TI 2023-083/IV
Tinbergen Institute Discussion Paper

Floods and financial stability: Scenario-based evidence from below sea level

Francesco G. Caloia[1]
Kees van Ginkel[2]
David-Jan Jansen[3]

1 Vrije Universiteit Amsterdam and De Nederlandsche Bank

2 Deltares, Delft

3 Vrije Universiteit Amsterdam, Tinbergen Institute and De Nederlandsche Bank

Floods and financial stability: Scenario-based evidence from below sea level

Francesco G. Caloia

Kees van Ginkel

David-Jan Jansen

This version: 20 December 2023
First draft: July 2022

Abstract

We study whether floods can affect financial stability through a credit risk channel. Our focus is on the Netherlands, a country situated partly below sea level, where insurance policies exclude property damages caused by some types of floods. Using geocoded data for close to EUR 650 billion in real estate exposures, we consider possible implications of such floods for bank capital. For a set of 38 adverse scenarios, we estimate that flood-related property damages lead to capital declines that mostly range between 30 and 50 basis points. We highlight how starting-point loan-to-value ratios are one important driver of capital impacts. Our estimates focus on property damages as the main transmission channel and are also subject to a number of assumptions. If climate change continues, more frequent floods or flood-related macrofinancial disruptions may have stronger implications for financial stability than our estimates so far indicate.

JEL codes: G21, Q54, R30

Keywords: floods, financial stability, real estate, credit risk, climate change

Acknowledgements: With thanks to Wouter Botzen, Karin de Bruijn, Maurice Bun, Hans Degryse, Mattia Girotti, Marco Hoogvliet, Hans de Moel, Remco van der Molen, Irene Monasterolo, David Rijsbergen, Stephan Rikkert, Kymo Slager, and Robert Vermeulen for constructive feedback over the course of this project. We also thank participants at an IMF/ECB workshop and a DNB seminar for useful comments. Van Ginkel acknowledges financial support under the Horizon Europe ACCREU project, grant agreement No. 101081358. Any errors and omissions remain our own responsibility. Views expressed in this paper do not necessarily coincide with those of de Vrije Universiteit, Deltares, de Nederlandsche Bank or the Eurosystem.

Affiliations: Caloia: De Nederlandsche Bank, Amsterdam; Vrije Universiteit Amsterdam; Van Ginkel: Deltares, Delft; Jansen (corresponding author): De Nederlandsche Bank, Amsterdam; Vrije Universiteit Amsterdam; Tinbergen Institute, Amsterdam; e-mail: d.jansen@dnb.nl.

1 Introduction

Floods can have severe socio-economic consequences. Examples include the devastating floods after hurricane Katrina in 2005 and the major impact of the 2021 summer floods in several European countries. So far, however, floods have not had major effects on the stability of the financial system. That is to say, flood-related disruptions have not impaired the financial system's ability to serve the real economy. This may change if climate change continues, as it would imply an increase in the frequency of extreme weather events (Auffhammer, 2018). A string of flood events and the associated macroeconomic disruptions may well start affecting financial institutions.

This paper focuses on the Netherlands, a densely-populated country situated partly below sea level. In the Dutch case, almost all insurance policies would exclude covering property damages from major floods. To be precise, insurers would not cover property damages if the flood were caused by a breach in primary flood defences.[1] These particular floods would, however, be the most impactful ones in terms of overall damages and socio-economic disruption. In the unlikely event of such a flood, one consequence is that banks need to consider accounting for increased credit risk for their real estate exposures. It is well known that these real estate exposures constitute a large share (roughly 45%) of banks' total assets. In addition, an institutional feature is that Dutch banks can grant relatively large loan amounts to borrowers, namely up to 100% of the collateral.[2] Also, nearly 50% of the mortgage portfolio of Dutch banks consist of interest-only mortgages. Such mortgages imply high exposure amounts as these loans do not amortize until maturity. This combination of large real estate exposures, high LTV ratios, and the inherent vulnerability to floods motivates our focus on banks' credit risk.

If and how climate change can affect financial stability remains a topic of debate. In a seminal speech, Mark Carney (2015) argued that climate change can affect financial stability through three main channels. One of these channels is physical risk, e.g. property damage and trade disruption due to floods and storms. Likewise, the European Systemic Risk Board (ESRB, 2016) has argued that a delay in the energy transition could affect systemic risk via three main channels, including more frequent natural catastrophes. The Network for Greening the Financial System (NGFS, 2019) has emphasized that climate change could have much larger impacts than other structural changes affecting the financial system. Bolton et al. (2020) suggest that climate change could even be the cause of the next systemic financial crisis. Others seem less convinced. Fed Governor Waller (2023) has argued that climate change is not a serious risk to the U.S. financial system, for instance because climate physical risks are not likely to have a large effect on economic growth. Hansen (2022) focuses

[1] The Appendix has further details on flood types relevant to the Netherlands.

[2] Dutch borrowers do indeed often take out loans close to the loan-to-value (LTV) limit of 100%. Household debt in the Netherlands is, therefore, the second-highest in the world (Caloia, 2022).

on climate analyses by central banks and issues a warning that there can be reputational costs to overselling insights from long-horizon climate stress tests. In a similar vein, Cochrane (2021) argues that climate change is highly unlikely to have a drastic and abrupt effect on financial institutions in the short term.

Our study on floods and financial stability has two distinctive features. First and foremost, we use scenario analysis rather than studying historical episodes. We make calculations for flood scenarios that are severe and unlikely from today's perspective, as is common in financial stress testing. We first focus on 32 scenarios for floods due to a single breach in flood defence systems. Information on these single-breach scenarios is available in an open-source information system. The single-breach flood scenarios we consider have a probability varying roughly between once in every 100 and once in every 10,000 years. Next, we consider even less likely events by studying six extreme multiple-breach scenarios. These six extreme multiple-breach scenarios were constructed by Dutch flood experts in 2007. In the end, all 38 flood scenarios we study are unlikely, which is reflective of the high level of flood protection in the Netherlands. By focusing on such unlikely—and, sometimes, very extreme scenarios—we aim to understand how tail events surrounding climate physical risk could have financial stability implications. This scenario-based approach sets our paper apart from work that has taken a backward-looking perspective on financial impacts of natural disasters. For instance, Blickle et al. (2022) consider data from the last 25 years and find that natural disasters have had insignificant or small effects on U.S. banks' performance. Taking a cross-country approach, Gramlich et al. (2023) also find little evidence that natural disasters have affected banks' Tier 1 capital ratio so far. Klomp (2014) focuses on large-scale natural disasters between 1997 and 2010 and finds that these disasters do increase the likelihood of bank defaults. Studying banks' lending decisions, Koetter et al. (2020) focus on a major flood of the river Elbe in Germany and find evidence for recovery lending (as opposed to less lending).

Second, our analysis uses various data sources that have granular information on real estate exposures. As the global financial crisis of 2007-09 has shown, real estate is an asset class with strong implications for financial stability. Real estate is also relevant from the perspective of climate physical risk, as floods can potentially lead to major damages to houses, offices and other properties. We use granular loan-level data on about 3 million residential properties and 175,000 commercial properties that underlie close to EUR 650 billion in exposures secured by real estate. In this way, we can include flood considerations in the credit risk estimation for a sample of eight Dutch banks.

In terms of analysis, we trace the possible impact of floods on bank capital via a number of steps. First, we use information on location-specific inundation depths under the 38 individual flood scenarios to compute property damages for each of the roughly 3 million properties in our data set. Second, we translate these property-specific damages into increases of the loan-to-value ratio of individual bank loans. In principle, we are able to analyze such LTV increases at the level of scenarios, banks, and types of exposure (i.e., either commercial or residential real estate loans). However, we will always present financial impacts that are aggregated across the banks in the sample. Third, using the shocks to LTV ratios, we compute implications for key credit risk parameters at the level of individual loans. Here, we consider changes in the probability of default (PD), the loss given default (LGD), and the risk weights of individual loans. Lastly, we consider the impact on bank capital, which we measure via the changes CET1 ratio. In the empirical section, we will discuss these financial impacts from the perspective of 38 individual flood scenarios and two main types of real estate exposures.

Turning to results, we estimate that flood-related property damages could cause a drop in system-wide bank capital that mostly ranges between 30 and 50 basis points. We show that flood-related capital depletions are a function of flood-related vulnerabilities and financial vulnerabilities. Concerning financial vulnerabilities, we highlight the role of the starting-point loan-to-value (LTV) ratio of the loan for which the damaged property serves as collateral. We show that the flood impact increases when properties are damaged that serve as collateral for loans with relatively high LTV ratios. Relative to current levels of bank capital, one could argue that the estimated levels of depletion seem manageable—especially since we already consider tail-risk events. At the same time, it is important to note that our analysis focuses on only one transmission channel and is subject to a number of assumptions. In particular, if climate change continues, more frequent floods or flood-related macrofinancial disruptions may have larger implications for financial stability than this paper's estimates indicate.

This paper proceeds as follows. Section 2 discusses related literature on climate change, property damages, and financial risks. Section 3 describes the various data sources, in particular the geocoded exposures data. Sections 4 and 5 discusses our methodology to link flood damages to banks' credit risk. Section 6 provides results for system-wide capital depletions. Section 7 provides a discussion of the results and also outlines a few key assumptions that could be explored in future research. Section 8 concludes.

2 Related literature

A first set of related papers studies how climate change could affect the financial system. As indicated in the survey by Acharya et al. (2023), floods have not received much explicit attention yet. In contrast, much of the focus has been on how disorderly transition paths could affect financial stability. Battiston et al. (2017) consider exposures of 50 European banks to climate policy relevant sectors. Using data on over EUR 1 trillion of equity holdings, they find that while direct equity exposures of these banks to fossil fuels are small, the overall exposures to climate policy relevant sectors are large as well as heterogeneous. Similarly, Vermeulen et al. (2021) find that financial losses under disruptive transition paths could be sizeable, suggesting that climate-transition shocks warrant close attention from a financial stability perspective. Jung et al. (2021) propose a new measure (CRISK) which is the expected capital shortfall of a financial institution in a climate stress scenario. This measure indicates a strong rise in climate vulnerabilities for banks in the U.S., U.K., Japan and France in recent years. Our paper contributes by extending the discussion to physical risk and focusing on the role of real estate as a key transmission channel. Also, our paper takes a innovative perspective by considering how a range of increasingly severe flood scenarios could affect financial stability. This approach sets our paper apart from work studying how natural disasters have so far impacted the financial system (Blickle et al., 2022; Gramlich et al., 2023; Koetter et al., 2020; Klomp, 2014).

