initial sales ratio (Sales ratio$_p^0$): Before the flood event, the sales ratio is 0.90.
- Probability of cure: The probability that a loan recovers after initial delinquency is assumed to be 0.15.
- Costs: Administrative costs incurred by the bank in the event of forced liquidation amount to €3,000. Relative to the exposure, this equals 0.012.

Using the just defined LGD methodology, the following values are obtained for loan i after the flood

1. Post-flood Loan-to-Value ratio: $LTV_i^S = 0.6 \cdot \frac{1}{1-0.23} = 0.779$
2. Post-flood sales ratio: $sales\ ratio_p^S = sales\ ratio_p^0 \cdot (1 - \phi_p^S) = 0.90 \cdot (1 - 0.23) = 0.693$
3. Post-flood Loss-Given-Loss: $LGL_i^S = \max\left(0, \frac{LTV_i^S - sales\ ratio_p^S}{LTV_i^S}\right) = \max\left(0, \frac{0.779-0.693}{0.779}\right) = 0.086$
4. Loss-Given-Default: $LGD_i^S = (1 - probability\ of\ cure) \cdot (LGL_i^S + costs) = (1 - 0.15) \cdot 0.086 + 0.012 = 0.0851$

Thus, the LGD for this individual loan i under flood scenario S is 0.0851. To compute the LGD multiplier, the initial LGD for loan i before the flood is required. Assume that $LGD_i = 0.04$. Now, extending the analysis to the banking system. Suppose a specific bank b is analyzed with the following characteristics

- w_b (Bank-level exposure share): 0.10 (this bank holds 10% of the total outstanding mortgage loans in the system)
- The bank has 1,000 loans in the affected postal code. For simplicity, assume that all loans have the same characteristics as loan i.

The entire financial system consists of multiple banks (assumed to be 8 in total). To simplify the example, it is assumed that none of the other banks have loan exposure in the affected postal code, meaning their LGD remains unchanged at the initial value. The scenario-specific LGD multiplier can be calculated as follows

$$m_{LGD}^S = \frac{\sum_b \sum_i w_b w_i LGD_{i,b}^S}{\sum_b \sum_i w_b w_i LGD_{i,b}} = \frac{(0.10 \cdot 1 \cdot 0.0851) + (0.90 \cdot 1 \cdot 0.04)}{(0.10 \cdot 1 \cdot 0.04) + (0.90 \cdot 1 \cdot 0.04)} = \frac{0.00851 + 0.036}{0.004 + 0.036} = \frac{0.04451}{0.04} = 1.113.$$

Thus, the scenario-specific LGD multiplier is 1.113. This implies that, on average across the financial system, the LGD has increased by 11.3% as a result of the flood scenario S.

Note that this example is intended purely for illustrative purposes, to provide an intuitive understanding of how flood scenarios might impact the LGD metric, and no particular significance should be attached to the hypothetical values used.

2. PD Impact

Following the discussion on how flood damage can reduce collateral value and consequently affect the Loss Given Default (LGD), attention should also be given to another key dimension of credit risk: the Probability of Default (PD). At first glance, PD might appear less directly relevant in the context of physical damage to real estate. After all, the primary impact of a flood is a reduction in collateral value, which immediately increases the potential losses in case of default (LGD). However, empirical research demonstrates that the likelihood of a borrower defaulting on their loan (PD) is often highly correlated with the Loan-to-Value (LTV) ratio [35]. This means that an increase in the LTV ratio caused by flood-related property damage not only amplifies the potential losses upon default but can also increase the probability of that default occurring. For this reason, the paper extends its analysis by incorporating a statistical model to estimate how changes in the LTV ratio, resulting from flood-induced real estate damage, influence the probability of borrowers defaulting on their loans.

The probability that a borrower defaults, conditional on the different flood scenarios, is estimated using the following regression model

$$y_{i,b,t} = c_b + \beta' \mathbf{Z}_{i,b,t} + \delta' \mathbf{X}_{i,b,t} + u_{i,b,t}, \tag{A.9}$$

where $y_{i,b,t}$ is the dependent variable, representing the default status of borrower i at bank b at time t. This is a binary variable that takes the value 1 in case of default and 0 otherwise.

The term c_b is the constant, representing the baseline default probability when all other variables in the model are set to zero. In practice, this serves as an intercept that captures the average probability of default, incorporating all unobserved influences not explicitly included in the model.

The term $\beta' \mathbf{Z}_{i,b,t}$ represents the key independent variables affecting the probability of default, where

- $\mathbf{Z}_{i,b,t}$ is a vector of primary independent variables that are expected to have a direct effect on the probability of default under flood scenarios. In this study, these include the Loan-to-Value (LTV) ratio, the mortgage interest rate and regional GDP growth. The LTV ratio plays a central role as it is directly impacted by flood-induced property devaluation.
- β' is a vector of coefficients that quantify the magnitude and direction of the effect of each independent variable in $Z_{i,b,t}$ on the probability of default $y_{i,b,t}$. A positive coefficient indicates that an increase in the corresponding independent variable raises the likelihood of default, while a negative coefficient implies the opposite.

