

UNIVERSITY OF TWENTE.

Faculty of Behavioural, Management and Social Sciences

Impact of Extreme Precipitation on the Dutch Social Housing Sector

Mart van den Hoek

M.Sc. Thesis

June 2026

Supervisors:
dr. B. Roorda
dr. R.A.M.G. Joosten
Finance Ideas
Faculty of Behavioural, Management and Social Sciences
Industrial Engineering & Management
Specialization of Financial Engineering & Management
University of Twente
P.O. Box 217
7500 AE Enschede
The Netherlands

Acknowledgements

My Master's programme in Industrial Engineering and Management at the University of Twente is concluded with my thesis. Looking back, I am grateful for the education and opportunities the university has offered me throughout my studies.

My research was conducted in collaboration with Finance Ideas in Utrecht. I would like to thank the organisation for proposing the topic, providing access to data, and sharing its knowledge of the social housing sector and its regulatory framework. In particular, I want to thank Yves Baljet and Victor Burger for their guidance and the many discussions that helped shape this research. The combination of academic methods and questions that matter to the sector made this project a rewarding one to work on.

I would also like to thank Dr. Berend Roorda for his supervision throughout the thesis process. His feedback sharpened both the methodology and the way its results are presented. My thanks also go to Dr. ir. Anouk Bomers for her feedback and for taking the time to answer my questions with regards to the hydraulic simulation. My thanks also go to Dr. Reinoud Joosten for taking the time to assess this work.

Finally, I want to thank my family and friends for their support during this final stage of my studies. Your encouragement made all the difference.

Mart van den Hoek
Ermelo, June 2026

Management Summary

Extreme rainfall is becoming more frequent and intense in the Netherlands; under the highest-emission scenarios its intensity could rise 10–30% by 2100. The resulting pluvial (rainfall-related) flooding can strike almost anywhere local drainage is overwhelmed, as the 2021 Limburg floods and the 2024 Enschede flash floods showed.

The Dutch social housing sector is highly exposed. Housing corporations own roughly 2.3 million dwellings, about 29% of the national stock, yet cannot raise rents to market levels, sell assets, or diversify geographically in response to risk. An estimated one third of their units face elevated pluvial risk. No prior study has quantified this risk for the sector.

My thesis closes that knowledge gap with a nationwide simulation that estimates flood depths from 5 m terrain data and KNMI rainfall scenarios, converts them to structural damage through a Dutch-calibrated European depth–damage curve, and expresses the result as Value at Risk (VaR), Expected Shortfall, and Expected Annual Damage (EAD), validated against an independent professional flood model. The shown figures cover the policy-relevant SSP5-8.5 2050 pathway, a high-end climate scenario at mid-century, and capture direct repair costs, a lower bound on total exposure.

Four findings stand out. First, the losses are material: a 100-year event produces an estimated EUR 457 million in direct structural damage (the 99% VaR). This assumes one storm strikes the whole country at once and is therefore an upper bound; the geographic spread of a nationwide portfolio places the realistic figure lower, with EUR 117 million as a conservative floor. Annualised, the Expected Annual Damage lies between roughly EUR 17 million and EUR 64 million per year. The EUR 17 million figure counts only the rarer events for which the model is best validated. It is therefore a conservative lower bound. the EUR 64 million figure also includes frequent events at face value, where the model is known to overestimate, and is an upper bound. Correcting that overestimation gives a central estimate of about EUR 34 million per year.

Second, the risk is driven by shallow water, not rare deep floods. As rainfall intensifies, losses grow almost entirely because more ground-floor dwellings cross the damage threshold, not because flooded homes flood deeper. This threshold-crossing mechanism is the central finding: modest measures raising the effective entry height, such as door barriers or street-level drainage, act precisely at the margin that dominates portfolio risk.

Third, the losses are absorbable nationally but tightly concentrated. The EUR 457 million shock equals about 18% of one year of the sector's aggregate operating cashflow, so it does not threaten solvency. Yet a few regions, led by Regio Rotterdam, Haaglanden, and Stedelijke Regio Breda–Tilburg, carry the highest exposure per dwelling, and under a one-off shock Breda–Tilburg is the only one whose Interest Coverage Ratio falls below the Aw/WSW norm of 1.4. The binding constraint is not solvency but the squeeze on cashflow already committed under the Nationale Prestatieafspraken. These are the national agreements in which the sector, the national government, municipalities, and tenant organisations have already committed corporations' investment capacity to new construction, sustainability, and affordability targets.

Fourth, inaction compounds: holding the 2050 distribution constant, thirty years without adaptation accumulate an expected EUR 505 million in direct damage (about EUR 377 million in present value at 2% real), and a 25-year event is more likely than not to strike at least once in that window. Because these losses are recurring, regionally concentrated, and near-certain to materialise, the case for adaptation rests on the avoidable, predictable cost that delay locks in.

I therefore argue for treating pluvial risk as a real item in long-term financial planning and supervisory stress testing. Adaptation should be prioritised on both aggregate exposure and damage per dwelling. Finally, early investment is warranted in shallow-water entry prevention and in the asset-level data that could replace today's conservative assumptions.

Contents

Acknowledgements i

Summary ii

List of Figures vii

List of Tables ix

List of Abbreviations xi

1 Introduction 1

1.1 Background . . . 1
1.2 Problem Context . . . 1
1.3 Research Objective and Questions . . . 2
1.4 Research Design and Methodological Scope . . . 4
1.5 Thesis Outline . . . 4

2 Literature Review 6

2.1 Climate Change and Extreme Precipitation in the Netherlands . . . 6
2.2 The Dutch Social Housing Sector and Its Financial Framework . . . 8
2.3 Flood Damage Modelling and Depth–Damage Functions . . . 9
2.4 Spatial Flood Simulation Approaches . . . 10
2.5 Financial Risk Quantification . . . 10
2.6 Flood Risk and Property Values . . . 12
2.7 Climate Adaptation in the Social Housing Sector . . . 13
2.8 Research Gap and Synthesis . . . 13

3 Data 15

3.1 Social Housing Portfolio . . . 15
3.2 KNMI Precipitation Data . . . 15
3.3 Geospatial Data . . . 16
3.4 Geocoding . . . 17

3.5 Descriptive Statistics . . . 17

4 Methodology **19**
4.1 Spatial Grid and Square Selection . . . 22
4.2 Terrain Data and Preprocessing . . . 22
4.3 Sink Mask Construction . . . 23
4.4 Routing Surface Construction . . . 24
4.5 Building Footprints and Roughness . . . 24
4.6 Rainfall Simulation . . . 25
4.6.1 Scenario rainfall depths and Gumbel extrapolation . . . 25
4.6.2 Design storm construction . . . 26
4.6.3 Simulation loop . . . 27
4.7 Damage Calculation . . . 28
4.7.1 Flood depth sampling per building . . . 28
4.7.2 Floor-level proxy . . . 29
4.7.3 Inundation threshold . . . 29
4.7.4 Depth–damage function . . . 29
4.8 Value at Risk Estimation . . . 30
4.9 Model Validation . . . 32
4.9.1 Aggregate validation statistics . . . 33
4.9.2 Cross-scenario damage ratios . . . 33
4.9.3 Sensitivity to sewer parameterisation . . . 34
4.9.4 Interpretation for VaR and ES . . . 34
4.10 Model Limitations . . . 35

5 Results **36**
5.1 Scope of the SSP5-8.5 2050 simulation run . . . 36
5.2 Portfolio-wide damage by return period . . . 37
5.3 Loss-exceedance curve, VaR, and Expected Shortfall . . . 38
5.4 Expected Annual Damage . . . 41
5.5 The risk of inaction over a 30-year horizon . . . 44
5.6 Spatial distribution of 100-year damage . . . 46
5.7 Spatial distribution of Expected Annual Damage . . . 47
5.8 Financial absorption capacity . . . 47
5.9 Depth distributions and threshold crossing . . . 51
5.10 Sensitivity to the inundation threshold . . . 52
5.11 Synthesis: Threshold crossing drives portfolio risk . . . 53

6 Conclusion **55**
6.1 Answer to the research questions . . . 55
6.2 Academic contribution . . . 57
6.3 Practical implications . . . 58
6.4 Limitations . . . 59
6.5 Recommendations . . . 60
6.6 Future research . . . 61

References **62**

A Use of Generative AI Tools **68**

B Methodology: Worked Examples and Derivations **70**
B.1 Flow routing equations and worked examples . . . 70
B.2 Routing-surface worked example . . . 73
B.3 Worked Value-at-Risk and Expected Shortfall calculations . . . 73
B.4 Loss-exceedance curve interpolation: linear versus $(\ln T, D)$ . . . 74
B.5 Micro-noise conditioning of the terrain . . . 74
B.6 Calibration of the depression-filling threshold . . . 75
B.7 Thermodynamic basis of the super-Clausius–Clapeyron scaling . . . 75
B.8 Choice of the inundation threshold . . . 76
B.9 Tail handling beyond the simulated range . . . 76
B.10 Truncated Expected Shortfall estimator . . . 76
B.11 Sink-mask detection parameters . . . 77

C Full Cross-Scenario Validation Table **78**

D Full Sewer-Parameterisation Sensitivity Table **81**

E Results: Robustness Checks **83**
E.1 Tail extrapolation robustness . . . 83

F Results: Spatial Damage and EAD Maps **85**
F.1 100-year direct structural damage . . . 86
F.2 Expected Annual Damage . . . 89

G Operating-cashflow buffer per Woondealregio **91**

H Interest Coverage Ratio per Woondealregio **95**
H.1 Stressed ICR under a one-off T=100 yr pluvial shock . . . 98

List of Figures

1.1 Thesis structure and research-question mapping. 5

3.1 Spatial distribution of social housing units across the simulation grid . . 18

4.1 Five-stage estimation pipeline from rainfall to portfolio risk metric. . . . 20
4.2 Gumbel Extrapolation 25
4.3 Example output of the simulation loop for a single 5 km simulation square. 28

5.1 Loss-exceedance curve and VaR/ES estimates for SSP5-8.5 2050 38
5.2 Spatial-dependence bracket for the SSP5-8.5 2050 loss-exceedance curve . 41
5.3 Trapezoidal decomposition used to derive Expected Annual Damage from the simulated loss-exceedance curve 42
5.4 The risk of inaction over a 30-year horizon under SSP5-8.5 2050 45
5.5 Per-Woondealregio damage absorption capacity 48
5.6 Impact of a T=100 yr pluvial shock on regional ICR 50
5.7 Distribution of maximum depth at the building for each return-period scenario 52
5.8 Damage and VaR sensitivity to the inundation threshold 53

B.1 Before and after a single routing sub-step 71
B.2 Top-down stencil illustrating Equations (B.1) and (B.2) 72

E.1 Tail extrapolation robustness check for the SSP5-8.5 2050 loss-exceedance curve. 83

F.1 Total 100-year direct structural damage by province under SSP5-8.5 2050 86
F.2 Total 100-year direct structural damage by Woondealregio under SSP5-8.5 2050 87
F.3 100-year direct structural damage per ground-floor unit by Woondealregio under SSP5-8.5 2050. 88
F.4 Expected Annual Damage by Woondealregio under SSP5-8.5 2050 89
F.5 Expected Annual Damage per ground-floor unit by Woondealregio under SSP5-8.5 2050 90

G.1 Per-Woondealregio damage absorption capacity, T=100 yr SSP5-8.5 2050 93

H.1 Per-Woondealregio ICR distribution, dVi2024 96
H.2 Impact of T=100 yr pluvial damage on regional ICR 100

List of Tables

2.1 KNMI point-rainfall quantiles (mm) for sub-daily storm durations and selected return periods . . . 7

3.1 Geospatial data sources used in the simulation pipeline. . . . 16
3.2 Descriptive statistics of the social housing portfolio. . . . 17

4.1 Climate-scaling parameters per KNMI'23 pathway . . . 26
4.2 Design storm rainfall totals . . . 26
4.3 European residential depth–damage curve . . . 30
4.4 Damage ratio (Tygron / local-inertial) by square and scenario, SSP5-8.5 2050 . . . 34

5.1 Portfolio-wide direct structural damage by return period, SSP5-8.5 2050. 37
5.2 Growth decomposition relative to the 5-year event. . . . 38
5.3 VaR and truncated ES of direct structural portfolio damage, SSP5-8.5 2050. 39
5.4 Expected Annual Damage under the two integration-range and calibration choices . . . 43
5.5 Exposure to pluvial events over a 30-year horizon of inaction . . . 44
5.6 Top-five provinces by total 100-year direct structural damage . . . 46
5.7 Top-five Woondealregio's by total 100-year direct structural damage . . . 46
5.8 Top-five Woondealregio's by 100-year direct structural damage per ground-floor unit . . . 46
5.9 Top-five Woondealregio's by aggregate Expected Annual Damage . . . 47
5.10 Top-five Woondealregio's by Expected Annual Damage per ground-floor unit 47
5.11 Top-three most-stressed Woondealregio's on each financial metric . . . 49
5.12 Sensitivity of VaR to the inundation threshold, SSP5-8.5 2050. . . . 52

C.1 Full per-run validation metrics . . . 79

D.1 Per-run damage ratio under baseline and raised sewer parameters . . . 82

G.1 Damage versus operating-cashflow buffer per Woondealregio . . . 92

H.1 Interest Coverage Ratio per Woondealregio, dVi2024 . . . 97

H.2 Pre- and post-shock ICR per Woondealregio 101

List of Abbreviations

Aw	Autoriteit Woningcorporaties (Dutch housing association supervisor)
WSW	Waarborgfonds Sociale Woningbouw (Social Housing Guarantee Fund)
DAEB	Diensten van Algemeen Economisch Belang (services of general economic interest)
NAP	Normaal Amsterdams Peil (Dutch elevation datum)
BAG	Basisregistratie Adressen en Gebouwen (Key Register of Addresses and Buildings)
BGT	Basisregistratie Grootschalige Topografie (Key Register of Large-Scale Topography)
AHN	Actueel Hoogtebestand Nederland (national LiDAR elevation dataset)
PDOK	Publieke Dienstverlening Op de Kaart (national geodata portal)
EGW	Eengezinswoning (single-family dwelling)
MGW	Meergezinswoning (multi-family dwelling, e.g. apartment)
KNMI	Koninklijk Nederlands Meteorologisch Instituut (Royal Netherlands Meteorological Institute)
STOWA	Stichting Toegepast Onderzoek Waterbeheer (Dutch applied water research foundation)
CBS	Centraal Bureau voor de Statistiek (Statistics Netherlands)
WOZ	Waardering Onroerende Zaken (Dutch statutory property valuation)
RIONED	Stichting RIONED (Dutch knowledge centre for urban water and sewerage)
VHE	Verhuureenheden (lettable/rental units)
VaR	Value at Risk
ES	Expected Shortfall
EAD	Expected Annual Damage

LEC	Loss-Exceedance Curve
ICR	Interest Coverage Ratio
LTV	Loan-to-Value ratio
DSCR	Debt Service Coverage Ratio
FRTB	Fundamental Review of the Trading Book (Basel market risk framework)
NGFS	Network for Greening the Financial System
CRISK	Climate-stress capital shortfall
ECB	European Central Bank
SSM	Standaardmethode Schade en Slachtoffers (Dutch standard flood damage method)
DDF	Depth–Damage Function
JRC	Joint Research Centre (European Commission)
FCF	Free Cash Flow
CPI	Consumer Price Index
IDF	Intensity–Duration–Frequency (curve)
GEV	Generalised Extreme Value (distribution)
GPD	Generalised Pareto Distribution
EVT	Extreme Value Theory
POT	Peaks-Over-Threshold
GAMLSS	Generalised Additive Model for Location, Scale, and Shape
AEP	Annual Exceedance Probability
CC	Clausius–Clapeyron
SSP	Shared Socioeconomic Pathway
RCP	Representative Concentration Pathway
IPCC	Intergovernmental Panel on Climate Change
AR6	(IPCC) Sixth Assessment Report
DTM	Digital Terrain Model
DSM	Digital Surface Model
LiDAR	Light Detection and Ranging

SWE	Shallow Water Equations
CFL	Courant–Friedrichs–Lewy (numerical stability condition)
EDT	Euclidean Distance Transform
RMSE	Root-Mean-Square Error
WCS	Web Coverage Service
WFS	Web Feature Service
OGC	Open Geospatial Consortium
PFR	Property-level Flood Resilience
NBS	Nature-Based Solutions
SUDS	Sustainable Urban Drainage Systems

Pand: Dutch term for a building or premises (the physical structure, which may contain one or more dwellings).
Pipeline: A chain of elements where the output of one element serves as the input to the next.
Woondealregio: A partnership between municipalities, provinces, housing corporations, and the central government.

Chapter 1

Introduction

1.1 Background

My research is conducted in collaboration with Finance Ideas, a financial consultancy in Utrecht. The firm advises organisations in the public and semi-public sectors, including housing associations, healthcare institutions, and institutional investors, on financial strategy, treasury management, and risk analysis.

In the Netherlands, social housing for low-income households is provided by non-profit organisations known as *woningcorporaties* (housing corporations), comparable to housing associations in the United Kingdom or Public Housing Agencies in the United States. They collectively own approximately 2.3 million dwelling units, around 29% of the total Dutch housing stock (Aedes, 2026), and operate under the supervision of the Autoriteit Woningcorporaties (Aw), which monitors financial management, governance, and adherence to social objectives (Autoriteit woningcorporaties, 2024). The anonymised Aw dataset used in my study was provided via Aw under a restricted-use agreement designed to prevent traceability to individual associations.

1.2 Problem Context

Climate change is intensifying extreme precipitation in the Netherlands. The KNMI'23 climate scenarios project a 10–30% increase in the intensity of heavy rainfall by 2100 under high-emission pathways, with the largest changes in summer convective events (IPCC, 2023; KNMI, 2023). Two types of flooding must be distinguished. Fluvial flooding originates from rivers exceeding their capacity, and is concentrated along well-defined floodplains where Dutch dike infrastructure has reduced annual exceedance probabilities to very low levels (Jonkman et al., 2008). Pluvial flooding occurs when rainfall intensity exceeds the local infiltration and drainage capacity of the urban surface. It depends on local terrain, soil permeability, and sewer capacity rather than on proximity to rivers, and

can therefore occur anywhere(KNMI, 2023). Recent events show what is at stake: the 2021 Limburg floods caused approximately €433 million in damage (Deltares, 2023), and the 2024 Enschede flash floods(No damage estimate found), a purely pluvial convective event, rendered dozens of homes uninhabitable (Tuijn, 2025).

The social housing sector faces a distinct form of exposure. About 33% of social housing units are estimated to face an elevated pluvial risk (Dutch Green Building Council, 2024), and several features of the sector amplify this risk. The Woningcorporaties cannot freely adjust their rents to market levels, they cannot sell their assets, and they cannot geographically diversify their portfolios in response to risk. The revenues of the Woningcorporaties come mostly from regulated rents that are subject to government-set caps, while extreme weather can cause sudden and large maintenance expenditures. In addition, a substantial share of the housing stock has been built before 1990, in neighborhoods that are served by aging combined sewer systems with a lower drainage capacity than modern separated systems. These older buildings also frequently sit at lower elevations relative to the street.

Despite this evident relevance, the literature does not provide a comprehensive quantification of the financial risk for the sector. Dutch flood-risk studies focus predominantly on fluvial scenarios (Caloia et al., 2023; Eichholtz et al., 2026). On the adaptation side, Roders and Straub (2015) and Boezeman and de Vries (2019) report that awareness has increased, while concrete actions remain limited. Recent methodological work has advanced as well (Allegri et al., 2024; Endendijk et al., 2023; Truong et al., 2024). However, none of these studies has combined flood simulation with portfolio-level financial risk measures for Dutch social housing specifically. Without this evidence base, the Woningcorporaties and the Aw lack the input that they need to prioritize adaptation investments and to integrate climate risk into their long-term financial planning.

1.3 Research Objective and Questions

The primary objective of my thesis is to quantify the financial risks that extreme pluvial precipitation poses to Dutch social housing portfolios. The focus is on the Value at Risk (VaR) at confidence levels from 90% to 99% (Hull, 2023). We trace the full range of confidence levels instead of reporting a single one, because this reveals how rapidly the potential losses escalate at higher confidence levels. This leads to the main research question:

> *To what extent do extreme pluvial precipitation events pose financial risks to Dutch social housing portfolios, and how can these risks be quantified using flood simulation and Value at Risk estimation?*

Four sub-questions structure this analysis:

1. **How can pluvial flood depths at social housing locations be simulated using high-resolution terrain data and extreme rainfall scenarios?** I answer this question in Chapter 4. I develop a 5 m raster simulation framework that combines AHN elevation data, BAG building footprints and a KNMI-derived design storm, and we validate this framework against the Tygron hydrodynamic platform.

2. **What is the magnitude and distribution of financial losses across social housing portfolios under extreme pluvial scenarios, as measured by Value at Risk?** I answer this question in Chapter 5. I link the simulated flood depths to the portfolio characteristics through the JRC European depth–damage curve, and we compute the VaR and the Expected Shortfall (ES) at confidence levels from 90% to 99%.

3. **How do financial risks from pluvial flooding vary across geographic locations, and which areas constitute spatial hotspots of portfolio vulnerability?** I asnwer this question by comparing the damage estimates across simulation squares that are ranked by portfolio density, and by aggregating these estimates to the gemeente, provincie and Woondealregio levels.

4. **Which mechanism drives portfolio damage growth across increasing return periods, and what does this imply for adaptation strategy?** This question is answered in Chapter 5. I decompose the scenario damage growth into an extensive margin (additional dwellings that cross the inundation threshold, the water depth at which a dwelling begins to register damage) and an intensive margin (deeper inundation of dwellings that are already damaged). We corroborate this decomposition with the building-level depth distributions and the threshold-sensitivity analysis.

Contributions

Academic. My research contributes to the literature on climate-related financial risk. It applies flood simulation and VaR estimation to nonprofit social housing, and in this way it bridges meteorological hazard modeling and portfolio-level risk management. My research combines high-resolution terrain-based flood modeling with the JRC European depth–damage curve (Huizinga et al., 2017). This curve is calibrated against SSM2023 replacement values (de Grave & Juch, 2023) and corroborated against recent Dutch empirical evidence (Endendijk et al., 2023). Finally, my research integrates these elements with portfolio-level risk measures (Hull, 2023).

Practical. For stakeholders in the Dutch social housing sector, my thesis provides actionable insights into the portfolio vulnerabilities. It quantifies the potential financial losses under extreme pluvial scenarios. In this way, it allows the Woningcorporaties and the Aw to prioritize adaptation measures, to allocate financial buffers and to integrate climate risk into their long-term financial planning.

1.4 Research Design and Methodological Scope

My research follows a quantitative and deductive approach in five phases. First, I identified the best-practice methods for pluvial flood simulation, depth–damage function calibration and financial risk quantification. Second, the data processing covers the Aw portfolio dataset, the KNMI radar precipitation, the AHN terrain, the BAG building footprints and the BGT land-use and water masks. Third, we run a 5 m raster flood simulation per simulation square, using a local-inertial routing scheme (Bates et al., 2010; Neal et al., 2012) with drainage through the sewer network and infiltration. Fourth, we compute the damage with the JRC depth–damage curve and SSM2023 replacement values, and we estimate the VaR and ES at confidence levels from 90% to 99%. Fifth, we validate the model against the Tygron hydrodynamic platform, and we complement it with parameter sensitivity analyses.

The simulation matrix consists of five climate pathways (a current-climate baseline, plus SSP2-4.5 and SSP5-8.5, each at a 2050 and a 2100 horizon), and these pathways are crossed with seven return periods. This results in 35 rainfall scenarios in total. Within the computational resources available for my thesis, the full national simulation run is delivered for the SSP5-8.5 mid-century pathway only, because this is the most policy-relevant case in the KNMI'23 ensemble. The generalization to the other four pathways requires no new methodology, and we identify it as future work. ES is reported alongside VaR throughout.

1.5 Thesis Outline

Figure 1.1 maps the six chapters to the research questions that they address. In Chapter 2 I review the relevant literature and identify the research gap. In Chapter 3 I describe the data sources. In Chapter 4 I explain the simulation framework and the VaR/ES estimation procedure, and I address SQ1. In Chapter 5 I present the numerical findings under SSP5-8.5 mid-century, and it addresses SQ2, SQ3 and SQ4. Fianlly in chapter 6 I synthesize the findings, answers the main research question and discuss the implications, limitations and future research.

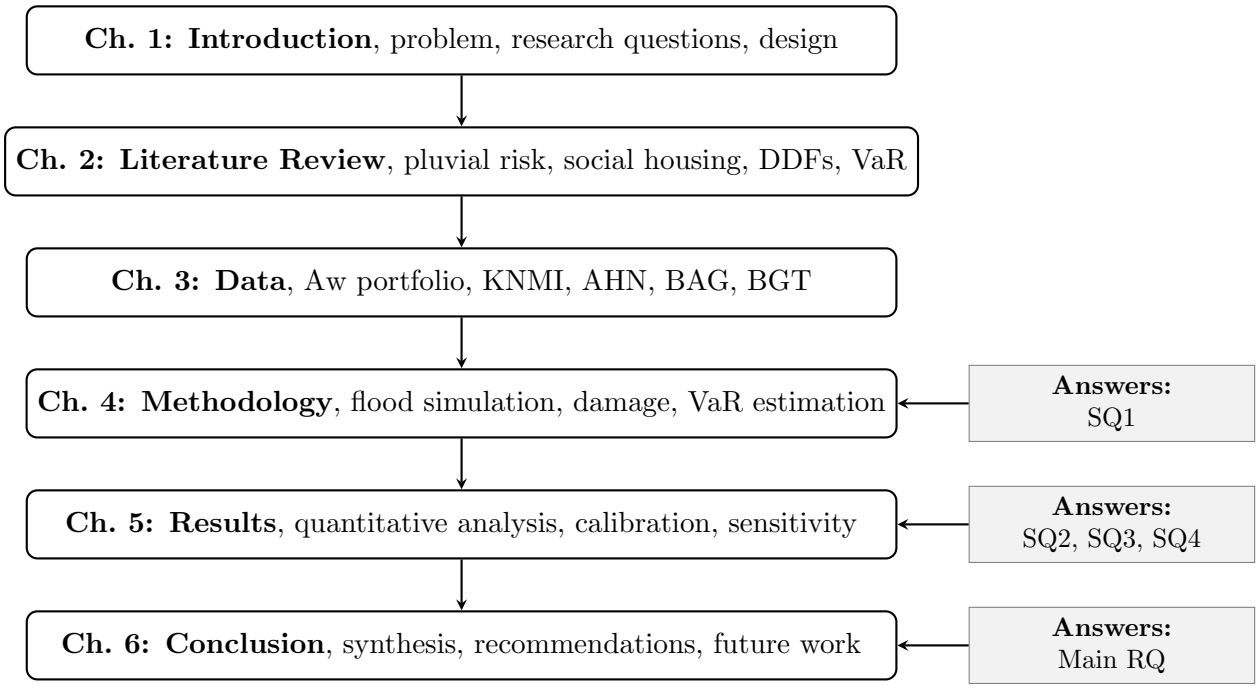


Figure 1.1: *Thesis structure and research-question mapping.*

Chapter 2

Literature Review

To quantify the financial risks that extreme pluvial precipitation poses to Dutch social housing portfolios, three things must be understood: the nature of the hazard itself, the institutional context of the sector that bears the risk, and the analytical tools available for converting physical exposure into monetary loss estimates. In this chapter I review seven strands of literature and synthesise the research gap that motivates my thesis. In Section 2.1 I review the climate science of extreme precipitation in the Netherlands. In Section 2.2 I describe the social housing sector and its financial structure. In Section 2.3 I examine depth–damage functions; in Section 2.4 I survey spatial flood simulation; In Section 2.5 I review financial risk quantification; In Section 2.6 I review flood risk and property values; in Section 2.7 I briefly examine climate adaptation; In Section 2.8 I synthesise the research gap.

2.1 Climate Change and Extreme Precipitation in the Netherlands

The IPCC AR6 (the most comprehensive, up-to-date scientific assessment of climate change) confirms that human-induced warming has already increased the frequency and intensity of heavy precipitation globally, with high confidence that this trend will continue under all emission pathways (IPCC, 2023). For the Netherlands, the KNMI'23 climate scenarios project a 10–30% increase in heavy hourly rainfall by 2100 under high-emission pathways, with the largest changes in summer convective events (KNMI, 2023). These projections are consistent with the Clausius–Clapeyron scaling, which predicts roughly a 7% increase in atmospheric moisture-holding capacity per degree of warming, although observed short-duration extremes in the Netherlands have frequently intensified faster than this rate (Lenderink & van Meijgaard, 2008).

KNMI's Dutch point-rainfall frequency statistics provide the quantitative basis for converting this intensification into design rainfall depths (Beersma et al., 2018; Koninklijk

Nederlands Meteorologisch Instituut, 2018). The statistics describe the rainfall depth expected to be exceeded once in T years for storm durations from 5 minutes to 10 days, fitted using the Generalised Extreme Value (GEV) distribution to annual maxima from Dutch rain-gauge networks. Critically, no statistically significant regional differences in extreme precipitation have been demonstrated for sub-24-hour durations: the statistics apply uniformly to any point in the country. In Table 2.1 I show the values for sub-daily durations and return periods up to $T = 100$ years; extrapolation beyond this range is handled in Chapter 4.

Table 2.1: *KNMI point-rainfall quantiles (mm) for sub-daily storm durations and selected return periods, valid at any location in the Netherlands. The simulation design storm spans a 120-minute window. Source: (Beersma et al., 2018; Koninklijk Nederlands Meteorologisch Instituut, 2018).*

Return period	T (yr)	5 min	15 min	30 min	60 min	**2 hr**	4 hr	8 hr	12 hr
2×/yr	0.5	6	9	10	13	**15**	19	22	25
1×/yr	1	7	11	14	16	**20**	23	28	31
1 in 2 yr	2	9	14	17	20	**24**	28	33	37
1 in 5 yr	5	11	17	21	26	**31**	36	42	45
1 in 10 yr	10	12	20	25	31	**37**	43	49	53
1 in 20 yr	20	14	24	30	37	**44**	51	58	62
1 in 50 yr	50	17	29	38	48	**57**	65	73	77
1 in 100 yr	100	20	35	46	58	**68**	78	86	90

The KNMI'23 climate scenarios are organised around two emission pathways used in this study (SSP2-4.5 and SSP5-8.5) at 2050 and 2100 horizons. Combined with the super-Clausius–Clapeyron scaling of Lenderink and van Meijgaard (2008) (14% per degree), local Netherlands warming of +1.3 to +4.4 °C above the 1991–2020 reference period yields rainfall intensity multipliers $\phi = 1.14^{\,\Delta T}$ from 1.19 (SSP2-4.5, 2050) to 1.78 (SSP5-8.5, 2100). A 50-year 2-hour design storm of 57 mm under present climate would rise to roughly 101 mm under the most severe pathway considered, motivating the climate-scenario stress tests in Chapter 4.

Stationarity and its consequences. Conventional hydrological frequency analysis assumes *stationarity*: the underlying distribution that generates extreme events does not drift over time. Under climate change this assumption is contested. Milly et al. (2008) declared "stationarity is dead", but Montanari and Koutsoyiannis (2014) and Serinaldi and Kilsby (2015) argued that the additional parameter uncertainty of non-stationary models often outweighs their predictive gain, even when the methods fit observed data well. The methodological literature has converged on a pragmatic middle ground. Non-stationary methods are warranted when three conditions are met: a physically plausible driver of non-stationarity exists, the observed record is long enough for the trend to be detectable against natural variability, and the added parameter uncertainty does not

overwhelm predictive accuracy (Cheng & AghaKouchak, 2014; François et al., 2019; Salas et al., 2018). For the present thesis this points to a scenario-based rather than frequency-based approach: a range of rainfall totals spanning present-day events to high-emission futures avoids committing to a specific non-stationary parameterisation, while still capturing the range of plausible loss outcomes. The VaR estimates in Chapter 5 are accordingly *conditional on a chosen climate pathway*, rather than a single climate-averaged figure.

2.2 The Dutch Social Housing Sector and Its Financial Framework

Housing corporations (*woningcorporaties*) collectively own approximately 2.3 million residential units, about 29% of the Netherlands' total housing stock (Aedes, 2026), and operate under the supervision of the Aw (Autoriteit woningcorporaties, 2024). Unlike commercial property owners, housing corporations cannot freely adjust rents to market levels, sell assets at will, or diversify their portfolios geographically: their core mission imposes constraints that make climate risk management fundamentally different from that of private landlords or institutional real-estate investors. Their revenue is dominated by regulated rents subject to government-set caps, while maintenance and renovation expenditures can spike unpredictably when extreme weather damages a portfolio. Dutch Green Building Council (2024) estimated that climate-related damage could cost the sector several hundred million euros annually under high-emission scenarios, with pluvial flooding the dominant driver. Boezeman and de Vries (2019) found that while awareness of climate risks has grown among housing corporations, concrete adaptation actions remain limited.

A structural feature of the framework is that the Interest Coverage Ratio (ICR, operating cashflow over interest expense, with a regulatory floor of 1.4) is the binding investment-capacity constraint for most corporations, rather than the Loan-to-Value ratio. The June 2024 Aw/WSW recalibration of beleidswaarde-based norms (beleidswaarde being the regulated, social-policy-adjusted valuation of the portfolio) left borrowing capacity limited by cashflow rather than collateral (Autoriteit woningcorporaties & Waarborgfonds Sociale Woningbouw, 2024; Finance Ideas, 2024). A one-off pluvial damage event that reduces operating cashflow can therefore push the ICR below its 1.4 threshold for the year of the event, and constrain subsequent investment capacity. This makes VaR a natural primary risk measure for the present analysis: it provides a threshold loss estimate at a specified confidence level, in line with the Aw's regulatory interest in ensuring that corporations maintain sufficient buffers to absorb shocks without compromising their social mission.