A second stream of literature considers how floods lead to property damage. For floods along the Meuse river in the Netherlands in the mid-1990s, Daniel et al. (2009) estimate a price decrease of up to 9%. Atreya and Ferreira (2015) combine a hedonic pricing model with geospatial information to study the 1994 flood in Albany, Georgia, and find damages up to 48%. Pistrika and Jonkman (2010) assess a database of 95,000 damage estimates for buildings in New Orleans after hurricane Katrina and report an average damage of 43.3%. Damage rates of more than 50% are reported for roughly a quarter of the assessed properties. Two further studies focus on floods in the United Kingdom. Using a repeat-sales model, Beltrán et al. (2019) estimate price declines in flooded areas between 21% (in case of coastal flooding) and 25% (in case of inland flooding). These price declines are estimated to be short-lived, as differences are no longer significant after a period of 5 years. Garbarino and Guin (2021) focus on one severe flood event flood in England in 2013-14. They find relatively small price effects related to this flood. Compared to unaffected properties in the same district, properties in affected areas shows decreases in sales prices of, at most, 4.2%.[3]

3 Data description

This paper uses information from three unique, proprietary data sources: loan-level data, administrative microdata, and supervisory data. The first two sources allow us to estimate the flood damages per individual property and map these into the credit risk parameters of the loans that these properties secure. In particular, we use the loan-level data to compute multipliers for credit risk parameters, while use the administrative microdata for information on property characteristics that are needed to compute flood-related property damages. We use the supervisory data to obtain starting points for capital ratios and credit risk parameters. Our analysis uses data for eight major Dutch banks, thus covering nearly all of the outstanding loans secured by real estate in the Netherlands.[4]

First, we use several vintages of loan level data. This data covers virtually all residential and commercial real estate properties that serve as collateral for loans provided by Dutch banks. The data

[3]There is also related work studying whether flood risk is already incorporated in property prices. For instance, Hino and Burke (2021) find little evidence that housing markets fully price information about flood risk in aggregate. See also Beltrán et al. (2018), Bosker et al. (2019), Mol et al. (2020), or Mutlu et al. (2023).

[4]For confidentiality reasons, the individual data sources have not been combined.

on exposures contain detailed information at a quarterly frequency. The data set has information on various aspects, such as the counterparty, the contract, the instrument and the protection related to each real estate loan. We use information as per end-2020.[5]

Second, we rely on administrative microdata at the level of the property available from Statistics Netherlands (in Dutch: CBS). This data set covers all objects registered in the basic property registry. This registry contains several types of information collected for administrative purposes. To begin with, the registry has information on the tax value of each property. In addition, the data set contains a range of property characteristics, such as floor surface (in square meters), location (at the postal-code level) and the type of real estate. We use postal codes to determine how much a property would be inundated under different flood scenarios, while we use the other characteristics in the computations for the flood-related damages to properties in case of a flood.

Third, we use supervisory data for individual banks from the Common (COREP) and Financial (FINREP) Reporting frameworks. The COREP and FINREP formats cover various aspects of banks' balance sheets, profitability, capital adequacy, and asset quality. The analysis in this paper uses starting points for credit risk parameters and capital positions, also at end-2020.

Table 1 shows descriptive statistics for the residential and commercial real estate data. On average, the collateralization of the portfolios (loan-to-value) is around 50%. In addition, a substantial amount of debt is non-amortizing (i.e. interest-only), meaning that outstanding debt remains constant over the lifetime of the loan. House prices have grown steadily in the Netherlands over the past decade, with a cumulative growth rate above 50%. On the one hand, this growth results in risky lending, as more-recent borrowers take high loan-to-values at origination (LTV-O), sometimes also with high interest-only shares or high loan-to-income (LTI) ratios. On the other hand, the increase in house prices has also increased the collateral value of less-recent exposures.

(insert Table 1 around here)

Figure 1 illustrates the distribution of real estate exposures across the municipalities in the Netherlands. Most of the real estate exposures are in the western half of the country, which is also the part of the country that is partly below sea level. In addition, there are also sizeable exposures in the vicinity of the four main rivers, which poses an additional risk in terms of flooding.

(insert Figure 1 around here)

[5]This vintage of the data was also a basis for Caloia and Jansen (2021). Presumably, the inherent vulnerability of real estate exposures to flood risk will not have changed materially since end-2020, although ideally one would like to use updated information at some point.

4 Methodology: Floods and damages

4.1 Selection of flood scenarios

The Netherlands is vulnerable to different types of floods. This paper considers scenarios with floods from either the sea or main rivers in areas protected by flood defences. There are two key reasons for this choice. First, these are the areas where the concentration of economic activity and real estate is the highest. If a flood event were to occur there, it is most likely to affect financial stability. Second, an important institutional consideration is that flood damages would, in that case, not be insured. Most Dutch insurance policies do not provide coverage for the specific type of floods that we consider in this paper. In principle, the government could step in and provide compensation for the flood damages. However, such government support would be discretionary and might well take some time to be available. In that case, in the short-run the increased risk would need to be factored in by the banking sector.

The analysis first focuses on 32 scenarios with single-breach flood events, i.e. cases where a local failure of flood defence systems leads to a flood in a specific part of the Netherlands. Figure 2 shows the location of the single-breach floods. These scenarios are available in an open-source information system called the LIWO.[6] In total, the LIWO contains more than 5,000 flood scenarios for the Netherlands.

We use two selection criteria. First, we want to cover all relevant parts of the Netherlands that could be vulnerable to breaches in primary flood defence systems (i.e., 'type B' floods).[7] Second, per area of the Netherlands, we then always select the scenario for which available calculations (as found in the LIWO) indicate the highest expected damages. Here, we do use a cut-off value of EUR 500mn. in terms of expected property damages. Based on these two criteria, we are left with 32 scenarios that are among the most impactful ones in terms of property damages. This particular selection procedure is in line with our objective of focusing on tail risks. The reported impacts are not necessarily representative of the average impact of single breaches in flood defence systems. We also note that one cannot simply combine results for our set of single breach scenarios to estimate the impact of multiple breaches occurring simultaneously. It is important to remember that we focus on the most impactful single-breach scenarios out of a possible set of 5,000 flood scenarios.

(insert Figure 2 around here)

[6]LIWO stands for Landelijk Informatiepunt Water en Overstromingen (Nationwide Information source Water and Floods). The LIWO-scenario's discussed in this paper can be viewed at `https://basisinformatie-overstromingen.nl/`.

[7]During the summer floods of 2021, a part of the Dutch province Limburg was also heavily affected. However, given that this paper focuses on large floods of 'type B', this region is not included in the scenario set.

During extreme weather conditions, multiple dike breaches may happen simultaneously. As with the 32 single-breach cases, we want to explore possible implications of the *extreme* versions of multiple-breach scenarios. To do so, we rely on a set of six extreme scenarios constructed by Dutch flood experts in 2007.[8] Each of these six extreme multiple-breach scenarios is composed of dozens of larger and smaller dike breaches occurring at the same time. We would emphasize that these are circumstances that the flood experts, at the time, considered to be at the edge of what could happen in a highly unlikely situation. Figure 3 shows the extent of floods under these six extreme multiple-breach scenarios.

(insert Figure 3 around here)

4.2 Inundation depths per postal-code area

For each flood scenario, a key metric is the inundation depth per location. Here, we do face a challenge concerning data granularity. The LIWO allows the user to assess inundation depth at the level of, at least, a 100x100 meter grid. However, the geocoded data on bank loans is only available at the level of four-digit postal codes. In practice, this could be either a small town or a neighbourhood in a larger city. While that level of geographical precision is still a major improvement compared to the nationwide supervisory data available in FINREP, it does mean we need to make assumptions when calculating the flood impact per scenario.

To obtain a proxy for the representative water depth per postal-code area, we proceed as follows. First, we select the part of the postal-code that contains residential or commercial real estate. This is necessary, because the inundation map also overlaps with other types of land use, such as agriculture or infrastructure. Given our focus on real estate, these other types of land use should be excluded from our calculations. Having excluded these other types, we then calculate the mean water depth over the built-up area only. The not-inundated parts of the built-up area per postal code are included as zero values when taking this mean for inundation depths. In other words: the mean water level is representative for all properties in the postal code area, including those properties that are not inundated.

4.3 Computing property damages

Our method for calculating property damage broadly follows the SSM-method, the standard method for calculating flood damage in The Netherlands (Slager and Wagenaar, 2017). This SSM method makes a distinction between direct damages (i.e. to the property) and indirect damages (i.e. to household goods and personal effects). We focus on the former, as it is only the direct damages to the property that would affect the collateral value.

[8]For details, see www.helpdeskwater.nl. URL last accessed on 24 November 2023.

The SSM method has two key inputs for determining property damages. First, a maximum amount of damage per property, for instance specified as a euro amount per m^2 for residential properties. Second, a damage factor $\theta(h)$ that increases non-linearly with inundation depth h, the so-called damage function. Slager and Wagenaar (2017) provide separate damage functions for different types of properties, such as single-family homes and apartments (for residential properties) or offices and industrial properties (for commercial real estate). We can apply almost all of these specific damage functions directly using the information in the loan-level data set. The exception concerns apartments, which constitute 14% of the properties in the data set (see Table 1). For apartments, flood damage factors differ between floor levels. As information on floor levels is not available in our data set, we use a weighted average of the damage functions for ground-floor and first-floor apartments. We place most weight (65%) on the damage function for first-floor apartments, given that we neither want to overestimate nor underestimate damages. Section 3.1 of the Appendix provides further details.

Next to this difference in damage function implementation, there are a few additional differences between this study and the most recent SSM-damage calculations (Slager and Wagenaar, 2017) for the 32 LIWO scenarios. Our study uses more recent exposure data (2020 instead of 2014), only includes properties that serve as collateral for a loan (instead of all properties), and uses a mean water depth per postal code (instead of per 100*100 meter grid cell). Despite these differences, there is a strong correlation ($\rho = 0.95$) between our damage estimates and those from the SSM method. Section 3.2 of the Appendix provides further details.

The key parameter linking floods to financial risk is ϕ. With this parameter ϕ, we denote the flood-induced decline in the collateral value of the individual property p. We compute this parameter at the property-level for each flood scenario. As described in the next section, this parameter ϕ will then affect the loan-to-value for each loan i for which the property serves as collateral. To compute ϕ, we first calculate the property-specific damages (in euros, price level 2020) as:

$$damage_p^S = \theta(h)_t^S \cdot max\ damage_t \cdot A_p \cdot \pi \tag{1}$$

Here, S indexes the scenario, p the property, t the type of property, while h denotes inundation depth (in meters). The parameter θ is the damage factor. This factor, which lies between 0 and 1, denotes the fraction of maximum damages that are incurred. The parameter θ is specific to property type and inundation depth. Furthermore, *max damage* is the maximum flood damages (per property type, in 2011 prices) as provided by Slager and Wagenaar (2017). The information on the size of the property (A, in m^2) is obtained from the administrative microdata available from Statistics Netherlands. The last term (π) is a correction factor to compute damages in terms of the 2020 price level, which is necessary as the collateral values are all observed at end-2020. The last section of the Appendix discusses how we calibrate this parameter π.