The term $\delta' \mathbf{X}_{i,b,t}$ accounts for so-called control variables, including

- $\mathbf{X}_{i,b,t}$ being a vector of additional variables that may influence the probability of default but are not the primary focus of this study. Including these variables allows for isolating the direct impact of flood-induced changes in the LTV ratio. Control variables include mortgage type, interest type, remaining loan term, initial LTV, property type and whether the mortgage benefits from the National Mortgage Guarantee (NHG) [1].

[1] The National Mortgage Guarantee (NHG, or Nationale Hypotheek Garantie) is a Dutch government-backed scheme that provides a safety net to both lenders and borrowers in case of payment difficulties due to circumstances beyond the borrower's

- δ' is a vector of coefficients that capture the effect of the control variables in $X_{i,b,t}$ on the probability of default.

Finally, $u_{i,b,t}$ is the error term, representing unobserved factors and random variation that affect the probability of default but are not explicitly included in the model.

Based on the estimated probability of default (A.9), a PD multiplier is derived to quantify the relative increase in default probabilities under flood scenario S. The system-wide PD multiplier for scenario S is given by

$$m_{PD}^{S} = \frac{\sum_b \sum_i w_b w_i \mathbb{E}(y_{i,b,T} | \mathbf{Z}_{i,b,T}^{S}, \mathbf{X}_{i,b,T})}{\sum_b \sum_i w_b w_i \mathbb{E}(y_{i,b,T} | \mathbf{Z}_{i,b,T}, \mathbf{X}_{i,b,T})} \tag{A.10}$$

where $y_{i,b,T}$ represents the default status of borrower i at bank b at the final reporting period T. Since it is a binary variable (1 for default, 0 otherwise), its expected value $\mathbb{E}(y_{i,b,T}|...)$ corresponds to the probability of default at a specified moment T and under scenario S.

$\mathbf{Z}_{i,b,T}^{S}$ is the vector of independent variables under the flood scenario S. The key difference compared to the baseline vector $\mathbf{Z}_{i,b,T}$ is that the LTV ratio has been adjusted to reflect flood-induced property devaluation.

As in the LGD case, it is again the case that w_b represents the bank-level weight, which reflects each bank's exposure as a share of total system-wide exposure, and w_i represents the loan-level weight, which accounts for each individual loan's exposure relative to the total exposure of the corresponding bank.

The PD multiplier provides a system-wide measure of how flood-induced changes in collateral values propagate into increased probabilities of default. This multiplier also plays an important role in the final quantification of the financial impact as a result of flood scenarios, which will become clear in the next part.

3. CET1 ratio Impact

To assess how flood scenarios could impact the financial health of banks, it is essential to examine their regulatory capital. A key metric in this regard is the Common Equity Tier 1 (CET1) ratio, which as discussed earlier, measures a bank's core equity capital relative to its risk-weighted assets (RWA), recall equation (2.8). A lower CET1 ratio signals a weaker capital position, potentially increasing financial vulnerability.

The denominator of the CET1 ratio, RWA, essentially quantifies the riskiness of a bank's assets. The RWA is computed as the product of Exposure at Default (EAD), a constant regulatory factor of 12.5% and the capital requirement factor K. This aligns with the Basel regulatory framework, as discussed in Section 2.2. Recall that the capital requirement factor K depends on both bank-specific and scenario-specific values of Loss Given Default (LGD) and Probability of Default (PD).

As discussed, the occurrence of a flood can change the LGD and PD for exposed loans. To capture this, the concept of scenario-specific RWA (RWA^S) is introduced, representing the RWA that would result if a particular flood scenario S occurs. To quantify the relative change in RWA due to a flood scenario, a scenario-specific RWA multiplier (m_{RW}^{S}) is defined as follows

$$\begin{aligned} m_{RW}^{S} &= \frac{\sum_b \sum_i w_b w_i RWA_{i,b}^{S}}{\sum_b \sum_i w_b w_i RWA_{i,b}} \\ &= \frac{\sum_b \sum_i w_b w_i K_{i,b}^{S}}{\sum_b \sum_i w_b w_i K_{i,b}}, \end{aligned}$$

where $RWA_{i,b}^{S}$ denotes the RWA for loan i of bank b under flood scenario S, and $RWA_{i,b}$ represents the initial RWA. Similarly, $K_{i,b}^{S}$ and $K_{i,b}$ denote the scenario-specific and initial capital requirement factors,

control.

respectively. The weights w_b and w_i account for the relative importance of each bank and loan. Essentially, this multiplier quantifies the average relative change in RWA across all banks and their real estate loan portfolios under scenario S.