2.3 Flood Damage Modelling and Depth–Damage Functions

Translating physical hazard into monetary damage requires depth–damage functions (DDFs), which express damage as a function of inundation depth. Their calibration has a decisive influence on computed losses.

The Dutch Standard Method. The *Standaardmethode Schade en Slachtoffers* (SSM), maintained by Deltares, is the standard approach for flood damage assessment in the Netherlands. Slager and Wagenaar (2017) documented the 2017 version, specifying maximum damage values per square metre with exponential depth–damage curves calibrated primarily on the 1953 North Sea and 1993/1995 Meuse floods. SSM has been widely adopted, including by Caloia et al. (2023) for banking-sector credit losses and by Kok et al. (2023) for the 2021 Limburg flood damage assessment. However, the SSM exponential curves were developed primarily for deep fluvial flooding and produce damage fractions that are too low at the shallow depths (0.05–0.50 m) characteristic of pluvial events (de Moel et al., 2015).

Empirical curves for shallow inundation. Recognising this limitation, Endendijk et al. (2023) estimated multivariate flood-vulnerability models using household survey data collected after the 2021 Limburg floods. Their dependent variable is a damage ratio: absolute damage relative to estimated replacement value. For structural building damage, inundation depth enters the model through a square-root term: across specifications, the coefficient on $\sqrt{d_{cm}}$ lies between roughly 0.014 and 0.017 and is statistically significant at the 1% level. Interpreted only as the marginal depth component, a coefficient of 0.015 implies depth-related increases of about $0.015\sqrt{10} = 4.7$ percentage points at 10 cm, 10.6 percentage points at 50 cm, and 15.0 percentage points at 1 m. The full Endendijk et al. model also includes an intercept and covariates such as flood-damage mitigation measures, homeownership, dwelling size, building year, location, flow velocity, and building type. The result should therefore be read as evidence for a concave shallow-depth damage relationship, rather than as the standalone damage function used in my thesis.

Uncertainty. Despite these advances, DDFs remain a major source of uncertainty in flood risk assessment. de Moel et al. (2015) showed that different published damage models applied to the same flood scenario produce estimates that vary by a factor of two to five even when the underlying inundation data are identical. A key driver is that most DDFs depend only on depth and ignore duration, velocity, sediment load, and contamination. For my thesis, this implies that point estimates of damage should be interpreted with caution and that a range of VaR estimates across confidence levels is more informative than any single damage figure.

2.4 Spatial Flood Simulation Approaches

Pluvial flood modelling requires simulating the accumulation and movement of water across the urban surface after intense rainfall. At portfolio scale, when the analysis spans millions of dwellings nationally, this involves a fundamental trade-off between physical realism and computational feasibility.

A standard is the full two-dimensional shallow water equations (SWE), implemented in software such as Tygron, which can be coupled with one-dimensional sewer network models. Neal et al. (2012) found that the full SWE produce the most accurate representation of flow through street networks, but they are computationally expensive, and student-licence implementations are restricted in the number of areas that can be simulated. This is my primary motivation for the simplified approaches.

Bates et al. (2010) introduced the local-inertial scheme. It retains the local acceleration term from the Saint-Venant equations while dropping the advective inertia. The resulting scheme is numerically stable at relatively large time steps governed by the Courant–Friedrichs–Lewy (CFL) condition, requires only a Manning friction term to close the momentum equation (Chow, 1959; Manning, 1891), and (unlike the simpler diffusive-wave approximation) carries momentum between timesteps. Neal et al. (2012) showed that without this momentum term water concentrates in local depressions rather than spreading through connected street channels, producing deeper but less extensive flooding. At portfolio level the local-inertial formulation offers a better trade-off than the diffusive-wave alternatives, and is the approach adopted in Chapter 4.

Simulation accuracy depends on terrain data quality. The Netherlands provides the AHN, a national LiDAR-derived elevation dataset at 0.5–5 m resolution. Raw LiDAR contains false depressions (measurement noise, vehicles, street furniture) that require filtering. Standard fill algorithms (Planchon & Darboux, 2002; Wang & Liu, 2006) are designed for fluvial applications and must be adapted for urban pluvial simulation, where some depressions are real retention features and others are artefacts; the depression-filling mask developed for this purpose is described in Chapter 4.

2.5 Financial Risk Quantification

Value at Risk (VaR) at confidence level α answers a single question: what is the loss level that the portfolio will not exceed in $\alpha \times 100\%$ of pluvial scenarios? Formally, VaR is the α-quantile of the loss distribution (Hull, 2023). Estimation can proceed parametrically (assuming a distributional form), through historical simulation, or through Monte Carlo methods. For natural-hazard risk, where loss distributions are typically heavy-tailed and asymmetric, the parametric approach is often inadequate, and simulation-based methods are preferred. The framework developed in Chapter 4 falls within the Monte Carlo tradi-

tion, with one important difference: the underlying scenarios are physically constrained by a hydrodynamic model rather than sampled from a statistical loss distribution.

Coherent risk measures. Artzner et al. (1999) proposed four axioms that any coherent risk measure should satisfy: monotonicity, translation invariance, positive homogeneity, and subadditivity. VaR satisfies the first three but fails subadditivity (the property that merging two portfolios never increases total risk) in general. Expected Shortfall (ES), defined as the conditional mean loss above VaR, is coherent in their sense (Acerbi & Tasche, 2002), and the 2019 revision of the Basel market-risk framework formally replaced VaR with ES at 97.5% as the primary measure for bank market risk capital (Basel Committee on Banking Supervision, 2019). However, Daníelsson et al. (2013) showed that VaR is in fact subadditive deep in the tail for distributions with sufficiently quickly decaying tails, and is more robust than ES in small samples. This is directly relevant to a setting where the loss distribution is reconstructed from only seven physically simulated scenarios. The present thesis takes a pragmatic position: VaR at confidence levels from 90% to 99% reveals the shape of the tail without requiring the stronger assumptions about tail behaviour that a stable ES estimate would need; ES is reported alongside VaR as a complementary view of tail severity.

Climate scenario analysis and physical-risk practice. The Network for Greening the Financial System (NGFS) publishes long-term climate scenarios that combine integrated assessment and macroeconomic models (Network for Greening the Financial System, 2024). The European Central Bank's 2022 climate stress test of 104 euro-area banks (European Central Bank, 2022) found that most banks lacked the data to identify which mortgage borrowers held flood-exposed properties, a gap that motivates the building-level approach adopted here. The closest Dutch analogue is Caloia et al. (2023), who estimated VaR for Dutch mortgage portfolios under fluvial flood risk; my thesis transposes that framing to pluvial risk and to the social housing portfolio. Commercial real-estate climate-risk frameworks (Morgan Stanley Capital International, 2020; United Nations Environment Programme Finance Initiative, 2023) continue to use VaR rather than ES as their headline metric even after the Basel shift to ES, confirming VaR's standing as the operative measure in real-estate climate risk communication. Abadie et al. (2016, 2019) applied VaR and ES to sea-level-rise damages for European coastal cities through a non-stationary stochastic diffusion process; the present thesis adopts a similar scenario-based logic but uses pluvial scenarios and empirical intensity–duration–frequency (IDF) mapping in place of stochastic diffusion.

2.6 Flood Risk and Property Values

The relationship between flood risk and property values is relevant because WOZ valuations and their derived measures (*marktwaarde*, the open-market value, and *beleidswaarde*, its social-policy-adjusted counterpart) serve as the denominator against which housing-corporation supervisory ratios are calculated. If flood risk is systematically capitalised into these valuations, built into them as the market prices the risk in, a re-rating of the portfolio's valuation basis can tighten regulatory ratios independently of direct repair costs.

Beltrán et al. (2018) conducted a meta-analysis of 37 studies and 364 point estimates and, after correcting for publication bias and heterogeneity, arrived at a best estimate of approximately −4.6% for inland-floodplain location. They identified flood history as a critical explanatory variable: discounts are systematically larger after recent floods, consistent with availability-heuristic interpretations of risk perception. More recent US evidence using property-level repeat-sales designs has produced larger estimates (Gourevitch et al., 2023; Sandoval Olascoga et al., 2025). Dutch evidence is more nuanced: Eichholtz et al. (2026) found a 1.1% average discount for fluvial-prone properties, and earlier work (Bosker et al., 2019; Daniel et al., 2009) found that historical flood events caused temporary price reductions whose decay depended on adaptation visibility. The most recent Dutch evidence comes from Mutlu et al. (2023), who analysed the 1993–1995 Meuse floods and found a pre-flood capitalised discount of 5.6%, rising sharply to 10.9% after the floods, decaying over 9–12 years and eventually vanishing as large-scale nature-based flood-management interventions were implemented.

Transmission to housing-corporation supervisory ratios. For commercial real-estate investors, capitalised flood discounts translate into realised losses through transactions. For Dutch social housing corporations, which do not sell into the open market, the effect operates through two channels of the Aw–WSW supervisory framework. The first is a *direct cashflow channel*: a pluvial event imposes immediate repair costs. Because rent increases are capped, these costs flow through to operating cashflow and therefore to the ICR, the binding investment-capacity constraint, tightening borrowing capacity precisely when reconstruction financing is most needed. The second is a slower *valuation channel*: if pluvial risk is capitalised into marktwaarde and beleidswaarde following events, as the hedonic evidence suggests (Beltrán et al., 2018; Mutlu et al., 2023), a geographically concentrated portfolio may face a systematic downward revision of its asset base that weakens LTV, Solvabiliteit (equity relative to total assets), and Dekkingsratio (loans relative to the value of the collateral). My thesis quantifies only the direct cashflow channel; the valuation channel lies beyond scope and is taken up in Chapter 6. The VaR and ES estimates in Chapter 5 should accordingly be interpreted as *lower bounds* on total financial risk.

2.7 Climate Adaptation in the Social Housing Sector

My thesis does not evaluate adaptation measures, but a brief review motivates why *quantifying* unadapted risk is a prerequisite for sector-level decision-making. Property-level flood resilience (PFR) measures are empirically effective: Kreibich et al. (2015) found that PFR reduced structural damage by 53% and contents damage by 48% on average across German households affected by the 2002 and 2013 Elbe floods, and Endendijk et al. (2023) produced the Dutch analogue, showing statistically significant damage reductions from household-level mitigation in the 2021 Limburg flood sample. Implementation lags behind technical feasibility: Seeteram et al. (2026) identify path dependencies in housing finance, insurance, and regulation, together with information asymmetries between residents, landlords, and insurers, as the dominant barriers. For non-profit social housing these are compounded by mission-driven constraints on rent increases and portfolio disposals.

2.8 Research Gap and Synthesis

The literature reviewed above is individually mature; what is missing is an integrated framework that combines these strands to estimate portfolio-level financial loss for Dutch social housing under pluvial flooding. Four specific gaps motivate my thesis. First, financial-risk studies for the Netherlands focus predominantly on fluvial flooding (Caloia et al., 2023; Eichholtz et al., 2026; Jonkman et al., 2008), despite pluvial flooding being spatially universal and affecting a far larger share of the housing stock (Dutch Green Building Council, 2024). Second, recent Dutch empirical work has shown that flood damage is concave and already material at the shallow depths (0.05–0.50 m) characteristic of pluvial events (Endendijk et al., 2023), but this shallow-depth sensitivity has not been propagated through a depth–damage function to portfolio-level VaR or ES estimates. Third, commercial climate-risk frameworks (Morgan Stanley Capital International, 2020; United Nations Environment Programme Finance Initiative, 2023) rely on proprietary models, coarse hazard grids, and commercial real-estate assumptions that do not transfer to the non-profit social housing context, while NGFS-style stress-testing (European Central Bank, 2022; Network for Greening the Financial System, 2024) operates at national and sectoral aggregation, leaving a scale gap for the building-level detail needed to prioritise adaptation. Fourth, how pluvial risk varies with urban morphology, sewer age, and terrain has been noted qualitatively (Dutch Green Building Council, 2024; Swinkels, 2023) but not quantified at portfolio level. The methodological foundations on which my thesis builds are Bates et al. (2010) (local-inertial routing), Huizinga et al. (2017) and de Grave and Juch (2023) (depth–damage curves and replacement values), KNMI (2023) and Lenderink and van Meijgaard (2008) (climate scenarios and CC scaling), Artzner

et al. (1999) and Acerbi and Tasche (2002) (coherent risk measures), and Hull (2023) (VaR estimation procedures), with Caloia et al. (2023) as the closest methodological predecessor. Chapter 3 describes the data used to implement this framework.

Chapter 3

Data

This chapter introduces the data sources I used in the pluvial flood risk simulation and portfolio risk analysis. It describes what each dataset contains, where it comes from, and what filtering is applied. How these data are transformed by me into simulation inputs (design storm construction, terrain preprocessing, grid assignment, and building-level matching) is deferred to Chapter 4.

3.1 Social Housing Portfolio

The primary dataset is an anonymised extract from the Aw administrative registry, provided via Finance Ideas under a restricted-use agreement. It contains 2,305,576 dwelling units (*verblijfsobjecten*) owned by Dutch housing associations, representing approximately 29% of the total Dutch housing stock (Aedes, 2026). Each record corresponds to a single unit and contains the fields used in the simulation: a unique identifier (*Eenheidscode*); the 6-digit postcode and street address used for geocoding; the housing type (*Bouwvorm*: *Eengezinswoning* (EGW), *Meergezinswoning* (MGW), or *Flexwoning*), which determines which units are treated as ground-floor exposed; the gross floor area (*AantalM2*, in m^2) used for per-m^2 damage; the construction year (*Bouwjaar*); the BAG *verblijfsobject* identifier linking the unit to building footprints; and the municipal code and province for spatial grouping.

3.2 KNMI Precipitation Data

The temporal structure of the design storm is derived from KNMI 5-minute radar precipitation data (dataset `nl_rdr_data_rfcor_5m`), retrieved via the KNMI Data Platform API. The product provides spatially explicit rainfall intensity at 1 km^2 resolution across the Netherlands, updated every 5 minutes and corrected against rain-gauge observations. The design storm uses data from the Enschede region during the July 2024 extreme rain-

fall event (Tuijn, 2025); the procedure for converting the raw radar time series into the normalised hyetograph (the storm's rainfall-intensity profile over time) is described in Section 4.6.2. The target rainfall totals to which the profile is scaled, covering return periods from 2 to 1,000 years, are drawn from the KNMI point-rainfall statistics shown in Table 4.2 of Section 4.6.1, with return periods beyond 100 years extrapolated through the Gumbel fit of Equation (4.2).

3.3 Geospatial Data

Four national geospatial datasets are used, all accessed via the Publieke Dienstverlening Op de Kaart (PDOK) platform (Table 3.1).

Table 3.1: *Geospatial data sources used in the simulation pipeline.*

Dataset	Source	Service	Purpose
AHN DTM (5 m)	Rijkswaterstaat	PDOK WCS	Terrain elevation
AHN DSM (5 m)	Rijkswaterstaat	PDOK WCS	Building heights
BAG Panden	Kadaster	PDOK WFS	Building footprints
BGT Begroeid terrein	Kadaster	PDOK OGC API	Green areas
BGT Waterdeel	Kadaster	PDOK OGC API	Water bodies

The Digital Terrain Model (DTM, layer `dtm_05m`) and Digital Surface Model (DSM, layer `dsm_05m`) come from the Actueel Hoogtebestand Nederland (AHN) at 5 m resolution. The DTM records bare-earth elevation in metres above the Dutch reference datum NAP; the DSM records the top surface including buildings and vegetation. Both products are based on airborne LiDAR surveys covering the full Netherlands, and are the national standard for hydrological and urban-drainage applications. Preprocessing (nodata handling, buffering, and depression filling) is described in Section 4.2.

Building footprint polygons are retrieved from the BAG (Basisregistratie Adressen en Gebouwen) via PDOK WFS. The BAG is the authoritative national register of Dutch buildings, providing polygon geometry and construction year. Records are deduplicated on the pand (building) identifier.

Two BGT (Basisregistratie Grootschalige Topografie) collections are used. *Begroeid terrein* polygons classified as `groenvoorziening` (managed green space) or `grasland overig` (grassland) with a minimum area of 500 m^2 contribute to the spatially varying Manning's roughness (surface-friction) grid, and to the identification of green retention features during terrain preprocessing. *Waterdeel* polygons (canals, rivers, harbours, and other open water) provide the authoritative classification of drainage sinks (points where surface water leaves the area).

3.4 Geocoding

Individual dwelling units are geocoded from their 6-digit postcode and house number to point coordinates using the BAG geocoding service. Coordinates are stored in the Dutch national coordinate reference system (EPSG:28992, Rijksdriehoekstelsel), yielding a GeoParquet file of 2,305,576 point geometries. How these points are assigned to the simulation grid, and joined to BAG building footprints for depth sampling, is described in Section 4.1.

3.5 Descriptive Statistics

Table 3.2 presents summary statistics for the portfolio.

Table 3.2: *Descriptive statistics of the social housing portfolio.*

Characteristic	**Value**
Total dwelling units	2,305,576
Housing associations represented	265
Housing type (Bouwvorm)	
MGW (apartment)	1,243,324 (53.9%)
EGW (single-family)	892,525 (38.7%)
Studenteneenheid (student housing)	48,488 (2.1%)
Extramuraal (care-linked)	9,590 (0.4%)
Flexwoning (modular)	8,586 (0.4%)
Other (remaining categories)	103,063 (4.5%)
Usable floor area (m^2)	
Mean / Median / interquartile range	78.6 / 77.8 / 61.6–93.8
Construction year	
Mean / Median	1976 / 1977
Pre-1945 / 1945–1990 / Post-1990	8.2% / 63.8% / 27.9%

Apartment buildings (MGW) dominate the portfolio in unit count but are concentrated in a smaller number of multi-storey buildings. Single-family homes (EGW), though fewer, are fully exposed to ground-level flooding and contribute disproportionately to portfolio flood risk relative to their share of units. The construction-year distribution is informative for vulnerability analysis: roughly 72% of units were built before 1990, and the post-war reconstruction period (1945–1990) alone accounts for nearly two-thirds. Older Dutch urban neighbourhoods more frequently feature combined sewer systems, which handle both rainwater and wastewater and have lower peak capacity than the separated systems in post-1990 developments (Stichting RIONED, 2004), and these neighbourhoods are also more likely to have ground floors at or near street level. Together these characteristics are expected to increase pluvial vulnerability, and motivate the structural treatment of housing type and construction year in Chapters 4 and 5.

The geographic distribution reflects Dutch urbanisation patterns: the largest concentrations are in the Randstad (Amsterdam, Rotterdam, The Hague, Utrecht), with secondary clusters in Eindhoven, Tilburg, Groningen, and Enschede. Figure 3.1 shows the spatial distribution of social housing density across the simulation grid.

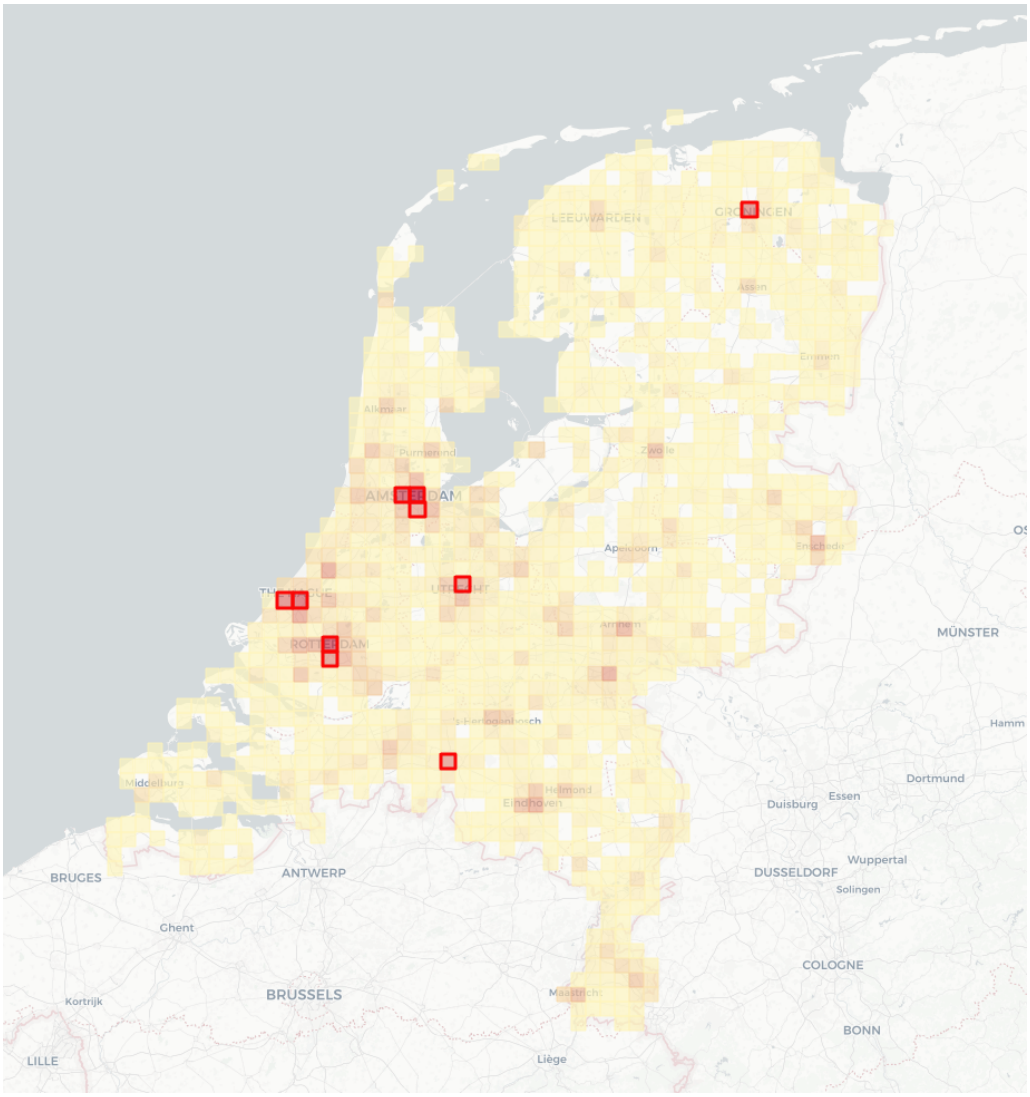


Figure 3.1: *Spatial distribution of social housing units across the 5 km simulation grid. Darker squares indicate higher concentrations of dwelling units; the four Randstad cities dominate the distribution, with each of the top-ranked squares containing tens of thousands of units.*

Chapter 4

Methodology

Chapter overview

In this chapter I develop the method used to answer the central research question: To what extent do extreme pluvial precipitation events pose financial risks to Dutch social housing portfolios, and how can these risks be quantified using flood simulation and Value at Risk estimation? Because no existing dataset records this risk at the building level, I cannot read it off directly and must instead estimate it. The approach taken here is a purpose-built physical simulation pipeline that turns extreme-rainfall scenarios into financial risk measures for the portfolio, while remaining reproducible from public Dutch data alone.

My pipeline follows one chain of reasoning. First, it translates an extreme rainfall event into water depths at building locations. Second, it converts these depths into structural damage per ground-floor dwelling. Third, it sums dwelling-level damages across the portfolio. Finally, it reports the resulting scenario losses as financial risk measures: Value at Risk (VaR), Expected Shortfall (ES), and Expected Annual Damage (EAD), respectively, the size of a rare bad-year loss, the average loss once such a bad year occurs, and the long-run average yearly damage. Figure 4.1 sketches this chain, with concrete numbers from the SSP5–8.5 2050 production run at the $T = 100$ year anchor.

The output of my pipeline is the empirical loss-exceedance curve (LEC) of pluvial flood damage to the Dutch social housing portfolio, the central risk quantity of my thesis. The LEC simply pairs each level of portfolio damage with the probability, in any given year, of damage at least that large. From this curve, I read VaR and ES off at supervisory-relevant confidence levels (90–99.9%) by interpolation. The remainder of this section restates the same chain in more technical detail, organised around four questions: what the estimator produces, why it is built as a bespoke physical simulation, how the construction is organised, and how the chapter is structured.



Figure 4.1: *Five-stage estimation pipeline from rainfall scenario to portfolio risk metric. Numerical examples correspond to the $T = 100$ yr scenario under SSP5-8.5 in 2050. The 99% VaR coincides with the $T = 100$ anchor by construction (AEP = 1%); ES at 99% is computed by composite trapezoidal rule over the simulated tail; EAD is the trapezoidal area under the loss-exceedance curve (LEC) over the simulated 25- to 1,000-year range.*

What

The estimator takes two inputs and returns a portfolio-aggregate damage estimate D_T for each of seven design return periods $T \in \{5, 10, 25, 50, 100, 500, 1{,}000\}$ years. A T-year

return period corresponds to an annual exceedance probability (AEP) of $1/T$: a $T = 100$ event has a 1% chance of being equalled or exceeded in any given year. The seven (loss, AEP) pairs form an empirical LEC, from which VaR_α and ES_α at any $\alpha \in [0.90, 0.999]$ follow by interpolation. I deliver the full national simulation run for the SSP5-8.5 mid-century pathway; the remaining KNMI'23 pathways are parameter swaps within the same pipeline (4.8).

Why a self-made physical simulation

I considered and rejected three alternatives. A purely statistical approach is infeasible: no public database records pluvial damage to housing associations at the building-level granularity required to fit a loss distribution. A national hazard-map approach (e.g. Klimaateffectatlas depths combined with a damage curve) gives a depth per dwelling but does not resolve, at building scale, how rainfall, terrain, and drainage interact at each individual return period. Scale rules out direct use of a commercial platform such as Tygron: 1,097 squares across seven return periods exceeds what per-square licensing supports within the computational resources of a thesis. The remaining route is a bespoke simulator, fast enough to scale, physically realistic enough to resolve rainfall–terrain interaction at building scale, and reproducible without proprietary dependencies. It uses the local-inertial formulation of the shallow water equations (Bates et al., 2010; Neal et al., 2012), a simplified but physically grounded scheme for moving water across a grid that, unlike the simplest routing rules, remembers the water's motion from one step to the next and therefore behaves well on flat terrain, and draws only on public Dutch input data (AHN, BAG, BGT, KNMI). I retain Tygron as the validation benchmark (4.9).

How

My pipeline runs five stages on each 5 km × 5 km simulation square independently, then aggregates to portfolio level. *(i) Terrain and infrastructure preparation*: my pipeline builds a routing surface from AHN DTM/DSM, depression-filled below 10 cm, with sink masks for canals, tunnel ramps, and large flat polder zones, and a spatially varying Manning's roughness grid. *(ii) Rainfall*: my pipeline scales an observed Enschede 2024 storm profile (Tuijn, 2025) to seven KNMI 2-hour return-period totals. *(iii) Hydrodynamic routing*: the solver advances the local-inertial equations at adaptive timesteps, with sewer abstraction (the removal of water by the sewer network) (Stichting RIONED, 2004), infiltration, and roof drainage as source/sink terms. *(iv) Building-level damage*: my pipeline samples the maximum depth in a 3 m buffer around each BAG pand, filters it at a 5 cm threshold, and converts it into a euro loss via the JRC depth–damage curve (Huizinga et al., 2017), a standard relationship between water depth and the fraction of a building's value lost, anchored at a full-replacement value of 1,300 €/m^2 (de Grave &

Juch, 2023). *(v) Portfolio aggregation*: my pipeline sums damage per scenario, forms a piecewise-linear LEC from the seven (loss, AEP) pairs, and computes VaR and ES at confidence levels from 90% to 99%.

Roadmap

Section 4.1 fixes the spatial grid. Sections 4.2–4.5 cover terrain preprocessing, sink masks, the routing surface, and the building and roughness layers. Section 4.6 specifies the rainfall and routing loop. Section 4.7 covers damage. Section 4.8 defines the LEC and the VaR/ES estimators. The chapter closes with the Tygron validation (4.9) and a catalogue of model limitations (4.10).

4.1 Spatial Grid and Square Selection

The 2.3 million social housing units are distributed nationally, making a single high-resolution simulation computationally infeasible. I therefore subdivide the portfolio into 5 km × 5 km squares (25 km^2 each). At a raster resolution of 5 m, each square is a 1,000 × 1,000 grid totalling one million simulation nodes, enough to resolve street widths and building footprints while remaining tractable.

This step delivers the spatial container in which all subsequent operations are performed: every later stage of my pipeline operates inside a single 5 km square at 5 m resolution, with dwelling units already linked to their hosting pand.

4.2 Terrain Data and Preprocessing

My simulation requires two terrain layers: the bare-earth Digital Terrain Model (DTM), which defines the surface over which water flows, and the Digital Surface Model (DSM), which includes buildings and vegetation and is used to raise building cells so they act as physical barriers in the routing scheme. My pipeline retrieves both from the national LiDAR dataset (AHN) at 5 m resolution via the PDOK Web Coverage Service, with a 250 m boundary buffer around each simulation square to avoid edge artefacts in the routing equations.

Micro-noise conditioning. My pipeline adds a negligible uniform random perturbation ($\epsilon_i \sim \mathcal{U}(0, 10^{-4}\,\text{m})$, at most 0.1 mm) to the DTM before routing to prevent static ponding on the perfectly flat LiDAR patches that are common in Dutch polder records; at four orders of magnitude below the depths of interest it cannot affect damage estimates (Appendix B.5).

Depression filling. Morphological reconstruction by erosion (Planchon & Darboux, 2002) removes false sub-decimetre holes in the LiDAR DTM, filling only depressions shal-

lower than $\delta_{\min} = 0.10\,\text{m}$. This preserves deeper features that genuinely store rainwater, such as canals, ditches, and sunken courtyards, while my pipeline excludes cells inside BGT *Begroeid terrein* polygons from filling so that the basin geometry of urban parks and green retention areas is retained. I calibrate the 10 cm threshold against the Tygron benchmark; the calibration and the two failure modes it steers between are detailed in Appendix B.6.

The output of this step is a noise-conditioned, depression-filled DTM and a raw DSM, which are combined in the next section into a single routing surface that lets buildings act as barriers while preserving natural retention features.

4.3 Sink Mask Construction

My pipeline identifies three categories of sink cells and zeroes the water at them after each routing sub-step, representing infrastructure that the surface model cannot physically resolve.

Open water cells (canals, rivers, harbours): my pipeline identifies these from the BGT `waterdeel` layer, rasterises them onto the simulation grid using the all-touched rule (every cell the polygon touches is included), and dilates them by one cell (5 m) to capture sloped banks; without this expansion the routing scheme cannot drain water into the canal fast enough during peak rainfall.

Local depressions are cells lying well below their surrounding neighbourhood (Appendix B.11 gives the exact detection rule), representing underpasses or tunnel ramps. They act as terminal sinks for uphill water that the local-inertial equations cannot drain (the underground storm-drain and tunnel-pumping infrastructure that prevents real-world flooding at these locations is not represented), and additionally force the CFL-adaptive timestep to impractically small values.

Large flat zones: my pipeline detects these automatically as large connected groups of cells sharing an essentially identical elevation (Appendix B.11 gives the detection thresholds). This general approach correctly identifies the Dutch polder landscape's engineered managed-drainage areas (Amsterdam's 0.644 m NAP polder level, Rotterdam's below-sea-level polders, etc.) without per-square calibration.

The result is a binary mask flagging the cells at which water is zeroed after each routing sub-step, capturing the drainage and water-management infrastructure that the surface model cannot explicitly resolve.

Together with the routing surface of the next section, this mask determines where water can pool, where it flows freely, and where it disappears from the simulated system.

4.4 Routing Surface Construction

The routing surface combines the depression-filled DTM with the building-elevated DSM:

$$z_i^{\text{route}} = \begin{cases} \max\left(d_i,\, z_i^{\text{fill}}\right) & i \in \mathcal{B} \\ z_i^{\text{fill}} & \text{otherwise} \end{cases} \tag{4.1}$$

where d_i is the DSM elevation, z_i^{fill} is the noise-conditioned depression-filled DTM elevation, and $\mathcal{B}$ is the set of building cells. The maximum guards against the rare case of the DSM dipping below ground elevation inside a building footprint. Sink cells retain their natural low elevation rather than receiving an artificial placeholder value, so the momentum-driven flow equations naturally direct water toward them. The scheme zeroes only the water depth after each routing sub-step, which lets the discharge arrays Q_x, Q_y continue carrying momentum toward the sink.

Appendix B.2 gives a worked three-cell example, a street cell, a normal building cell, and a rare LiDAR-artefact building cell, illustrating Equation (4.1).

The construction is therefore conservative in a precise sense: when $d_i > z_i^{\text{fill}}$ (the usual case), it preserves the full DSM-derived building height as a barrier; when $d_i \leq z_i^{\text{fill}}$ (a LiDAR artefact), it defaults to the ground elevation and still prevents water ingress at this resolution.

4.5 Building Footprints and Roughness

My pipeline retrieves building footprint polygons from the BAG via PDOK WFS and constructs the building mask by rasterising BAG *pand* polygons onto the simulation grid using pixel-centre membership: a cell is classified as a building cell only if its centre falls inside a building polygon.

Surface roughness. How fast water spreads also depends on what it flows over, rough grass slows it more than smooth asphalt, a friction effect captured by *Manning's roughness* coefficient n. My pipeline constructs a spatially varying roughness grid for non-building ground cells using the BGT green-area mask: paved urban surfaces (streets, parking) receive $n = 0.025$, while green areas receive $n = 0.050$. These values lie within the ranges tabulated for smooth asphalt/concrete and short-grass/dense-urban surfaces in standard open-channel and overland-flow references (Chow, 1959; Manning, 1891). The roughness grid governs the friction term in the momentum equation (B.1).