Lastly, we obtain ϕ by expressing the property-specific damages as a fraction of the observed collateral value in the loan-level data, where we restrict the value of ϕ to be in the [0,1) interval.

$$\phi_p^S = min\left(\frac{damage_p^S}{property\ value_p};\ 1\right) \tag{2}$$

It is important to emphasize again that we are able to compute ϕ at the level of individual properties. At the level of the property, we observe the collateral value as well as the size of the property. The only restriction remains that the scenario-specific inundation depths can only be approximated at the level of the four-digit postal-code areas.

5 Methodology: Credit risk impact

The credit risk impact for banks will be driven by a few key parameters. We discuss, in turn, our models for the loss-given-default and the probability-of-default parameters. We close by linking the credit risk impact to banks' capital adequacy, which we measure by the CET1 ratio. The CET1 ratio, a key measure in financial supervision, is the ratio between the high-quality regulatory capital and banks' risk weighted assets. We measure the credit risk impact over a one-year horizon. The starting points are taken as per end-2020.

5.1 Loss given default

The loss given default (LGD) indicates the part of the loan that the lender cannot recover in case the borrower were to default. Using the loan-level data, we use a model that estimates LGD paths as a function of flood-related collateral damages. To compute the LGD paths, a number of other metrics are important: the loan-to-value ratio and the loss-given-loss.

First and foremost, the estimated decline in the collateral value (ϕ_p^S) will lead to an increase in the loan-to-value ratio of the loans for which the property serves as collateral. This means those loans will become riskier from the perspective of the lender. The LTV of loan i under flood scenario S is given by:

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

Here, p indexes properties, LTV_i^0 represents the starting-point loan-to-value ratio (per end-2020), and LTV_i^S is the LTV ratio for loan i per flood scenario. The parameter ϕ_p^S is computed based on equations 1 and 2 above.

Second, the increase in the LTV ratio under the flood scenario will increase the loss-given-loss (LGL) of the loan. The LGL indicates whether a lender will be able to cover its exposures when selling the property that serves as collateral. A positive LGL indicates that the lender will not be able to recover the exposure in full. The link between the LTV and LGL is given by:

$$LGL_i = \frac{max[0\ ;\ exposure - liquidation\ value]}{exposure} \tag{4}$$

Here, the liquidation value is a prudent estimate of the price of the asset in the event of a forced sale. If we divide both the numerator and the denominator by the property value and then condition on the flood scenario S, we obtain the following expression:

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

Here, the sales ratio denotes the ratio between the liquidation and current value of the property. In line with global standards on valuations[9] Our calculations for the liquidation value take into account the costs for getting the property in a condition ready for sale. So, the sales ratio in each flood scenario S is reduced by an amount equal to the property damage.[10]

Third, an increase in the LGL will increase the LGD parameter of loan i. This relationship is given by:

$$LGD_i^S = ((1 - probability\, of\, cure) * LGL_i^S) + costs \tag{6}$$

Here, LGD indicates the loss-given-default for loan i per flood scenario S, while the LGL is the loss-given-loss as computed in equation 4. The probability of cure denotes the percentage of loans that previously reported arrears and, post restructuring, does not present arrears. The last term (*costs*) indicates the administration costs that the bank would incur when moving ahead with selling the property.

For each of the single-breach flood scenarios, we make calculations for LGD paths based on equations (1) - (6). Given the granularity of the loan-level data, we do so separately for residential and commercial real estate exposures. The starting-point data is a cross-section of exposures backed by real estate at end-2020. When available (i.e. for commercial real estate), we use the reported information on liquidation values (for sales ratios) and administration costs. For residential real estate, this information is not available to us. Therefore, we use values in line with those reported for commercial real estate exposures.

Turning to system-wide impacts, the scenario-specific LGD multiplier for banks (relatively to the starting point) is given by:

$$m_{LGD}^S = \sum_b \sum_i w_b w_i \frac{LGD_{i,b}^S}{LGD_{i,b}} \tag{7}$$

Here, the weights w_b and w_i represent the bank-level and loan-level exposures, as a share of the total system exposure and the total bank exposure, respectively. For a given flood scenario S, the final impact on the banks reflects their relative exposure to the flood scenarios, the associated property damages, as well as their starting point for the LGD.

[9]See, for instance, Red Book on Global Standards `https://www.rics.org/profession-standards/rics-standards-and-guidance/sector-standards/valuation-standards/red-book/red-book-global`.

[10]In other words, we calculate as follows $sales\, ratio_p^S = sales\, ratio_p^0 * (1 - \phi_p^S)$.

5.2 Probability of default

Usually, a second key parameter for credit risk is the probability of default (PD). At first sight, the PD may seem less relevant to this paper, given our focus on how property damages map into loss-given-default parameters. However, PD paths can be relevant, given that empirically PDs are found to comove with loan-to-value ratios. Therefore, we estimate the probability that a borrower defaults on the loan, conditional on each of the single-breach flood scenarios, using a panel data approach that includes LTV ratios as one of the explanatory variables. We estimate the following equation:

$$y_{i,b,t} = c_b + \beta' \mathbf{Z}_{i,b,t} + \delta' \mathbf{X}_{i,b,t} + u_{i,b,t} \tag{8}$$

Here, the dependent variable $y_{i,b,t}$ is the default status of the counterparty i of bank b at time t. The set of control variables $\mathbf{X}_{i,b,t}$ includes the mortgage type (amortizing, deferred-amortization, or interest-only), the interest rate type (variable, fixed, or fixed with reset), the residual maturity, the LTV at inception, the type of real estate that serves as collateral, and whether or not a National Mortgage Guarantee (NHG) is attached to residential loans.[11] We estimate a pooled OLS model using the loan-level data.[12] The main independent variables $\mathbf{Z}_{i,b,t}$ are the current loan-to-value, as well as the mortgage interest rate and regional GDP growth. The parameters β then capture the effect on the default probability of collateral, funding and income risk, respectively. We obtain the system-wide PD multiplier in each scenario as:

$$m_{PD}^S = \sum_b \sum_i w_b \, w_i \frac{E(y_{i,b,t} | \mathbf{X}_{i,b,t} \, , \, \mathbf{Z}_{i,b,t}^S \, , t = T \, , s = S)}{E(y_{i,b,t} | \mathbf{X}_{i,b,t} \, , \, \mathbf{Z}_{i,b,t} \, , t = T)} \tag{9}$$

Here, the numerator represents the predicted value of eq. (3) conditional on the values of the LTV under each scenario, as of the last reporting period. The weights w_b and w_i are again the bank-level and loan-level exposure shares.

In principle, we could extend the analysis further by also projecting PD paths in line with the implied change for each main independent variable $\mathbf{Z}_{i,b,t}$ under the various flood scenarios. In particular, we could include considerations such as growth slowdowns or increases in risk premia.[13] However, this paper focuses on the role of property damages. The effect on the PD that we show, therefore, comes purely from the flood-related damages and their effect on the LTV ratio via ϕ_p^S.

5.3 The impact on the CET1 ratio

We measure the impact of flood damages by computing how changes in credit risk parameters affect system-wide bank capital. The metric we use is the CET1 ratio, i.e. the ratio between banks'

[11] The availability of such an NHG guarantee would serve as a key risk mitigant, as the bank would be more likely to recoup a sizeable part of its outstanding exposure in case of a default.

[12] We estimate the model at annual frequency. The sample is 2012 - 2018 (residential real estate) or 2015-2020 (commercial real estate).

[13] See also Caloia and Jansen (2021) for further discussion on this point.

high-quality regulatory capital (CET1 capital) and their risk weighted assets (RWA). Our calculations take as a starting point a system-wide CET1 ratio of 16.7%.

First, we consider the flood-related changes in expected loss (EL) for banks, which will affect the numerator of the CET1 ratio. Expected loss can be computed by multiplying the probability of default (PD), the loss given default (LGD), and the exposure at default (EAD, measured in euros). Hence, the expected loss under flood scenario S is then given by the product of starting-point risk parameters (as given by the supervisory data) and multipliers (as computed by us):

$$EL^S = PD \cdot LGD \left(m_{PD}^S \, m_{LGD}^S \right) \cdot EAD \tag{10}$$

As expected loss increases under the flood scenarios, banks are assumed to book additional provisions, thus lowering the available CET1 capital to hold against their risk weighted exposures. Although we compute changes in EL at the level of individual banks—as well as portfolios and flood scenarios—we will only report system-wide impacts.

Turning to the denominator of the CET1 ratio, we compute risk weighted assets as per the requirements from the Basel Framework. We compute risk weighted assets as the product of exposures at default, a constant factor of 12.5% (which represents the inverse of the 8% Basel minimum capital requirement) and a factor K. This factor K is the relevant capital requirement from the Basel Framework that depends, among other things, on the bank-specific and scenario-specific paths for LGDs and PDs that we compute.[14] Having computed scenario-specific RWAs, we can then also derive the implied multipliers for risk weights (m_{RW}^S) by comparing to the starting-point RWAs (as given in the supervisory data):

$$\begin{aligned} m_{RW}^S &= \sum_b \sum_i w_b w_i \frac{RWA_{i,b}^S}{RWA_{i,b}} \\ &= \sum_b \sum_i w_b w_i \frac{K_{i,b}^S}{K_{i,b}} \end{aligned} \tag{11}$$

Having derived bank-specific, portfolio-specific, and scenario-specific paths for expected loss and risk weights, we can then consider implications for capital positions. As noted above, we only report system-wide impacts. Comparing the paths for EL and RWAs to the starting points (as given by the supervisory data), we can compute changes in the CET1 ratio per flood scenario as:

$$\Delta CET1\, ratio^S = \frac{CET1}{RWA} - \frac{CET1 - \Delta EL^S}{RWA + \Delta RWA^S} \tag{12}$$

Here, ΔEL^S represents the difference between the EL under the flood scenario and starting-point expected loss, while ΔRWA^S denotes the scenario-specific change in the risk weighted assets. Both

[14] We refer to the documentation by the Basel Committee for the details of RWA computations: `https://www.bis.org/basel_framework/index.htm?m=2697`. On a technical note, the RWA formula differs between residential and commercial real estate exposures concerning the full maturity adjustment. We take this difference into account in the calculations.

the increase in the expected loss and the higher risk weights will lead to a lower CET1 ratio under the flood scenario. Hence, we will often use the term capital depletions when discussing changes in the CET1 ratio.