Since the capital requirement factor K depends on the LGD and PD, the previously calculated scenario-specific multipliers for LGD (m^S_{LGD}) and PD (m^S_{PD}) are the key drivers behind variations in K and, consequently, in RWA. More specifically, for residential mortgages, the scenario-specific capital requirement factor is given by

$$K^S = LGD \cdot m^S_{LGD} \cdot \Phi\left(\frac{\Phi^{-1}(PD \cdot m^S_{PD})}{\sqrt{1-\rho}} + \sqrt{\frac{\rho}{1-\rho}} \cdot \Phi^{-1}(0.999)\right) - PD \cdot m^S_{PD} \cdot LGD \cdot m^S_{LGD}, \quad \text{(A.11)}$$

where Φ is the cumulative standard normal distribution function, and ρ represents the asset correlation parameter as defined in the Basel framework. A higher LGD and/or PD under scenario S leads to a higher K^S, which in turn increases RWA^S, resulting in an RWA multiplier greater than one.

Finally, the impact of these scenario-induced RWA changes on the CET1 ratio is examined. The formula is given by:

$$\Delta CET1\ ratio^S = \frac{CET1}{RWA} - \frac{CET1 - \Delta EL^S}{RWA + \Delta RWA^S} \quad \text{(A.12)}$$

where ΔEL^S represents the difference between the expected loss under the flood scenario and the starting-point expected loss and ΔRWA^S denotes the scenario-specific change in risk-weighted assets. CET1 represents a bank's core capital, previously denoted as CET1-capital in Equation (2.8), but abbreviated as CET1 in this formula for conciseness. To better understand this formula, its derivation is as follows

$$\Delta CET1\ ratio^S = CET1\ ratio^0 - CET1 ratio^S = \frac{CET1}{RWA} - \frac{CET1^S}{RWA^S} \quad \text{(A.13)}$$

$$= \frac{CET1}{RWA} - \frac{CET1 - \Delta EL^S}{RWA \cdot m^S_{RW}} = \frac{CET1}{RWA} - \frac{CET1 - \Delta EL^S}{RWA + \Delta RWA^S}. \quad \text{(A.14)}$$

This derivation clarifies how the CET1 ratio is affected by both expected credit losses (ΔEL^S) and changes in risk-weighted assets (ΔRWA^S). The first term, $\frac{CET1}{RWA}$, represents the initial CET1 ratio in the absence of a flood scenario, while the second term incorporates the effects of flood scenario S. Specifically, a higher ΔEL^S leads to a reduction in CET1 capital, and a higher ΔRWA^S increases the denominator, amplifying the decline in the CET1 ratio. Alternatively, the final expression in the denominator can also be conveniently rewritten as $RWA^S = RWA \cdot m^S_{RW}$, utilizing the previously defined scenario-specific RWA multiplier.

In summary, the entire methodology essentially builds up to this final dependency, the flood-scenario-specific CET1 ratio (A.12). It begins with the construction of a parameter that represents the general impact of flooding on collateral values, ϕ^S_p (A.1). This fraction is then incorporated into the scenario-specific Loan-to-Value ratio, LTV^S_i (A.3). The adjusted loan-to-value subsequently influences both the LGD and PD modeling each in its own way, ultimately resulting in the scenario-specific multipliers m^S_{LGD} (A.8) and m^S_{PD} (A.10), respectively. In the final step these multipliers are used, in line with Basel regulatory formulas, to compute the capital requirement K^S (A.11), which feeds into the adjusted risk-weighted assets RWA^S, and finally determines the overall impact on the $CET1\ ratio^S$.

The objective of this section was to provide a detailed overview of how a Dutch state-of-the-art, publicly available flood risk integration methodology is structured. From the initial construction of the flood-induced collateral depreciation parameter to the final impact assessment on the CET1 capital requirements, it is clear that numerous assumptions are made throughout the entire process. Many of these assumptions could benefit from further investigation, ensuring that the overall modeling approach remains as precise, robust, and comprehensive as possible.

It is also important to acknowledge certain limitations of the methodology. First, the model does not incorporate potential mitigating factors, such as insurance payouts, government relief schemes, or other compensatory mechanisms that may reduce actual credit losses. Second, the approach is based on scenario analysis and does not attempt to estimate the probability of the flood event occurring. As a result, it does not quantify the full expected impact, but rather the conditional impact under the assumption that a specific flood scenario takes place.

B

Income growth parameters

To construct a realistic income forecast over the life cycle of Dutch mortgage holders, age-specific income statistics were used. Specifically, the model relies on the 2022 average gross annual salary per age group in the Netherlands, as reported by Statista based on data from the Centraal Bureau voor de Statistiek (CBS) [59].

Figure B.1 displays the average annual salary by age. A clear pattern emerges: income rises steeply during the early stages of one's career, continues to grow at a more modest pace during mid-career, and then gradually declines toward retirement. Based on this pattern, the following four income segments were defined, each associated with an empirically derived annual growth or decline factor.

Note that only data from the age of 20 onwards was used, as younger age groups are likely to include part-time student jobs or other non-representative income sources not reflective of individuals who are eligible for a mortgage.

- **Early-career** ($a < \chi_e = 35$)
  This segment spans ages 20 to 34. Income grows from an average of €18,470 (20–24) to €42,430 (30–34). The compound annual growth rate over 15 years is:
  $$\left(\frac{42{,}430}{18{,}470}\right)^{1/15} \approx 1.0570.$$
  This corresponds to an approximate yearly increase of **6%**, leading to the chosen parameter value: $r_e = \mathbf{1.06}$.