Roof drainage precomputation. Each building cell drains rainfall to the nearest non-building ground cell, representing gutters and downpipes. My pipeline precomputes

the target cell once per square using the Euclidean distance transform (Maurer et al., 2003) and reuses it at every timestep.

These three artefacts, a building mask, a roughness grid, and a roof-drainage lookup table, close the static-input set required by the rainfall simulation: the next section adds rainfall in time and routes the resulting water across the surface.

4.6 Rainfall Simulation

4.6.1 Scenario rainfall depths and Gumbel extrapolation

I derive target rainfall depths over the 2-hour simulation window from the KNMI 2-hour point-rainfall statistics (Table 2.1), matched to the 120-minute simulation window. For $T \leq 100$ years, I read depths directly from the tabulated quantiles (Koninklijk Nederlands Meteorologisch Instituut, 2018). For $T > 100$ years (relevant for the far tail of the loss distribution), I extrapolate the tabulated quantiles using a Gumbel fit to the KNMI 2-hour quantiles:

$$Q(T) \approx 16.4 + 10.3 \ln T \quad \text{(mm, 2-hour duration)}. \tag{4.2}$$

The extrapolated quantiles (≈ 80 mm at $T = 500$, ≈ 88 mm at $T = 1{,}000$) are treated as indicative stress-test inputs rather than precise point estimates; their role in the LEC tail is discussed in 4.8. Because the simple Gumbel form underestimates the upper tail relative to a full GEV with non-zero shape, these quantiles represent lower-bound estimates.

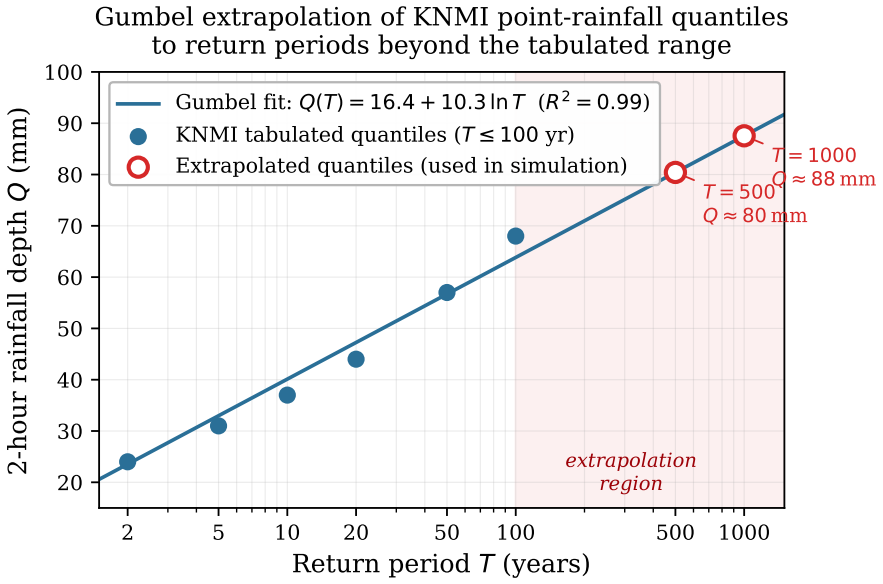


Figure 4.2: *Gumbel Extrapolation*

The thermodynamic basis for this scaling, the Clausius–Clapeyron relation and the

empirically observed "super-CC" rate of roughly 14% per °C for short Dutch convective extremes (Lenderink & van Meijgaard, 2008), is explained in Appendix B.7. I scale baseline depths to climate-changed conditions by the super-Clausius–Clapeyron factor $\phi = 1.14^{\Delta T}$, where ΔT is the local Netherlands warming relative to the 1991–2020 reference period (KNMI, 2023). Table 4.1 lists ΔT and the resulting multiplier for each pathway; the multipliers range from 1.19 (SSP2-4.5, 2050) to 1.78 (SSP5-8.5, 2100).

Table 4.1: *Climate-scaling parameters per KNMI'23 pathway. Local warming ΔT relative to the 1991–2020 reference period gives the super-Clausius–Clapeyron rainfall multiplier $\phi = 1.14^{\Delta T}$, applied to the baseline design-storm depths in Table 4.2.*

Pathway	Horizon	ΔT (°C)	$\phi = 1.14^{\Delta T}$
SSP2-4.5	2050	1.3	1.19
SSP2-4.5	2100	2.5	1.40
SSP5-8.5	2050	1.6	1.24
SSP5-8.5	2100	4.4	1.78

This step therefore produces the seven climate-scaled rainfall depths that anchor the LEC: each depth is one P_{target} value that the design storm of the next subsection scales to.

Table 4.2: *Design storm rainfall totals (mm over 120 minutes) for seven return periods under four climate conditions plus a current-climate baseline. Values for $T > 100$ years (marked †) use Equation (4.2). Climate-scaled values are obtained by multiplying the baseline by $\phi = 1.14^{\Delta T}$.*

		Baseline	SSP2-4.5		SSP5-8.5	
T (yr)	AEP (%)	current (mm)	2050 ($\phi = 1.19$)	2100 ($\phi = 1.40$)	2050 ($\phi = 1.24$)	2100 ($\phi = 1.78$)
5	20.0	31	37	43	38	55
10	10.0	37	44	52	46	66
25	4.0	47	56	66	58	84
50	2.0	57	68	80	71	101
100	1.0	68	81	95	84	121
500 †	0.2	80	95	112	99	142
1 000 †	0.1	88	105	123	109	157

4.6.2 Design storm construction

I derive the temporal structure of the design storm from KNMI 5-minute radar precipitation data (dataset `nl_rdr_data_rfcor_5m`). Rather than imposing a synthetic hyetograph, I extract an observed storm profile from the July 2024 Enschede event (Tuijn, 2025), a recent, well-documented convective storm representative of summer pluvial extremes. A 3×3 neighbourhood maximum around the Pathmos weather station isolates the spatial peak of the convective cell over Enschede, capturing the storm core rather than the regional average. From this time series, a rolling sum identifies the 120-minute window (24

timesteps of 5 minutes) with the highest cumulative rainfall. I normalise this window, including its sharp peak and asymmetric decay, to sum to unity and scale it to each target total P_{target}:

$$p_t^* = \widetilde{p}_t \cdot P_{\text{target}}, \quad \widetilde{p}_t = p_t / \textstyle\sum_\tau p_\tau. \tag{4.3}$$

4.6.3 Simulation loop

At each 5-minute timestep, the model adds rainfall uniformly to all cells and transfers water on building cells to the precomputed nearest ground cell to conserve mass.

Sewer abstraction and infiltration. The model represents the sewer system as a capacity-limited drain following the RIONED Leidraad (Stichting RIONED, 2004), using a baseline capacity of 8 mm/hr and an 8 mm initial buffer. It applies a uniform infiltration rate of 1.5 mm/hr to non-building cells to reflect low urban soil permeability. Appendix B.1 provides the exact mathematical formulations.

Flow routing. The model routes water using the local-inertial formulation of the shallow water equations (Bates et al., 2010). This scheme retains the local acceleration term while dropping the advective inertia, carrying momentum between timesteps. The solver chooses the timestep adaptively using the Courant–Friedrichs–Lewy (CFL) condition to maintain numerical stability. Appendix B.1 details the full momentum and depth-update equations, along with two worked examples tracing the cell-to-cell flux. The model caps water depths at 2 m to prevent isolated deep pits from forcing the CFL timestep to impractically small values.

Figure 4.3 shows this depth field for one example square. Most paved surfaces carry only a shallow sheet of water, while the local-inertial routing concentrates deeper accumulation in local depressions and along connected low-lying street corridors, the spatial pattern that the damage calculation of the next section samples around each building footprint.

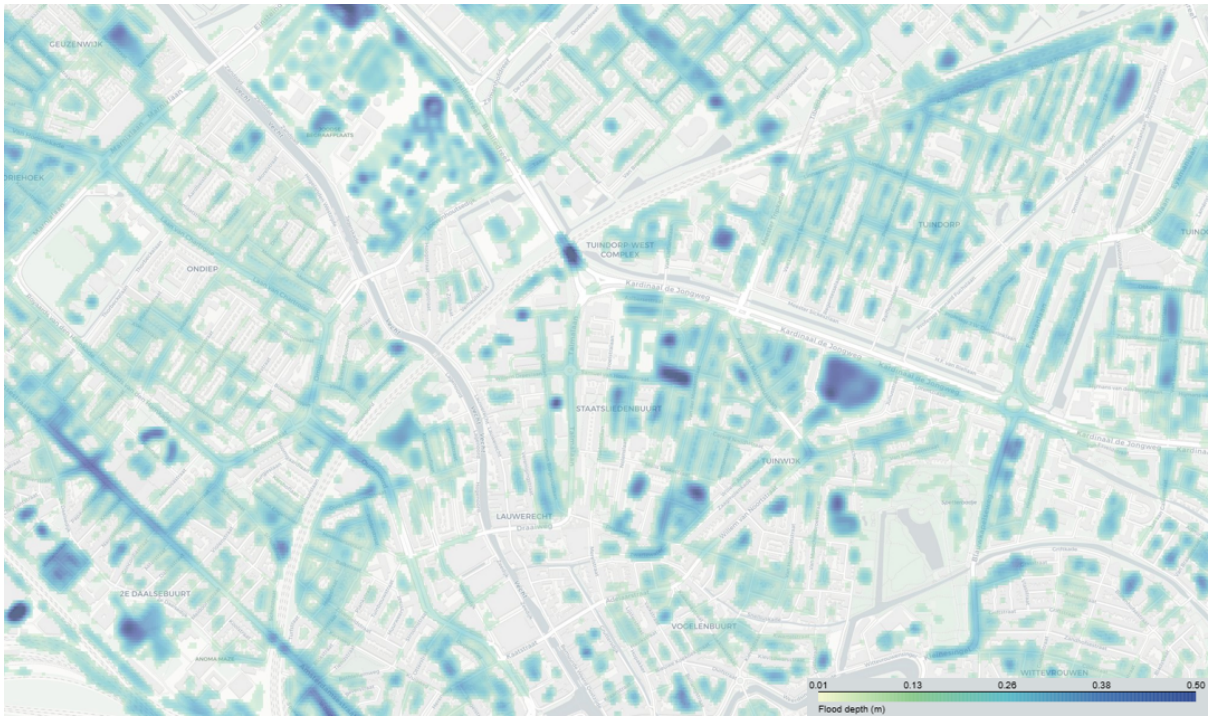


Figure 4.3: *Example output of the simulation loop for a single 5 km simulation square covering part of northern Utrecht (Ondiep, Lauwerecht, Staatsliedenbuurt, Tuindorp, and adjacent neighbourhoods). The colour field is the per-cell maximum water depth w_i^{sample} reached during the 120-minute design storm; lighter cells are shallow surface water and darker cells mark deeper accumulation in local depressions and low-lying street corridors. Building footprints carry no water and appear as gaps in the field, reflecting the building-barrier routing surface of Section 4.4.*

Post-rainfall drainage. After 24 rainfall timesteps, the model executes an additional 5 drainage steps (25 minutes) without new rainfall to let water redistribute from roof-drainage injection points and shallow ponding areas into connected street-level pools. The output of this loop is the per-cell maximum water depth w_i^{sample}, which feeds into the damage calculation of the next section.

4.7 Damage Calculation

4.7.1 Flood depth sampling per building

For each BAG pand, my pipeline samples the maximum flood depth within a 3 m perimeter around the building footprint:

$$d_{\text{pand}} = \max_{i \in \mathcal{A}(\text{pand}, 3\,\text{m})} w_i^{\text{sample}}, \tag{4.4}$$

where $\mathcal{A}(\text{pand}, 3\,\text{m})$ denotes raster cells within 3 m of the polygon boundary. The buffer is needed because at 5 m resolution water in a flooded street is represented in cells adjacent

to the building; 3 m corresponds roughly to the nearest half-cell around the footprint. Taking the maximum (rather than the mean) reflects that flood damage is driven by the peak water level at any point of entry.

4.7.2 Floor-level proxy

The Aw dataset records individual dwelling units (*verblijfsobjecten*) without explicit floor-level information, so I derive a proxy from the *Bouwvorm* (building-form) classification: my pipeline treats single-family houses (EGW) and flex units as ground-floor exposed; for multi-family buildings (MGW), it flags only the first registered unit per building (*pand*) as ground-floor and assigns upper-floor units zero damage. This is a conservative lower bound for MGW damage, since some buildings contain multiple ground-floor units.

4.7.3 Inundation threshold

A central parameter is the *inundation threshold* d_{th}: the water depth below which no structural damage is assumed to occur, representing the doorstep or finished-floor height that water must overcome to enter the building. Two reference points give a wide range, the Dutch building decree's 20 mm at the main entrance (Ministerie van Binnenlandse Zaken en Koninkrijksrelaties, 2022) and the *Waterschadeschatter*'s implicit 15 cm above terrain (Nelen & Schuurmans, 2024), but neither is fully appropriate for a mixed social housing portfolio (Appendix B.8).

My thesis adopts $d_{\text{th}} = 5\,\text{cm}$ as the baseline. It sits between the two reference points and is intended to represent a plausible weighted average across a mixed social housing portfolio. My pipeline sets damage to zero below this threshold:

$$D_i = \begin{cases} D_i^{\text{full}} & d_i \geq d_{\text{th}} \\ 0 & d_i < d_{\text{th}} \end{cases} \tag{4.5}$$

where D_i^{full} is the unconditional structural damage from the depth–damage function. Chapter 5 presents a sensitivity analysis varying $d_{\text{th}} \in \{5, 10, 15\}\,\text{cm}$.

4.7.4 Depth–damage function

I quantify flood damage to building structures using the European residential depth–damage curve from Huizinga et al. (2017), combined with maximum damage values from the Dutch national standard SSM2023 (de Grave & Juch, 2023). I include only direct damage to building structure and exclude household contents damage, since it falls to the tenant rather than the housing association.

The SSM2023 standard specifies a maximum structural damage of 1,295 €/m^2 for residential buildings (2022 prices, ex-VAT) (de Grave & Juch, 2023), closely consistent with SSM2017's 1,000 €/m^2 (2011 prices) adjusted by CBS CPI to 2024 (1,300 €/m^2). The simulation uses 1,300 €/m^2. Table 4.3 gives the fractional curve from Huizinga et al. (2017), which I interpolate linearly. I selected the JRC curve over the SSM2023 built-in functions because the latter produce damage fractions that are too low at the shallow depths characteristic of pluvial flooding. A recent Deltares evaluation (de Moel et al., 2024) found that SSM2023 very likely underestimates structural damage to single-family homes for water depths of 0–2 m, and explicitly advises the financial sector against using SSM2023's built-in structural functions for per-building risk calculations.

Table 4.3: *European residential depth–damage curve (Huizinga et al., 2017).*

Depth (m)	0.0	0.5	1.0	1.5	2.0	3.0	4.0	5.0
Damage fraction	0.00	0.25	0.40	0.50	0.60	0.75	0.85	0.95

The damage per dwelling unit u is therefore

$$D_u = \begin{cases} 0 & d_{\text{pand}(u)} < 0.05\,\text{m} \\ f(d_{\text{pand}(u)}) \cdot 1{,}300 \cdot A_u & d_{\text{pand}(u)} \geq 0.05\,\text{m} \end{cases} \cdot \mathbf{1}[u \in \text{ground floor}], \tag{4.6}$$

where A_u is the usable floor area (in m^2) and $f(\cdot)$ is the JRC damage fraction. For a representative 80 m^2 ground-floor dwelling, this yields damage estimates from €2,600 at 5 cm depth to €26,000 at 50 cm.

Repeating this calculation for every ground-floor unit in the portfolio and summing yields the portfolio damage $D(T)$ for a single return-period scenario. Across the seven simulated return periods, this produces the seven (loss, AEP) anchors that the VaR estimation of the next section converts into a loss-exceedance curve.

4.8 Value at Risk Estimation

In plain terms, this section converts the seven simulated scenarios into one curve and then reads two numbers off it. The curve is the loss-exceedance curve (LEC): it pairs every level of portfolio damage with the annual probability that damage of at least that size occurs, so it captures *how bad* an event can be together with *how often* it happens. Value at Risk (VaR) is a single point on this curve, the damage level that is exceeded only in a chosen small fraction of years. It answers the question a supervisor or treasurer asks when sizing a buffer: *how large is a rare-but-plausible shock that the portfolio must be able to absorb without breaching its interest-coverage ratio?* Expected Shortfall (ES) looks past that point and averages the curve beyond it: *given* that such a rare shock

occurs, how severe is it on average. The distinction matters because two portfolios can share the same VaR yet differ sharply in how catastrophic their worst outcomes are. VaR is the primary metric here because it maps directly onto the stress-test logic of Aw/WSW supervision; ES is reported alongside it as a complementary view of tail severity. The remainder of this section makes each step precise: first the LEC, then VaR, then ES.

The simulation produces seven (loss, AEP) pairs per simulation square. Sorting by damage so that $D_1 < D_2 < \cdots < D_7$, the empirical loss-exceedance curve is the piecewise-linear function $\widehat{G} : [D_1, D_7] \to [\pi_7, \pi_1]$ connecting the seven pairs:

$$\widehat{G}(d) = \pi_k + \frac{\pi_{k+1} - \pi_k}{D_{k+1} - D_k}(d - D_k), \quad d \in [D_k, D_{k+1}], \tag{4.7}$$

for $k = 1, \ldots, 6$. I choose linear interpolation for simplicity, imposing no functional form on the LEC between the seven anchors; an alternative $(\ln T, D)$ interpolation agrees with it to within 2 % (Appendix B.4), and Appendix B.3 gives a worked example of reading the curve.

For confidence level $\alpha \in [0.90, 0.99]$, Value at Risk is the α-quantile of the damage distribution:

$$\mathrm{VaR}_\alpha = \inf\{d \in R : \widehat{G}(d) \leq 1 - \alpha\}. \tag{4.8}$$

Because $\widehat{G}$ is continuous and strictly decreasing, the infimum lies on the interior of the curve and can be evaluated by linear interpolation. I estimate VaR at confidence levels from 90% to 99% in steps of 1%, tracing the shape of the tail rather than committing to a single threshold. Appendix B.3 gives worked examples for $\mathrm{VaR}_{0.95}$ ($\approx$ €344 M) and $\mathrm{VaR}_{0.99}$ ($=$ €457 M).

Tail handling beyond $T = 1{,}000$ years. The highest simulated scenario corresponds to $T = 1{,}000$ years ($\pi_7 = 0.001$), so I restrict reported VaR values to $\alpha \leq 0.999$ and interpret them as the in-sample tail of the seven-scenario distribution. Appendix B.9 sets out the treatments considered for the deeper tail ($\alpha > 0.999$); I report one as a robustness check in Chapter 5 and identify another as a methodological extension in Chapter 6.

Pathway-conditional interpretation. Within the computational resources available for my thesis I deliver the full 1,097-square national simulation run for the SSP5-8.5 mid-century pathway only, the most policy-relevant case in the KNMI'23 ensemble; generalisation to the other pathways requires no new methodology (5.1). The VaR estimates are therefore *pathway-conditional*: they assume the return-period frequencies are unchanged while the rainfall at each return period shifts to its future-climate value, rather than being single climate-averaged estimates (which would also require modelling the frequency shift under a non-stationary EVT).

Expected Shortfall. Expected Shortfall is a *coherent* risk measure (Artzner et al., 1999) and the standard tail metric adopted by banking supervisors after the Fundamental Review of the Trading Book (FRTB). It is not part of the core VaR framework, but $\widehat{G}$ already contains the information needed to compute it:

$$\mathrm{ES}_\alpha = E[D \mid D \geq \mathrm{VaR}_\alpha] = \frac{1}{1-\alpha}\int_\alpha^1 \mathrm{VaR}_u \, \mathrm{d}u. \tag{4.9}$$

Because $\widehat{G}$ is piecewise linear, I evaluate the integral by the composite trapezoidal rule, truncated at the upper edge of the simulated range ($u_m = 0.999$); the simulation does not constrain the deep tail beyond $T = 1{,}000$ years, which I therefore exclude from the reported figure. Appendix B.10 gives the exact truncated estimator and its justification. Appendix B.3 provides a fully worked numerical example of this calculation. I present ES alongside VaR as a complementary view of tail severity rather than as a co-equal headline metric: I formulate the primary conclusions in terms of VaR.

This completes the estimation pipeline. From the spatial container of 4.1 through to the VaR and ES of this section, every step transforms one well-defined object into the next: terrain into routing surface, routing surface into depth field, depth field into per-dwelling damage, per-dwelling damage into seven (loss, AEP) anchors, and the anchors into the loss-exceedance curve from which VaR and ES are read.

4.9 Model Validation

I validate the simulation against Tygron, a professional hydrodynamic platform that solves the full 2D shallow water equations with a coupled sewer network model. The validation covers five simulation squares (indices 0, 4, 5, 7, 8) across the seven return-period scenarios under SSP5-8.5 2050, yielding 35 paired runs. These squares correspond to high-density social-housing areas in Amsterdam, Den Haag, Groningen, Tilburg, and Utrecht; the set therefore combines three Randstad cities with a northern (Groningen) and a southern (Tilburg) urban centre, so that the validation is not confined to a single regional terrain or drainage regime. I apply the same damage function to both engines, so any difference reflects differences in simulated flood depths rather than damage methodology. The Tygron project matches the local-inertial simulation in rainfall time series, Manning roughness, and the 120-minute storm window followed by 25-minute drainage. I run Tygron in its operational configuration with its spatially explicit sewer network, rather than the simplified uniform 8 mm/hr capacity of this study. I do not treat Tygron's sewer network as ground truth either, so I report residual bias as a calibration ratio that bounds the VaR and ES estimates upward (rather than as an absolute error), and 4.9.3 establishes that the bias is not driven by sewer parameterisation in the first place.

4.9.1 Aggregate validation statistics

Across the 35 paired runs, the ratio of Tygron-simulated portfolio damage to local-inertial-simulated portfolio damage ranges from 0.30 to 0.89 (mean 0.64, SD 0.18). The local-inertial model produces higher portfolio damage than Tygron in all 35 runs, so the direction of bias is conservative throughout. Grid-level RMSE on the depth field ranges from 0.051 m to 0.139 m (mean 0.091 m), and mean depth bias is mildly positive at low rainfall and turns slightly negative at the highest return periods, consistent with the convergence of the two models at the upper anchors. Pearson correlation (agreement on the absolute damage levels) between the two models' estimates is 0.918, and Spearman rank correlation (agreement on the ordering of scenarios from least to most damaging) is 0.905. Computed separately for each square, the within-square Spearman correlation is $\rho = 1.00$ in every case: the two models agree perfectly on which scenarios produce the largest damages. The modest cross-square drop in aggregate Spearman ($\rho = 0.905$) reflects differences in the absolute scaling of damage between squares, not disagreement on scenario ranking.

4.9.2 Cross-scenario damage ratios

The damage ratio rises with rainfall intensity in every square, with inter-model agreement tightest at the highest return periods (Table 4.4); the only minor reversal is in square 5 between $T = 500$ and $T = 1{,}000$ (0.86 to 0.83), which does not change the qualitative pattern. At $T = 5$ (38 mm), the ratio ranges from 0.30 (square 7) to 0.63 (square 0), a spread of 0.33; at $T = 1{,}000$ (109 mm), the ratio ranges from 0.78 (square 4) to 0.89 (square 8), a spread of 0.11. The physical interpretation: at low totals a large fraction of precipitation is absorbed by parameterised subsurface processes (sewer abstraction, infiltration) whose configuration differs between the two models; at high totals these absorption mechanisms are exceeded regardless of parameter choice, so both models route similar overland volumes governed primarily by topography. This is favourable for the VaR/ES estimates of Chapter 5: confidence levels of 95% and above draw on the upper end of the LEC, where the mean damage ratio across the five validation squares rises to 0.84 at $T = 1{,}000$, with cross-square spread of 0.11.

Table 4.4: *Damage ratio (Tygron / local-inertial) by square and scenario, SSP5-8.5 2050. Values below 1.0 indicate that the local-inertial model produces higher portfolio damage than Tygron.*

Square	T=5 (38 mm)	T=10 (46 mm)	T=25 (58 mm)	T=50 (71 mm)	T=100 (84 mm)	T=500 (99 mm)	T=1000 (109 mm)
0	0.63	0.67	0.72	0.76	0.79	0.82	0.84
4	0.33	0.42	0.51	0.58	0.65	0.72	0.78
5	0.33	0.43	0.54	0.64	0.73	0.86	0.83
7	0.30	0.38	0.49	0.60	0.68	0.79	0.87
8	0.43	0.50	0.62	0.71	0.79	0.86	0.89

4.9.3 Sensitivity to sewer parameterisation

To test whether the residual bias is attributable to the uniform sewer abstraction, I reran the model using upper-bound parameters from the RIONED design range ($q_{\text{sewer}} = 11\,\text{mm/hr}$, buffer $= 10\,\text{mm}$). As detailed in Appendix D, this 37% increase in capacity shifted the cross-run mean damage ratio by only 4 percentage points. This falsifies the hypothesis that sewer parameterisation is the dominant bias source and confirms that the residual bias stems instead from structural model-scope differences, such as spatial sewer heterogeneity and the floor-level proxy.

4.9.4 Interpretation for VaR and ES

Three implications follow. First, the local-inertial model produces systematically higher damage estimates than Tygron across all 35 paired runs, and the sewer-sensitivity analysis establishes that this bias reflects intrinsic model-scope differences rather than parameter choice. The resulting VaR and ES estimates should therefore be read as upper bounds relative to the Tygron benchmark, a direction of bias preferable to underestimation for capital planning and Aw supervisory review. Second, the bias is strongly rainfall-intensity-dependent, with inter-model agreement tightest at the highest return periods (cross-square damage ratio rising into the 0.78–0.89 range at $T = 1{,}000$, mean 0.84). This is favourable for the VaR and ES framework because confidence levels of 95% and above are populated by these extreme events: the implied upper-bound multiplier $1/r$ that relates the local-inertial estimate to a Tygron-calibrated central value shrinks from approximately $1/0.55 \approx 1.8$ at the low-rainfall end to $1/0.84 \approx 1.2$ at the upper anchors. Third, within-square Spearman rank correlation is $\rho = 1.00$ in every square, so the scenario-based risk *ranking* is robust to the choice of hydrodynamic engine: the disagreement is a calibration effect on absolute damage levels, not a structural disagreement on which events drive tail risk. The full 35-run per-metric table is in Appendix C.

4.10 Model Limitations

The validation establishes that the local-inertial model produces conservative damage estimates whose bias is not attributable to sewer capacity within the RIONED design range. The remainder of this section catalogues the limitations that shape the residual bias. These cover both choices internal to the hydrodynamic model (routing scheme, sewer abstraction, building-depth sampling) and choices in the surrounding methodology (floor-level proxy, JRC damage curve, tail extrapolation); together they account for the directional conservatism documented in 4.9.

- **Routing scheme.** The local-inertial formulation neglects sub-grid channel flow, backwater effects, and advective inertia. For shallow depths and moderate velocities these omissions are expected to be secondary relative to the sewer and floor-level limitations (Bates et al., 2010).

- **Sewer model.** The uniform 8 mm/hr capacity does not resolve spatial variation, in particular the difference between modern separated systems and pre-war combined systems. The sensitivity analysis (4.9.3) shows that absolute-rate changes have at most four-percentage-point effects on the damage ratio; the missing spatial heterogeneity is the dominant disagreement source, and is intrinsic to the simplified surface-routing scheme.

- **Floor-level proxy and damage curve uncertainty.** The first-unit-per-pand proxy for MGW ground-floor units (Section 4.7.2) is the leading candidate explanation for the residual MGW-Randstad bias. The JRC European depth–damage curve (Huizinga et al., 2017) is a continent-wide average that may not perfectly reflect Dutch building practices; the recent Deltares evaluation (de Moel et al., 2024) confirms that steeper damage curves than SSM2023's defaults are warranted at shallow depths.

- **Tail extrapolation.** The two stress-test scenarios ($T = 500, 1{,}000$) use the simple Gumbel special case of GEV, which underestimates the upper tail relative to a non-zero shape; the resulting VaR values should be regarded as lower-bound estimates beyond $T = 100$.

Chapter 5

Results

Chapter overview

This chapter presents the numerical results of the SSP5-8.5 2050 simulation run. The chapter has four purposes. First, it quantifies the portfolio-level damage generated by the seven simulated return-period events and translates these losses into VaR and ES. Second, it describes the spatial distribution of damage at the *gemeente*, *provincie*, and Woondealregio levels. Third, it sets the modelled losses against the operating-cashflow base that has to absorb them, both at sector level and per Woondealregio. Fourth, it tests how strongly the results depend on the assumed inundation threshold. The local-inertial model is directionally conservative relative to Tygron, so the absolute damage figures represent upper-bound estimates within the modelling framework, while the spatial ranking and relative scenario ordering are more robust.

Headline finding. One result dominates the remainder of this chapter and answers sub-question 4. Across all simulated return periods, additional dwellings crossing the 5 cm damage threshold primarily drive portfolio loss growth, rather than deeper inundation of dwellings already affected. This *threshold-crossing mechanism* recurs throughout the chapter: it shapes the empirical loss-exceedance curve (Section 5.3), is visible directly in the building-level depth distributions (Section 5.9), and explains the strong sensitivity of VaR to the assumed threshold (Section 5.10); Section 5.11 consolidates the evidence under SQ4. It also carries the practical implication taken up in Chapter 6: relatively modest measures that prevent shallow water entry can produce disproportionately large reductions in portfolio VaR.

5.1 Scope of the SSP5-8.5 2050 simulation run

The simulation run covers the full national set of 1,097 priority squares under the SSP5-8.5 2050 pathway. Across these squares, the spatial damage calculation includes 2,305,576

dwelling units. For each square, the pipeline simulates the same seven 120-minute rainfall scenarios: $T = 5$, 10, 25, 50, 100, 500, and 1,000 years. Under SSP5-8.5 in 2050 these correspond to rainfall totals of 38, 46, 58, 71, 84, 99, and 109 mm respectively. The resulting damage estimates are pathway-conditional: they represent the loss distribution implied by the SSP5-8.5 2050 rainfall intensities, not an unconditional forecast of future annual losses.

5.2 Portfolio-wide damage by return period

Table 5.1 reports the portfolio-wide damage estimates by return period. The total damage increases from EUR 246.3 million for the 5-year event to EUR 525.2 million for the 1,000-year event. The 100-year event, corresponding to 84 mm of rainfall, produces EUR 457.1 million in direct structural damage. Since the annual exceedance probability of the 100-year event is 1%, this scenario is also the 99% VaR anchor of the simulated loss distribution.

Table 5.1: *Portfolio-wide direct structural damage by return period, SSP5-8.5 2050.*

T (years)	AEP (%)	Rainfall (mm)	Total damage (EUR M)	Damaged units (#)	Share damaged (%)	Mean damage (EUR/damaged unit)
5	20.0	38	246.3	37,147	1.62	6,631
10	10.0	46	296.0	43,377	1.89	6,825
25	4.0	58	354.1	50,775	2.21	6,973
50	2.0	71	410.2	58,265	2.54	7,040
100	1.0	84	457.1	64,053	2.79	7,137
500	0.2	99	500.4	69,089	3.01	7,242
1,000	0.1	109	525.2	71,978	3.14	7,297

The most important pattern is that the extensive margin mainly drives portfolio loss growth (an increase in damage from additional dwellings crossing the 5 cm barrier rather than from greater maximum depth per house). From $T = 5$ to $T = 1{,}000$, total damage increases by a factor of 2.13. Over the same interval, the number of damaged units increases by a factor of 1.94, while mean damage per damaged unit increases by only a factor of 1.10. In other words, higher rainfall totals mainly cause additional dwellings to cross the 5 cm damage threshold; they do not produce a proportional increase in damage severity among dwellings that are already affected. This is the *threshold-crossing mechanism*: higher rainfall mainly pushes additional ground-floor dwellings above the damage threshold rather than deepening the inundation of those already affected. This mechanism also drives three further results presented later in the chapter: the flattening of the loss-exceedance curve beyond $T = 100$ (Section 5.3), the building-level depth distributions (Section 5.9), and the strong VaR response to the threshold assumption (Section 5.10).

Table 5.2: *Growth decomposition relative to the 5-year event.*

T (years)	Rainfall (mm)	Damaged units (#)	Growth in damaged units ($\times T5$)	Mean damage (EUR/unit)	Growth in total damage ($\times T5$)
5	38	37,147	1.00	6,631	1.00
10	46	43,377	1.17	6,825	1.20
25	58	50,775	1.37	6,973	1.44
50	71	58,265	1.57	7,040	1.67
100	84	64,053	1.72	7,137	1.86
500	99	69,089	1.86	7,242	2.03
1,000	109	71,978	1.94	7,297	2.13

This result is relevant for adaptation policy. Measures that prevent shallow water from entering dwellings, such as local threshold protection, door barriers, adjusted ground-floor detailing, or street-level drainage improvements, may reduce portfolio loss materially even if they do not eliminate pluvial flooding in the surrounding public space. The reason is that much of the simulated loss arises at the transition from non-damaged to damaged status.

5.3 Loss-exceedance curve, VaR, and Expected Shortfall

The seven return-period losses define the empirical loss-exceedance curve introduced in Section 4.8. Figure 5.1 shows the resulting curve and the VaR/ES estimates on the 90–99% confidence grid. The 90%, 99%, and 99.9% VaR values coincide with simulated scenario anchors. The 95% VaR lies between the 10-year and 25-year anchors, and I obtain it by interpolation in the piecewise-linear (D, π) LEC.