6 Results

This section discusses how flood-related property damages would affect the capital position of the Dutch banking system. To start with, we show the relationship between our calculated property damages and capital depletion. We find that this relationship is, within the set of scenarios we study, non-linear. Figure 4 illustrates this non-linearity. The figure shows estimated property damages (in EUR bn., horizontal axis) and the system-wide decline in the CET1 ratio (in basis points, vertical axis). Here, the horizontal axis shows combined damages for residential and commercial real estate. The gray line plots an estimated fractional polynomial of degree 2. The figure shows results for the single-breach scenarios (blue diamonds) as well as extreme multiple-breach floods (red circles).

(insert Figure 4 around here)

The figure indicates that the 32 single-breach flood scenarios would all have similar implications for bank capital. Three of the multiple-breach scenarios have comparable implications for CET1 capital as well. Overall, this cluster of 35 scenarios would be characterized by property damages between EUR 0.5bn and EUR 3bn. (in price level of 2020). Capital depletions are estimated to be between 30 and 50 basis points. There is an indication that capital depletions would taper off quickly as property damages move toward zero, which is a first indication of non-linearity. Turning to the other end of the spectrum, the remaining three scenarios—all of them multiple-breach cases—stand out in terms of total damages and capital depletion. For the two scenarios with property damages around EUR 6 bn., we estimate a CET1 decline of around 75 basis points. Non-linearity is most pronounced when moving to the most extreme scenario, for which we estimate a capital depletion of close to 110 basis points. This scenario considers the extremely unlikely case in which a major multiple-breach event leads to a flooding of most of the western part of the Netherlands.

In terms of composition, we find that most of the capital depletion is due to an increase in the risk weights. The top panel of Figure 5 provides a breakdown across scenarios between the two key components of the change in the CET1 ratio (as per Equation (12)). For the single-breach scenarios, the average contribution of expected losses to capital depletion is around 6 basis points, i.e. around 15% of the total depletion. The average contribution of risk weighted assets is, in turn, 34 basis point. When moving to the multiple-breach scenarios, the RWA effect is even more prominent. On average, the contribution to the overall depletion is 66 basis points. Another decomposition (Figure 5, bottom panel) indicates that the role of PDs already strongly increases when considering multiple-breach scenarios. Across the 38 adverse scenarios, most of the capital depletion is due to increase in LGDs. However, the contribution of the PD becomes pronounced for multiple-breach scenarios. It should be noted again that this increase in PDs is purely due to damage-related increases in loan-to-value ratios. Adding further macrofinancial shocks, such as growth declines, will further amplify the capital impact through higher PDs.

(insert Figure 5 around here)

The key linking pin underlying such aggregate results is Equation 3, which connects flood-related damages (the ϕ parameter) to loan-to-value ratios. To explore this mechanism, we use outcomes for two of the 38 flood scenarios. Figure 6 shows results for one single-breach scenario with a relatively low overall impact (top panel) and results for one multiple-breach scenario with a fairly high outcome (bottom panel). Each red circle corresponds to a postal code area that is affected under the respective scenario. The size of the circle indicates how many properties within that area would be affected by a flood. The horizontal axis indicates the average property damage in that area (i.e. the mean value for ϕ), while the vertical axis indicates the average starting point of the loan-to-value ratio. Of these four metrics, the LTV ratio is a key indication of financial vulnerabilities, while the other three indicate how vulnerable the property itself would be to floods.

(insert Figure 6 around here)

Starting with the top panel of Figure 6, we sketch the reasons why the overall impact remains low in the single-breach scenarios. One factor for the low impact is that single-breach events impact only few postal codes areas. In addition, in some of those areas relatively few properties are damaged. Third, the average damages remain contained, in this case to 20% or less. Lastly, this flood affects also affects a few areas with fairly low starting-point LTV ratios. In other words, the flood affects areas where the existing level of financial vulnerabilities (as measured via LTV ratios) was low. The bottom panel of Figure 6 shows the other end of the spectrum. In this extreme multiple-breach scenario, many postal codes areas are affected. Within those areas, the number of damaged properties is high, while average damages can run up to 70%. Lastly, this extreme flood also impacts areas where financial vulnerabilities are larger, as seen by the relatively high starting-point LTV ratios.

7 Discussion

Taken at face value, our estimates for tail-risk events would suggest that the Dutch banking system is resilient in the face of flood risk. We find that the system-wide capital depletion under adverse flood scenarios would be mostly between 30 and 50 basis points. Relative to recent levels of bank capital, such depletions seem manageable. In this section, we provide some further intuition for this finding. At the same time, we highlight how three elements of our research design are relevant for the interpretation of our findings.

Starting with flood-related vulnerabilities, although a sizeable portion of the Netherlands is below sea level, our findings reflect the fact that the level of flood protection is extremely high. Considering the occurrence of a flood—let alone a multiple-breach event—already implies that one is imagining situations with low levels of probability of ever occurring. More importantly, given the set-up of its flood protection mechanisms, a local flood in one part of the Netherlands would not automatically also affect the rest of the country. This level of containment makes it unlikely that a large share of banks' real estate exposures would be damaged at the same point in time. Figure 7 provides an illustration of this point. The top panel focuses on banks' exposures to residential real estate, while the bottom panel focuses on commercial real estate. In both panels, the horizontal axis shows (as a fraction) how many of the properties in the loan-level data are damaged in the various scenarios. The vertical axis shows the average value loss (as a fraction of current value) of damaged properties. Hence, the values on the vertical axis correspond to the parameter ϕ defined above. In nearly all of the flood scenarios we consider, more than 95% of the properties would not incur damages. The main exception is the extreme scenario for the western part of the Netherlands. Both for residential and commercial real estate, this particular extreme multiple-breach flood would affect close to 25% of the properties that serve as loan collateral.

(insert Figure 7 around here)

Perhaps most importantly, our paper highlights that loan characteristics are highly relevant also. In this paper, given our focus on real estate, we have focused on loan-to-value ratios as a key aspect of financial vulnerabilities. In some flood scenarios, the damaged property serves as collateral to loans with fairly low starting-point LTV ratios. From a financial stability perspective, this means there is a degree of resilience to flood events. In such cases, mild to moderate floods would not immediately lead to large declines in bank capital.

Turning to assumptions, our paper has abstracted from various other types of socioeconomic impacts that a flood may have. Even though property damages will be a relevant component of the overall effect on bank capital, there are other ways in which floods could wreak havoc. For instance, if a major flood were to lead to a decline in economic growth, the increased probability-of-default would lead to a further increase in banks' credit losses (Caloia and Jansen, 2021). Also, if climate change were to continue, more frequent floods could mean that banks have less time available to restore their capital position. We leave the consideration of such macrofinancial channels for future research.

A second important assumption concerns the damage calculations that we use. Even though we use the most current damage curves available in the Dutch context, there is some debate on whether these curves may be on the conservative side. For instance, Endendijk et al. (2023) compare actual damages in the wake of the 2021 summer floods to the numbers suggested by the method of Slager and Wagenaar (2017). They find indications that actual damages may have been higher than suggested by the SSM-method for inundation depths below 175 cm. Any upward revision of damages curves would have direct implications for our credit risk calculations, naturally.

A last important assumption to highlight remains our modelling of inundation depth at the level of four-digit postal code areas. In reality, water stress will not be uniformly distributed across postal code areas. Given the granularity of the exposures data available to us, there is no immediate way to assess the extent to which this affects our estimates for property damages. If such more detailed data were readily available, this would be one clear dimension along which the precision of our estimates could be assessed.

8 Conclusions

This paper uses scenario analysis to study how property damages caused by severe and unlikely floods would affect credit risk for a sample of Dutch banks. We focus on specific floods causing property damages that would not be covered by current insurance policies. We calculate the implications for banks' capital positions using geocoded data on EUR 650bn. in exposures backed by real estate. We find that the system-wide capital depletion would mostly range between 30 and 50 basis points. In terms of drivers, we show that the overall financial impact is a combination of flood-related vulnerabilities and financial vulnerabilities. In terms of the latter factor, we highlight how in some cases floods would have mild effects as they would affect properties that serve as collateral for loans with fairly low loan-to-value ratios.

Compared to current levels of bank capital, the estimated degrees of flood-related capital depletion seem manageable. This finding is not unexpected, given the investment that the Netherlands continues to make in an adequate system of flood protection. The probability of a major flood is very low indeed. Also, in case a breach in flood defence system were to occur, effects in terms of damaged properties would often still be contained. However, it is also important to note that our estimations focus on one particular channel (i.e. property damages) and are subject to a number of assumptions (for instance concerning damage calculations). In particular, an open question is how our conclusions on implications for financial stability would be affected if climate change continues. More frequent floods and associated macroeconomic disruptions may then have stronger implications for financial stability than these present-day estimates on increased credit risk indicate. For instance, if floods were to happen more often, banks would have less time available to replenish capital. In addition, if floods were to lead to recessions, banks would face a further increase of credit losses. A third effect could be that a flood event makes this type of risk more salient, which could lead to a revaluation of properties located in flood zones.[15] We leave the exploration of such additional considerations for future research.

[15]Based on two surveys of Dutch homeowners, Jansen (2023) finds that flood risk awareness has indeed increased since the 2021 floods.

References

Acharya, V.A., R. Berner, R. Engle, H. Jung, J. Stroebel, X. Zeng, and Y. Zhao (2023). Climate stress testing. *Annual Review of Financial Economics* 15: 291-326.

Atreya, A. and Ferreira, S. (2015). Seeing is believing? Evidence from property prices in inundated areas. *Risk Analysis* 35(5): 828-848.

Auffhammer, M. (2018). Quantifying economic damages from climate change. *Journal of Economic Perspectives* 32(4): 33-52.

Battiston, S., Mandel, A., Monasterolo, I., Schütze, F., and Visentin, G. (2017). A climate stress-test of the financial system. *Nature Climate Change* 7: 283-288.

Beltrán, A., Maddison, D., and Elliott, R. (2018). Is flood risk capitalised into property vales? *Ecological Economics* 146: 668-685.

Beltrán, A., Maddison, D., and Elliott, R. (2019). The impact of flooding on property prices: A repeat-sales approach. *Journal of Environmental Economics and Management* 95: 62-86.

Blickle, K.S., Hamerling, S.N., Morgan, D.P. (2022) How bad are weather disasters for banks? Federal Reserve Bank of New York Staff Reports No. 990.

Bolton, P., Després, M., Pereira da Silva, L. A., Samama, F. and Svartzman, R. (2020). The green swan: Central banking and financial stability in the age of climate change. Bank for International Settlements.