- **Mid-career** ($\chi_e \leq a < \chi_m = 55$)
  This segment spans ages 35 to 54. Income increases from €46,890 (35–39) to €52,090 (50–54). The calculated average growth rate over 20 years is:
  $$\left(\frac{52{,}090}{46{,}890}\right)^{1/20} \approx 1.0053.$$
  This corresponds to an approximate yearly increase of **1%**, leading to the chosen parameter value: $r_m = \mathbf{1.01}$.

- **Late-career** ($\chi_m \leq a < \chi_l = 65$)
  From age 55 to 64, income declines from €50,330 (55–59) to €47,350 (60–64). The computed decline over 10 years is:
  $$\left(\frac{47{,}350}{50{,}330}\right)^{1/10} \approx 0.9939.$$
  This corresponds to an approximate yearly decrease of **1%**, leading to the chosen parameter value: $r_l = \mathbf{0.99}$.

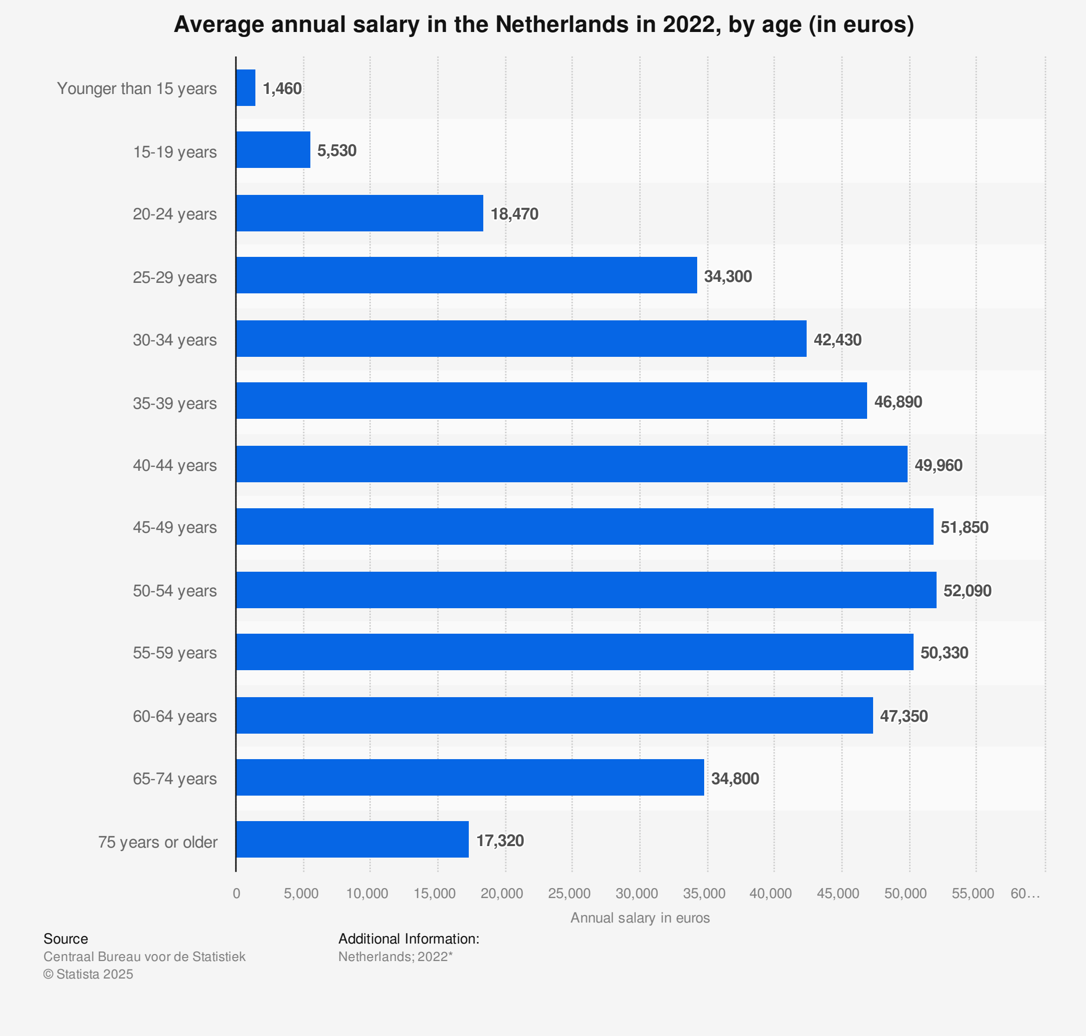

Figure B.1: Average gross annual salary in the Netherlands by age group (2022), based on data from Statista/CBS [59].

- **Post-career / Retirement** ($a \geq \chi_l = 65$)
  In retirement (ages 65–74), income drops from €47,350 to €34,800. The implied annual decline over 10 years is:
  $$\left(\frac{34{,}800}{47{,}350}\right)^{1/10} \approx 0.9697.$$
  This corresponds to an approximate yearly decrease of **3%**, leading to the chosen parameter value: $r_p = \mathbf{0.97}$.