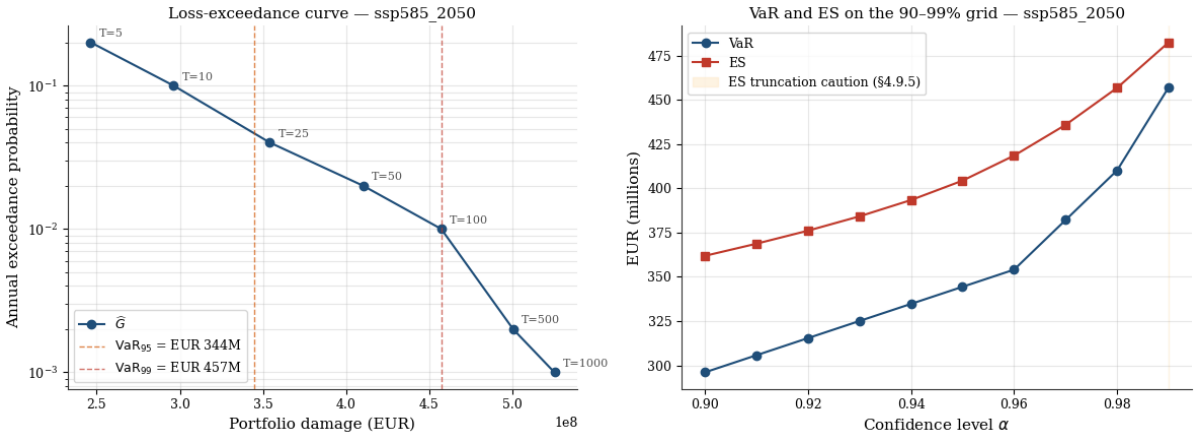


Figure 5.1: *Loss-exceedance curve and VaR/ES estimates for SSP5-8.5 2050. The left panel shows the simulated loss-exceedance curve. The right panel reports VaR and truncated ES across the 90–99% confidence grid.*

Table 5.3 reports the main risk metrics. The 99% VaR is EUR 457.1 million and

the 99.9% VaR is EUR 525.2 million. The corresponding truncated ES at 99% is approximately EUR 482.5 million, calculated over the simulated interval from $\alpha = 0.99$ to $u_m = 0.999$. I therefore interpret ES as the average severity within the simulated tail, not as an extrapolation beyond the 1,000-year event.

Table 5.3: *VaR and truncated ES of direct structural portfolio damage, SSP5-8.5 2050.*

Confidence level α	Equivalent AEP	VaR (EUR M)	Truncated ES (EUR M)
90.0%	10.0%	296.0	361.9
95.0%	5.0%	344.4	404.4
99.0%	1.0%	457.1	482.5
99.9%	0.1%	525.2	–

Two conclusions follow. First, the simulated tail is material: under the baseline 5 cm threshold, the 99% loss approaches half a billion euros. Second, the curve becomes flatter beyond the 100-year event. Damage increases by EUR 43.3 million between the 100-year and 500-year events and by EUR 24.8 million between the 500-year and 1,000-year events, compared with EUR 210.8 million between the 5-year and 100-year events. This shape is the threshold-crossing mechanism showing up in the tail of the LEC: by the time rainfall reaches the 100-year total of 84 mm, it has already pushed most ground-floor dwellings exposed at this resolution across the 5 cm boundary. Further rainfall therefore produces only modest incremental damage, by widening the affected set rather than by deepening existing inundation.

The loss-exceedance curve treats each return period as a single event that falls uniformly across the entire portfolio: the T-year loss is the sum of every square's T-year damage, and that sum receives the annual exceedance probability $\pi = 1/T$. This embeds an assumption of perfect positive (*comonotonic*) spatial dependence between square-level losses: in plain terms, it presumes that all 1,097 squares realise their T-year rainfall in the same year. For a country-scale portfolio this is the most conservative dependence structure available: under comonotonicity the aggregate VaR and ES are additive across squares and attain their no-diversification maximum (McNeil et al., 2015), so the figures in Table 5.3 represent an *upper bound* on portfolio tail risk. The assumption overstates dependence on physical grounds, because the dominant pluvial extremes in the Netherlands are convective summer storms with footprints of the order of tens of square kilometres; squares separated by more than the storm's decorrelation length, the distance beyond which two locations rarely flood in the same storm, rarely realise their extreme rainfall in the same event, so the true portfolio loss benefits from a spatial diversification that the uniform-event curve ignores.

To bound the magnitude of this effect, I computed the opposite extreme by random sampling (a Monte Carlo simulation). Each square draws its annual return period independently from the same piecewise-linear marginal (D, π) curve used to construct the

LEC, and the simulation sums the square-level losses into a portfolio annual loss over 10^6 simulated years. This independence case maximises spatial diversification and yields a *lower bound*. The two bounds are far apart in the tail. At the 99% level, portfolio VaR ranges from EUR 117 million under independence to EUR 457 million under comonotonicity, and ES from EUR 125 million to EUR 483 million; at the 99.9% level VaR ranges from EUR 136 million to EUR 525 million. This band, shown in Figure 5.2, is wider than the inundation-threshold sensitivity of Section 5.10 and the tail-extrapolation spread, which makes the spatial-dependence assumption the single largest source of uncertainty in the reported tail metrics. The realistic value lies inside the band, and for a portfolio dispersed across the whole country it lies closer to the lower bound than to the upper one, because no single convective event covers the national footprint at T-year intensity. Two qualifications keep this in perspective. First, the lower bound is itself conservative as a floor: complete independence between neighbouring squares is as unphysical as complete dependence, and a realistic decorrelation length would place the estimate above EUR 117 million. Second, even under maximal diversification the simulated tail remains material, with a 99% VaR of EUR 117 million and a 99% ES of EUR 125 million, so the qualitative conclusion that pluvial risk is financially non-trivial survives even the most optimistic dependence assumption.

I do not model the intermediate case explicitly. A spatial-dependence model, for instance a copula coupling the square-level losses with a correlation that decays beyond the storm decorrelation length, would narrow the EUR 117–457 million band to a central estimate, but calibrating such a structure requires data on the joint occurrence of rainfall extremes across squares that lies beyond the present scope. I retain the distribution-free bracket because the comonotonic bound holds under any dependence structure (McNeil et al., 2015); I leave an explicit copula model as the natural refinement for future work (Section 6.6).

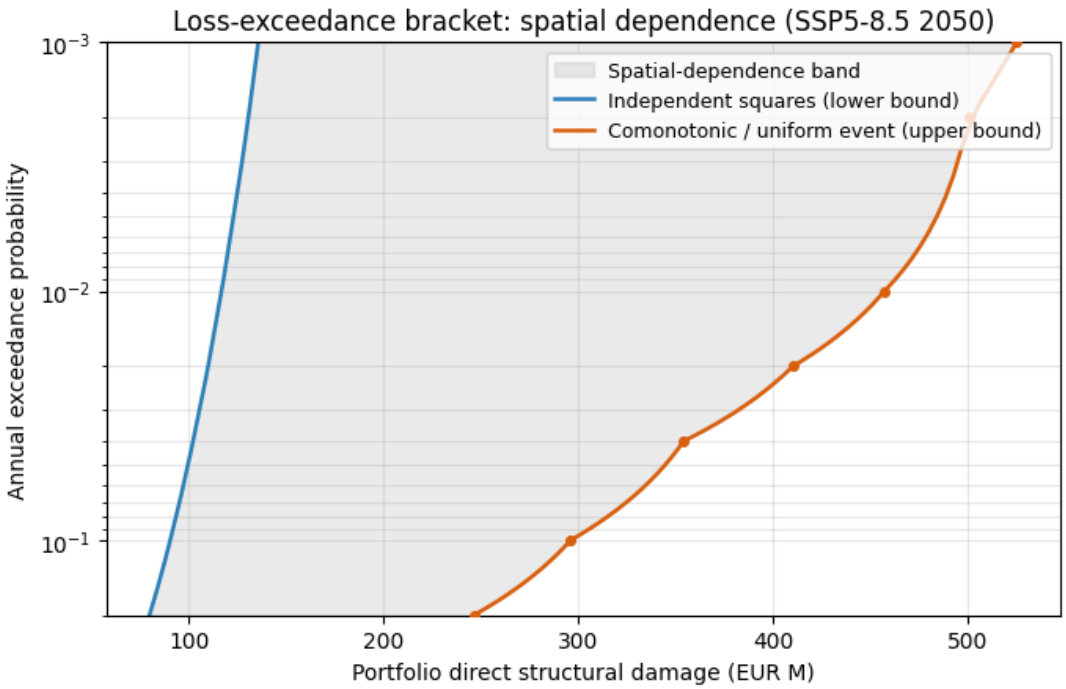


Figure 5.2: *Spatial-dependence bracket for the SSP5-8.5 2050 loss-exceedance curve. My thesis LEC of Figure 5.1 assumes a spatially uniform event and coincides with the upper (comonotonic) bound; the lower bound assumes square-level losses are mutually independent. The shaded band is the range of portfolio VaR attributable to the spatial-dependence assumption alone. Both curves share identical per-square marginal (D, π) distributions and differ only in the assumed spatial dependence.*

Unlike the tail metrics, Expected Annual Damage is unaffected by this assumption. Because EAD is an expectation, it is invariant to the spatial dependence structure: a Monte Carlo drawing each square's return period independently reproduces, to within rounding, the same EAD as the comonotonic trapezoidal integration over the identical return-period range. The spatial-dependence bound therefore qualifies the VaR and ES estimates only; the EAD reported in Section 5.4 is unchanged.

5.4 Expected Annual Damage

Expected Annual Damage (EAD) summarises the annualised financial burden implied by the LEC. In my thesis, I calculate EAD over the simulated exceedance-probability range from the *25-year* to the 1,000-year event. I deliberately exclude the frequent end of the curve. The cross-scenario validation in Section 4.9 shows that the local-inertial model overestimates damage most strongly at low rainfall totals, where the damage ratio relative to Tygron falls to 0.30–0.63 at $T = 5$, against 0.78–0.89 at $T = 1{,}000$. Because EAD weights the most frequent events most heavily, retaining the $T = 5$ and $T = 10$ anchors would let exactly the least-reliable part of the curve dominate the annualised figure; restricting the integration to $T \geq 25$ removes this band. Integrating the remaining

scenario damages by trapezoidal decomposition gives a national scenario-based EAD of approximately EUR 16.8 million per year. This is not a full climatological EAD: it excludes both rainfall more frequent than the 25-year event and the deep tail beyond the 1,000-year event, and it represents the annualised damage contribution of the 25- to 1,000-year band rather than a complete expected-loss figure. Because the excluded $T = 5$–25 band would have added roughly EUR 47 million per year at face value, the restriction is deliberately conservative: it trades the inclusion of a large but poorly-validated frequent-event contribution for an annualised figure drawn only from the return periods where the model agrees most closely with the Tygron benchmark.

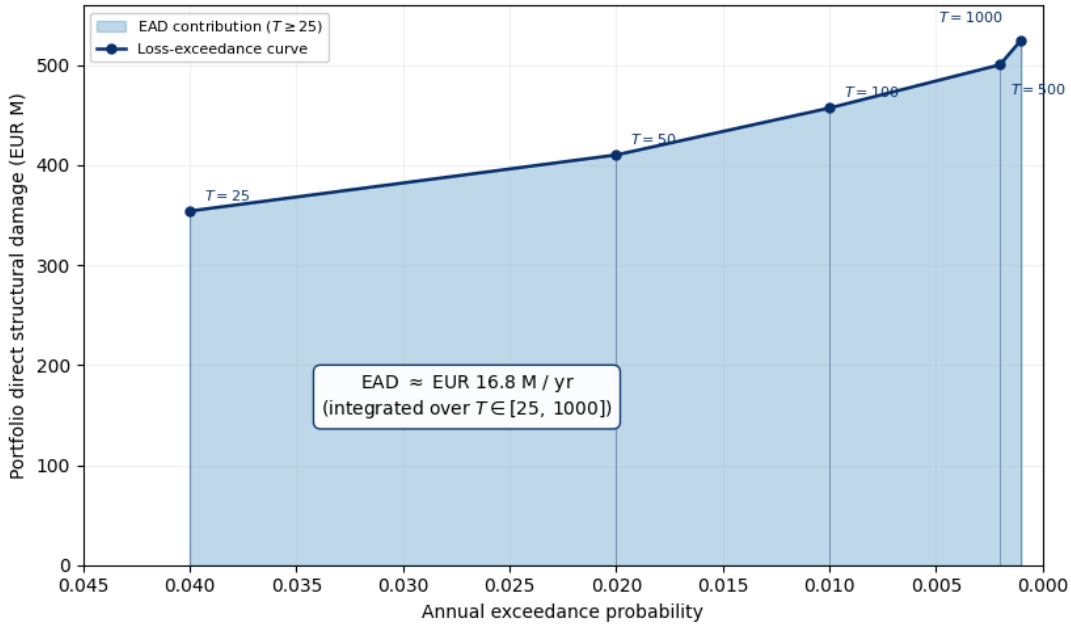


Figure 5.3: *Trapezoidal decomposition used to derive Expected Annual Damage from the simulated loss-exceedance curve. The EAD reported in this chapter is the annualised contribution over the simulated 25- to 1,000-year event range; the $T = 5$ and $T = 10$ anchors are excluded because the validation shows the model is least reliable at these low rainfall totals.*

The $T \geq 25$ restriction is one of two distinct choices embedded in the headline figure; the other is whether to apply the validation calibration before integrating. The two choices are separable, and Table 5.4 reports EAD under all four combinations. The upper-right cell is the headline figure of approximately EUR 16.8 million per year. Reading across the upper row, integrating the uncalibrated curve over the full $T \geq 5$ range gives approximately EUR 63.5 million per year, of which the $T = 5$–25 band alone contributes roughly EUR 47 million. This large contribution is a direct consequence of the annualised weighting: because the trapezoidal integral weights each band by its width $\Delta\pi$ in exceedance-probability space, and $\Delta\pi$ is two orders of magnitude larger between the frequent anchors than between the rare ones, the most frequent events dominate the figure. Excluding them is therefore not a marginal adjustment but the single largest

driver of the reported EAD.

Table 5.4: *Expected Annual Damage under the two integration-range and calibration choices, SSP5-8.5 2050 (EUR million per year).*

	$T \geq 5$ (full range)	$T \geq 25$ (validated range)
Uncalibrated (model damage)	63.5	16.8
Validation-calibrated (Tygron-adjusted)	33.6	11.1

The validation-calibrated column applies the cross-square mean damage ratio to each return-period loss before integration, so that it removes the known overestimation at low rainfall totals rather than accepting it at face value or excluding it wholesale. This is the most complete use of the validation data: it retains the heavily-weighted frequent events, which are, after all, simulated outputs rather than missing data, while correcting them for the bias the validation quantifies. On this basis the calibrated full-range EAD is approximately EUR 34 million per year, of which roughly EUR 22 million still originates in the $T = 5$–25 band even after bias correction. The calibrated figure represents a central estimate, bracketed by the uncalibrated full-range EAD (≈EUR 63.5 million, an upper bound that lets the low-return-period overestimation through unadjusted) and the uncalibrated validated-range EAD (EUR 16.8 million, a conservative lower anchor that discards the frequent band entirely).

Two qualifications bound the interpretation of the calibrated estimate. First, calibration corrects the central tendency of the low-return-period bias but not its dispersion: the $T = 5$ damage ratio varies across the validation squares from 0.30 to 0.63, so the single mean ratio applied to the frequent band carries a wider relative uncertainty than the ratios at the upper anchors, where the cross-square spread narrows to 0.11. Second, the calibration assumes that the mean ratio of the five validation squares is representative of the full portfolio; this is the same representativeness assumption that underlies the upper-bound interpretation of the VaR and ES estimates in Section 5.3. I retain the headline EAD of this chapter at the conservative EUR 16.8 million per year, reporting the calibrated EUR 34 million and uncalibrated EUR 63.5 million figures alongside so that the sensitivity to both choices is visible rather than concealed in a single number.

EAD and VaR answer different questions. VaR is most useful for stress testing and assessing the size of an extreme one-off shock. EAD is more useful for long-run budgeting and prioritising recurring adaptation expenditure. The two metrics are complementary: a region can have high VaR because it is vulnerable to rare extreme events, while also having high EAD if a substantial part of the damage is already present at the more frequent end of the retained range (from $T = 25$ upward).

5.5 The risk of inaction over a 30-year horizon

The risk metrics reported so far describe a single year or a single event. For a sector that plans maintenance and investment over decades, the more relevant question is what the SSP5-8.5 2050 loss distribution implies over a multi-year horizon if the sector undertakes no adaptation. This section answers that question for a 30-year horizon, treating the SSP5-8.5 2050 loss distribution as a stationary approximation held constant over the period. The resulting figures are a scenario projection of the consequences of inaction, not a forecast of realised losses, and, like every estimate in this chapter, they are conditional on the modelling assumptions of Chapter 4.

Over a 30-year horizon, even moderately rare events become likely. The 4% annual exceedance probability of the 25-year event compounds to a 71% chance of at least one occurrence within 30 years; the 50-year event reaches 45%, and the 100-year event reaches 26%. Table 5.5 and Figure 5.4(A) report these probabilities together with the expected number of occurrences. One clarification matters for reading these numbers. The probabilities are per-event annual exceedance probabilities, whereas the associated loss (the EUR 457.1 million 100-year shock) is the spatially uniform, comonotonic upper bound on the portfolio loss (Section 5.3); paired this way they describe a stylised stress horizon rather than the probability of a portfolio-wide loss of that exact size in a given year, which under realistic spatial dependence is lower. I report the cumulative cost below on the dependence-invariant EAD basis, unaffected by this distinction. The headline is that a 25-year pluvial event is more likely than not to strike the portfolio at least once in the next three decades, and a one-in-four chance of a 100-year event is not a remote contingency for a long-horizon planner.

Table 5.5: *Exposure to pluvial events over a 30-year horizon of inaction, SSP5-8.5 2050. Occurrence probabilities and expected counts are per-location (point) statistics; single-event damages are the spatially uniform (comonotonic) portfolio upper bound.*

T (years)	Annual AEP (%)	$P(\geq 1$ in 30 yr) (%)	Expected # in 30 yr (#)	Single-event damage (EUR M)
25	4.0	71	1.20	354.1
50	2.0	45	0.60	410.2
100	1.0	26	0.30	457.1
500	0.2	6	0.06	500.4
1,000	0.1	3	0.03	525.2

Aggregating the annual risk over the horizon gives the cumulative expected damage of inaction. Using the EAD of Section 5.4 (EUR 16.8 million per year on the conservative $T \geq 25$ basis), 30 years of inaction accumulate an expected EUR 505 million in direct structural damage, shown as the rising line in Figure 5.4(B). Discounting at a 2% real rate lowers the present value to approximately EUR 377 million.

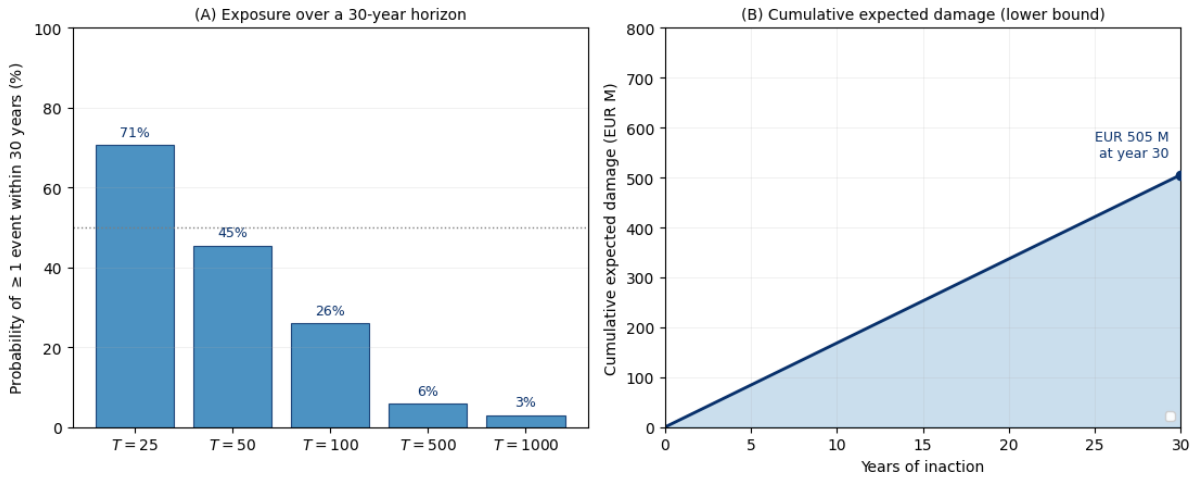


Figure 5.4: *The risk of inaction over a 30-year horizon under SSP5-8.5 2050. Panel (A): probability of at least one event of each return period within 30 years. Panel (B): cumulative expected direct structural damage accumulating over 30 years of inaction, on the conservative $T \geq 25$ EAD basis, against the sector closing cash position.*

Three qualifications make this figure a firm lower bound on the true cost of inaction. First, it rests on the $T \geq 25$ EAD and therefore omits the annual damage of the more frequent $T < 25$ events entirely, which at face value would have added on the order of EUR 47 million per year (Section 5.4). Second, it captures only the direct repair cost, and not the capitalised mark-down on property values that would also erode beleidswaarde and LTV headroom. Third, it holds the climate at its 2050 intensity rather than allowing exceedance probabilities to keep shifting through the second half of the horizon.

Against these three qualifications, all of which raise the true cost above the headline figure, one consideration works in the opposite direction, and on a different dimension: not the size of the loss but the sector's capacity to bear it. In aggregate cashflow terms the amount is absorbable: EUR 505 million over 30 years is about one-fifth of a single year of aggregate DAEB free cashflow (Section 5.8). The binding constraint is therefore not sector solvency but the regional concentration of the losses and the opportunity cost of diverting cashflow from reinvestment already committed under the Nationale Prestatieafspraken (the National Performance Agreements on social-housing investment). The practical implication is that the case for adaptation rests less on an existential sector-wide threat than on the recurring, regionally concentrated, and near-certain-to-materialise nature of the losses that inaction locks in.

5.6 Spatial distribution of 100-year damage

The spatial distribution of damage is highly uneven. As Tables 5.6 and 5.7 demonstrate, the largest absolute 100-year damages concentrate heavily in the major urbanised provinces (Zuid-Holland, Noord-Brabant, Noord-Holland) and Woondealregio's (such as Metropoolregio Amsterdam and Regio Rotterdam). These figures primarily reflect portfolio scale: regions with the largest social housing footprints mechanically face the highest aggregate exposure.

Table 5.6: *Top-five provinces by total 100-year direct structural damage, SSP5-8.5 2050. The full choropleth map is in Appendix F (Figure F.1).*

Rank	Province	Total damage (EUR M)
1	Zuid-Holland	122.6
2	Noord-Brabant	81.7
3	Noord-Holland	80.5
4	Gelderland	35.6
5	Utrecht	33.5

Table 5.7: *Top-five Woondealregio's by total 100-year direct structural damage, SSP5-8.5 2050. The full choropleth map is in Appendix F (Figure F.2).*

Rank	Woondealregio	Total damage (EUR M)
1	Metropoolregio Amsterdam	75.0
2	Regio Rotterdam	68.7
3	Stedelijke Regio Breda–Tilburg	50.7
4	Haaglanden	35.5
5	U16/Regio Utrecht	29.5

Because absolute damage is partly a scale effect, I report damage per ground-floor unit as a complementary intensity measure. This normalisation shifts the risk profile substantially (Table 5.8). While the Metropoolregio Amsterdam leads in aggregate damage, Regio Rotterdam, Haaglanden, and Stedelijke Regio Breda–Tilburg emerge as the clear hotspots for per-unit vulnerability.

Table 5.8: *Top-five Woondealregio's by 100-year direct structural damage per ground-floor unit, SSP5-8.5 2050. The full choropleth map is in Appendix F (Figure F.3).*

Rank	Woondealregio	Damage per unit (EUR k/unit)
1	Regio Rotterdam	1.5
2	Haaglanden	1.2
3	Stedelijke Regio Breda–Tilburg	1.0
4	Groningen–Assen	0.81
5	Metropoolregio Amsterdam	0.75

5.7 Spatial distribution of Expected Annual Damage

The Expected Annual Damage (EAD) maps show a similar geographic split between scale and vulnerability. The largest annualised financial burdens fall on the high-density regions in the Randstad and Brabant (Table 5.9).

Table 5.9: *Top-five Woondealregio's by aggregate Expected Annual Damage, SSP5-8.5 2050 ($T \geq 25$ basis). The full choropleth map is in Appendix F (Figure F.4).*

Rank	Woondealregio	EAD (EUR M/yr)
1	Metropoolregio Amsterdam	2.7
2	Regio Rotterdam	2.5
3	Stedelijke Regio Breda–Tilburg	1.8
4	Haaglanden	1.3
5	U16/Regio Utrecht	1.1

However, when corrected for portfolio size (Table 5.10), the recurring annualised risk per exposed unit again points strongly to Regio Rotterdam, Haaglanden, and Stedelijke Regio Breda–Tilburg.

Table 5.10: *Top-five Woondealregio's by Expected Annual Damage per ground-floor unit, SSP5-8.5 2050 ($T \geq 25$ basis). The full choropleth map is in Appendix F (Figure F.5).*

Rank	Woondealregio	EAD per unit (EUR/unit/yr)
1	Regio Rotterdam	53
2	Haaglanden	42
3	Stedelijke Regio Breda–Tilburg	37
4	Groningen–Assen	29
5	Metropoolregio Amsterdam	27

These EAD results reinforce that large urban regions dominate the aggregate burden, but per-unit EAD highlights a smaller set of regions where adaptation expenditure may yield the highest risk reduction per treated dwelling.

5.8 Financial absorption capacity

Aggregate damage figures gain financial meaning only when compared to the cashflow base that has to absorb them. This section reports the ratio of T=100 yr modelled damage to one year of DAEB-tak operating cashflow (FCF) per Woondealregio. Cash items: I take these from the dVi2024 release and allocate them to Woondealregio's *pro rata* to each corporation's regional dwelling-count share; Appendix G reports the full aggregation table. I report the complementary regulatory metric, the Interest Coverage Ratio (ICR), in summary form here and in full in Appendix H.

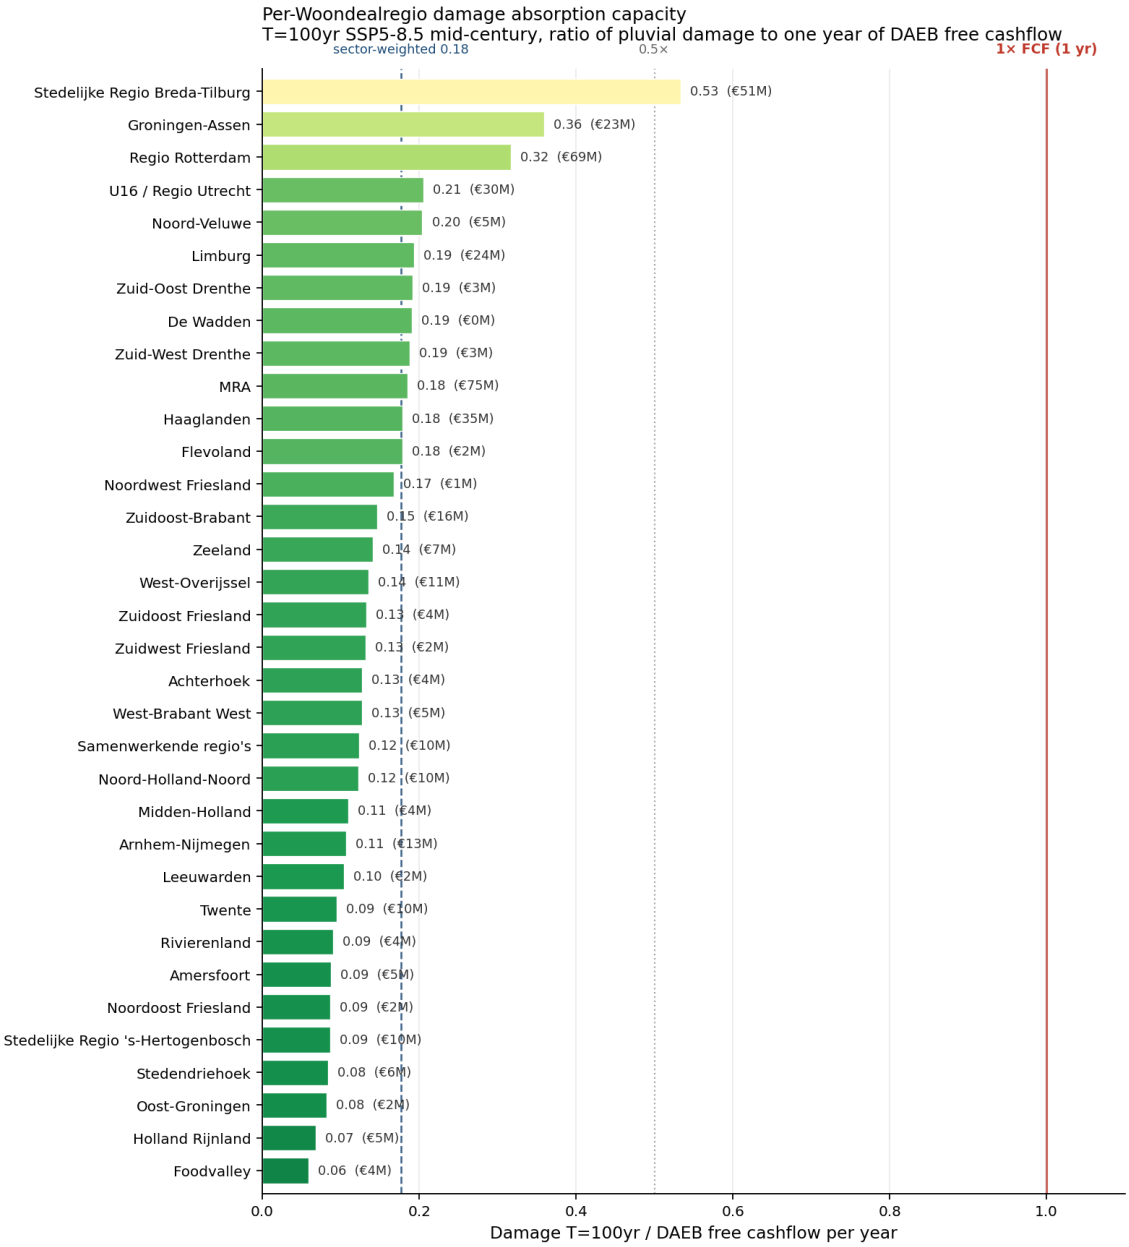


Figure 5.5: *Damage-to-free-cashflow ratio per Woondealregio at the T=100 yr SSP5-8.5 mid-century anchor. The dashed blue line marks the sector-weighted ratio (0.18); the dotted grey line the half-year-of-FCF reference (0.5); the solid red line at 1.0 the full-year reference. One region crosses 0.5, none crosses 1.0; the full table is reported in Appendix G.*

Three patterns stand out. First, at sector level the modelled damage is absorbable: across all 34 reporting Woondealregio's, the T=100 yr damage of EUR 457.1 million is 0.18 times one year of aggregate DAEB-tak free cashflow (EUR 2,581 million). Including the closing cash position (EUR 727 million) brings the ratio to 0.14. Second, this sectoral capacity masks substantial regional dispersion: 2 Woondealregio's absorb more than one third of their annual FCF, and 1 crosses the half-year-of-FCF reference. Stedelijke Regio Breda-Tilburg sits highest at 0.53, followed by Groningen-Assen at 0.36 and Regio Rotterdam at 0.32. Third, the five regions with the highest absorption ratios jointly carry

EUR 177 million of damage at $T = 100$, about 39% of the national total. Financial stress therefore concentrates in a small subset of regions rather than spreading evenly across the portfolio.

The headroom shown in Figure 5.5 represents potential rather than uncommitted capacity. All 32 out of 34 reporting regions record net investing outflows that exceed their operating cashflow in the 2024 reporting year; sector-wide, $|\mathrm{CF}_{\mathrm{inv}}|$ is 2.4 times FCF. Pluvial damage absorbed from operating cashflow would therefore come at the expense of reinvestment under the *Nationale Prestatieafspraken*.

Convergence with the supervisory metric. Table 5.11 reports the top three most-stressed regions on each metric. The two rankings overlap on Stedelijke Regio Breda-Tilburg, which is the most stressed region on both the D/FCF ratio (0.53) and the VHE-weighted ICR (weighted by lettable units, *verhuureenheden*; 1.72). Before the modelled shock, all 34 Woondealregio's clear the Aw/WSW supervisory norm of 1.4 on the regional weighted ICR, with Leeuwarden (1.75) and De Wadden (1.81) closest to the threshold; the sector-weighted ICR is 2.33 (Appendix H).

Table 5.11: *Top-three most-stressed Woondealregio's on each financial metric, dVi2024 cashflow inputs and T=100 yr SSP5-8.5 2050 damage.*

Buffer pressure (D/FCF)		Interest coverage ($\mathrm{ICR_w}$)	
Woondealregio	D/FCF	Woondealregio	$\mathrm{ICR_w}$
Stedelijke Regio Breda-Tilburg	0.53	Stedelijke Regio Breda-Tilburg	1.72
Groningen-Assen	0.36	Leeuwarden	1.75
Regio Rotterdam	0.32	De Wadden	1.81
Sector-weighted aggregate	0.18	Sector-weighted aggregate	2.33

The two metrics measure different financial constraints. ICR is a good-year measure of whether operating cashflow covers interest payments; the D/FCF buffer ratio is a shock-absorption measure of whether the same cashflow can absorb a one-off pluvial event. In an ordinary year, before any pluvial damage, all 34 regions clear the ICR norm while Breda-Tilburg's D/FCF reaches 0.53, so for most regions the financial strain from a T=100-year pluvial event shows up not as an inability to cover interest, but as less cashflow left for the investment programme. In Breda-Tilburg specifically, absorbing EUR 51 million of damage would consume just over half of one year of operating cashflow, with the remainder still required to service interest.