Bosker, M., Garretsen, H., Marlet, G. and Van Woerkens, C. (2019). Nether Lands. Evidence on the price and perception of rare natural disasters. *Journal of the European Economic Association* 17: 413-453.

Caloia, F. G. (2022). Borrower-based measures, house prices and household debt. De Nederlandsche Bank Working Paper No. 738, Available at SSRN: `http://dx.doi.org/10.2139/ssrn.4034335`.

Caloia, F. G. and D. Jansen (2021). Flood risk and financial stability: Evidence from a stress test for the Netherlands. De Nederlandsche Bank Working Paper No. 730.

Carney, Mark J. (2015). Breaking the tragedy of the horizon—climate change and financial stability. Speech at Lloyd's of London, 29 September 2015.

Cochrane, J. H. (2021). A convenient myth: Climate risk and the financial system. *National Review*, 17 November 2021.

Daniel, V. E., Florax, R.J.G.M., and Rietveld, P. (2009). Floods and residential property values: A hedonic price analysis for the Netherlands. *Built Environment* 35(4): 563-576.

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

European Systemic Risk Board (2016). Too late, too sudden: Transition to a low-carbon economy and systemic risk, Reports of the Advisory Scientific Committee. No 6 / February 2016.

Garbarino, N., and Guin B. (2021). High water, no marks? Biased lending after extreme weather. *Journal of Financial Stability* 54(100874).

Gramlich, D., T. Walker, Y.Zhao, and M. Bitar (2023). After the storm: Natural disasters and bank solvency. *International Journal of Central Banking* 19(2): 199-249.

Hansen, L. P. (2022). Central banking challenges posed by uncertain climate change and natural disasters. *Journal of Monetary Economics* 125: 1-15.

Hino, M. and Burke, M. (2021). The effect of information about climate risk on property values. *PNAS* 118(17).

Jansen, David-Jan (2023). Homeowners and flood risk: A disconnect between awareness and actions? DNB Working Paper No. 791.

Jung, H., Engle, R. and Berner, R. (2021). Climate stress testing. Federal Reserve Bank of New York Staff Reports No. 977.

Klomp, J. (2014). Financial fragility and natural disasters: An empirical analysis. *Journal of Financial Stability* 13: 180-192.

Koetter, M., Noth, F., Rehbein, O. (2020). Borrowers under water! Rare disasters, regional banks, and recovery lending. *Journal of Financial Intermediation* 43(100811).

Mol, J. M., Botzen, W.J., Blasch, J.E. and de Moel, H. (2020). Insights into flood risk misperceptions of homeowners in the Dutch river delta. *Risk Analysis* 40(7): 1450-1468.

Mutlu, A., D. Roy, and T. Filatova (2023). Capitalized value of evolving flood risks discount and nature-based solution premiums on property prices. *Ecological Economics* 205(107682).

Network for Greening the Financial System (2019). A call for action: Climate change as a source of financial risk.

Pistraka, A. and Jonkman, S. N. (2010). Damage to residential buildings due to flooding of New Orleans after hurricane Katrina. *Natural Hazards* 54: 413-434.

Slager, K. (2019). Handboek overstromingsrisico's op de kaart: Over de methode van kaartproductie.

Slager, K., and Wagenaar, D. (2017). Standaardmethode 2017. Schade en slachtoffers als gevolg van overstromingen.

Vermeulen, R., Schets, E. Lohuis, M., Kölbl, B., Jansen, D., and Heeringa, W. (2021). The heat is on: A framework for measuring financial stress under disruptive energy transition scenarios. *Ecological Economics* 190(107205).

Waller, C. J. (2023). Climate change and financial stability. Speech given at the IE University-Banco de España-Federal Reserve Bank of St. Louis Conference Current Challenges in Economics and Finance, 11 May 2023.

Figure 1: Distribution of real estate exposures

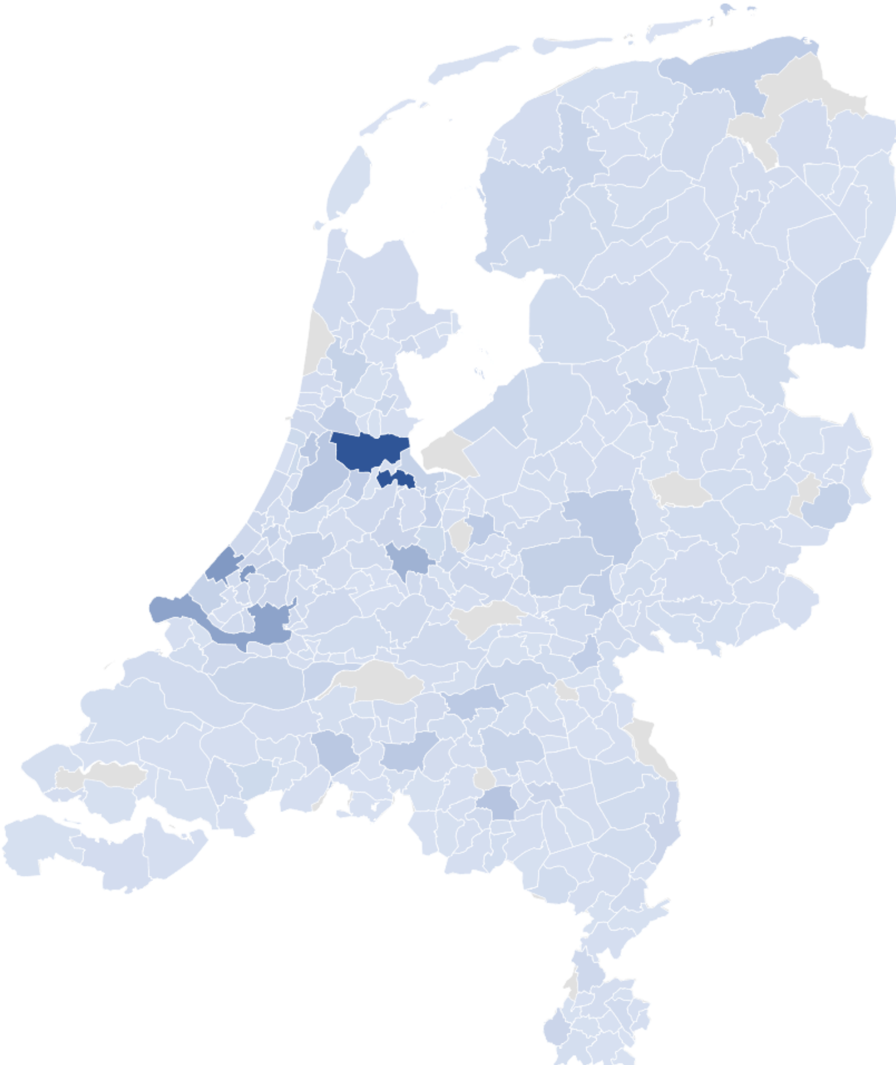

Note: This figure shows the geographical distribution of EUR 650bn. in Dutch banks' real estate exposures by municipality at end-2020. Darker colours indicate a higher level of exposure. This figure combines information for exposures to residential and commercial real estate.

Figure 2: Scenario set for single-breach floods

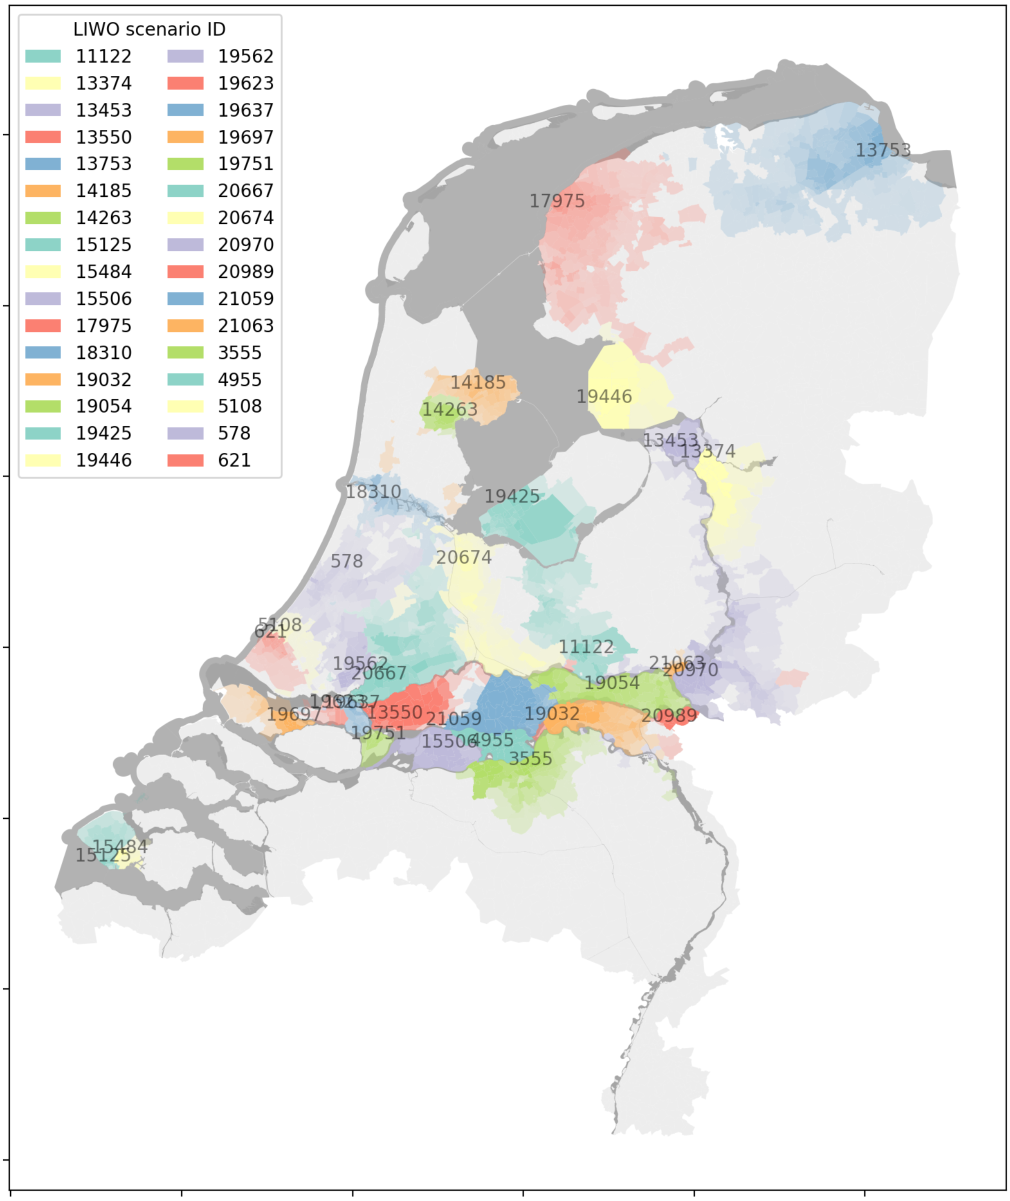

Note: This figure gives an overview of 32 scenarios with single-breach floods. In each scenario, a local breach in the system of flood defence leads to a flood in a specific part of the country. The source for the flood scenarios is the LIWO, which is an open-source information system. From the LIWO, we select the scenarios that have the largest impact in each compartment of the flood protection system. In addition, we use a cut-off value of EUR 500mn. for overall estimated damages.