It is important to emphasize that these income values represent the average gross salary from formal employment per age group. Pension benefits, such as AOW or employer pension schemes, are not included in these figures. Therefore, the reported income for individuals aged 75 or older — which still averages €17,320 — reflects a small subset of elderly individuals who continue to engage in paid work beyond retirement age.

C

Damage-dependent behavioural probabilities

Expert consultations highlighted that in practice the likelihood of households choosing between repair, loan financing or non-repair would primarily depend on the size of the incurred damage relative to key financial quantities such as property value and household income. Based on this feedback, an additional analysis was carried out in which behavioural probabilities were no longer treated as fixed scenario inputs, but were instead modelled as functions of damage-to-value and damage-to-income ratios. Acknowledging that several alternative specifications are possible, this section presents one possible approach. The results that follow provide insight into the type of shifts in credit risk outcomes that emerge once dynamic behavioural probabilities are introduced.

In this approach, the modelling of behavioural responses is approached as a compositional framework: four probabilities ($P_{\text{no-repair}}$, P_{self}, P_{AAB}, P_{ext}) are jointly defined such that they always sum to one. The design combines a set of boundary conditions with flexible functional forms that allow calibration to expert judgement and baseline assumptions. Below, the different steps of the proposed method are presented.

Step 1: Repair vs. No-Repair The first split is defined as a function of the damage-to-value ratio $r_{DV} = \frac{\text{damage}}{\text{property value}} \in [0, 1]$. A power function is applied:

$$P(\text{repair}) = \left(1 - \frac{r_{DV}}{\kappa}\right)_+^{\xi}, \qquad P(\text{no-repair}) = 1 - P(\text{repair}),$$

where $(x)_+ = \max(x, 0)$.

- This specification enforces the logical constraint that if damage equals property value (i.e. $r_{DV} \geq \kappa$), then $P(\text{repair}) = 0$.
- At $r_{DV} = 0$, the probability of repair equals 1.
- The parameter κ sets the cut-off point (e.g. $\kappa = 0.70$ ensures that when damage reaches 70% of collateral value, no-repair becomes certain).
- The parameter ξ calibrates the curvature, e.g. to match a baseline repair probability of 98% at a typical ratio r^*_{DV}.

Step 2: Allocation within Repair Conditional on repair, the probabilities are split into self-funding and loan-based repair as a function of the damage-to-income ratio $r_{DI} = \frac{\text{damage}}{\text{annual income}}$:

$$q_{\text{self}|\text{repair}} = 1 - g(r_{DI}), \qquad q_{\text{loan}|\text{repair}} = g(r_{DI}),$$

with

$$g(r_{DI}) = 1 - e^{-\zeta r_{DI}},$$

a monotone function that ensures $q_{\text{loan|repair}} = 0$ at $r_{DI} = 0$ and increases with the burden of damage relative to income. The slope parameter ζ is calibrated to match a targeted loan-share at a reference r^*_{DI}.

Within the loan option, a constant split is applied between ABN AMRO and external financing:

$$q_{\text{AAB|repair}} = \pi_{\text{AAB}} \cdot g(r_{DI}), \qquad q_{\text{ext|repair}} = (1 - \pi_{\text{AAB}}) \cdot g(r_{DI}),$$

where π_{AAB} represents the constant share of ABN AMRO among loan-based solutions (e.g. $\pi_{\text{AAB}} = 0.60$).

Final unconditional probabilities The unconditional probabilities follow as:

$$\begin{aligned}
P(\text{no-repair}) &= 1 - \left(1 - \tfrac{r_{DV}}{\kappa}\right)_+^{\xi}, \\
P(\text{self}) &= \left(1 - \tfrac{r_{DV}}{\kappa}\right)_+^{\xi} \left[1 - g(r_{DI})\right], \\
P(\text{AAB}) &= \left(1 - \tfrac{r_{DV}}{\kappa}\right)_+^{\xi} \pi_{\text{AAB}}\, g(r_{DI}), \\
P(\text{ext}) &= \left(1 - \tfrac{r_{DV}}{\kappa}\right)_+^{\xi} (1 - \pi_{\text{AAB}})\, g(r_{DI}).
\end{aligned}$$

Key properties

- All four probabilities sum to one by construction.
- Boundary conditions are respected: $P(\text{repair}) = 1$ at $r_{DV} = 0$, and $P(\text{repair}) = 0$ at $r_{DV} \geq \kappa$; if damage $= 0$, then $P(\text{self}) = 1$.
- The functional forms are parsimonious, interpretable, and allow calibration to expert judgement.
- Parameters $(\kappa, \xi, \zeta, \pi_{\text{AAB}})$ can be adjusted to test different scenarios or perform sensitivity analysis.

In Table C.1, a set of impact results is presented, obtained under the functional behavioural extension. For this sample, the parameters were set to a repair cut-off of $\kappa = 0.70$, a curvature of $\xi = 0.277$, a loan probability slope of $\zeta = 2.0$, and an ABN loan share of $\pi_{\text{AAB}} = 0.875$.