This changes when I charge the modelled damage in full against a single year's cashflow. Figure 5.6 recomputes the regional ICR under a one-off T=100 yr shock, allocating each region's damage to its corporations *pro rata* to their dwelling share and deducting it from that year's operating cashflow (Appendix H.1 details the construction). The average regional ICR falls by 0.18, and exactly one region crosses the supervisory norm:

Stedelijke Regio Breda-Tilburg, which moves from 1.72 to 1.36.

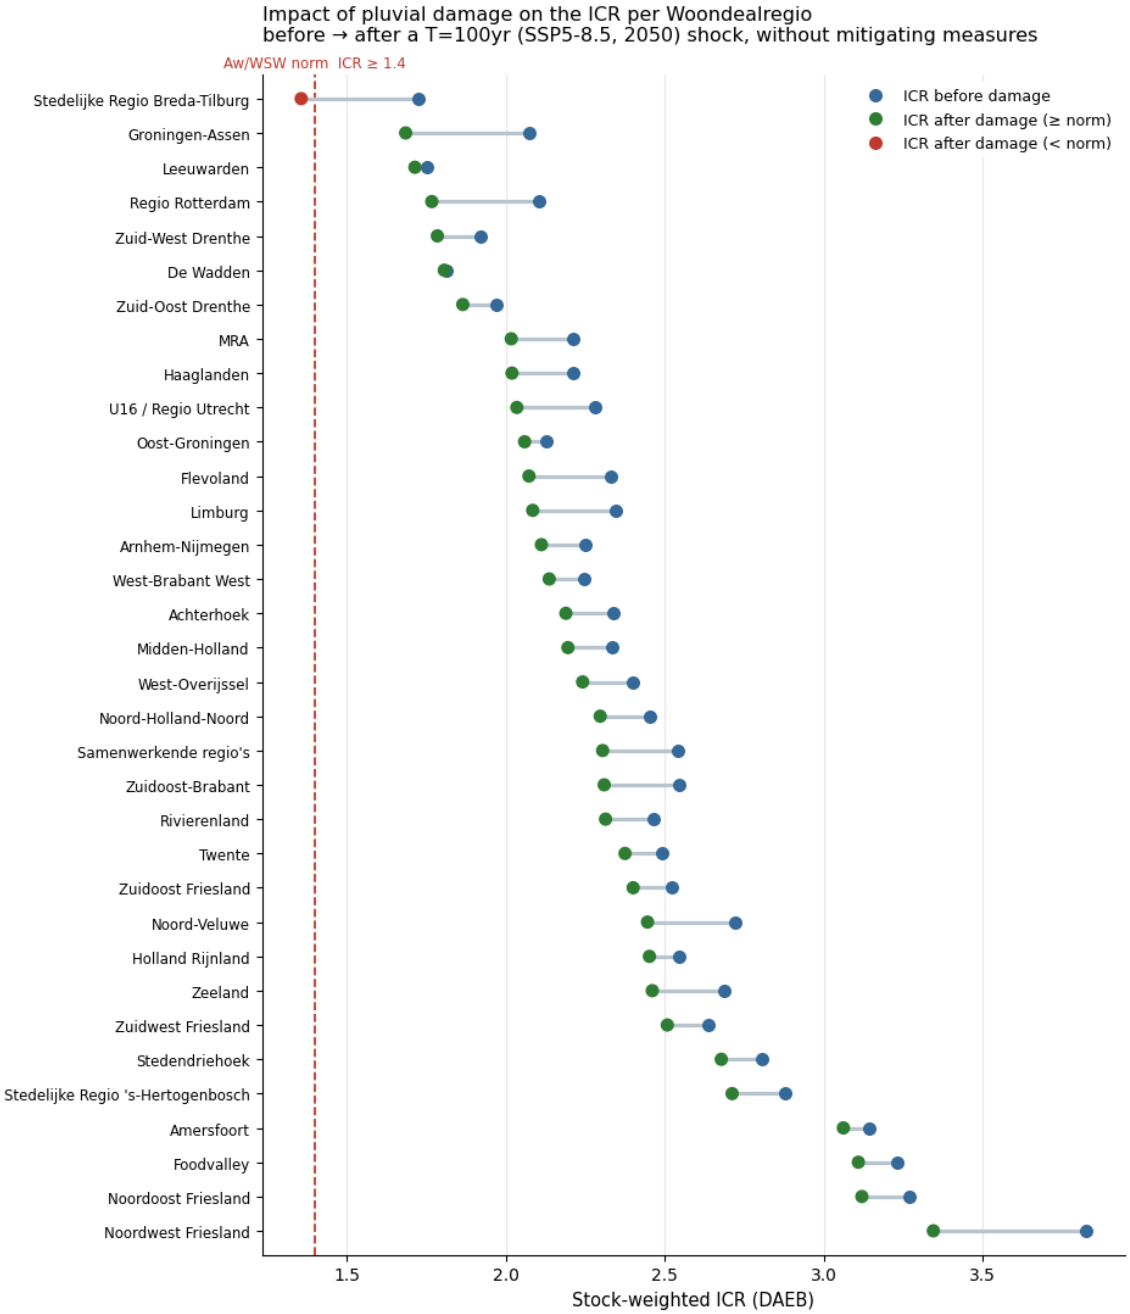


Figure 5.6: *Impact of a one-off T=100 yr pluvial shock (SSP5-8.5, 2050) on the VHE-weighted ICR per Woondealregio, without mitigating measures. Each region's pre-shock ICR (blue) is linked to its post-shock value; green markers remain at or above the Aw/WSW norm of 1.4, the single red marker falls below it. Regions are sorted by post-shock ICR. The full table and construction are reported in Appendix H.1.*

The ranking carries one important nuance. Noordwest Friesland records the largest absolute drop (−0.48, from 3.83 to 3.35), yet it remains comfortably above the norm; what matters for supervision is the post-shock distance to 1.4, not the size of the move. The regions sitting closest to the norm after the shock, Breda-Tilburg (1.36), Groningen-Assen

(1.68), Leeuwarden (1.71), and Regio Rotterdam (1.77), are a mix of high-damage regions and regions that simply began near the threshold, and Breda-Tilburg is the single case where high pre-shock stress, large absolute damage, and the resulting breach coincide.

The financial signature of pluvial risk for the social housing sector is therefore an investment-capacity story rather than a solvency story. Even under the deliberately adverse assumption that a 1%-annual-probability event lands in full on a single year's cashflow, only one region breaches the supervisory norm, and only narrowly, which bounds the interest-coverage dimension of the risk while leaving the investment-capacity dimension as the binding concern.

5.9 Depth distributions and threshold crossing

The distribution of building-level depths shows the threshold-crossing mechanism directly. Figure 5.7 plots the per-scenario depth histograms on a logarithmic scale: the overwhelming majority of ground-floor units remain dry or experience only shallow water, with the unit count falling by roughly an order of magnitude over the first decimetres of depth, while a long but thin right tail extends beyond the 5, 10, and 15 cm thresholds. As rainfall increases from $T = 5$ to $T = 1{,}000$, the shape of this distribution is remarkably stable; the main shift is not a uniform deepening of dwellings already flooded, but additional mass crossing the lower thresholds. The annotated counts make this explicit: the number of units at or above the 5 cm threshold grows from 37,147 at $T = 5$ to 64,053 at $T = 100$ and 71,978 at $T = 1{,}000$. The bottom-right panel summarises this threshold-crossing growth across all three thresholds and shows that it decelerates at the highest return periods, the same saturation that flattens the loss-exceedance curve. This is precisely what generates the extensive-margin growth pattern of Section 5.2 and the flattening of the LEC in Section 5.3.

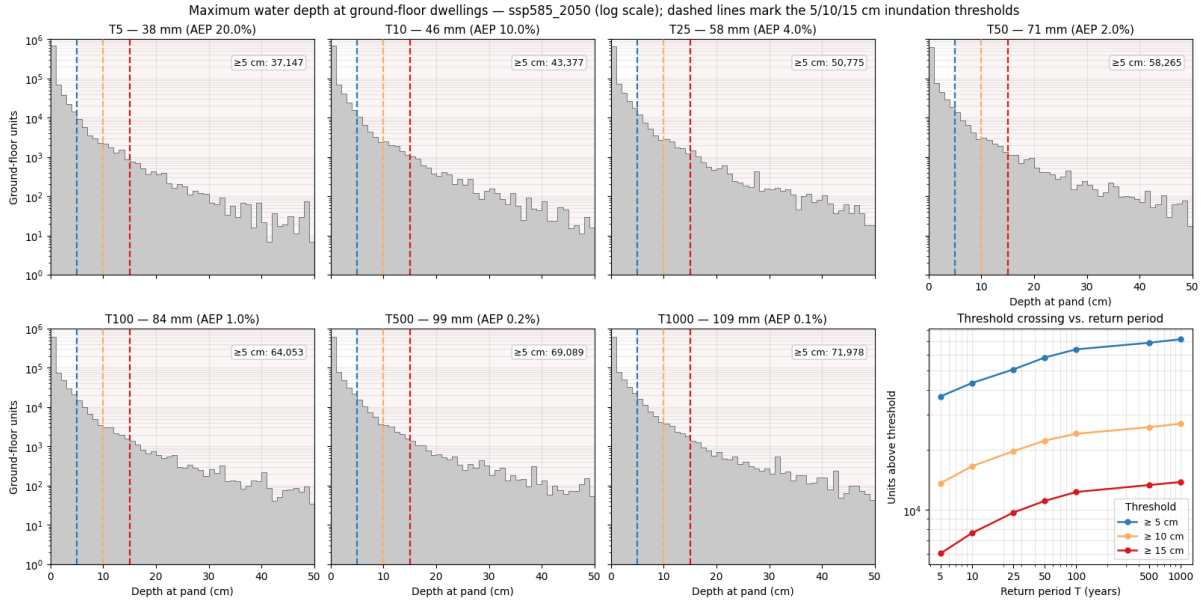


Figure 5.7: *Distribution of maximum water depth at ground-floor dwellings for each return-period scenario under SSP5-8.5 2050, on a logarithmic count scale. Dashed vertical lines mark the 5, 10, and 15 cm inundation thresholds used in the sensitivity analysis; each panel reports the number of units at or above the 5 cm threshold. The bottom-right panel plots the number of units above each threshold against the return period.*

This depth pattern is consistent with the damage-growth decomposition in Table 5.2. The loss distribution is sensitive to the threshold because the histograms concentrate dwellings in the shallow-depth range where a small change in assumed entry height changes whether damage occurs at all.

5.10 Sensitivity to the inundation threshold

The baseline uses a 5 cm damage threshold. This is intentionally conservative: it assumes that shallow water reaching the building perimeter can generate structural damage once the depth exceeds 5 cm. To test the importance of this assumption, I repeat the damage calculation with 10 cm and 15 cm thresholds. Table 5.12 reports the effect on VaR at confidence levels that coincide with scenario anchors.

Table 5.12: *Sensitivity of VaR to the inundation threshold, SSP5-8.5 2050.*

Confidence level α	5 cm (EUR M)	10 cm (EUR M)	Reduction vs. 5 cm	15 cm (EUR M)	Reduction vs. 5 cm
90.0%	296.0	187.6	-36.6%	124.7	-57.9%
99.0%	457.1	294.9	-35.5%	210.8	-53.9%
99.9%	525.2	340.7	-35.1%	246.1	-53.1%

Raising the threshold from 5 cm to 10 cm reduces VaR by approximately 35–37%. Raising it to 15 cm reduces VaR by approximately 53–58%. The relative reduction

is stable across the tail, which indicates that the threshold assumption affects the entire simulated loss distribution rather than only one specific return period. If deep inundation of dwellings already affected drove loss growth, raising the activation threshold by 5–10 cm would barely change the VaR. The 35–58% reductions observed instead reveal that a large share of dwellings sits in the shallow-depth band just above 5 cm, and the sensitivity result therefore constitutes a third, independent line of evidence for the threshold-crossing mechanism alongside the growth decomposition (Section 5.2) and the depth distributions (Section 5.9).

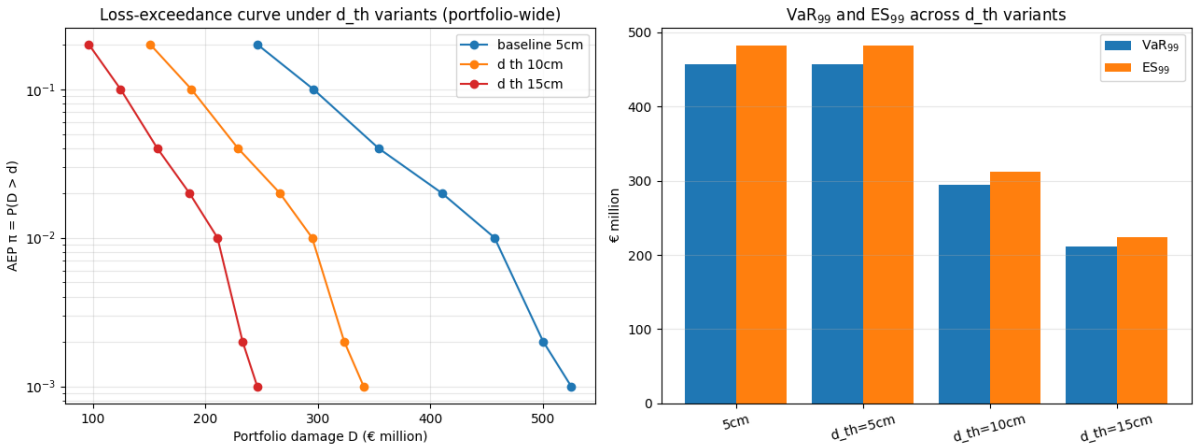


Figure 5.8: *Damage and VaR sensitivity to the inundation threshold. The baseline uses a 5 cm threshold; 10 cm and 15 cm thresholds are shown as robustness checks.*

This sensitivity result strengthens the interpretation that the social housing portfolio is exposed mainly through shallow inundation. The baseline loss estimates are therefore not precise predictions of realised repair costs. They are scenario-based financial risk estimates conditional on a conservative threshold assumption. At the same time, the sensitivity analysis is informative for adaptation: relatively small increases in effective entry height or local protection can materially reduce portfolio VaR.

5.11 Synthesis: Threshold crossing drives portfolio risk

Four main findings emerge from the simulation run, answering the empirical sub-questions of my thesis.

First, the *extensive margin*, additional ground-floor dwellings crossing the 5 cm inundation threshold, drives portfolio loss growth almost entirely, rather than the *intensive margin* of deeper inundation among already-affected buildings. The growth decomposition (Section 5.2), the shallow-skewed depth distributions (Section 5.9), and the high sensitivity of VaR to the assumed threshold (Section 5.10) confirm this mechanism. The

financial signature of pluvial risk for social housing is therefore a problem of *shallow inundation at scale*.

Second, the absolute losses are material. Under SSP5-8.5 in 2050, a 100-year event generates an estimated EUR 457.1 million in direct structural damage, representing the 99% VaR of the simulated loss distribution. Over a 30-year horizon, inaction accumulates an expected damage of at least EUR 505 million.

Third, when set against the sector's cashflow base, these losses are absorbable at the national level (consuming roughly 18% of annual DAEB free cashflow) but tightly concentrated. Financial stress generally shows up not as a breach of the Aw/WSW Interest Coverage Ratio, but as compression of the residual cashflow available for regional investment programmes. Stedelijke Regio Breda–Tilburg is the distinct exception, facing the highest pressure on both the cashflow buffer and the regulatory ICR metric.

Finally, the spatial distinction between aggregate exposure and damage intensity directly informs adaptation strategy. While regions like the Metropoolregio Amsterdam face the largest absolute financial shocks due to their size, per-unit vulnerability hotspots like Regio Rotterdam and Haaglanden offer higher potential returns on adaptation per exposed dwelling. Because the threshold-crossing mechanism dominates the portfolio VaR, relatively modest measures that raise the effective entry height, such as door barriers, ground-floor detailing, or street-level drainage improvements, act exactly at the margin where they can provide the greatest risk reduction.

Chapter 6

Conclusion

In this chapter I will synthesise the results of my thesis and answer the main research question. I answer the research questions by summarising the empirical findings, then set out the academic and practical contributions, the main limitations, recommendations, and directions for future research. My central conclusion is that extreme pluvial precipitation constitutes a material financial risk for the Dutch social housing sector, but not primarily a problem of rare catastrophic depths: it is a problem of many ground-floor dwellings sitting close to the point at which shallow water becomes damaging, and of that damage concentrating in a small number of financially exposed regions.

6.1 Answer to the research questions

My main research question was:

> *To what extent do extreme pluvial precipitation events pose financial risks to Dutch social housing portfolios, and how can these risks be quantified using flood simulation and Value at Risk estimation?*

The answer is that extreme pluvial precipitation poses a material, spatially concentrated, and threshold-sensitive financial risk to the Dutch social housing sector, but one that the sector can absorb out of cashflow at the national level. Under the SSP5–8.5 2050 pathway, the simulated 99% VaR of direct structural damage is EUR 457.1 million, the 99.9% VaR is EUR 525.2 million, and the truncated 99% Expected Shortfall is EUR 482.5 million. These headline figures assume a spatially uniform event across the portfolio and therefore represent the comonotonic (perfect-dependence) *upper bound* on tail risk; the spatial diversification of a country-wide portfolio places the realistic 99% VaR lower, with EUR 117 million as the independence-case floor and the true value closer to that floor (Section 5.3). The financial significance of these numbers emerges only against the cashflow base that has to absorb them: the EUR 457.1 million 100-year shock equals

roughly 18% of one year of aggregate DAEB free cashflow, so at sector level the loss is absorbable and pluvial risk is not, in the first instance, a solvency problem. The binding constraint is instead distributional. The loss concentrates in a small set of regions, and absorbing it diverts cashflow away from reinvestment already committed under the *Nationale Prestatieafspraken*. I quantify this risk by constructing a physically constrained loss-exceedance curve from a small set of physically simulated return-period events and reading VaR and ES off the resulting curve; Chapter 4 sets out the full method.

Sub-question 1: Simulating pluvial flood depths

A national 5 m raster simulation pipeline makes this quantification possible at portfolio scale. It combines AHN terrain data, BAG building footprints, BGT land-use and water masks, and KNMI-derived rainfall scenarios in a local-inertial surface-flow model, and resolves local terrain and building effects efficiently enough to cover the entire social housing portfolio in a single reproducible framework (Chapter 4).

Sub-question 2: Magnitude and distribution of financial losses

My simulated loss distribution is material and strongly right-tailed (dominated by rare, large losses). A rare but plausible two-hour pluvial event under SSP5–8.5 2050 generates several hundred million euros in direct structural damage alone (EUR 457.1 million for the 100-year event under the comonotonic upper bound). Annualised, the scenario-based Expected Annual Damage (EAD) spans a bracket rather than a single number: EUR 16.8 million per year on the conservative validated range ($T \geq 25$, uncalibrated), EUR 34 million per year as the validation-calibrated central estimate, and EUR 63.5 million per year as the uncalibrated upper bound (Section 5.4). The cumulative consequence of inaction follows directly: held at the 2050 intensity and the conservative EAD, thirty years of doing nothing accumulate an expected EUR 505 million in direct structural damage, EUR 377 million in present value at a 2% real discount rate (Section 5.5). These magnitudes place pluvial flooding outside the category of marginal operational issues and inside long-term financial planning.

Sub-question 3: Spatial variation, hotspots, and financial stress

Risk is spatially concentrated, and the concentration is sharper in financial terms than in damage terms alone. In absolute damage, the largest losses fall on highly urbanised, portfolio-dense regions such as the Metropoolregio Amsterdam and Regio Rotterdam; normalised by the number of ground-floor units, the vulnerability hotspots shift toward Regio Rotterdam, Haaglanden, and Stedelijke Regio Breda–Tilburg. Setting the modelled damage against each region's operating cashflow sharpens the picture further. At sector

level the 100-year damage is 0.18 times one year of aggregate DAEB free cashflow, but this masks substantial dispersion: the five most exposed regions jointly carry EUR 177 million of the 100-year damage (about 39% of the national total), and Stedelijke Regio Breda–Tilburg alone reaches a damage-to-cashflow ratio of 0.53. The supervisory dimension confirms where the stress binds. Before any shock, all 34 Woondealregio's clear the Aw/WSW Interest Coverage Ratio (ICR) norm of 1.4 (sector-weighted ICR 2.33); under a one-off 100-year damage shock charged in full against a single year's cashflow, exactly one region falls below the norm, Stedelijke Regio Breda–Tilburg, which moves from 1.72 to 1.36 (Section 5.8). Breda–Tilburg is therefore the single region where high absolute damage, the thinnest cashflow buffer, and a supervisory-norm breach coincide, which makes it the clearest joint priority for adaptation and supervisory attention. For the sector as a whole, the financial signature is an investment-capacity story rather than a solvency story: 32 of the 34 reporting regions already record investing outflows larger than their operating cashflow, so damage absorbed from cashflow competes directly with committed reinvestment rather than threatening interest cover.

Sub-question 4: Mechanism behind portfolio loss growth

The *extensive margin*, additional ground-floor dwellings crossing the 5 cm inundation threshold, drives portfolio loss growth almost entirely, rather than the *intensive margin* of deeper inundation among already-affected buildings. The shallow-skewed building-level depth distributions and the high sensitivity of VaR to the assumed threshold confirm this: raising the threshold from 5 cm to 10 cm reduces VaR by 35–37% (Section 5.10). The financial implication is direct, and it is what makes the mechanism matter for practice. Because the portfolio VaR concentrates in a thin band of dwellings sitting just above the damage threshold, relatively modest measures that raise the effective entry height at the ground floor act exactly at the margin that dominates the VaR, and can therefore deliver disproportionate reductions in tail risk per euro spent.

6.2 Academic contribution

My academic contribution of my thesis is the integration of three literatures: pluvial flood simulation, depth–damage modelling, and financial risk quantification. Dutch flood-risk studies have focused mainly on fluvial flooding or on commercial mortgage and real-estate exposures. My thesis applies a spatially explicit physical-risk framework to the non-profit social housing sector, where regulated revenues, social obligations, and limited geographical flexibility make risk management different from that of commercial property owners.

My thesis demonstrates that a physically constrained loss-exceedance curve can be

constructed from a limited number of hydrodynamic design events. This is useful in settings where historical loss data are unavailable and where purely statistical loss modelling is infeasible. The approach does not require fitting an unconditional loss distribution from observed claims. Instead, it uses physically simulated return-period events as control points and reads VaR and ES off the resulting LEC, which makes the method transparent, reproducible, and directly linked to meteorological return periods.

My second contribution is the decomposition of loss growth into extensive and intensive components. My finding that additional dwellings crossing a shallow damage threshold dominate loss escalation adds an interpretable mechanism to the VaR results. It therefore does not only report how large the tail losses are, but also explains why they grow: the key process is threshold crossing among ground-floor dwellings rather than rapidly increasing severity among already-damaged dwellings.

6.3 Practical implications

My results carry four practical implications. First, pluvial flood risk is large enough to belong in long-term financial planning and stress testing rather than to be treated as a marginal operational risk; and because the binding constraint is investment capacity rather than solvency, it belongs specifically in reinvestment and *Nationale Prestatieafspraken* planning, not only in solvency buffers. Second, adaptation prioritisation should combine aggregate and per-unit metrics, since the two identify different things: where the largest financial shocks fall, and where treatment is most efficient per exposed dwelling. Third, shallow-water entry prevention is the most promising early adaptation theme, because it acts on the threshold-crossing margin that dominates the portfolio VaR. Fourth, asset-level data collection (ground-floor status, entry height, basement presence, installed mitigation) is itself a risk-management measure, because it would let future models replace conservative proxies with building-specific estimates.

Ultimately, my thesis shows that pluvial flood risk can be translated from a local physical hazard into an actionable portfolio-level financial metric. The most important insight for practice is the meaning behind the numbers: the risk is not primarily a problem of rare catastrophic depths, nor of sector-wide insolvency, but of many dwellings sitting close to the point at which shallow water becomes damaging, and of that damage falling on a few regions whose cashflow is already fully committed to reinvestment. For housing associations and supervisors, this provides a clear path forward: if the first centimetres of water entry can be prevented in the right places, a substantial part of the portfolio's pluvial VaR may be reduced, and the cashflow that would otherwise absorb the damage stays available for the investment programme. Quantifying the risk is only the first step; the next is to use these estimates to prioritise investments that protect tenants, preserve financial capacity, and make the social housing stock more resilient to a wetter climate.

6.4 Limitations

Several limitations qualify these results. They are ordered here from the most material for the financial interpretation to the most technical.

Damage scope. My thesis estimates direct structural damage only. It does not include contents damage, tenant displacement, temporary accommodation, rental-income losses, emergency response costs, insurance deductibles, indirect operational disruption, or the possible valuation channel through *WOZ*, *marktwaarde*, or *beleidswaarde*. For direct structural damage under the baseline physical assumptions, the 5 cm threshold is conservative. For total economic and financial impact, however, the estimates are incomplete and represent lower bounds; in particular, a valuation mark-down would also erode *beleidswaarde* and LTV headroom beyond the repair cost captured here.

Spatial-dependence assumption. The headline VaR and ES figures assume a spatially uniform event across all 1,097 squares (comonotonic dependence), which maximises aggregate tail risk and yields an upper bound. The opposite extreme of mutually independent square-level losses gives a 99% VaR of EUR 117 million; the realistic value lies inside this band and, for a country-wide portfolio, closer to the lower bound (Section 5.3). This is the single largest source of uncertainty in the reported tail metrics, and it qualifies VaR and ES only: the EAD is invariant to the dependence structure.

Threshold sensitivity. The baseline 5 cm threshold is a modelling assumption, and the sensitivity analysis shows that the results respond strongly to it. Any conclusion about absolute damage levels therefore depends on reading the 5 cm baseline together with the 10 cm and 15 cm robustness checks.

Damage-function uncertainty. The damage calculation uses a depth–damage function anchored to replacement values. As the literature shows, depth–damage functions are a major source of uncertainty in flood-risk modelling. They typically omit duration, flow velocity, contamination, building materials, repair-cost inflation, and behavioural response. The reported values therefore represent scenario-based estimates rather than precise predictions of realised repair bills.

Ground-floor exposure proxy. The Aw register does not contain unit-level floor information. I treat EGW units as ground-floor exposed and represent MGW exposure through a first-unit-per-pand proxy. This is a reasonable portfolio-level approximation, but it can misclassify individual apartment buildings, so the per-ground-floor-unit metrics represent upper-bound estimates for MGW-dominated regions (e.g., Dutch *galerijflats*, gallery-access apartment blocks).

Tail estimation. I build the LEC from seven scenario points. Simulated losses directly determine VaR at the scenario anchors, but interior quantiles require interpolation, and I truncate ES at the 1,000-year event. Deeper-tail risk beyond 99.9% therefore remains outside the scope of the main estimates.

Climate-scenario scope. I deliver the full simulation run for SSP5–8.5 in 2050 only. I designed the method so that parameter changes can simulate other pathways and horizons, but my thesis does not provide full national results for the current climate, SSP2–4.5, or the 2100 horizons. The results are therefore pathway-conditional and do not represent single climate-averaged estimates of future climate risk.

Hydrodynamic and rainfall simplification. The model uses a local-inertial surface-routing scheme rather than a fully coupled surface-water and sewer-network model, and it represents sewer capacity, infiltration, open-water drainage, and tunnel sinks through simplified approximations. The rainfall forcing uses a single design storm profile scaled to return-period totals. Both choices are necessary for national scalability, but they mean that highly local drainage features and the dependence of real damage on storm track and spatial rainfall gradients are not fully resolved.

6.5 Recommendations

Based on the results, five recommendations follow.

1. **Use VaR and EAD together in sector-level climate-risk planning.** Use VaR for stress testing and buffer assessment, and EAD for long-run budgeting and adaptation prioritisation. The two metrics answer different questions and should not be treated as substitutes.

2. **Prioritise regions with both high aggregate exposure and high per-unit intensity.** Aggregate hotspots such as the Metropoolregio Amsterdam and Regio Rotterdam matter for sector-level loss concentration. Per-unit hotspots such as Regio Rotterdam, Haaglanden, Stedelijke Regio Breda–Tilburg, and Groningen–Assen matter for efficient adaptation targeting. Stedelijke Regio Breda–Tilburg warrants particular attention as the one region where high damage, the thinnest cashflow buffer, and a supervisory-norm breach coincide.

3. **Focus early adaptation on shallow-water entry prevention.** The threshold sensitivity results indicate that modest increases in effective entry height can substantially reduce VaR. This makes property-level flood resilience and local threshold protection promising first-order interventions.

4. **Collect asset-level exposure data.** Housing associations should prioritise machine-readable data on ground-floor status, entrance height, basement presence, building-level adaptations, renovation status, and known drainage problems. These variables would materially improve the next generation of risk models.

5. **Integrate pluvial scenarios into financial supervision and investment planning.** The Aw and WSW could use pluvial stress scenarios to test whether associations in exposed regions, and Stedelijke Regio Breda–Tilburg in particular, retain enough cashflow flexibility to absorb repair costs without breaching the ICR norm or crowding out required investment.

6.6 Future research

Several extensions would strengthen and broaden the analysis.

First, future work should extend the simulation run to the full climate matrix: the current-climate baseline, SSP2–4.5 in 2050 and 2100, and SSP5–8.5 in 2100. This would allow my thesis framework to quantify not only a single pathway-conditional risk distribution, but also the incremental effect of climate change across pathways and time horizons.

Second, future work could refine the model for priority hotspots using more detailed hydrodynamic information. A practical next step would be a two-stage modelling strategy: use the national 5 m model to identify high-risk regions, then simulate those regions with a coupled surface-water and sewer model or a calibrated municipal drainage model.

Third, future work should develop a Dutch pluvial depth–damage function. The current approach relies on a general residential depth–damage relationship. Future work could use claims data, repair invoices, and post-event inspection records from housing associations to calibrate a more appropriate function for shallow pluvial inundation, social-housing building types, and Dutch repair-cost conditions.

Fourth, future work could extend the financial analysis from portfolio damage to association-level balance-sheet and cashflow stress testing. Linking simulated losses to operating cashflow, ICR, debt service, insurance coverage, and investment plans would show which associations are most likely to face binding financial constraints after a severe pluvial event.

Fifth, future work should evaluate adaptation measures explicitly. The threshold sensitivity analysis suggests large potential benefits from increasing effective entry heights, but a full policy recommendation requires comparing risk reduction with measure costs. Future research could therefore estimate benefit–cost ratios for a number of adaptation measures.

Sixth, future research could improve tail estimation. A peaks-over-threshold or non-stationary extreme-value model could be combined with additional physically simulated scenarios to estimate risk beyond the 1,000-year event. Such an extension would be especially relevant for supervisory stress testing, but it would require more scenario nodes than were feasible in my thesis.

Bibliography

Abadie, L. M., de Murieta, E. S., & Galarraga, I. (2016). Climate risk assessment under uncertainty: An application to main European coastal cities. *Frontiers in Marine Science*, *3*, 265. https://doi.org/10.3389/fmars.2016.00265

Abadie, L. M., Galarraga, I., & de Murieta, E. S. (2019). Risk measures and the distribution of damage curves for 600 European coastal cities. *Environmental Research Letters*, *14*(6), 064021. https://doi.org/10.1088/1748-9326/ab185c

Acerbi, C., & Tasche, D. (2002). On the coherence of expected shortfall. *Journal of Banking & Finance*, *26*(7), 1487–1503. https://doi.org/10.1016/S0378-4266(02)00283-2

Aedes. (2026). Aedes: Dutch association of housing corporations. https://aedesdatacentrum.nl/mosaic/aedesdatacentrum/voldoende-huurwoningen?scrollid=888

Allegri, E., Zanetti, M., Torresan, S., & Critto, A. (2024). Pluvial flood risk assessment for 2021–2050 under climate change scenarios in the metropolitan city of Venice. *Science of the Total Environment*, *914*, 169925. https://doi.org/10.1016/j.scitotenv.2024.169925

Artzner, P., Delbaen, F., Eber, J.-M., & Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, *9*(3), 203–228. https://doi.org/10.1111/1467-9965.00068

Autoriteit woningcorporaties. (2024). Het werk van de aw. https://www.ilent.nl/onderwerpen/over-de-autoriteit-woningcorporaties

Autoriteit woningcorporaties & Waarborgfonds Sociale Woningbouw. (2024, June). Aw en WSW publiceren nieuwe grenswaarden financiële normen (bestuurlijke brief).

Basel Committee on Banking Supervision. (2019). *Minimum capital requirements for market risk* (tech. rep. No. d457). Bank for International Settlements. https://www.bis.org/bcbs/publ/d457.htm

Bates, P. D., Horritt, M. S., & Fewtrell, T. J. (2010). A simple inertial formulation of the shallow water equations for efficient two-dimensional flood inundation modelling. *Journal of Hydrology*, *387*(1–2), 33–45. https://doi.org/10.1016/j.jhydrol.2010.03.027

Beersma, J., Versteeg, H., & Hakvoort, H. (2018). *Neerslagstatistieken voor korte duren—actualisatie 2018* (tech. rep. No. 2018-12). Stichting Toegepast Onderzoek Waterbeheer (STOWA).