Figure 3: Scenario set for extreme multiple-breach floods

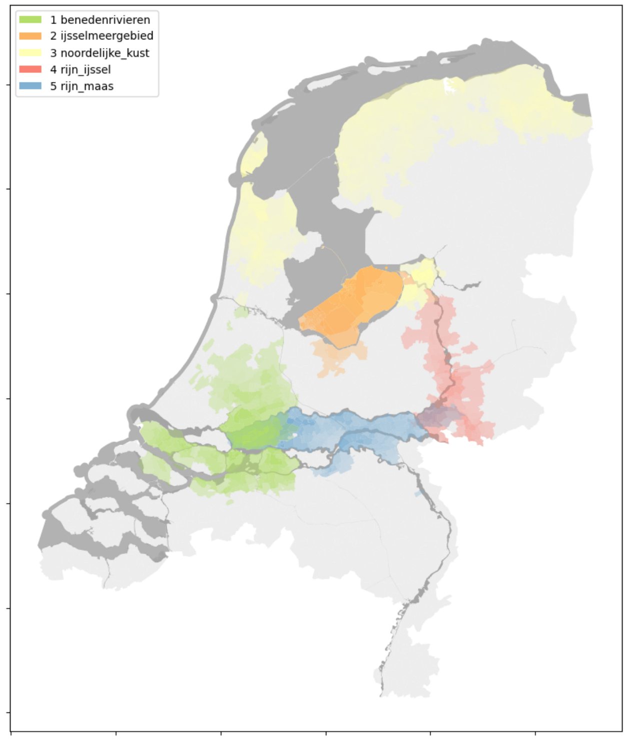

(a) Extreme scenarios 1 to 5

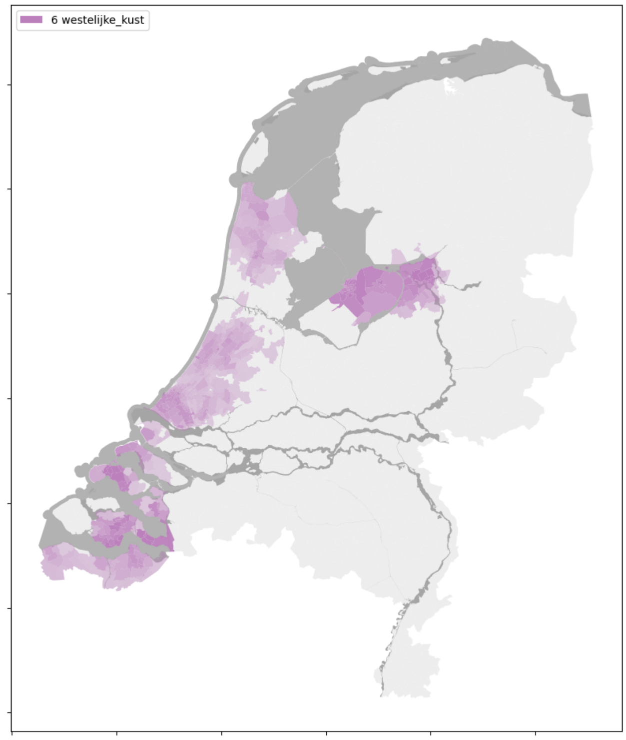

(b) Extreme scenario 6

Note: This figure gives an overview of six scenarios for extreme multiple-breach floods. Dutch flood experts constructed these six extreme multiple-breach scenarios in 2007. These six scenarios are intended to represent extreme impacts that are still theoretically conceivable, yet very unlikely.

Figure 4: Property damage and capital depletion under 38 flood scenarios

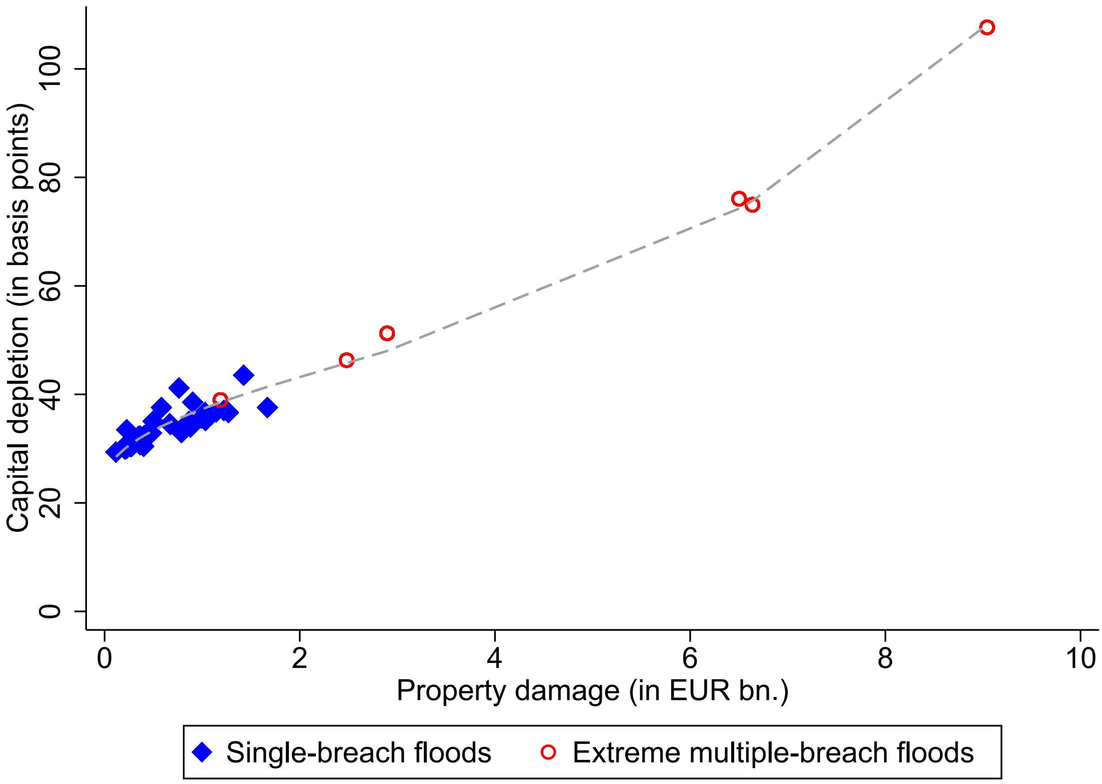

Note: This figure presents a scatter plot for estimated property damages (in EUR bn., horizontal axis) and banks' capital depletion (in basis points, vertical axis). The blue diamonds show results for 32 single-breach flood scenarios; the red circles indicate results for six extreme multiple-breach scenarios. The property damages (horizontal axis) are the sum total of direct damages to residential properties and commercial real estate that serve as collateral for bank loans. Insurance policies would not cover these property damages. We use the CET1 ratio to measure bank capital. The capital depletion is a weighted average across a sample of eight Dutch banks. The dashed gray line plots an estimated fractional polynomial of degree 2.

Figure 5: Decomposing the capital depletion

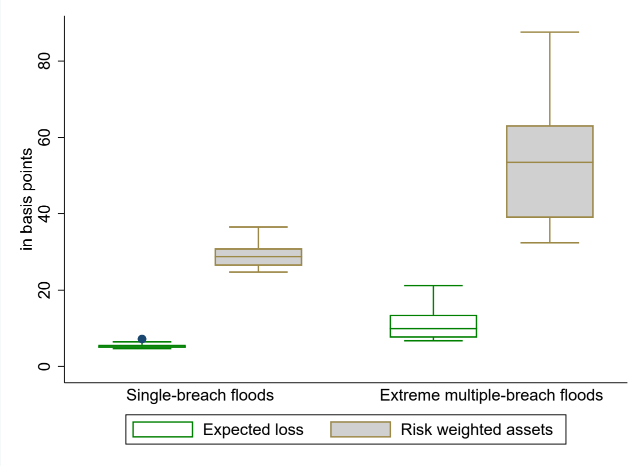

(a) Components of the capital ratio

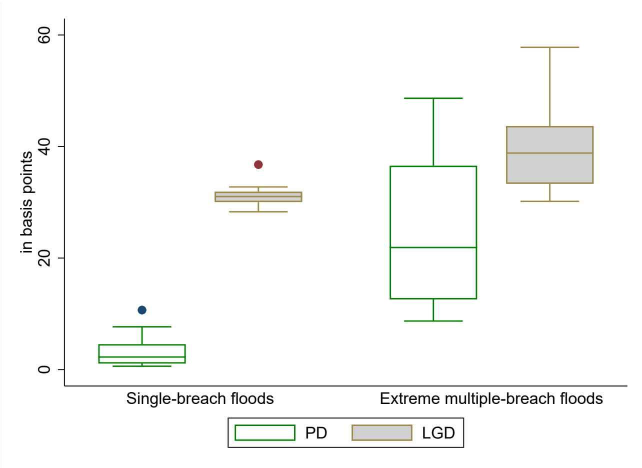

(b) Credit risk parameters

Note: This figure shows a decomposition for capital depletions across the 32 single-breach and the 6 multiple-breach scenarios. In the top panel, the contributions to the overall depletion are split between expected losses and risk weights, as per Equation (12) in the main text. The bottom panel shows a decomposition across credit risk parameters, i.e. the probability of default (PD) and the loss given default (LGD).