	Flood Physical	**Flood Perception**	**Foundation**	**Earthquake**
LGD % change (performing)	0.00222	0.00331	0.07388	0.00062
LGD % change (in-default)	0.00088	0.00137	0.04536	0.00008
PD % change	0.00437	0.00517	0.04182	0.00103
EAD % change	0.00536	0.00775	0.20971	0.00069
RWA % change	0.01012	0.01405	0.30926	0.00073
EL % change	0.00865	0.01203	0.27512	0.00047

Table C.1: **Behavioural extension with dynamic repair probabilities:** CER impact estimation results on LGD, PD, EAD, RWA and EL under defined behavioural parameters: cut-off for repair $\kappa = 0.70$, curvature $\xi = 0.277$, loan probability slope $\zeta = 2.0$, ABN loan share $\pi_{\text{AAB}} = 0.875$.

Table C.2 summarises the average behavioural probability distributions across the three datasets (EAD/PD, LGD in-default, and LGD performing), based on the same functional parameterisation of the repair and loan decision process.

The results indicate that the behavioural extension produces intuitive probability distributions across risk types and datasets. For low-impact risks such as earthquake, the distribution is heavily skewed towards self-repair (above 90%), with negligible no-repair or external loan probabilities. For more material risks such as foundation, the loan channel dominates, with ABN AMRO accounting for the majority of repair

	Flood Physical	Flood Perception	Foundation	Earthquake
EAD / PD dataset:				
Self-repair	0.63	0.56	0.23	0.96
ABN loan	0.31	0.37	0.62	0.03
External loan	0.04	0.05	0.09	0.00
No repair	0.02	0.02	0.06	0.00
LGD In-Default:				
Self-repair	0.59	0.51	0.17	0.95
ABN loan	0.35	0.41	0.67	0.04
External loan	0.05	0.06	0.10	0.01
No repair	0.02	0.02	0.06	0.00
LGD Performing:				
Self-repair	0.62	0.56	0.21	0.94
ABN loan	0.32	0.37	0.63	0.05
External loan	0.05	0.05	0.09	0.01
No repair	0.02	0.02	0.06	0.00

Table C.2: Average behavioural probability distributions per dataset (EAD/PD, LGD in-default, LGD performing), under functional specification of repair and loan decision probabilities.

financing. Flood-related risks fall in between, with balanced shares between self-repair and loans, and only a small fraction of no-repair outcomes. These patterns align with expectations: higher average damages (relative to income or property value) shift probability mass away from self-repair towards loan financing, while extreme non-repair remains rare except under high damage-to-value ratios.

$$\text{damage}_{i,\text{CER}} = V_i \times D_i^{\text{CER}} \times M_i$$

- V_i: initial collateral value
- D_i^{CER}: damage function under CER
- M_i: mitigation factor

$$\begin{aligned} LGD_{i,E(\text{CER})} = {} & LGD_{i,\text{non-CER}} \\ & + p_{i,\text{CER}} \left(LGD_{i,\text{CER}} - LGD_{i,\text{non-CER}} \right) \end{aligned}$$

$$\begin{aligned} LGD_{i,E(\text{CER})} = {} & LGD_{i,\text{non-CER}} \\ & + p_{i,\text{CER}} \Big(LGD_{i,\text{CER}} \\ & \quad - LGD_{i,\text{non-CER}} \Big) \end{aligned}$$

$$PD_{\text{CER},i}^{\text{behaviour}} = \lambda_{\text{CER},i}^{\text{behaviour}} \cdot PD_{\text{non-CER},i}$$

$$\begin{aligned} RWA_{\text{CER}} = {} & 12.5 \cdot LGD_{E(\text{CER})} \\ & \cdot \Phi\Bigg(\frac{\Phi^{-1}(PD)}{\sqrt{1-\rho}} \\ & \quad + \sqrt{\frac{\rho}{1-\rho}} \cdot \Phi^{-1}(0.999) \Bigg) \\ & \cdot EAD \\ & - PD \cdot LGD_{E(\text{CER})} \cdot EAD \end{aligned}$$

$$\begin{aligned} EL_{\text{CER}} = {} & PD \\ & \cdot LGD_{E(\text{CER})} \\ & \cdot EAD \end{aligned}$$

$$\begin{aligned} RWA_{\text{CER}} = {} & 12.5 \cdot LGD_{E(\text{CER})} \\ & \cdot \Phi\left(\frac{\Phi^{-1}(PD)}{\sqrt{1-\rho}} + \sqrt{\frac{\rho}{1-\rho}} \cdot \Phi^{-1}(0.999) \right) \\ & \cdot EAD - PD \cdot LGD_{E(\text{CER})} \cdot EAD \end{aligned}$$

$$EL_{\text{CER}} = PD \cdot LGD_{E(\text{CER})} \cdot EAD$$

$$\begin{aligned} LGD_{i,E(\text{CER})} = {} & LGD_{i,\text{non-CER}} \\ & + p_{i,\text{CER}} \left(\Delta LGD_{i,\text{CER}} \right) \end{aligned}$$

Bibliography

[1] W. J. W. Botzen, J. C. J. M. van den Bergh and L. M. Bouwer, *Climate change and increased risk for the insurance sector: a global perspective and an assessment for the Netherlands*. Springer Nature, (2009), https://doi.org/10.1007/s11069-009-9404-1.