Beltrán, A., Maddison, D., & Elliott, R. J. R. (2018). Is flood risk capitalised into property values? *Ecological Economics*, *146*, 668–685. https://doi.org/10.1016/j.ecolecon.2017.12.015

Boezeman, D., & de Vries, T. (2019). Climate proofing social housing in the Netherlands: Toward mainstreaming? *Journal of Environmental Planning and Management*, *62*(8), 1446–1464. https://doi.org/10.1080/09640568.2018.1510768

Bosker, M., Garretsen, H., Marlet, G., & Van Woerkens, C. (2019). Nether Lands: Evidence on the price and perception of rare natural disasters. *Journal of the European Economic Association*, *17*(2), 413–453. https://doi.org/10.1093/jeea/jvy002

Caloia, F., van Ginkel, K., & Jansen, D.-J. (2023). *Floods and financial stability: Scenario-based evidence from below sea level* (Working Paper No. 796). De Nederlandsche Bank. https://www.dnb.nl/en/publications/research-publications/working-paper-2023/796-floods-and-financial-stability-scenario-based-evidence-from-below-sea-level/

Cheng, L., & AghaKouchak, A. (2014). Nonstationary precipitation intensity-duration-frequency curves for infrastructure design in a changing climate. *Scientific Reports*, *4*, 7093. https://doi.org/10.1038/srep07093

Chow, V. T. (1959). *Open-channel hydraulics*. McGraw-Hill.

Daniel, V. E., Florax, R. J., & Rietveld, P. (2009). Floods and residential property values: A hedonic price analysis for the Netherlands. *Built Environment*, *35*(4), 563–576. https://doi.org/10.2148/benv.35.4.563

Daníelsson, J., Jorgensen, B. N., Samorodnitsky, G., Sarma, M., & de Vries, C. G. (2013). Fat tails, VaR and subadditivity. *Journal of Econometrics*, *172*(2), 283–291. https://doi.org/10.1016/j.jeconom.2012.08.011

de Grave, P., & Juch, S. (2023). *Standaardmethode 2023: Schade en slachtoffers als gevolg van overstromingen* (tech. rep.). Deltares.

de Moel, H., Endendijk, T., van Ederen, D., Juch, S., & van Ginkel, C. (2024). *Berekenen van overstromingsschade aan woningen: Nieuwe inzichten voor financiële toepassingen* (tech. rep. No. 11211522-004-0001). Deltares.

de Moel, H., Jongman, B., Kreibich, H., Merz, B., Penning-Rowsell, E., & Ward, P. J. (2015). Flood risk assessments at different spatial scales. *Mitigation and Adaptation Strategies for Global Change*, *20*(6), 865–890. https://doi.org/10.1007/s11027-015-9654-z

Deltares. (2023). *Schadeberekening overstromingen Limburg 2021* (tech. rep.). https://publications.deltares.nl/11207700_007_0011.pdf

Dutch Green Building Council. (2024). Effect van klimaatverandering op woningcorporaties. https://www.dgbc.nl/publicaties/effect-van-klimaatverandering-op-woningcorporaties/

Eichholtz, P., Kok, N., & Weenink, P. (2026). Dutch dilemma: Housing prices and flood risk exposure. *Real Estate Economics*, 1–35. https://doi.org/10.1111/1540-6229.70027

Endendijk, T., Botzen, W. J. W., de Moel, H., Aerts, J. C. J. H., Slager, K., & Kok, M. (2023). Flood vulnerability models and household flood damage mitigation measures: An econometric analysis of survey data. *Water Resources Research*, *59*(8). https://doi.org/10.1029/2022WR034192

European Central Bank. (2022). *2022 climate risk stress test* (tech. rep.). https://www.bankingsupervision.europa.eu/ecb/pub/pdf/ssm.climate_stress_test_report.20220708~2e3cc0999f.en.pdf

Finance Ideas. (2024, July). Nieuwe grenswaarden financiële normen: Wat zijn de consequenties? Retrieved April 14, 2026, from https://finance-ideas.nl/nieuwe-grenswaarden-financiele-normen/

François, B., Schlef, K. E., Wi, S., & Brown, C. M. (2019). Design considerations for riverine floods in a changing climate – a review. *Journal of Hydrology*, *574*, 557–573. https://doi.org/10.1016/j.jhydrol.2019.04.068

Gourevitch, J. D., Kousky, C., Liao, Y., Nolte, C., Pollack, A. B., Porter, J. R., & Weill, J. A. (2023). Unpriced climate risk and the potential consequences of overvaluation in US housing markets. *Nature Climate Change*, *13*(3), 250–257. https://doi.org/10.1038/s41558-023-01594-8

Huizinga, J., de Moel, H., & Szewczyk, W. (2017). *Global flood depth-damage functions: Methodology and the database with guidelines* (EUR 28552 EN). Joint Research Centre, European Commission. https://doi.org/10.2760/16510

Hull, J. C. (2023). *Risk management and financial institutions* (6th). Wiley.

IPCC. (2023). Sections. In Core Writing Team, H. Lee & J. Romero (Eds.), *Climate change 2023: Synthesis report. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change* (pp. 35–115). https://doi.org/10.59327/IPCC/AR6-9789291691647

Jonkman, S. N., Kok, M., & Vrijling, J. K. (2008). Flood risk assessment in the Netherlands: A case study for dike ring south Holland. *Risk Analysis*, *28*(5), 1357–1374. https://doi.org/10.1111/j.1539-6924.2008.01103.x

KNMI. (2023). *KNMI'23 klimaatscenario's voor Nederland* (tech. rep.). KNMI. https://cdn.knmi.nl/system/ckeditor/attachment_files/data/000/000/357/original/KNMI23_klimaatscenarios_gebruikersrapport_23-03.pdf

Kok, M., Slager, K., de Moel, H., Botzen, W., de Bruijn, K., Wagenaar, D., Rikkert, S., Koks, E., & van Ginkel, K. (2023). Rapid damage assessment caused by the flooding event 2021 in Limburg, Netherlands. *Journal of Coastal and Riverine Flood Risk*, *2*, 10. https://doi.org/10.59490/jcrfr.2023.0010

Koninklijk Nederlands Meteorologisch Instituut. (2018, April). *Neerslagstatistiek* (tech. rep.). https://cdn.knmi.nl/system/downloads/files/000/000/042/original/Neerslagstatistiek_mrt2018.pdf

Kreibich, H., Bubeck, P., Van Vliet, M., & de Moel, H. (2015). A review of damage-reducing measures to manage fluvial flood risks in a changing climate. *Mitigation and Adaptation Strategies for Global Change*, *20*(6), 967–989. https://doi.org/10.1007/s11027-014-9629-5

Lenderink, G., & van Meijgaard, E. (2008). Increase in hourly precipitation extremes beyond expectations from temperature changes. *Nature Geoscience*, *1*(8), 511–514. https://doi.org/10.1038/ngeo262

Manning, R. (1891). On the flow of water in open channels and pipes. *Transactions of the Institution of Civil Engineers of Ireland*, *20*, 161–207.

Maurer, C. R., Qi, R., & Raghavan, V. (2003). A linear time algorithm for computing exact Euclidean distance transforms of binary images in arbitrary dimensions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *25*(2), 265–270. https://doi.org/10.1109/TPAMI.2003.1177156

McNeil, A. J., Frey, R., & Embrechts, P. (2015). *Quantitative risk management: Concepts, techniques and tools* (Revised). Princeton University Press.

Milly, P. C. D., Betancourt, J., Falkenmark, M., Hirsch, R. M., Kundzewicz, Z. W., Lettenmaier, D. P., & Stouffer, R. J. (2008). Stationarity is dead: Whither water management? *Science*, *319*(5863), 573–574. https://doi.org/10.1126/science.1151915

Ministerie van Binnenlandse Zaken en Koninkrijksrelaties. (2022). Besluit bouwwerken leefomgeving. https://wetten.overheid.nl/BWBR0041297/
Artikel 4.27: Toegankelijkheid. In werking getreden 1 januari 2024 met de Omgevingswet.

Montanari, A., & Koutsoyiannis, D. (2014). Modeling and mitigating natural hazards: Stationarity is immortal! *Water Resources Research*, *50*(12), 9748–9756. https://doi.org/10.1002/2014WR016092

Morgan Stanley Capital International. (2020). *Climate value-at-risk: Powering better investment decisions* (tech. rep.). MSCI ESG Research. https://www.msci.com/documents/1296102/16985724/MSCI-ClimateVaR-Introduction-Feb2020.pdf

Mutlu, A., Roy, D., & Filatova, T. (2023). Capitalized value of evolving flood risks discount and nature-based solution premiums on property prices. *Ecological Economics*, *205*, 107682. https://doi.org/10.1016/j.ecolecon.2022.107682

Neal, J., Villanueva, I., Wright, N., Willis, T., Fewtrell, T., & Bates, P. (2012). How much physical complexity is needed to model flood inundation? *Hydrological Processes*, *26*(15), 2264–2282. https://doi.org/10.1002/hyp.8339

Nelen & Schuurmans. (2024, June). *WaterSchadeSchatter (WSS): Gebruikershandleiding* (tech. rep. No. STOWA 2013-11). Stichting Toegepast Onderzoek Waterbeheer (STOWA). Amersfoort. https://www.waterschadeschatter.nl/static_media/lizard_damage/help/gebruikers%20handleiding.pdf

Network for Greening the Financial System. (2024). *NGFS climate scenarios for central banks and supervisors – phase V* (tech. rep.). NGFS. https://www.ngfs.net/en/publications-and-statistics/publications/ngfs-climate-scenarios-central-banks-and-supervisors-phase-v

Planchon, O., & Darboux, F. (2002). A fast, simple and versatile algorithm to fill the depressions of digital elevation models. *CATENA*, *46*(2–3), 159–176. https://doi.org/10.1016/S0341-8162(01)00164-3

Roders, M., & Straub, A. (2015). Assessment of the likelihood of implementation strategies for climate change adaptation measures in Dutch social housing. *Building and Environment*, *83*, 168–176. https://doi.org/10.1016/j.buildenv.2014.07.014

Salas, J. D., Obeysekera, J., & Vogel, R. M. (2018). Techniques for assessing water infrastructure for nonstationary extreme events: A review. *Hydrological Sciences Journal*, *63*(3), 325–352. https://doi.org/10.1080/02626667.2018.1426858

Sandoval Olascoga, S., Shu, E., & Porter, J. (2025). Flood risk and property value changes: Understanding the impact of climate event exposure in the context of population change. *Frontiers in Environmental Economics*, *4*, 1615802. https://doi.org/10.3389/frevc.2025.1615802

Seeteram, N. A., Shi, L., Mach, K. J., Carrère, A., Chatterjee, T., Garner, A., & Horton, R. M. (2026). Challenges and opportunities in scaling climate-resilient housing solutions in the United States. *Nature Communications*, *17*, 2032. https://doi.org/10.1038/s41467-026-68595-x

Serinaldi, F., & Kilsby, C. G. (2015). Stationarity is undead: Uncertainty dominates the distribution of extremes. *Advances in Water Resources*, *77*, 17–36. https://doi.org/10.1016/j.advwatres.2014.12.013

Slager, K., & Wagenaar, D. (2017). *Standaardmethode 2017: Schade en slachtoffers als gevolg van overstromingen* (tech. rep. No. 11200580-004). Deltares.

Stichting RIONED. (2004). *Leidraad riolering, module C2100: Rioleringsberekeningen, hydraulisch functioneren* (tech. rep.). Stichting RIONED. Ede, the Netherlands. https://www.vanluijtelaar.nl/wp-content/uploads/2019/12/C2100-hydraulisch-functioneren.pdf

Swinkels, M. W. A. (2023). *Quantifying social vulnerability for decision-making on pluvial flood adaptation* [Master's thesis]. University of Twente. https://essay.utwente.nl/fileshare/file/97109/Swinkels_MA_ET.pdf

Truong, C., Malavasi, M., & Goldstein, M. A. (2024). Timing is (almost) everything: Real options, extreme value theory, climate adaptation, and flood risk management.

Journal of Environmental Management, *370*, 122621. https://doi.org/10.1016/j.jenvman.2024.122621

Tuijn, J. (2025). Structurele aanpak wateroverlast Enschede. https://www.waterforum.net/52373-structurele-aanpak-wateroverlast-enschede/

United Nations Environment Programme Finance Initiative. (2023). *Climate risks in the real estate sector: Sectoral risk briefings* (tech. rep.). UNEP FI. https://www.unepfi.org/wordpress/wp-content/uploads/2023/03/Real-Estate-Sector-Risks-Briefing.pdf

Wang, L., & Liu, H. (2006). An efficient method for identifying and filling surface depressions in digital elevation models for hydrologic analysis and modelling. *International Journal of Geographical Information Science*, *20*(2), 193–213. https://doi.org/10.1080/13658810500433453

Appendix A

Use of Generative AI Tools

In accordance with the guidelines of the Industrial Engineering and Management programme, this appendix discloses the generative AI tools used during this thesis project, the purposes for which they were used, and the manner in which their output was reviewed and verified. The intended use of AI tools was discussed with the supervisors at the start of the project.

Tools used

- **Claude (Anthropic)**, accessed via claude.ai; models Claude Opus 4.6, 4.7 and 4.8, used between February 2026 and June 2026.

Purpose of use

Generative AI was used in a supporting role only. The research design, the formulation of the research questions, the modelling choices, the analysis, and the conclusions of this thesis are my own. Specifically, AI tools were used for:

- **Text support:** suggesting grammar, style, and consistency corrections to text I had written; rewriting passages drafted by me to improve clarity and academic tone; suggesting reformulations in response to supervisor feedback; and performing a terminology consistency audit (e.g., standardising terms and checking that acronyms are introduced at first use) across chapters. The arguments, structure, findings, and interpretations in all chapters were written and led by me.
- **Coding support:** assisting in debugging the flood simulation code and analysis notebooks; suggesting implementations of components I had specified; and suggesting ways to visualise results. The simulation model, the methodological choices embedded in the code, and all transformations of data affecting the analysis were designed by me.

Review and verification

All AI-suggested output was reviewed before being incorporated. Text suggestions were checked against the underlying results and either edited or rejected where they did not accurately reflect my findings; all reported figures were verified against the simulation outputs. Code suggestions were tested on the actual data, and automated edits to thesis source files were applied through scripts with assertion-based checks and verified afterwards by inspecting diffs. Factual claims and references suggested by AI were never accepted without verification against the original sources.

Declaration

I declare that generative AI tools were used ethically, responsibly, and in line with the guidelines of the Industrial Engineering and Management programme. I take full intellectual responsibility for the content and the results of this thesis.

Appendix B

Methodology: Worked Examples and Derivations

B.1 Flow routing equations and worked examples

Flow Routing Equations. Water is routed using the local-inertial formulation of the shallow water equations (Bates et al., 2010), which retains the local acceleration term while dropping the advective inertia, yielding a momentum equation per cell face:

$$Q^{n+1} = \frac{Q^n - gh_f \Delta t\, \partial\eta/\partial x}{1 + g\Delta t n^2 |Q^n| / h_f^{7/3}}, \tag{B.1}$$

where Q is discharge per unit width (m^2/s), $\eta = w + z$ is water surface elevation, h_f is the flow depth at the face, n is Manning's roughness, and $g = 9.81\,\mathrm{m/s}^2$. Persistent flux arrays Q_x, Q_y store the discharge at each cell face and carry momentum between timesteps.

Depths update from flux divergence:

$$w_i^{n+1} = w_i^n + (\Delta t/\Delta x)(Q_{x,i-1} - Q_{x,i} + Q_{y,i-1} - Q_{y,i}). \tag{B.2}$$

Sewer abstraction and infiltration. The sewer system is modelled as a capacity-limited drain following the RIONED Leidraad (Stichting RIONED, 2004): $\Delta w_{i,t}^{\text{sewer}} = \min(w_{i,t}, q_{\text{sewer}}\Delta t)$ for non-building cells. A uniform infiltration rate is applied to all non-building cells as $w_{i,t} = \max(0, w_{i,t} - 1.5\Delta t/1000)$.

Numerical example. A single face between two adjacent street cells illustrates how Equations (B.1) and the depth update operate together (Figure B.1). Take cell i holding $w_i = 0.10\,\mathrm{m}$ of water over terrain $z_i = 1.00\,\mathrm{m\,NAP}$, and dry cell $i{+}1$ at $z_{i+1} = 0.95\,\mathrm{m\,NAP}$, both on a paved surface ($n = 0.025$), and start the simulation from rest ($Q^n = 0$). The

water-surface elevations are $\eta_i = 1.10$ and $\eta_{i+1} = 0.95\,\mathrm{m}$, the gradient at the face is $\partial\eta/\partial x = (0.95 - 1.10)/5 = -0.030$, and the flow depth at the face is $h_f = \max(\eta_i, \eta_{i+1}) - \max(z_i, z_{i+1}) = 1.10 - 1.00 = 0.10\,\mathrm{m}$. With a sub-step of $\Delta t = 1\,\mathrm{s}$, Equation (B.1) gives

$$Q^{n+1} = \frac{0 - 9.81 \cdot 0.10 \cdot 1 \cdot (-0.030)}{1 + 0} = +0.029\,\mathrm{m}^2/\mathrm{s},$$

positive (directed from the wet cell toward the dry one) as the head gradient requires, and well below the multi-face clamp $w_i \Delta x/(4\Delta t) = 0.125\,\mathrm{m}^2/\mathrm{s}$. The depth update then transfers $(\Delta t/\Delta x)\, Q^{n+1} = 0.029/5 = 5.9\,\mathrm{mm}$ of water from cell i to cell $i+1$, conserving mass exactly: cell i drops 0.006 m and cell $i+1$ rises 0.006 m. At the following sub-step two changes appear simultaneously: the head gradient has shrunk because the cells have partially equalised, and Q^n is now non-zero, so the friction denominator in Equation (B.1) exceeds unity and damps further growth of the flux. The persistent memory of Q^n between sub-steps is what distinguishes the local-inertial scheme from the diffusive-wave approximation (Neal et al., 2012), which would recompute the flux from scratch at every step with $\Delta x = 5\,\mathrm{m}$.

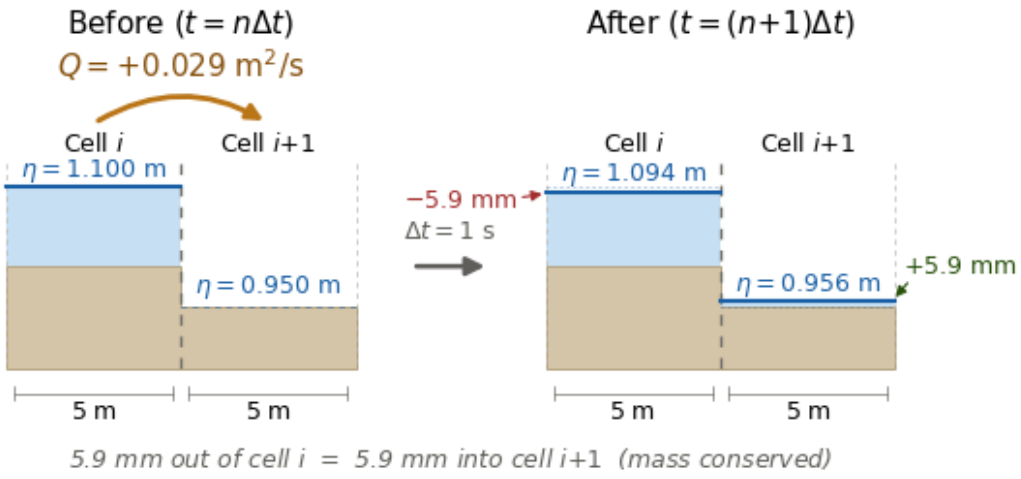


Figure B.1: *Before and after of a single 1-second routing sub-step for the worked example.* Left: *the initial state with $w_i = 0.10\,m$ of water on cell i and a dry cell i+1; the water-surface gradient drives a discharge $Q = +0.029\,m^2/s$ across the face.* Right: *the state one second later, after the depth update of Equation (B.2): cell i has lost 5.9 mm of water depth and cell i+1 has gained the same 5.9 mm, exactly conserving mass. Faint dashed lines mark the original water-surface elevations for reference. Vertical scale exaggerated relative to horizontal.*

The single-face example demonstrates Equations (B.1) and (B.2) for one cell pair, but in production every cell exchanges discharge with all four neighbours simultaneously. The depth update of Equation (B.2) sums the four face contributions, and the sign of each is set by the per-face momentum equation of Equation (B.1). Figure B.2 visualises a representative five-cell stencil and shows how the four face discharges combine to a net

depth change for the central cell.

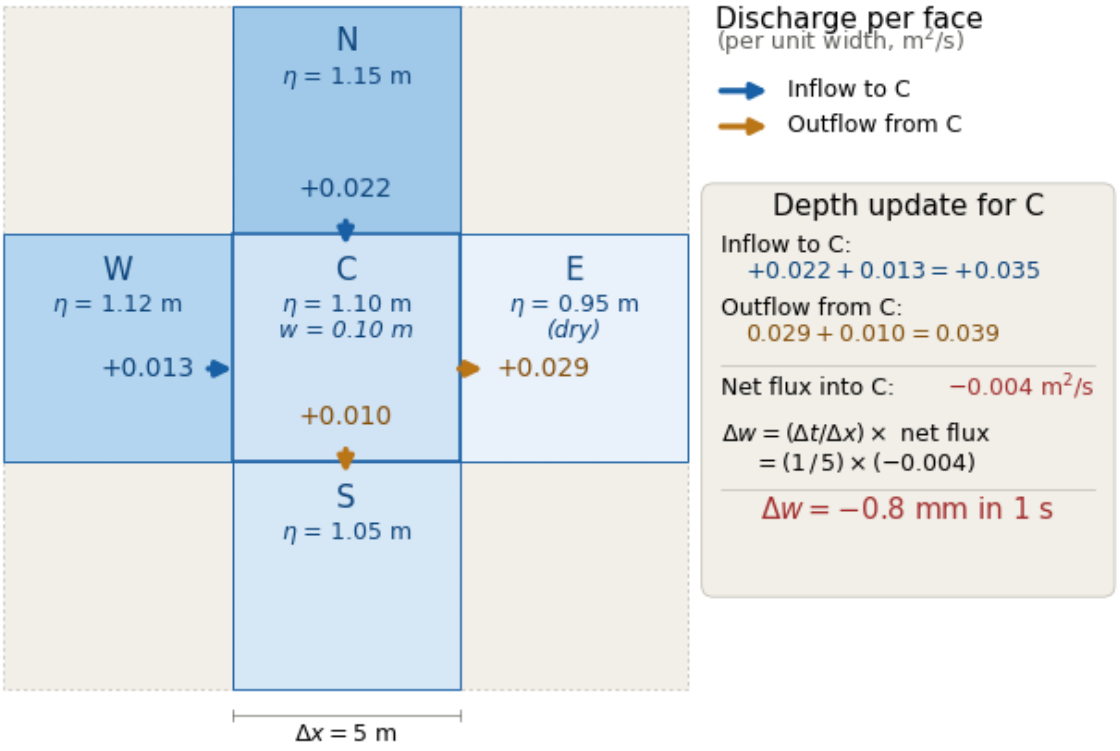


Figure B.2: *Plan view of the five-cell stencil that the local-inertial scheme operates on at every timestep, illustrating how Equation (B.1) (per-face momentum) and the depth update of Equation (B.2) (cell-level continuity) interact.*

The central cell C exchanges discharge with four neighbours N, E, S, W across its four faces, and the depth update of Equation (B.2) sums the four face contributions. Each cell label reports the water-surface elevation $\eta = w + z$; for cell C the standing water depth $w = 0.10\,\mathrm{m}$ is also shown, and cell E is marked as dry. Arrows mark the discharge per unit width Q at each face, computed from the local water-surface gradient and friction-conditioned momentum carried over from the previous sub-step via Equation (B.1): blue arrows indicate inflow into C (from higher-head neighbours N and W), amber arrows indicate outflow from C (toward the lower-head neighbours E and S). The numerical values are schematic rather than exactly back-derived from Equation (B.1); they are chosen to balance to a net inflow of $-0.004\,\mathrm{m}^2/\mathrm{s}$ that propagates through the depth update to a depth change of $\Delta w = -0.8\,\mathrm{mm}$ in the $\Delta t = 1\,\mathrm{s}$ sub-step. The corner cells (dashed grey) are not directly coupled to C in the five-point stencil of Equation (B.2).

B.2 Routing-surface worked example

This example accompanies the routing-surface construction of Section 4.4 and illustrates how Equation (4.1) treats three adjacent cells.

Numerical example. Three adjacent 5 m cells in a typical Dutch street scene illustrate how Equation (4.1) operates. All elevations are in metres above NAP.

- *Street cell* (not in $\mathcal{B}$): the filled DTM gives $z_i^{\text{fill}} = 1.20\,\text{m}$. The second branch applies, so $z_i^{\text{route}} = 1.20\,\text{m}$. Water routes across this cell freely, governed only by terrain and friction.
- *Typical building cell* ($i \in \mathcal{B}$): $z_i^{\text{fill}} = 1.20\,\text{m}$ and the DSM returns $d_i = 9.50\,\text{m}$, capturing a roughly 8 m-tall row house including its roof. The first branch gives $z_i^{\text{route}} = \max(9.50,\, 1.20) = 9.50\,\text{m}$, placing the cell 8.30 m above the adjacent street. The momentum equation (4.6.3) cannot push water uphill against such a gradient, so the building acts as an impermeable barrier and the routing scheme redirects flow around its perimeter.
- *Artefact building cell* ($i \in \mathcal{B}$): $z_i^{\text{fill}} = 1.20\,\text{m}$ but the DSM returns the anomalously low value $d_i = 1.05\,\text{m}$, because the LiDAR pulse penetrated a roof opening (skylight, ventilation shaft, courtyard cutout) and returned a ground-level echo. Without the maximum, the cell would sit 0.15 m below the street and water would route *into* the building rather than around it. With the maximum, $z_i^{\text{route}} = \max(1.05,\, 1.20) = 1.20\,\text{m}$: the cell is restored to street level and inflow is prevented. This case is rare (1% of building cells in the Amsterdam benchmark) but, if left uncorrected, produces visible artefact-driven damage clusters at the locations of glass-roofed and partially-covered buildings.

B.3 Worked Value-at-Risk and Expected Shortfall calculations

These worked calculations accompany Section 4.8 and show, step by step, how the loss-exceedance curve is read and how VaR and ES are computed at the reported anchors.

Reading a probability off the curve

As a numerical illustration, between the $T = 100$ anchor ($D = €457\,\text{M}$, $\pi = 0.01$) and the $T = 500$ anchor ($D = €500\,\text{M}$, $\pi = 0.002$), a portfolio damage of €478 M corresponds to $\widehat{G}(478) = 0.01 + \frac{0.002-0.01}{500-457}(478-457) \approx 0.006$, i.e. an annual exceedance probability of $\approx 0.6\%$ or a return period of ≈ 165 years.

Worked Value-at-Risk calculation

As a numerical example, $\mathrm{VaR}_{0.95}$ solves $\widehat{G}(d) = 0.05$, which falls between the $T = 10$ anchor ($\pi = 0.10$, $D = €296\,\mathrm{M}$) and the $T = 25$ anchor ($\pi = 0.04$, $D = €354\,\mathrm{M}$); inverting Equation (4.7) gives

$$\mathrm{VaR}_{0.95} = 296 + \frac{354 - 296}{0.10 - 0.04}\,(0.10 - 0.05) \approx €344\,\mathrm{M}.$$

At $\alpha = 0.99$ the target $\widehat{G}(d) = 0.01$ coincides with the $T = 100$ anchor, so $\mathrm{VaR}_{0.99} = €457\,\mathrm{M}$ directly.

Worked Expected Shortfall calculation

As a numerical illustration, $\widehat{\mathrm{ES}}_{0.99}$ uses VaR values at the three anchors in the truncated range, $u = 0.99$ ($€457\,\mathrm{M}$), $u = 0.998$ ($€500\,\mathrm{M}$), and $u = 0.999$ ($€525\,\mathrm{M}$):

$$\int_{0.99}^{0.999} \mathrm{VaR}_u\,\mathrm{d}u \approx \tfrac{0.008}{2}(457 + 500) + \tfrac{0.001}{2}(500 + 525) \approx 4.34\ €\mathrm{M},$$

yielding $\widehat{\mathrm{ES}}_{0.99} \approx 4.34/0.009 \approx €483\,\mathrm{M}$.

B.4 Loss-exceedance curve interpolation: linear versus $(\ln T, D)$

This note expands the comparison of the two interpolation schemes for the loss-exceedance curve referenced in Section 4.8.

A natural alternative is to interpolate in $(\ln T, D)$ space and recover $\pi = 1/T$, equivalent to piecewise-Gumbel patches between control points; this implicitly imposes a thin-tailed shape on the LEC between simulated anchors. The two schemes coincide at every anchor by construction and differ only in their interior. Under the SSP5-8.5 mid-century production LEC, the two schemes agree closely: the maximum discrepancy across the 90–99% confidence grid is below 2 % (e.g. $\mathrm{VaR}_{0.95}^{(\ln T, D)} = €340\,\mathrm{M}$ versus $\mathrm{VaR}_{0.95}^{(D,\pi)} = €344\,\mathrm{M}$, a relative deviation of 1.3 %); the $\mathrm{ES}_{0.99}$ deviation is around 1 %. Both are an order of magnitude below the validation cross-square spread (4.9.1) and the threshold sensitivity (5.10), and do not change any conclusion of the thesis.

B.5 Micro-noise conditioning of the terrain

This note expands the micro-noise safeguard applied to the DTM in Section 4.2.

The local-inertial routing formulation (4.6.3) carries momentum between timesteps and avoids the well-known flat-terrain failure mode of steepest-descent algorithms (see Wang & Liu, 2006). It nevertheless requires a non-zero surface gradient to set an initial flow direction. Perfectly flat patches, common in Dutch polder LiDAR records, where elevation is constant to centimetre precision over large areas, yield a zero gradient, so the routing arrays have no direction in which to push water and it ponds statically. A small uniform random perturbation $\epsilon_i \sim \mathcal{U}(0, 10^{-4}\,\mathrm{m})$ added to the DTM before routing breaks these ties. At 0.1 mm maximum it lies four orders of magnitude below the depths of interest and cannot affect the damage estimates.

B.6 Calibration of the depression-filling threshold

This note expands the choice of the $\delta_{\min} = 0.10\,\mathrm{m}$ depression-filling threshold in Section 4.2.

Depression filling raises low-lying cells until water can drain to the edge of the domain, eliminating "holes" in the terrain. In the Netherlands this is delicate: many real depressions (sunken courtyards, ditches, retention basins) genuinely store rainwater and should be preserved, while LiDAR measurement noise produces spurious sub-decimetre holes that must be removed. The threshold $\delta_{\min}$ is the dividing line between these two cases.

The 10 cm threshold was calibrated against the Tygron reference simulation for the Amsterdam benchmark square: with higher thresholds, real retention features began to be filled and the model under-predicted local ponding; with lower thresholds, residual LiDAR noise left enough micro-depressions to trap rainfall everywhere, producing a "wet everywhere" artefact absent in the Tygron benchmark. The 10 cm value is the point at which both failure modes were avoided.

B.7 Thermodynamic basis of the super-Clausius–Clapeyron scaling

This note expands the rainfall-intensification factor applied in Section 4.6.1.

The Clausius–Clapeyron (CC) relation states that the atmosphere can hold approximately 7% more water vapour per degree of warming, providing a thermodynamic lower bound on rainfall intensification under climate change. The "super-CC" factor of 14% per °C adopted here reflects the empirical finding of Lenderink and van Meijgaard (2008) that short-duration convective rainfall extremes in the Netherlands have intensified at roughly twice the basic CC rate, attributed to dynamic feedbacks (latent-heat release, stronger updrafts) inside developing storms.

B.8 Choice of the inundation threshold

This note expands why neither normative reference point is adopted directly for the inundation threshold d_{th} of Section 4.7.3.

The Dutch building decree (*Besluit bouwwerken leefomgeving*, art. 4.27) caps the height between finished floor and adjacent terrain at 20 mm at the main entrance (Ministerie van Binnenlandse Zaken en Koninkrijksrelaties, 2022), but this norm applies only to new-build front entrances and is therefore not representative of an existing, mixed-age portfolio. The *Waterschadeschatter* (WSS) implicitly assumes 15 cm above terrain when reconstructing building heights from filtered AHN data (Nelen & Schuurmans, 2024), but this overstates the effective inflow height, because the lowest inflow opening, which governs damage during a pluvial event, is rarely the front door. The 5 cm baseline sits between the two and is intended as a plausible weighted average across the portfolio; the $\{5, 10, 15\}$ cm sensitivity analysis of Chapter 5 maps the consequences of this choice.

B.9 Tail handling beyond the simulated range

This note expands the treatment of confidence levels beyond the simulated range ($\alpha > 0.999$) for the VaR estimator of Section 4.8.

The highest simulated scenario corresponds to $T = 1{,}000$ years ($\pi_7 = 0.001$), so VaR at confidence levels up to $\alpha = 0.999$ lies within the simulated range. For $\alpha > 0.999$, three treatments were considered: (1) clamping the curve at D_7; (2) extrapolating linearly in $(\ln T, D)$ space, consistent with the Gumbel extrapolation of 4.6.1; and (3) fitting a Generalised Pareto Distribution, the standard tool of extreme-value theory (EVT) for describing the rare upper tail of a distribution. The primary analysis adopts (1); reported VaR values are restricted to $\alpha \leq 0.999$ and interpreted as the in-sample tail of the seven-scenario distribution. Treatment (2) is reported as a robustness check in Chapter 5, and (3) is identified as a methodological extension in Chapter 6.

B.10 Truncated Expected Shortfall estimator

This note expands the truncated Expected Shortfall estimator of Section 4.8.