Figure 6: Average damages and starting-point loan-to-value ratios

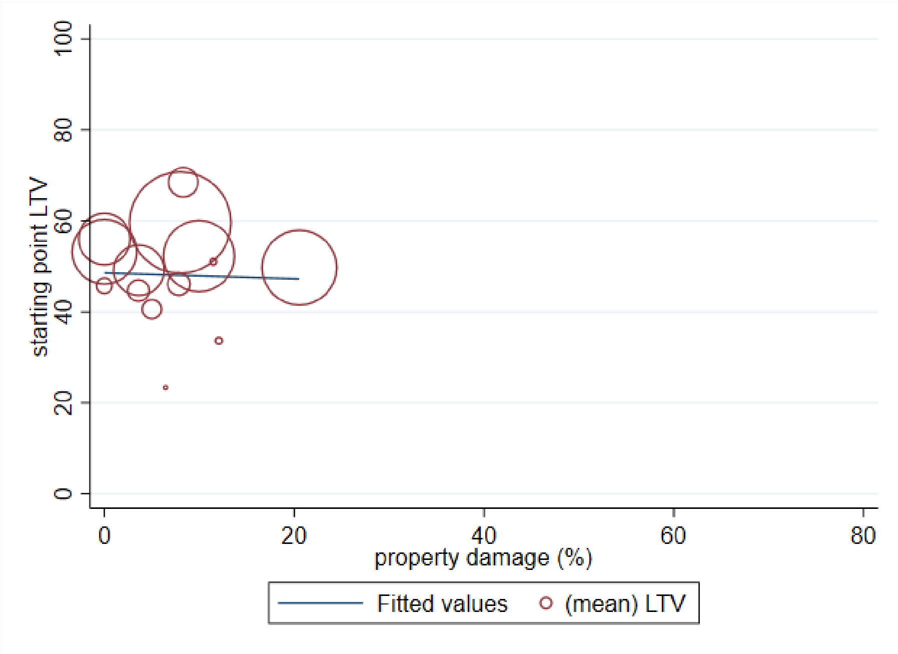

(a) Single-breach flood scenario

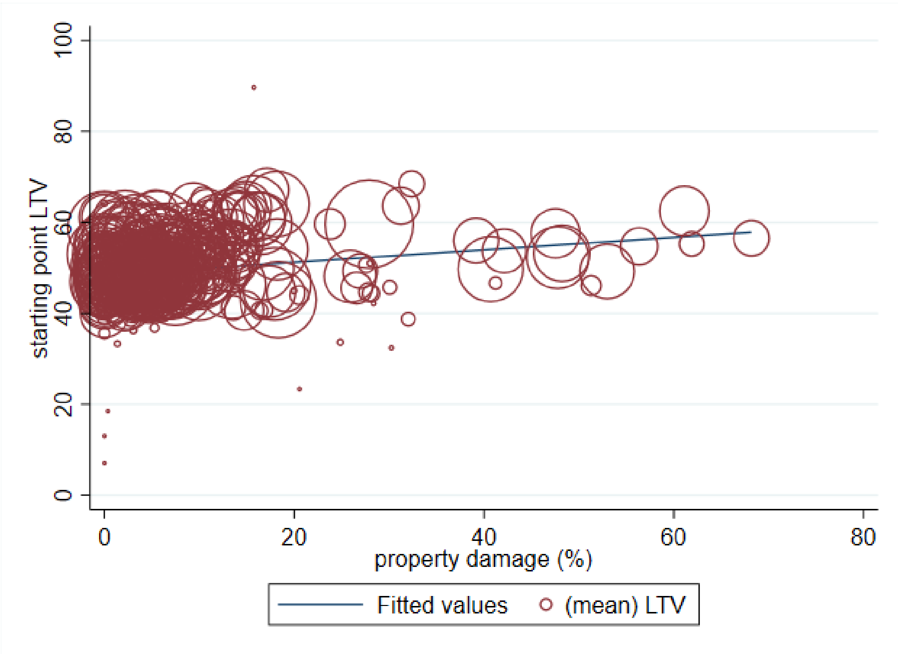

(b) Extreme multiple-breach flood scenario

Note: This figure illustrates four drivers of capital depletion based on two scenarios, one with a single-breach scenario (top panel) and one with an extreme multiple-breach scenario (bottom panel). Each red circle corresponds to a postal code area that is affected under the respective scenario. The size of the circle indicates the number of properties within that area that are affected by the flood. The horizontal axis indicates the average property damage in that area (or ϕ), while the vertical axis indicates the average starting point of the loan-to-value ratio.

Figure 7: Damaged properties and value losses across flood scenarios

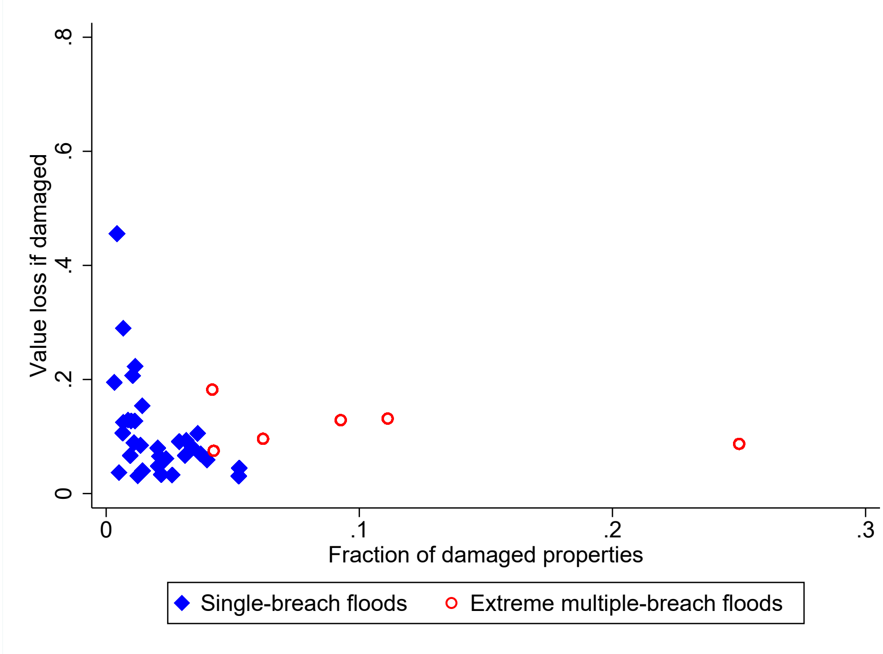

(a) residential

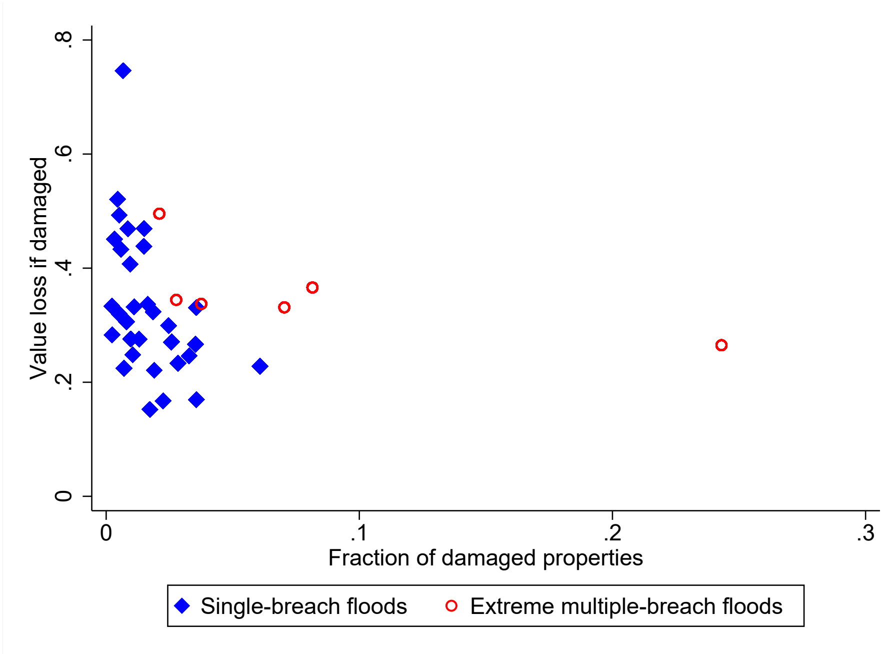

(b) commercial

Note: This figure show scatter plots for the fraction of damaged properties (horizontal axis) and associated property damages (as fraction of value, vertical axis) The blue diamonds show results for 32 single-breach flood scenarios; the red circles indicate results for six extreme multiple-breach scenarios. The average property damages correspond to mean of the parameter ϕ in equation 2 of the main text. This parameter ϕ indicates the percentage value loss for affected properties in a given flood scenario. Panel a) shows this relationship for residential properties; panel b) focuses on commercial real estate.

Table 1: Descriptive statistics

Loan characteristics	Residential Real Estate *mean*	Residential Real Estate *median*	Commercial Real Estate *mean*	Commercial Real Estate *median*
Interest rate (%)	2.60	2.35	2.23	2.25
LTV (%)	52.3	54.5	53.4	49.0
LTV-O (%) New business	76.7	83.0	66.0	65.0
LTV-O (%) Starters, New business	86.6	95.0	-	-
Interest-only (%)	50.9	50.7	22.4	0.00
Outst.amount (€)	182,672	152,750	1,677,100	357,707
Protection value (€)	384,371	321,000	4,914,289	605,810

	Residential Real Estate		Commercial Real Estate	
Type of Property	Detached (%)	32.8	Residential (%)	52.7
	Apartment (%)	14.3	Offices (%)	26.9
	Terraced (%)	42.4	Retail (%)	7.0
	Commercial (%)	2.3	Industrial (%)	4.9
Location	Large 4 cities (%)	14.3	Large 4 cities (%)	24.1
Total Portfolio	Instruments (N)	6,129,774	Instruments (N)	46,744
	Protections (N)	5,352,319	Protections (N)	222,928
	of which properties (N)	3,030,383	of which properties (N)	175,721
	Contracts (N)	3,124,998	Contracts (N)	40,979
	Exposure (€m)	577,217	Exposure (€m)	70,730
	In default (%)	0.90	In default (%)	3.62

Note: Summary statistics for loan-level data related to residential and commercial real estate for eight Dutch banks at end-2020. The top panel presents information on risk characteristics, such as the loan-to-value (LTV) ratio and the LTV ratio at origination (LTV-O). The bottom panel presents information on the type of property, the location, and the portfolio characteristics. The labels *Starters* and *New business* refer to first-time buyers in the housing market and to newly originated loans, respectively.

APPENDIX

1. Flood types relevant to the Netherlands

As described in Slager (2019), the at-risk areas in the Netherlands are distributed across four catchment areas, i.e. that of the rivers Rhine, Meuse, Scheldt, and Ems. For each of these four catchment areas, a further distinction is made based on two dimensions: whether or not there is protection against water, and whether the water system is primary or regional. All four combinations of these two dimensions are indicated by the letters A - D. Table A.1 has an overview. This paper focuses on areas at risk from flood type B. Property damages by floods of this type are almost always excluded from coverage by insurance policies.

Table A.1 Taxonomy of flood risks in the Netherlands

This table describes the basic taxonomy for areas with a potentially significant flood risk. The columns indicate whether or not there is protection against water; the rows indicate whether the water system is main or regional. The resulting four classes are indicated by the letters A - D. See Slager (2019).