[2] L. Jiang, T. Gerkema, D. Idier, A.B.A. Slangen and K. Soetaert, *Effects of sea-level rise on tides and sediment dynamics in a Dutch tidal bay*. Ocean Sci., (2020), https://doi.org/10.5194/os-16-307-2020.

[3] S. Hauswirth, *Hydrological Droughts in the Netherlands: from Simulations to Projections*. Utrecht Studies in Earth Sciences, volume 305, (2024), https://dspace.library.uu.nl/handle/1874/437427.

[4] Smith, Adam B. and Katz, Richard W., *U.S. Billion-dollar weather and climate disasters: Data sources, trends, accuracy and biases*. Natural Hazards, (2013), https://doi.org/10.1007/s11069-013-0566-5.

[5] Asian Development Bank, *Climate-Related Disasters in Asia and the Pacific*. Asian Development Bank, 2013.

[6] R. Vermeulen, E. Schets, M. Lohuis, B. Kölbl, D.-J. Jansen and W. Heeringa, *The heat is on: A framework for measuring financial stress under disruptive energy transition scenarios*. Ecological Economics, (2021), https://doi.org/10.1016/j.ecolecon.2021.107205.

[7] W. Salet, *Public Norms in Practices of Transitional Planning - The Case of Energy Transition in The Netherlands*. Sustainability, (2021), https://doi.org/10.3390/su13084454.

[8] European Banking Authority, *Final Guidelines on the Management of ESG Risks*. European Banking Authority, 2025.

[9] KPMG AG Wirtschaftsprüfungsgesellschaft, *Climate and Environmental Risks in Credit Risk Models: Market Survey 2024*. KPMG, 2024.

[10] Basel Committee on Banking Supervision, I, *International Convergence of Capital Measurement and Capital Standards*. Bank for International Settlements, 1988.

[11] Basel Committee on Banking Supervision, II, *International Convergence of Capital Measurement and Capital Standards: A Revised Framework*. Bank for International Settlements, 2004.

[12] Basel Committee on Banking Supervision, III, *Basel III: A Global Regulatory Framework for More Resilient Banks and Banking Systems*. Bank for International Settlements, 2011.

[13] Basel Committee on Banking Supervision, IV, *Basel III: Finalising post-crisis reforms*. Bank for International Settlements, 2017.

[14] European Central Bank, I, *Guide on climate-related and environmental risks: Supervisory expectations relating to risk management and disclosure*. European Central Bank, 2020.

[15] European Banking Authority (EBA), *Report on the Role of Environmental and Social Risks in the Prudential Framework*. European Banking Authority, 2023.

[16] McNeil, Alexander J. and Frey, Rudiger and Embrechts, Paul, *Quantitative Risk Management: Concepts, Techniques and Tools*. Princeton University Press, 2005.

[17] Bank for International Settlements, *IRB approach: risk weight functions*. Bank for International Settlements, 2023.

[18] Bank of England, *Modelling Credit Risk*. Bank of England, 2025.

[19] Basel Committee on Banking Supervision, *An Explanatory Note on the Basel II IRB Risk Weight Functions*. Bank for International Settlements, 2005.

[20] M.-S. Tsai and L.-C. Chen, *The calculation of capital requirement using Extreme Value Theory*. Economic Modelling, (2011), https://doi.org/10.1016/j.econmod.2010.08.010.

[21] European Central Bank, *ECB Banking Supervision Newsletter – May 2024*. European Central Bank, 2024.

[22] Basel Committee on Banking Supervision, II, *Basel II: International Convergence of Capital Measurement and Capital Standards: A Revised Framework - Comprehensive Version*. Bank for International Settlements, 2006.

[23] Bank for International Settlements, *History of the Basel Committee and Its Membership*. Bank for International Settlements, 2014.

[24] European Central Bank, *The Single Supervisory Mechanism (SSM) Explained*. European Central Bank, 2020.

[25] De Nederlandsche Bank (DNB), *Toezicht op financiële instellingen*. De Nederlandsche Bank, 2024.

[26] Bank of England, *Modelling Credit Risk*. Bank of England, 2015.

[27] Basel Committee on Banking Supervision, *Calculation of RWA for Credit Risk – Standardised Approach: Individual Exposures (CRE20)*. Bank for International Settlements, 2025.

[28] European Central Bank, *Supervisory Guide on Internal Models*. European Central Bank, 2024.

[29] BDO, *Final EBA Guidelines on ESG Risk Management – Implications for Governance and Risk Management*. BDO, 2025.

[30] O. Bin and S. Polasky, *Effects of flood hazards on property values: Evidence before and after Hurricane Floyd*. Land Economics, (2004), https://doi.org/10.2307/3655805.

[31] C. Kousky, *Learning from extreme events: Risk perceptions after the flood*. Land Economics, (2010), https://doi.org/10.3368/le.86.3.395.