Because $\widehat{G}$ is piecewise linear, the integral in Equation (4.9) is evaluated by composite trapezoidal rule. Integration is truncated at the upper edge of the simulated range $u_m = 0.999$, and the normalisation is set to $1/(u_m - \alpha)$, so that the reported estimator is the mean VaR over the simulated quantile range $[\alpha, u_m]$:

$$\widehat{\mathrm{ES}}_\alpha = \frac{1}{u_m - \alpha} \int_\alpha^{u_m} \mathrm{VaR}_u \, \mathrm{d}u. \tag{B.3}$$

This estimator coincides with the theoretical ES of Equation (4.9) when the contribution from $u > u_m$ is negligible; for the present application, the very deep tail beyond $T =$ 1,000 years is unconstrained by simulation and is excluded from the reported figure on those grounds.

B.11 Sink-mask detection parameters

This note records the exact detection thresholds for the two automatically detected sink categories of Section 4.3.

Local depressions. A cell is flagged as a local depression if it lies more than 2.5 m below the mean elevation of its 55 m neighbourhood (computed with an 11-cell uniform filter). These cells represent underpasses or tunnel ramps.

Large flat zones. The DTM is rounded to 1 mm precision, elevations appearing in more than 500 cells are identified, and connected components exceeding 500 cells (12,500 m^2) are flagged. This correctly identifies the engineered managed-drainage areas of the Dutch polder landscape without per-square calibration.

Appendix C

Full Cross-Scenario Validation Table

Table C.1 reports the complete set of per-run validation metrics for the 35-run SSP5-8.5 2050 cross-scenario comparison against Tygron summarised in Section 4.9. Each row corresponds to one (square, return period) pair.

Table C.1: *Full per-run validation metrics for the 35-run SSP5-8.5 2050 cross-scenario comparison. Columns: damage ratio (Tygron / local-inertial), grid-level RMSE (m), mean depth bias (local minus Tygron, m), number of ground-floor units flagged by both models ("Both"), by the local-inertial model only ("Ours"), by Tygron only ("Tyg"), and by neither ("None"). Concordant percentage is $(n_{Both}+n_{None})/n_{GF}$; conditional agreement is $n_{Both}/(n_{Both}+n_{Ours}+n_{Tyg})$.*

Sq	T	mm	Ratio	RMSE	Bias	Both	Ours	Tyg	None	Conc.%	Cond.%
0	T5	38	0.63	0.064	0.0024	713	1,182	665	2,210	61.3	27.9
0	T10	46	0.67	0.075	0.0026	936	1,160	693	1,981	61.2	33.6
0	T25	58	0.72	0.089	0.0023	1,124	1,186	745	1,715	59.5	36.8
0	T50	71	0.76	0.103	0.0018	1,312	1,181	763	1,514	59.2	40.3
0	T100	84	0.79	0.116	0.0008	1,516	1,184	738	1,332	59.7	44.1
0	T500	99	0.82	0.129	-0.0007	1,679	1,162	725	1,204	60.4	47.1
0	T1000	109	0.84	0.137	-0.0019	1,777	1,119	736	1,138	61.1	48.9
4	T5	38	0.33	0.051	0.0064	202	728	129	1,798	70.0	19.1
4	T10	46	0.42	0.060	0.0063	276	739	184	1,658	67.7	23.0
4	T25	58	0.51	0.072	0.0053	357	765	244	1,491	64.7	26.1
4	T50	71	0.58	0.085	0.0040	432	799	308	1,318	61.3	28.1
4	T100	84	0.65	0.096	0.0021	509	823	369	1,156	58.3	29.9
4	T500	99	0.72	0.109	-0.0005	567	808	391	1,091	58.0	32.1
4	T1000	109	0.78	0.117	-0.0025	652	771	403	1,031	58.9	35.7
5	T5	38	0.33	0.051	0.0030	334	1,187	323	2,967	68.6	18.1
5	T10	46	0.43	0.060	0.0021	448	1,178	485	2,700	65.4	21.2
5	T25	58	0.54	0.072	0.0002	615	1,163	677	2,356	61.8	25.1
5	T50	71	0.64	0.083	-0.0020	701	1,156	837	2,117	58.6	26.0
5	T100	84	0.73	0.094	-0.0046	1,019	904	973	1,915	61.0	35.2
5	T500	99	0.86	0.106	-0.0078	1,142	839	1,188	1,642	57.9	36.0
5	T1000	109	0.83	0.109	-0.0014	1,087	927	1,206	1,591	55.7	33.8
7	T5	38	0.30	0.060	0.0051	411	2,339	633	4,595	62.7	12.1
7	T10	46	0.38	0.070	0.0045	718	2,298	763	4,199	61.6	19.0
7	T25	58	0.49	0.084	0.0024	1,042	2,262	1,037	3,637	58.6	24.0
7	T50	71	0.60	0.099	-0.0002	1,367	2,107	1,214	3,290	58.4	29.2
7	T100	84	0.68	0.113	-0.0037	1,616	2,044	1,369	2,949	57.2	32.1
7	T500	99	0.79	0.128	-0.0085	1,941	1,921	1,584	2,532	56.1	35.6
7	T1000	109	0.87	0.138	-0.0121	2,130	1,763	1,662	2,423	57.1	38.3
8	T5	38	0.43	0.052	0.0052	325	928	379	2,849	70.8	19.9
8	T10	46	0.50	0.061	0.0050	470	1,075	424	2,512	66.5	23.9
8	T25	58	0.62	0.074	0.0039	659	1,192	580	2,050	60.5	27.1
8	T50	71	0.71	0.087	0.0025	846	1,086	591	1,958	62.6	33.5
8	T100	84	0.79	0.098	0.0004	929	1,034	701	1,817	61.3	34.9
8	T500	99	0.86	0.110	-0.0024	1,037	981	785	1,678	60.6	37.0
8	T1000	109	0.89	0.118	-0.0046	1,080	991	791	1,619	60.2	37.7

The mean-bias column reports the mean difference (local-inertial minus Tygron) across non-sink cells in the 5 m grid. Positive values indicate that the local-inertial model is wetter on average than Tygron. The pattern across the 35 runs is non-monotonic: bias is mildly positive at low rainfall totals (averaging +4.4 mm at $T = 5$ across the five squares) and becomes slightly negative at the highest return periods (averaging −4.5 mm at $T = 1000$). The same sign reversal is visible within each square individually. This is consistent with the convergence of the two models toward similar inundation extents under the most extreme rainfall scenarios, while the residual disagreement at intermediate return periods is dominated by structural model-scope differences (spatial sewer heterogeneity, building-perimeter depth-sampling conventions, and the MGW first-unit-per-pand proxy) rather than by a single systematic over- or under-routing tendency. The mean-bias values shown here are grid-cell averages; at building perimeters, where the depth-damage function

is sampled, typical biases are larger because water preferentially ponds against urban frontages.

Appendix D

Full Sewer-Parameterisation Sensitivity Table

This appendix documents the full per-run results of the sewer-parameterisation sensitivity experiment summarised in Section 4.9.3. The experiment raises the sewer abstraction rate q_{sewer} from 8 to 11 mm/hr and the initial sewer buffer from 8 to 10 mm, the upper end of the RIONED design range (Stichting RIONED, 2004), raising the total absorbable rainfall over the 120-minute event from ∼21 mm to ∼28 mm. The local-inertial simulation is re-run on validation squares 4, 5, 7, and 8 across all seven SSP5-8.5 2050 return-period scenarios, holding all other parameters and the Tygron benchmark fixed.

Aggregate result. Across the 28 paired runs the per-run damage ratio r = Tygron/local-inertial shifts by either +0.00 or +0.01, with the cross-run aggregate ratio rising from $\bar{r} = 0.486$ at the baseline parameterisation to $\bar{r} = 0.529$ under the raised parameterisation, a 4-percentage-point shift in response to a 37% increase in rate-limited sewer capacity and a 25% increase in initial absorption capacity. The within-run shifts are an order of magnitude smaller than the cross-square spread of damage ratios documented in Table 4.4, falsifying the hypothesis that sewer parameterisation within the RIONED design range is the dominant driver of the residual bias against Tygron. Instead, the bias points to structural differences in model scope: the spatial heterogeneity of the sewer network, the building-depth sampling rule, and the floor-level proxy for MGW units, none of which are tunable through sewer parameters.

Per-run table. Table D.1 reports the per-run damage ratio under the two parameter settings. Square 0 is excluded because its baseline $T = 5$ damage ratio of 0.57 already sits at the high end of the observed range, leaving little headroom for further upward movement.

Table D.1: *Per-run damage ratio (Tygron / local-inertial) under the baseline sewer parameterisation ($q_{sewer} = 8\,mm/hr$, $buffer = 8\,mm$) and the raised parameterisation ($q_{sewer} = 11\,mm/hr$, $buffer = 10\,mm$), for the four validation squares and seven SSP5-8.5 2050 return periods that constitute the sensitivity matrix.*

Square	T (yr)	mm	Ratio (baseline)	Ratio (raised)
4	5	38	0.29	0.30
4	10	46	0.36	0.37
4	25	58	0.43	0.44
4	50	71	0.48	0.49
4	100	84	0.53	0.54
4	500	99	0.58	0.58
4	1000	109	0.62	0.62
5	5	38	0.30	0.31
5	10	46	0.39	0.40
5	25	58	0.48	0.49
5	50	71	0.55	0.55
5	100	84	0.61	0.61
5	500	99	0.71	0.71
5	1000	109	0.68	0.68
7	5	38	0.28	0.29
7	10	46	0.34	0.35
7	25	58	0.43	0.44
7	50	71	0.50	0.50
7	100	84	0.56	0.56
7	500	99	0.64	0.64
7	1000	109	0.69	0.69
8	5	38	0.39	0.40
8	10	46	0.45	0.46
8	25	58	0.55	0.55
8	50	71	0.61	0.61
8	100	84	0.66	0.66
8	500	99	0.70	0.70
8	1000	109	0.72	0.72
Cross-run aggregate			**0.486**	**0.529**

Appendix E

Results: Robustness Checks

This appendix collects a robustness check referenced in Chapter 5 that is not required for the main 90–99.9% results.

E.1 Tail extrapolation robustness

The primary analysis avoids extrapolating beyond the 1,000-year event. This is consistent with the limited number of physically simulated scenarios and the uncertainty in the far tail of rainfall extremes. Figure E.1 reports the robustness check using a log-return-period extrapolation. The check does not change the main conclusions within the reported 90–99.9% confidence range: VaR at the scenario anchors is identical by construction, and the difference at interior quantiles is small relative to both the validation spread and the threshold sensitivity.

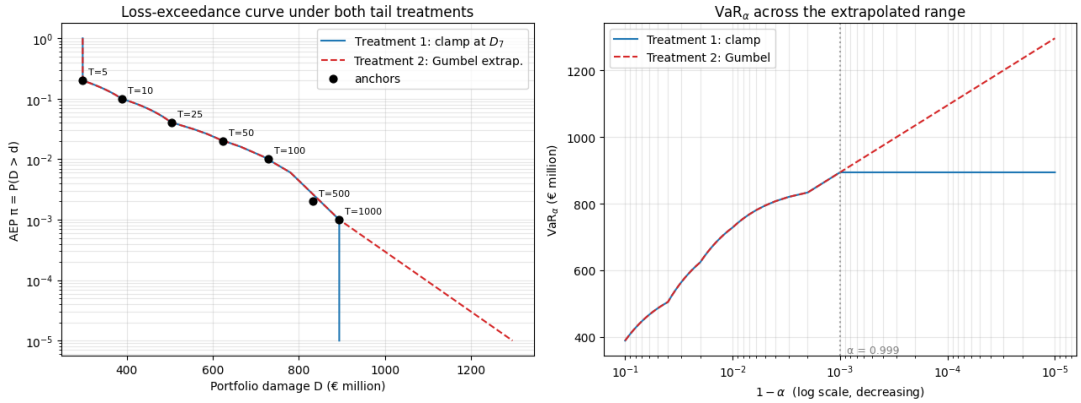


Figure E.1: *Tail extrapolation robustness check for the SSP5-8.5 2050 loss-exceedance curve. The main analysis reports VaR only within the simulated range up to* $\alpha = 0.999$.

The robustness check supports the choice made in Section 4.8: extrapolation beyond the simulated range is not required for the main 90–99.9% results. A peaks-over-threshold or non-stationary EVT extension could be useful in future work, but would require more

than seven physically simulated control points to be statistically defensible.

Appendix F

Results: Spatial Damage and EAD Maps

This appendix collects the full choropleth maps for the spatial 100-year damage and Expected Annual Damage results of Chapter 5. The body of that chapter reports the top-five regions for each map in tabular form; the maps below show the complete national distribution.

F.1 100-year direct structural damage

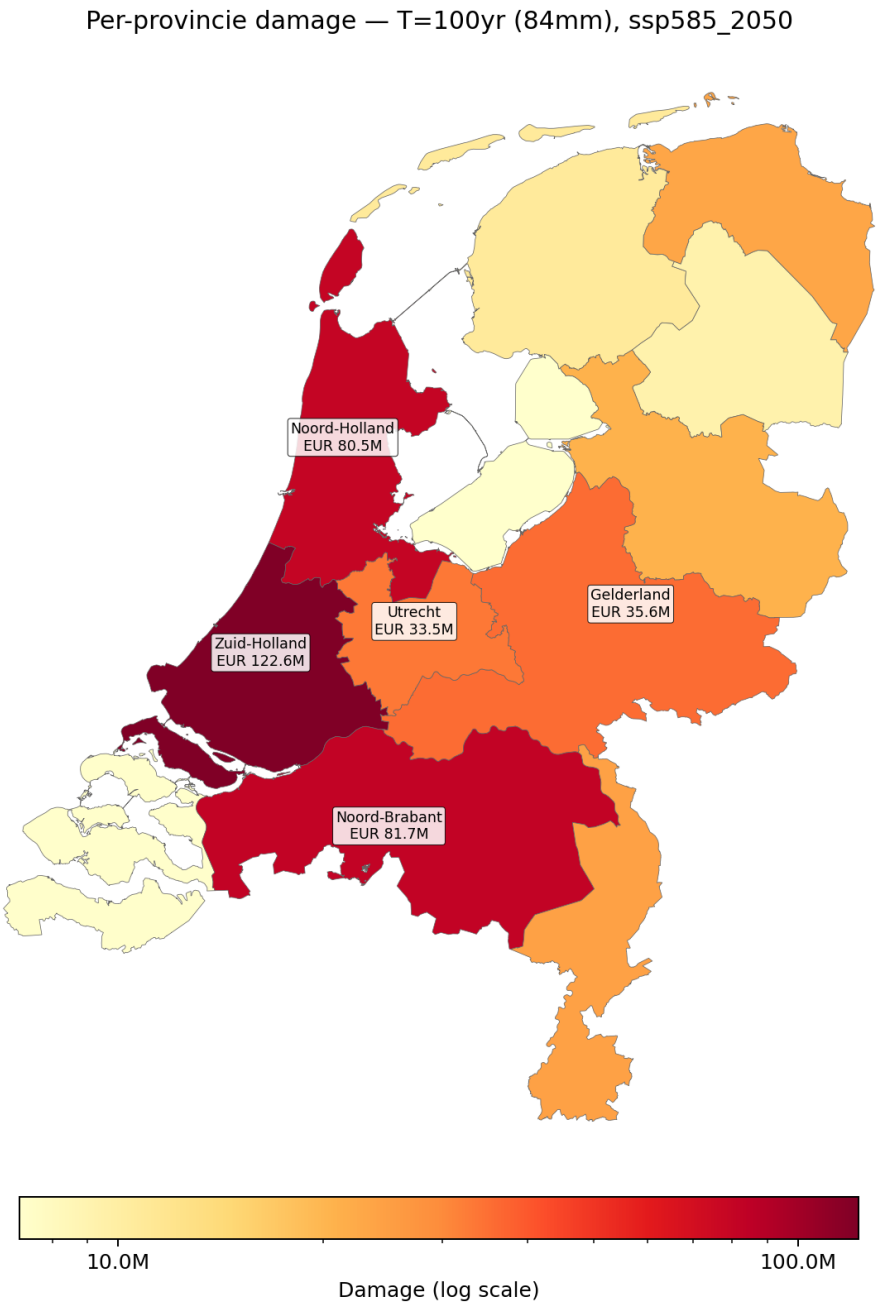


Figure F.1: *Total 100-year direct structural damage by province under SSP5-8.5 2050. Colours use a logarithmic scale.*

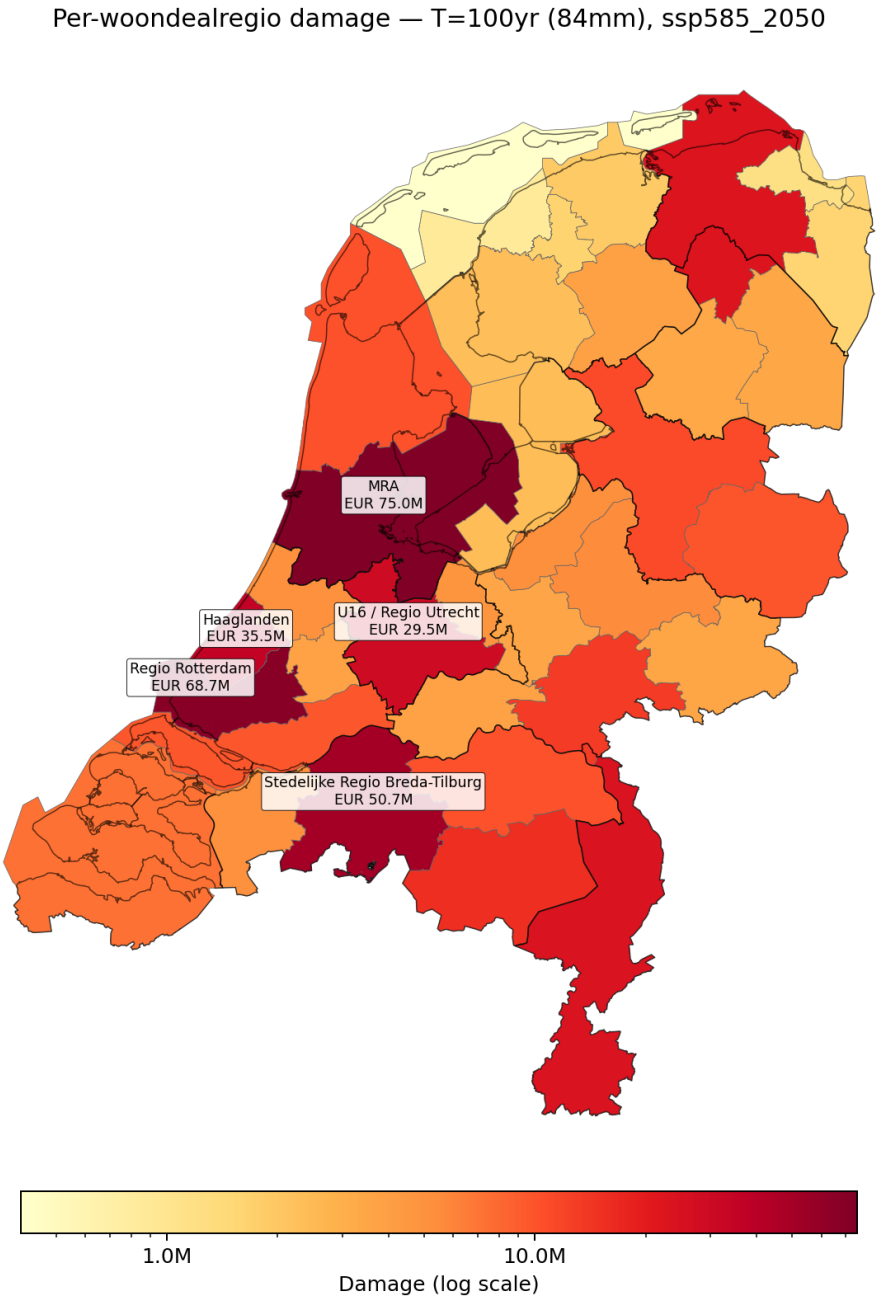


Figure F.2: *Total 100-year direct structural damage by Woondealregio under SSP5-8.5 2050. Colours use a logarithmic scale.*

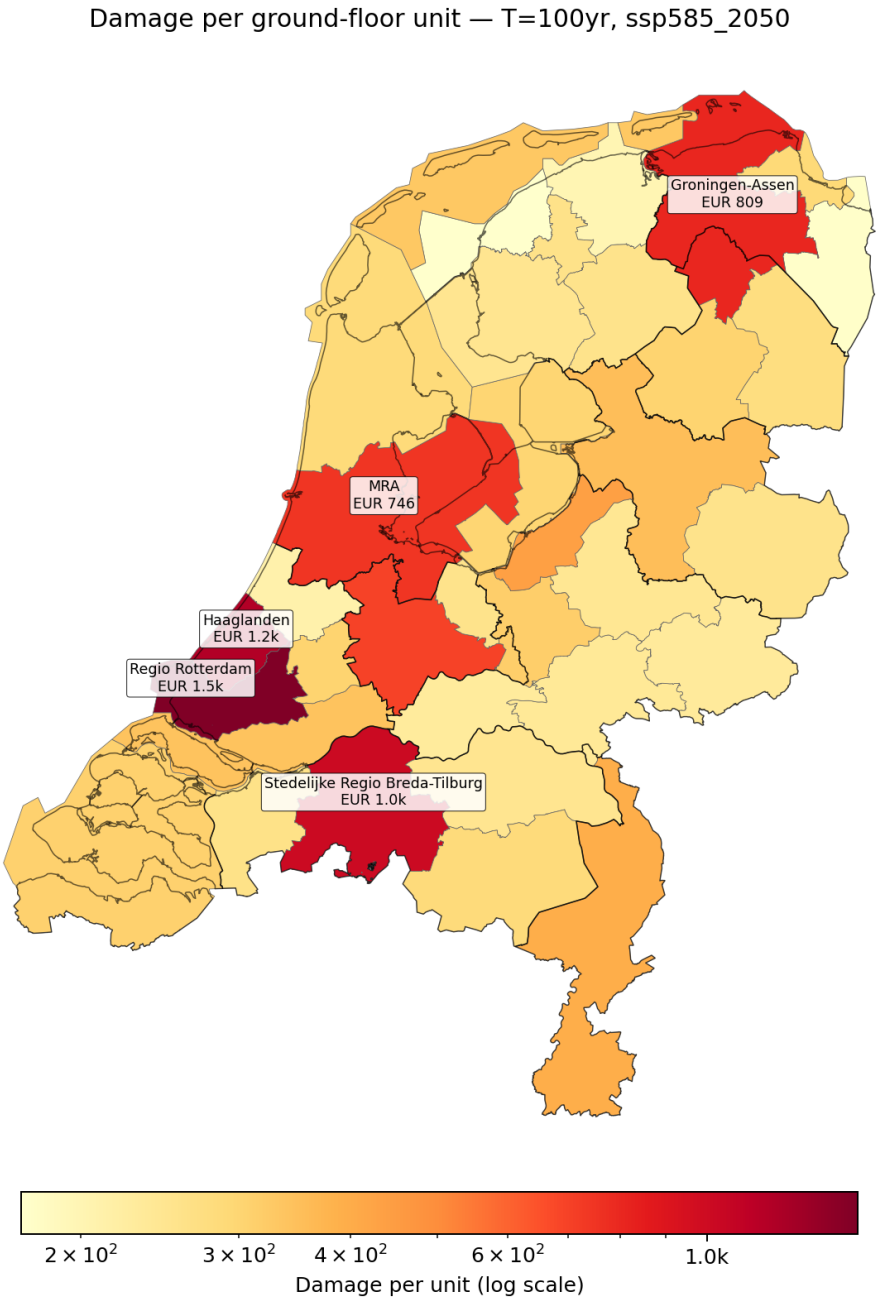


Figure F.3: *100-year direct structural damage per ground-floor unit by Woondealregio under SSP5-8.5 2050. Colours use a logarithmic scale.*

F.2 Expected Annual Damage

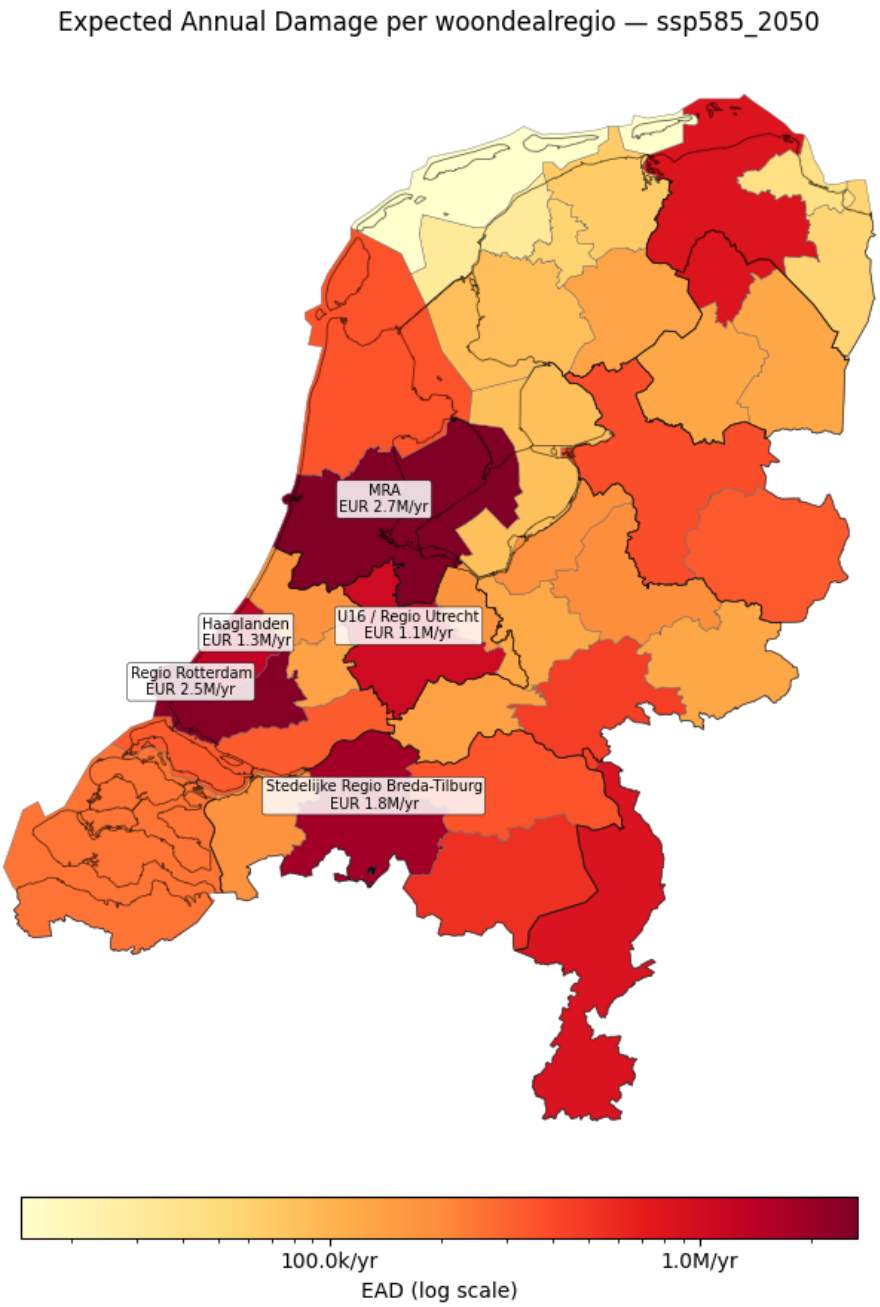


Figure F.4: *Expected Annual Damage by Woondealregio under SSP5-8.5 2050. Colours use a logarithmic scale.*

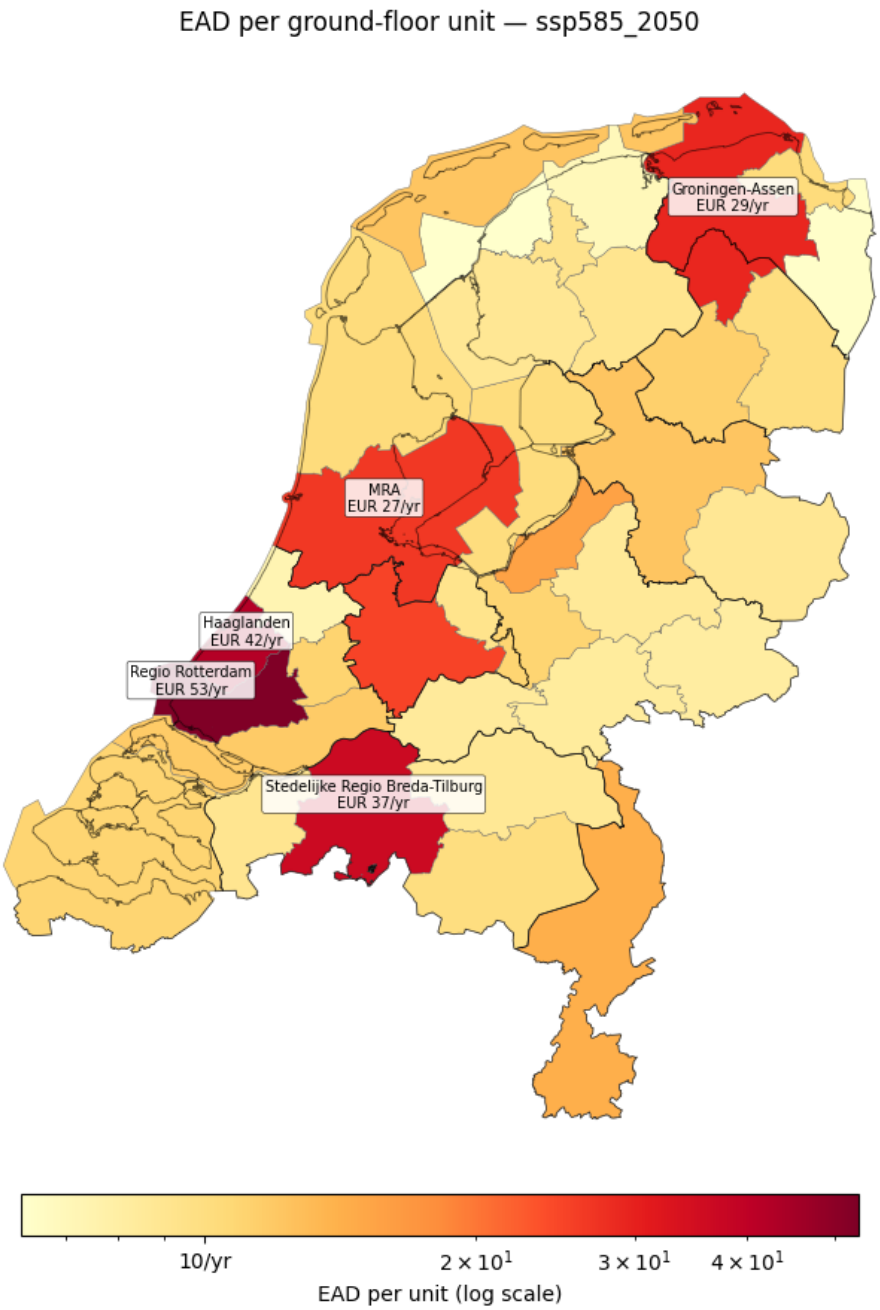


Figure F.5: *Expected Annual Damage per ground-floor unit by Woondealregio under SSP5-8.5 2050. Colours use a logarithmic scale.*

Appendix G

Operating-cashflow buffer per Woondealregio

This appendix reports the full damage-versus-buffer table. Each row aggregates the DAEB-tak cash items of the corporations operating in that Woondealregio, allocated *pro rata* to the region's share of the corporation's national dwelling count. All cash and damage figures are expressed in millions of euros and refer to verslagjaar 2024; the T=100 yr damage estimate is taken from the SSP5-8.5 mid-century pathway. Regions are sorted by descending damage-to-FCF ratio, so that the regions with the tightest absorption headroom appear first. The data covers 34 of the 35 Woondealregio's defined under the 2022 *Nationale Prestatieafspraken*; the omitted region holds a small share of the corporation footprint and was not matched to a DAEB-reporting toegelaten instelling under the dVi2024 filter conventions.

Table G.1: *T=100 yr pluvial damage and operating-cashflow buffer per Woondealregio, verslagjaar 2024.* VHE *is the DAEB-tak self-contained and shared-housing dwelling count from the dVi2024* Gemeente Aantal Ultimo *table.* FCF *is the operating cashflow (*Totaal van kasstroom uit operationele activiteiten*),* CF_{inv} *the standalone cashflow from investing activities (negative = net investment outflow, so that* $FCF + CF_{inv}$ *equals the residual after the year's investment programme), and* Cash *the closing cash position.* Damage *is the per-square pluvial damage at the T=100 yr SSP5-8.5 2050 anchor, summed over corporation-owned panden within the region.* D/FCF *is the damage-to-free-cashflow ratio;* D/(F+C) *the damage-to-(FCF+Cash) ratio. Negative Cash values reflect working-capital overdrafts on the balance-sheet reporting date.*

Woondealregio	VHE	FCF (€M)	CF_{inv} (€M)	Cash (€M)	Damage (€M)	D/FCF	D/(F+C)
Stedelijke Regio Breda-Tilburg	96,204	95.1	-385.8	19.5	50.7	0.53	0.44
Groningen-Assen	52,950	64.4	-178.4	6.9	23.2	0.36	0.33
Regio Rotterdam	201,439	216.4	-243.3	217.9	68.7	0.32	0.16
U16 / Regio Utrecht	110,938	143.4	-469.6	20.7	29.5	0.21	0.18
Noord-Veluwe	18,872	25.6	-56.4	8.2	5.2	0.20	0.15
Limburg	128,446	124.8	-310.8	14.6	24.1	0.19	0.17
Zuid-Oost Drenthe	19,980	18.0	-65.2	5.8	3.5	0.19	0.14
De Wadden	998	0.8	-2.8	0.4	0.2	0.19	0.13
Zuid-West Drenthe	18,028	18.1	-49.6	5.1	3.4	0.19	0.15
Metropoolregio Amsterdam	365,045	405.6	-1,119.1	24.8	75.0	0.18	0.17
Haaglanden	154,734	198.1	-333.9	362.6	35.5	0.18	0.06
Flevoland	12,165	13.5	-27.0	2.2	2.4	0.18	0.15
Noordwest Friesland	2,498	5.2	-4.8	0.0	0.9	0.17	0.17
Zuidoost-Brabant	96,874	106.9	-406.5	-57.5	15.7	0.15	0.32
Zeeland	39,955	51.4	-125.0	-17.4	7.3	0.14	0.21
West-Overijssel	59,908	81.6	-209.1	-1.2	11.1	0.14	0.14
Zuidoost Friesland	22,971	28.8	-37.3	-2.8	3.8	0.13	0.15
Zuidwest Friesland	15,353	18.0	-21.5	0.2	2.4	0.13	0.13
Achterhoek	24,202	28.2	-60.4	0.2	3.6	0.13	0.13
West-Brabant West	31,831	40.7	-107.8	2.7	5.2	0.13	0.12
Samenwerkende regio's	50,034	78.0	-202.6	-14.4	9.6	0.12	0.15
Noord-Holland-Noord	61,155	83.1	-226.3	26.8	10.2	0.12	0.09
Midden-Holland	26,288	34.4	-120.6	13.2	3.8	0.11	0.08
Arnhem-Nijmegen	108,852	123.6	-287.7	30.4	13.2	0.11	0.09
Leeuwarden	17,425	15.7	-42.2	7.5	1.6	0.10	0.07
Twente	75,380	103.7	-291.9	7.2	9.8	0.09	0.09
Rivierenland	19,925	41.9	-46.5	-1.3	3.8	0.09	0.09
Amersfoort	39,885	55.2	-93.1	20.4	4.8	0.09	0.06
Noordoost Friesland	8,701	22.3	-22.0	3.3	1.9	0.09	0.08
Stedelijke Regio 's-Hertogenbosch	52,647	117.3	-183.8	-3.0	10.2	0.09	0.09
Stedendriehoek	39,982	66.3	-90.3	3.6	5.5	0.08	0.08
Oost-Groningen	13,199	19.5	-107.6	13.8	1.6	0.08	0.05
Holland Rijnland	65,561	75.2	-207.9	2.4	5.1	0.07	0.07
Foodvalley	40,965	59.8	-89.4	4.1	3.5	0.06	0.05
Sector aggregate	**2,093,390**	**2,580.6**	**-6,226.1**	**726.7**	**456.0**	**0.18**	**0.14**

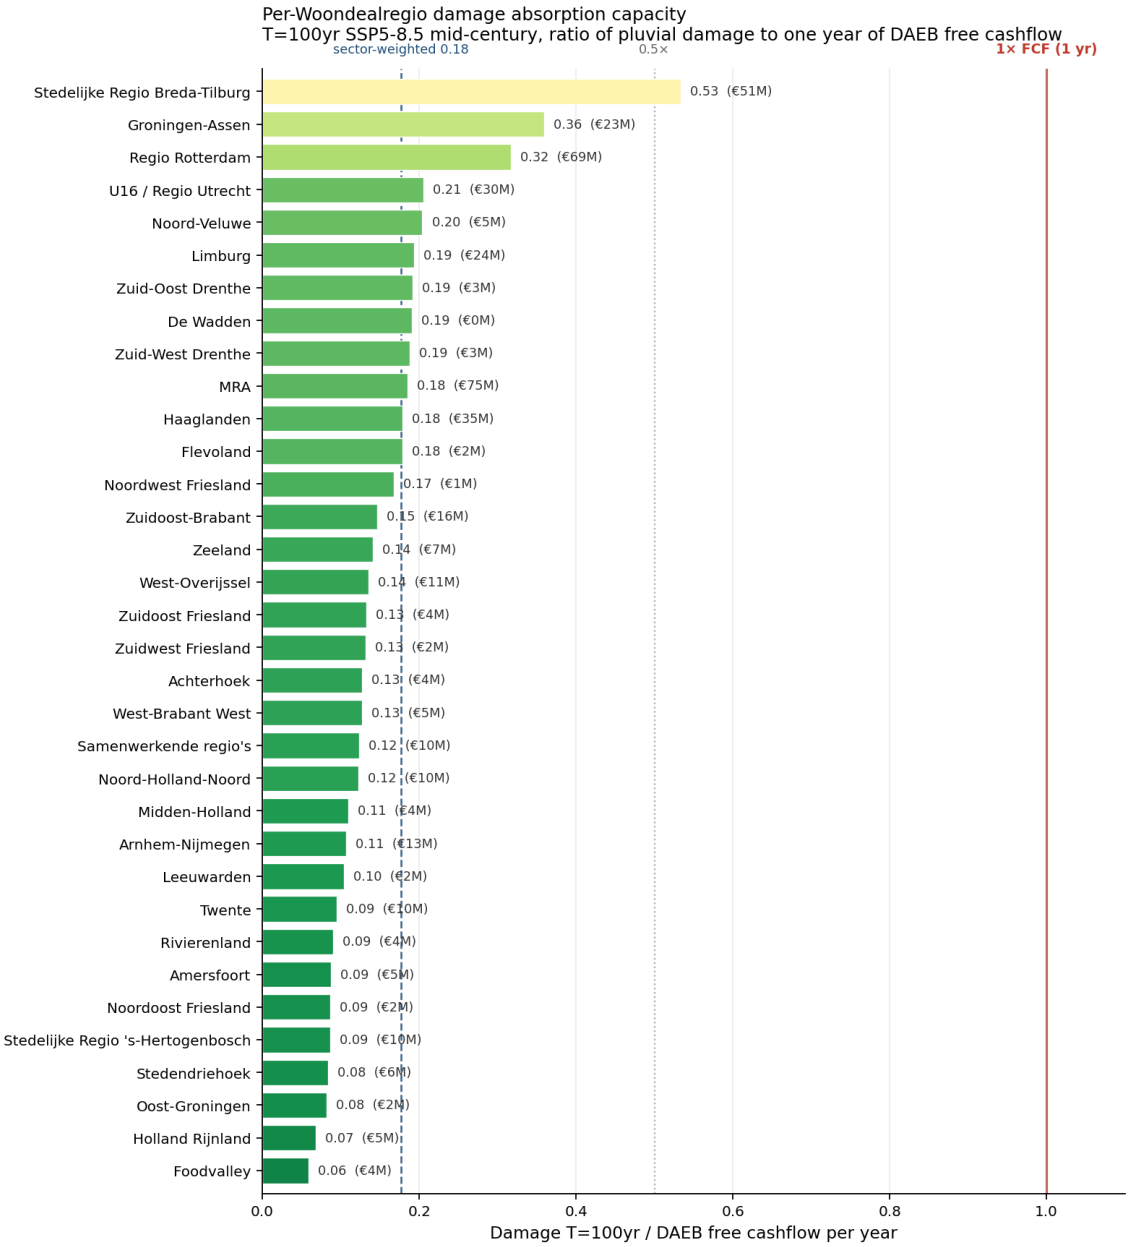


Figure G.1: *Damage-to-free-cashflow ratio per Woondealregio at the T=100 yr SSP5-8.5 mid-century anchor. The dashed blue line marks the sector-weighted ratio (0.18); the dotted grey line marks the half-year-of-FCF threshold (0.5); the solid red line at 1.0 marks the one-year-of-FCF reference. Stedelijke Regio Breda-Tilburg (0.53) is the only region above the half-year reference; Groningen-Assen (0.36) and Regio Rotterdam (0.32) follow at the upper end of the remaining distribution.*

Reading the table. A row with D/FCF < 1 indicates that the T=100 yr damage estimate is smaller than one year of free cashflow generated by the corporations operating in that region. Every row in the table satisfies this condition. The most stressed region is Stedelijke Regio Breda-Tilburg at 0.53, which combines the second-largest absolute damage (€51 M) with an operating cashflow base of €95 M; this is also the only region whose D/FCF exceeds the half-year-of-FCF reference. Groningen-Assen (0.36, €23 M damage

on a €64 M FCF base) and Regio Rotterdam (0.32, €69 M damage on a €216 M FCF base) sit at the upper end of the remaining distribution. Three regions with small absolute denominators deserve a separate reading: Noord-Veluwe (0.20 on a €26 M FCF base), De Wadden (0.19 on a €0.8 M FCF base), and Noordwest Friesland (0.17 on a €5 M FCF base) show ratios driven by a small denominator rather than by exceptionally high pluvial exposure. The least stressed regions are Foodvalley (0.06), Holland Rijnland (0.07), and Oost-Groningen (0.08). The CF_{inv} column makes clear that this absorptive capacity is not idle headroom: all 34 regions report net investing outflows that exceed operating cashflow in verslagjaar 2024. The buffer figures in this table should accordingly be read as the financial headroom that *would* be available for shock absorption if reinvestment were paused, rather than as currently uncommitted cash.

Sector aggregate. The sector aggregates in the final row are obtained by summing the absolute cash items across the 34 reported regions and recomputing the ratios on the totals. The portfolio damage-to-FCF ratio of 0.18 implies that absorbing a T=100 yr SSP5-8.5 pluvial event would consume approximately 18% of the sector's annual DAEB free cashflow. Combining FCF with the closing cash position brings this to 0.14, well below any threshold that would mechanically threaten ICR continuity but, with material implications for the sector's contemporaneous investment programme under the *Nationale Prestatieafspraken*: the CF_{inv} column sums to €-6,226 M against €2,581 M of FCF, a ratio of approximately 2.4. The sector in aggregate is investing more than twice its annual operating cashflow, so the buffer headroom shown here represents potential rather than actual absorption capacity.

Appendix H

Interest Coverage Ratio per Woondealregio

This appendix complements the cashflow-buffer analysis of §5.8 with a regulatory-norm benchmark: the Interest Coverage Ratio (ICR), defined as

$$\text{ICR} \;=\; \frac{\text{CF}_{\text{op}} + |\text{Saldo rente}|}{|\text{Saldo rente}|},$$

where CF_{op} is the operating cashflow and the saldo rente is interest paid net of interest received. Under the *Gezamenlijk Beoordelingskader Aw-WSW* the DAEB-tak ICR must exceed 1.4 (Autoriteit woningcorporaties, 2024); corporations falling below this threshold trigger heightened supervisory scrutiny. ICR is computed at corporation level from the dVi2024 release and aggregated to the Woondealregio level by VHE-weighting within each region; the within-region 25th and 75th percentiles (ICR_{P25}, ICR_{P75}) summarise corporation-level heterogeneity within the region.

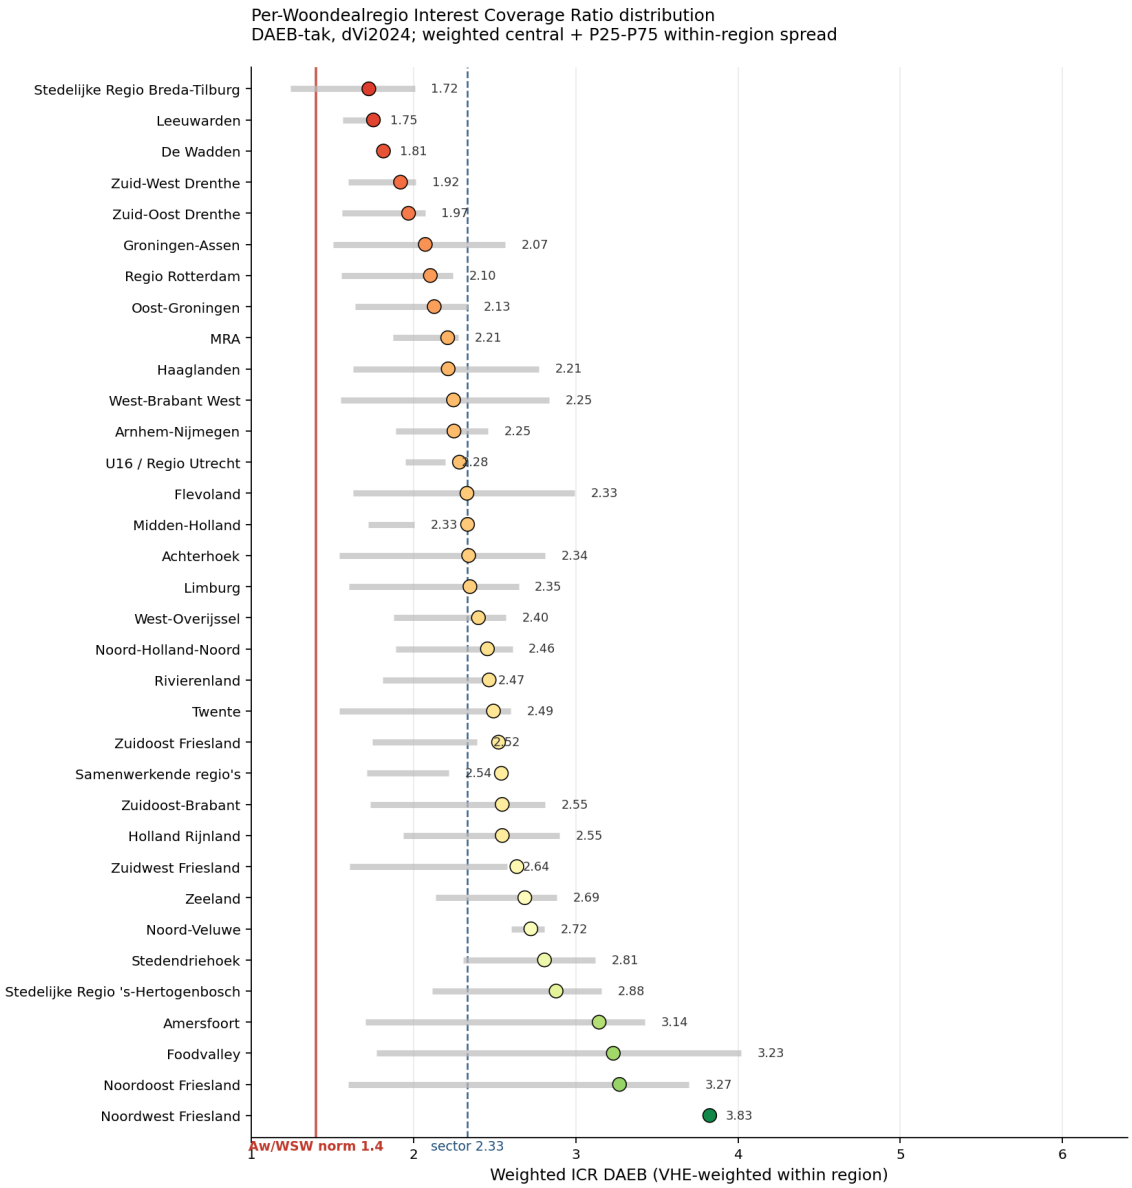


Figure H.1: *Weighted Interest Coverage Ratio per Woondealregio, sorted ascending so the most-stressed regions appear at the top. Filled markers are the VHE-weighted regional ICR; horizontal grey bars span the within-region P25–P75 of corporation-level ICR. The solid red line at 1.4 marks the Aw/WSW supervisory norm (Autoriteit woningcorporaties, 2024); the dashed blue line is the sector-weighted average (2.33). All 34 Woondealregio's sit above the supervisory norm on the regional weighted ICR; the most-stressed regions are Stedelijke Regio Breda-Tilburg (1.72), Leeuwarden (1.75), and De Wadden (1.81).*

Table H.1: *Per-Woondealregio Interest Coverage Ratio under the* Gezamenlijk Beoordelingskader Aw-WSW *norm of 1.4, verslagjaar 2024. Regions are sorted by ascending weighted ICR. Damage and D/FCF columns are reproduced from Table G.1 for cross-reference.*

Woondealregio	VHE	n_{corp}	ICR_{P25}	ICR_{w}	ICR_{P75}	Damage €M	D/FCF
Stedelijke Regio Breda-Tilburg	96,132	13	1.24	1.72	2.01	50.7	0.53
Leeuwarden	17,396	4	1.56	1.75	1.75	1.6	0.10
De Wadden	998	1	1.81	1.81	1.81	0.2	0.19
Zuid-West Drenthe	13,242	4	1.60	1.92	2.01	3.4	0.19
Zuid-Oost Drenthe	19,980	7	1.56	1.97	2.07	3.5	0.19
Groningen-Assen	44,451	12	1.50	2.07	2.56	23.2	0.36
Regio Rotterdam	201,211	23	1.55	2.10	2.24	68.7	0.32
Oost-Groningen	13,199	6	1.64	2.13	2.33	1.6	0.08
Metropoolregio Amsterdam	364,893	33	1.87	2.21	2.28	75.0	0.18
Haaglanden	154,734	21	1.63	2.21	2.77	35.5	0.18
West-Brabant West	30,759	7	1.55	2.25	2.84	5.2	0.13
Arnhem-Nijmegen	108,645	22	1.89	2.25	2.46	13.2	0.11
U16 / Regio Utrecht	110,652	15	1.95	2.28	2.19	29.5	0.21
Flevoland	12,147	6	1.63	2.33	2.99	2.4	0.18
Midden-Holland	25,245	7	1.72	2.33	2.01	3.8	0.11
Achterhoek	24,197	6	1.54	2.34	2.81	3.6	0.13
Limburg	105,014	21	1.60	2.35	2.65	24.1	0.19
West-Overijssel	59,906	18	1.88	2.40	2.57	11.1	0.14
Noord-Holland-Noord	52,880	13	1.89	2.46	2.61	10.2	0.12
Rivierenland	19,884	6	1.81	2.47	2.42	3.8	0.09
Twente	74,463	16	1.54	2.49	2.60	9.8	0.09
Zuidoost Friesland	20,773	5	1.75	2.52	2.39	3.8	0.13
Samenwerkende regio's	49,990	10	1.71	2.54	2.22	9.6	0.12
Zuidoost-Brabant	96,657	14	1.73	2.55	2.81	15.7	0.15
Holland Rijnland	65,561	16	1.94	2.55	2.90	5.1	0.07
Zuidwest Friesland	15,353	5	1.61	2.64	2.58	2.4	0.13
Zeeland	39,953	10	2.14	2.69	2.88	7.3	0.14
Noord-Veluwe	18,853	7	2.60	2.72	2.80	5.2	0.20
Stedendriehoek	39,976	12	2.30	2.81	3.12	5.5	0.08
Stedelijke Regio 's-Hertogenbosch	52,647	10	2.11	2.88	3.16	10.2	0.09
Amersfoort	39,842	10	1.70	3.14	3.42	4.8	0.09
Foodvalley	35,257	11	1.77	3.23	4.02	3.5	0.06
Noordoost Friesland	5,752	4	1.60	3.27	3.70	1.9	0.09
Noordwest Friesland	2,426	1	3.83	3.83	3.83	0.9	0.17
Sector-weighted aggregate	**2,033,068**	**265**		**2.33**		**456.0**	**0.18**

Reading the table. All 34 reporting Woondealregio's sit above the Aw/WSW norm of 1.4 on the VHE-weighted regional ICR. The most-stressed regions are Stedelijke Regio Breda-Tilburg (1.72), Leeuwarden (1.75), De Wadden (1.81), and the two Drenthe regions (Zuid-West Drenthe at 1.92 and Zuid-Oost Drenthe at 1.97), all of which are within 0.5 of the supervisory threshold. Stedelijke Regio Breda-Tilburg is also the only region whose within-region 25th percentile sits below the norm (P25 = 1.24), indicating that at least a quarter of the corporations operating in this region report ICRs below 1.4 even though the VHE-weighted regional aggregate clears the threshold. At the other end of the distribution, Noordwest Friesland records the highest weighted ICR at 3.83, a small region with a single reporting corporation, and Noordoost Friesland follows at 3.27. The

sector-weighted ICR across the 265 reporting corporations and 2,033,068 VHE is 2.33.

Implication for the chapter narrative. Read alongside Table G.1, the ICR distribution confirms that in an ordinary year, before any pluvial shock, the binding direct-cashflow constraint for most Woondealregio's is the FCF buffer (Section 5.8) rather than the ICR norm. Two configurations are visible. In one group, Stedelijke Regio Breda-Tilburg, Groningen-Assen, and Regio Rotterdam, both indicators point the same way: ICR sits in the 1.7–2.1 band and the T=100 yr damage absorbs between 32% and 53% of one year of free cashflow. Breda-Tilburg is the most concentrated case: weighted ICR of 1.72, P25 of 1.24, and a damage-to-FCF ratio of 0.53. In a second group, Leeuwarden, De Wadden, and the two Drenthe regions, the regional ICR sits close to the norm but the absolute T=100 yr damage is small relative to operating cashflow (D/FCF between 0.10 and 0.19), so the dominant supervisory concern is corporation-level interest coverage rather than pluvial-event absorption. For the policy-significant high-damage regions outside Breda-Tilburg (Metropoolregio Amsterdam, Haaglanden, U16 / Regio Utrecht), the weighted ICR sits between 2.21 and 2.28, well above the norm, and the corresponding damage-to-FCF ratios are 0.18, 0.18, and 0.21 respectively, so the financial transmission of a T=100 yr pluvial event in these regions is to modest FCF compression rather than ICR breach. This good-year reading changes once the modelled damage is charged against a single year's cashflow, as the stressed ICR of Section H.1 shows: under that one-off shock, Breda-Tilburg is the one region that slips below the supervisory norm.

H.1 Stressed ICR under a one-off T=100 yr pluvial shock

The ICR reported above is a steady-state measure: it asks whether operating cashflow covers interest payments in an ordinary year. To translate the modelled pluvial damage into the same supervisory metric, this section recomputes the ICR under the assumption that the full regional T=100 yr damage materialises as a one-off charge against a single year's operating cashflow, with no mitigating measures and no insurance recovery. This is a deliberately conservative stress: it concentrates an event with a 1% annual exceedance probability into one reporting year and charges it in full to that year's cashflow.

Construction. The pre-shock ICR is the regional figure of Table H.1. Regional damage is first allocated to the corporations active in each Woondealregio *pro rata* to their share of regional dwelling units (VHE), so that a corporation carrying a fraction $w_{c,r}$ of region r's dwellings is charged $w_{c,r}\,D_r$ of that region's damage D_r. The stressed corporation-level ICR then charges this allocated damage against operating cashflow while retaining

the appendix definition of ICR,

$$\mathrm{ICR}_c^{\mathrm{shock}} \;=\; \frac{\left(\mathrm{CF}_{\mathrm{op},c} - w_{c,r}\,D_r\right) + |\mathrm{Saldo\ rente}_c|}{|\mathrm{Saldo\ rente}_c|} \;=\; \mathrm{ICR}_c \;-\; \frac{w_{c,r}\,D_r}{|\mathrm{Saldo\ rente}_c|},$$

which is exactly the pre-shock ICR of this appendix reduced by the allocated damage expressed in units of the corporation's annual interest bill. Stressed corporation-level ICRs are then aggregated to the region by the same VHE-weighting used for the pre-shock figure. Only corporations reporting at least 50 VHE in a region are included, consistent with the aggregation elsewhere in this appendix.

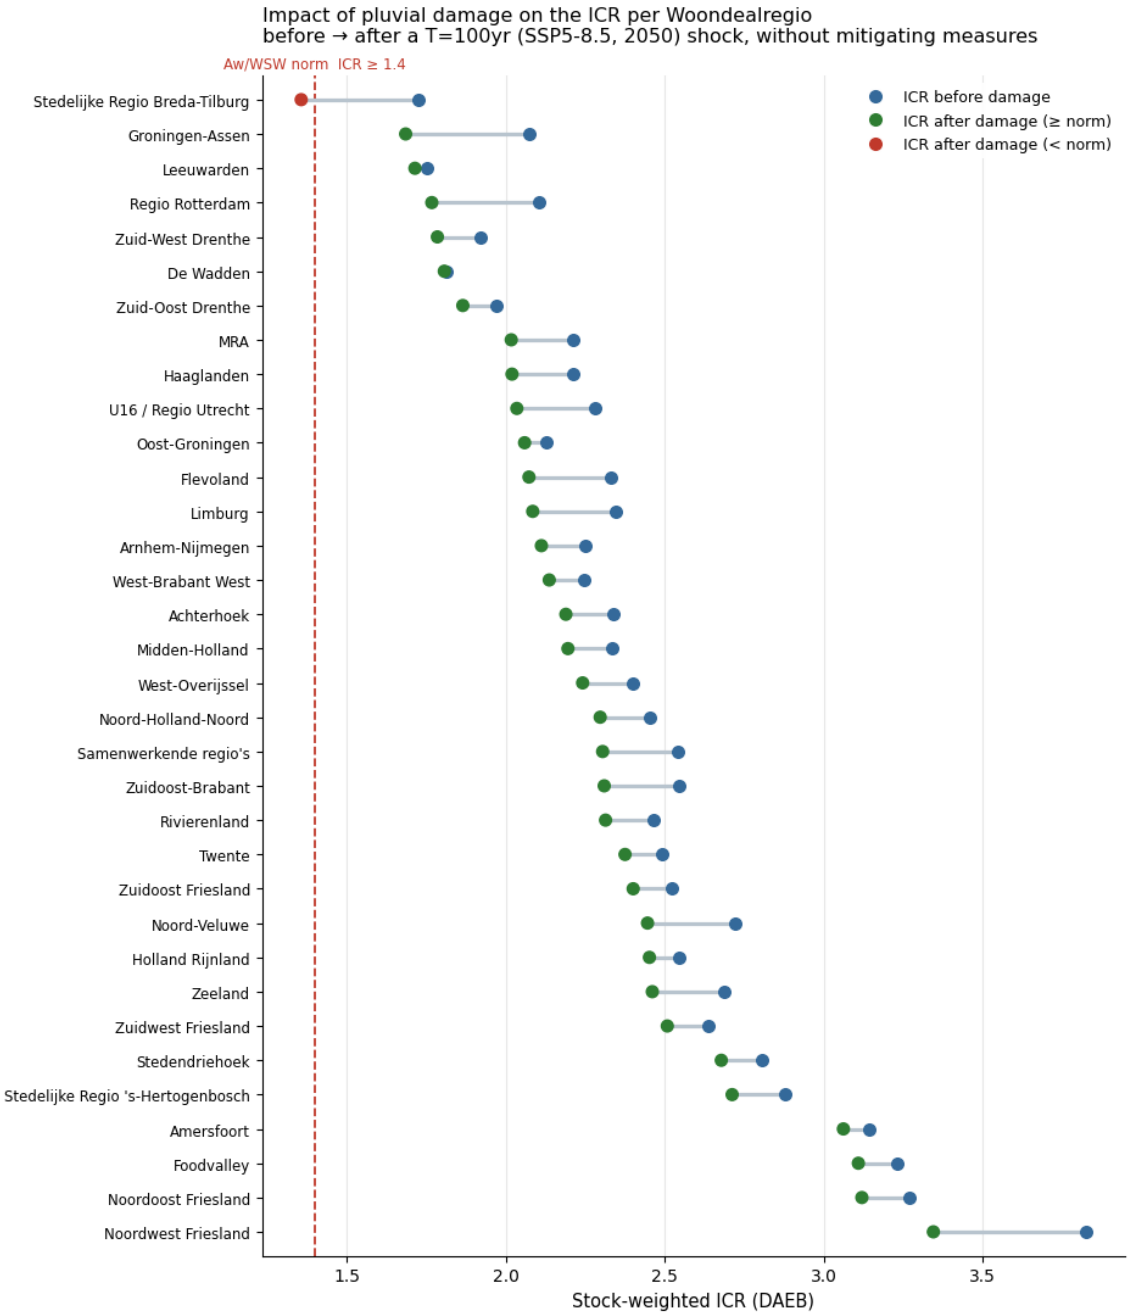


Figure H.2: *Impact of a one-off T=100 yr pluvial shock (SSP5-8.5, 2050) on the VHE-weighted ICR per Woondealregio, without mitigating measures. Each region is shown as a pre-shock value (blue) linked to its post-shock value; green markers remain at or above the Aw/WSW norm of 1.4 after the shock, the single red marker falls below it. The dashed red line marks the norm. Regions are sorted by post-shock ICR, so the most-stressed regions appear at the top.*

Table H.2: *VHE-weighted Interest Coverage Ratio per Woondealregio before and after a one-off T=100 yr SSP5-8.5 2050 pluvial shock, without mitigating measures. Δ is the change in ICR. Regions are sorted by ascending post-shock ICR. An asterisk (∗) marks the single region that falls below the Aw/WSW norm of 1.4 after the shock.*

Woondealregio	VHE	ICR_{pre}	ICR_{post}	Δ	Damage €M	<norm
Stedelijke Regio Breda-Tilburg	96,132	1.72	1.36	-0.37	50.7	∗
Groningen-Assen	44,451	2.07	1.68	-0.39	23.2	
Leeuwarden	17,396	1.75	1.71	-0.04	1.6	
Regio Rotterdam	201,211	2.10	1.77	-0.34	68.7	
Zuid-West Drenthe	13,242	1.92	1.78	-0.13	3.4	
De Wadden	998	1.81	1.81	-0.01	0.2	
Zuid-Oost Drenthe	19,980	1.97	1.86	-0.10	3.5	
Metropoolregio Amsterdam	364,893	2.21	2.02	-0.19	75.0	
Haaglanden	154,734	2.21	2.02	-0.19	35.5	
U16 / Regio Utrecht	110,652	2.28	2.03	-0.25	29.5	
Oost-Groningen	13,199	2.13	2.06	-0.07	1.6	
Flevoland	12,147	2.33	2.07	-0.26	2.4	
Limburg	105,014	2.35	2.08	-0.26	24.1	
Arnhem-Nijmegen	108,645	2.25	2.11	-0.14	13.2	
West-Brabant West	30,759	2.25	2.14	-0.11	5.2	
Achterhoek	24,197	2.34	2.19	-0.15	3.6	
Midden-Holland	25,245	2.33	2.20	-0.14	3.8	
West-Overijssel	59,906	2.40	2.24	-0.16	11.1	
Noord-Holland-Noord	52,880	2.46	2.30	-0.16	10.2	
Samenwerkende regio's	49,990	2.54	2.30	-0.24	9.6	
Zuidoost-Brabant	96,657	2.55	2.31	-0.24	15.7	
Rivierenland	19,884	2.47	2.31	-0.15	3.8	
Twente	74,463	2.49	2.38	-0.12	9.8	
Zuidoost Friesland	20,773	2.52	2.40	-0.12	3.8	
Noord-Veluwe	18,853	2.72	2.45	-0.28	5.2	
Holland Rijnland	65,561	2.55	2.45	-0.09	5.1	
Zeeland	39,953	2.69	2.46	-0.22	7.3	
Zuidwest Friesland	15,353	2.64	2.51	-0.13	2.4	
Stedendriehoek	39,976	2.81	2.68	-0.13	5.5	
Stedelijke Regio 's-Hertogenbosch	52,647	2.88	2.71	-0.17	10.2	
Amersfoort	39,842	3.14	3.06	-0.08	4.8	
Foodvalley	35,257	3.23	3.11	-0.12	3.5	
Noordoost Friesland	5,752	3.27	3.12	-0.15	1.9	
Noordwest Friesland	2,426	3.83	3.35	-0.48	0.9	

Reading the table. Before the shock, every Woondealregio clears the Aw/WSW norm of 1.4. After a one-off T=100 yr shock charged in full to a single year's cashflow, the average regional ICR falls by 0.18 (unweighted), and exactly one region crosses the norm: Stedelijke Regio Breda-Tilburg, which moves from 1.72 to 1.36. Two features of the ranking are worth stressing. First, the size of the ICR drop is not by itself a measure of vulnerability. The largest absolute fall is recorded by Noordwest Friesland (−0.48, from 3.83 to 3.35), yet it remains far above the norm; a large drop from a high starting point is comfortably absorbed. What matters for supervision is the post-shock distance to 1.4, not the size of the move. Second, the regions closest to the norm after the shock, Breda-Tilburg (1.36), Groningen-Assen (1.68), Leeuwarden (1.71), and Regio Rotterdam

(1.77), are a mixture of high-damage regions (Rotterdam at €68.7 M, Breda-Tilburg at €50.7 M) and regions that simply began close to the threshold (Leeuwarden, whose €1.6 M damage barely moves its ICR). Breda-Tilburg is the one region where a high pre-shock stress, a large absolute damage, and the resulting norm breach coincide, which is consistent with its position as the most-stressed region on both the D/FCF buffer ratio and the pre-shock ICR (Section 5.8).

Interpretation and caveats. This stressed ICR should be read as a worst-case single-year sensitivity rather than as an expected outcome. Charging a 1%-annual- probability event in full to one reporting year, with no insurance recovery, no spreading of repair expenditure across years, and no mitigation, is the most adverse plausible timing for the cashflow impact. Under that assumption only one of the 34 regions breaches the supervisory norm, and even that breach is modest (1.36 against a 1.4 threshold). The result reinforces the chapter's central financial finding: a T=100 yr pluvial event does not threaten the sector's interest-coverage position at large, and its supervisory significance is concentrated in a single region whose stress is already visible on the unstressed metrics. As with the rest of this analysis, the figure captures only Channel 1 (direct repair cost affecting operating cashflow) and not the Channel 2 capitalised mark-down on property values (Section 2.6); it is therefore a lower bound on the full balance-sheet impact.