[32] A. Atreya and S. Ferreira and W. Kriesel, *Forgetting the flood? An analysis of the flood risk discount over time*. Land Economics, (2013), https://doi.org/10.3368/le.89.4.577.

[33] A. Atreya and S. Ferreira, *Seeing is believing? Evidence from property prices in inundated areas*. Risk Analysis, (2015), https://doi.org/10.1111/risa.12307.

[34] Rijkswaterstaat, *Landelijk Informatiesysteem Water en Overstromingen (LIWO)*. Rijkswaterstaat, 2025.

[35] Francesco G. Caloia and Kees van Ginkel and David-Jan Jansen, *Floods and financial stability: Scenario-based evidence from below sea level*. Tinbergen Institute Discussion Paper, 2023.

[36] T. Endendijk and W. W. Botzen and H. de Moel and J. C. Aerts and S. J. Duijndam and K. Slager and B. Kolen and M. Kok, *Experience from the 2021 floods in the Netherlands: Household survey results on impacts and responses*. Journal of Coastal and Riverine Flood Risk, (2023), https://doi.org/10.59490/jcrfr.2023.0009.

[37] P. Eichholtz and N. Kok and P. Weenink, *Overstromingsrisico beïnvloedt woningwaarde, maar impact is beperkt*. ESB (Economisch-Statistische Berichten), 2025.

[38] A. Mutlu and D. Roy and T. Filatova, *Capitalized value of evolving flood risks discount and nature-based solution premiums on property prices*. Ecological Economics, (2023), https://doi.org/10.1016/j.ecolecon.2022.107682.

[39] A. Beltrán and D. J. Maddison and R. J. R. Elliott, *Is flood risk capitalised into property values?* Ecological Economics, (2018), https://doi.org/10.1016/j.ecolecon.2017.12.015.

[40] Steven Hommes and Sandra Phlippen and Jeannine van Reeken-van Wee and Christiaan Schreuder and Fonger Ypma, *Gemelde funderingsschade leidt tot forse prijskorting bij woningverkoop*. ESB (Economisch Statistische Berichten), 2023.

[41] Raad voor de leefomgeving en infrastructuur (Rli), *Goed gefundeerd: advies om te komen tot een nationale aanpak van funderingsproblematiek*. Raad voor de leefomgeving en infrastructuur, 2024.

[42] Kennis Centrum Aanpak Funderingsproblematiek (KCAF), *Impact van droogte op funderingen*. KCAF, 2020.

[43] P. M. Maurenbrecher and A. Den Outer and H. J. Luger, *Review of Geotechnical Investigations Resulting from the Roermond April 13, 1992 Earthquake*. Third International Conference on Recent Advances in Geotechnical Earthquake Engineering and Soil Dynamics, 1995.

[44] Koninklijk Nederlands Meteorologisch Instituut (KNMI), *Aardbeving Roermond 1992*. KNMI, 1992.

[45] berekenhet.nl, *Inflatie berekenen - bereken hoeveel geld minder waard is geworden*. berekenhet.nl, 2025.

[46] Kadaster, *Basisregistratie Adressen en Gebouwen (BAG)*. PDOK, 2024.

[47] Global Facility for Disaster Reduction and Recovery, *ThinkHazard! Netherlands - Earthquake Risk Report*. Global Facility for Disaster Reduction and Recovery, 2024.

[48] Global Earthquake Model Foundation, *GEM Global Hazard Map*. Global Earthquake Model Foundation, 2023.

[49] US Geological Survey, *Earthquake Hazards Program*. US Geological Survey, 2023.

[50] Carin A.B. van der Cruijsen, Jakob de Haan, David-Jan Jansen, and Robert H.J. Mosch, *Households' Decisions on Savings Accounts After Negative Experiences with Banks During the Financial Crisis*. The Journal of Consumer Affairs, (2016), https://doi.org/10.1111/joca.12106.

[51] De Nederlandsche Bank, *Borrower-Based Measures, House Prices and Household Debt*. De Nederlandsche Bank, 2022.

[52] Basel Committee on Banking Supervision, *Studies on the Validation of Internal Rating Systems*. Bank for International Settlements, 2005, Available online.

[53] European Banking Authority, *Discussion Paper on Machine Learning for IRB Models*. European Banking Authority, 2021.

[54] ABN AMRO Economisch Bureau, *ESG & Economie – Welke EU-landen zullen het meest te lijden hebben van extreme klimaatrampen?* ABN AMRO Economisch Bureau, 2024.

[55] European Central Bank, *Good practices for climate-related and environmental risk management: Observations from the 2022 thematic review*. European Central Bank, 2022.

[56] Klimaateffectatlas, *Plaatsgebonden overstromingskans*, 2025.

[57] Informatiepunt Leefomgeving (IPLO), *Crisismanagement water*, 2023.

[58] Kymo Slager and Dennis Wagenaar, *Standaardmethode 2017: Schade en slachtoffers als gevolg van overstromingen*. Deltares, 2017.

[59] Statista, *Average annual salary in the Netherlands in 2023, by age group*. Statista, 2024.