	Protection against water	
	yes	*no*
Water system		
Main	B	A
Regional	C	D

2. Return periods for single-breach flood scenarios

Figure A.1 provides information on the return periods for the 32 single-breach flood events. The source for these return periods is the LIWO (variable ‘Ref’). The vertical axis reports (on a log scale) the return periods (in years). The Netherlands is divided in a number of areas when it comes to flood protection (in Dutch: dijkring). The numbers on the horizontal axis correspond to those geographical areas. Most of the return periods are between once every 100 and once every 10,000 years. The highest probability in our scenario set is once every 100 years, for example for LIWO scenario no. 19562. The lowest probability is less than once every 10,000,000 years. This is for LIWO scenario 5108—a scenario with a coastal flood affecting Scheveningen.

Figure A.1 Return periods of single-breach flood scenarios

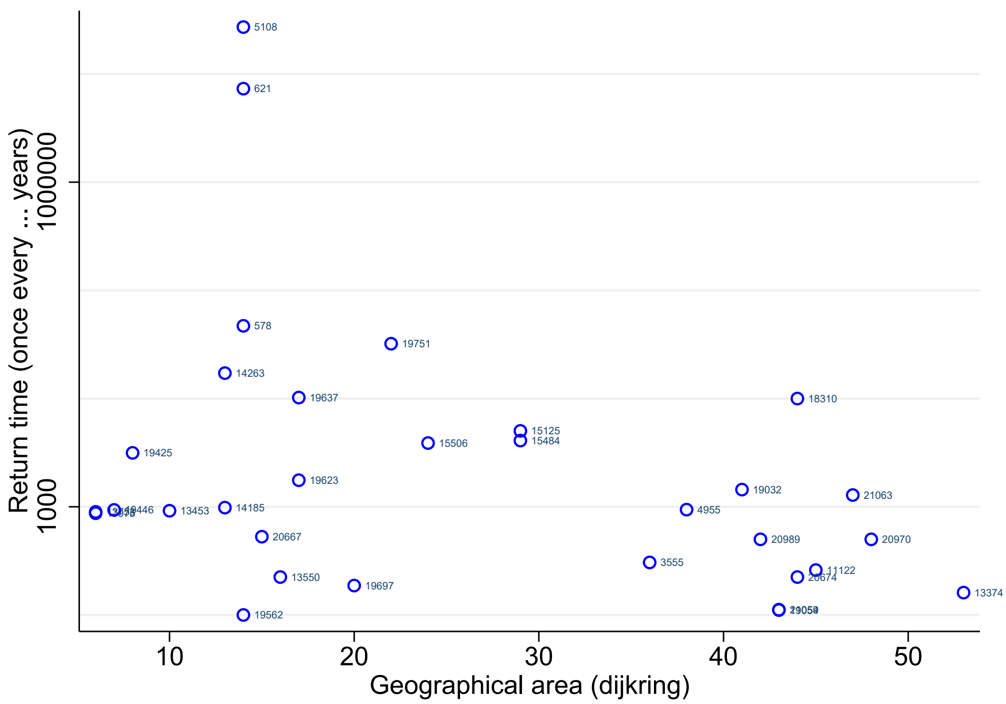

Note: The figure shows on the vertical axis the return frequencies (in years on a log scale) for the 32 flood scenarios with a single breach in flood protection. The labels in the graph correspond to the number of each scenario in the LIWO. The horizontal shows the number of the area of the Netherlands (in Dutch: dijkring) in which the flood would occur.

3. Background on damage calculations

3.1 Damage curves from the SSM methodology

Our method for calculating property damage broadly follows the SSM-method, the standard method for calculating flood damage in The Netherlands (Slager and Wagenaar, 2017). Figure A.2 shows the damage curves that we apply. Panel A shows two damage functions for residential real estate, while panel B shows three damage functions for commercial real estate. In all cases, the curves indicate how to map inundation depth (horizontal axis) into a fraction of the maximum damage amounts (vertical axis). This latter fraction corresponds to the parameter θ in the main text.

Starting with the top panel, the red curve shows the relevant damage function for single-family homes as provided by Slager and Wagenaar (2017). For floods up to 2.5 meters, the damage factor remains below 0.2. Between 2.5 meters and 5 meters, the damage fraction quickly increases to 1.

The blue line in panel A denotes the damage function we use for apartments. Slager and Wagenaar (2017) provide, in fact, three separate damage functions for apartments. These three separate curves are for apartments that lie, respectively, on the ground floor, the first floor, or above the first floor. Of these three curves, the one for the ground floor is the most severe, whereas the damage function for apartments above the first floor would be least severe. In fact, that particular damage function only has positive values for floods above 5 meters.

Given that our data set does not have information on the floor level of the apartment, we need to make some assumptions on the relevant damage curve. Our estimates use a damage curve that gives most weight to the first-floor damage curve. We see that as a reasonable middle ground to avoid both an overestimation and an underestimation of damages. Next to that, we give some weight to the damage function for the ground floor. The blue curve in panel A weights the curve for ground-floor apartments with 35% and the curve for first-floor apartments with 65%. This means that the curve remains relatively flat for floods with inundation depths below 2.5 meters. With this particular curve, we find a strong correlation between our damage estimates and the number reported in the LIWO, as discussed in Section 3.2 of the Appendix. We considered four additional weighting schemes that, for instance, also assign some weight to the curve for apartments above the first floor. In those other cases, the correlation between our estimates and the LIWO numbers would be lower, though still at least 0.85.[16]

The bottom panel of Figure A.2 shows four separate damage curves for commercial real estate suggested by Slager and Wagenaar (2017). These functions are, respectively, for retail properties, industrial properties, offices, and residential properties. Of these, the curve for retail properties is the steepest, reaching the maximum damage factor for floods of three meters. For residential properties that serve as collateral for commercial real estate loans, we use the same damage curves as for apartments.

[16] Results available upon request.

Figure A.2 Flood damage functions

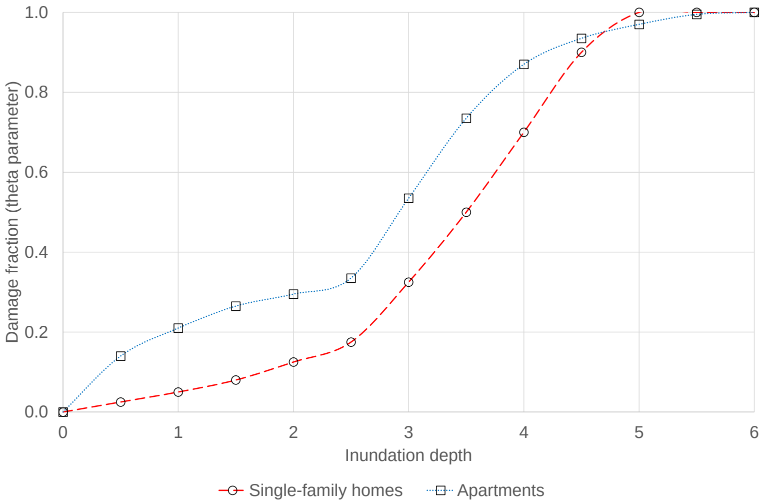

(a) residential real estate

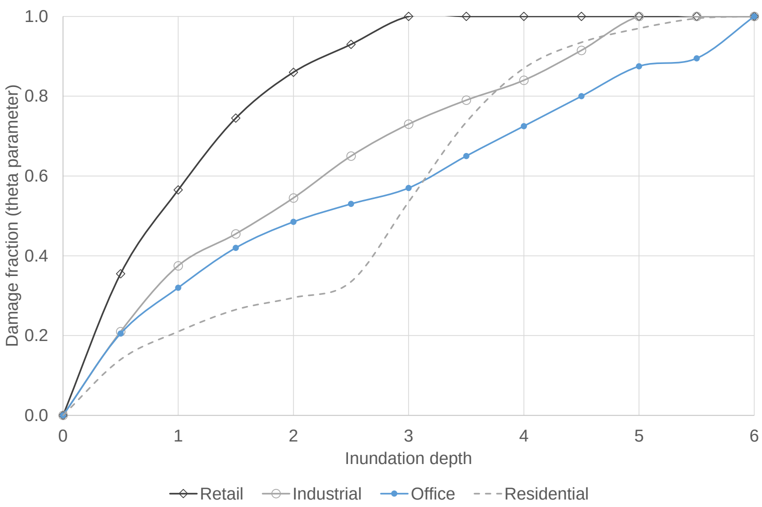

(b) commercial real estate

Note: The figure shows the various damage functions used in this paper. The top panel focuses on residential real estate (either single-family homes or apartments), the bottom panel shows damage curves for four types of properties (retail properties, industrial properties, offices, residential properties) that serve as collateral for commercial real estate loans. The horizontal shows the inundation depth (in meters), while the vertical axis shows the damage fraction (parameter θ in the main text). The damage functions are based on Slager and Wagenaar (2017).

3.2 Comparing our damage estimates to the LIWO

The LIWO reports estimated property damages for residential real estate against which we can compare our computations based on the loan-level data. To make the comparison, we can compute damages at the four-digit postal codes and then aggregate per scenario. In doing so, two methodological points are important.

First, to compare against the LIWO numbers, we need to calculate damages in 2011 prices. This means we will no longer use the π parameter to account for price level differences between 2011 and 2020.

Secondly, we need to account for the different coverage between our loan-level data and the LIWO. Our damage calculations only cover a part of the properties used to calculate damages for the LIWO. In our case, the data set only has information on properties that serve as loan collateral. In contrast, the LIWO uses the property registry and, therefore, has information on all properties that would be affected under each of the flood scenarios. To make a fair comparison, therefore, we apply a rescaling of our damage estimates as follows. Per postal code, we collect information from Statistics Netherlands on the number of residential properties. We then compute a scaling factor, per postal code, by dividing the total number of properties (according to the CBS) by the number of residential properties in our data set. We then multiply our damage estimates, per postal code, with this scaling factor and again compute aggregate damages per scenario.

Figure A.3 shows a scatter plot comparing damages estimates (in EUR bn., price level 2011) as reported in the LIWO (horizontal axis) and our calculations following the SSM method, with an adjustment for the differences in the number of properties (vertical axis). The correlation between the two sets of estimates is 0.95.

Measured in 2011 prices, the damages in the 32 single-breach scenarios we study would range between EUR 0.5 bn and EUR 3.5 bn. In closing, it is important to emphasize that these are not the numbers that we use in our credit risk calculations. These numbers reported on the vertical axis in Figure A.6 are only calculated for the purposes of comparing against the LIWO damage estimates. In our credit risk calculations, we measure damages in 2020 prices and only consider direct damages to properties when these serve as loan collateral.

Figure A.3 Comparing damage estimates (in EUR bn.)

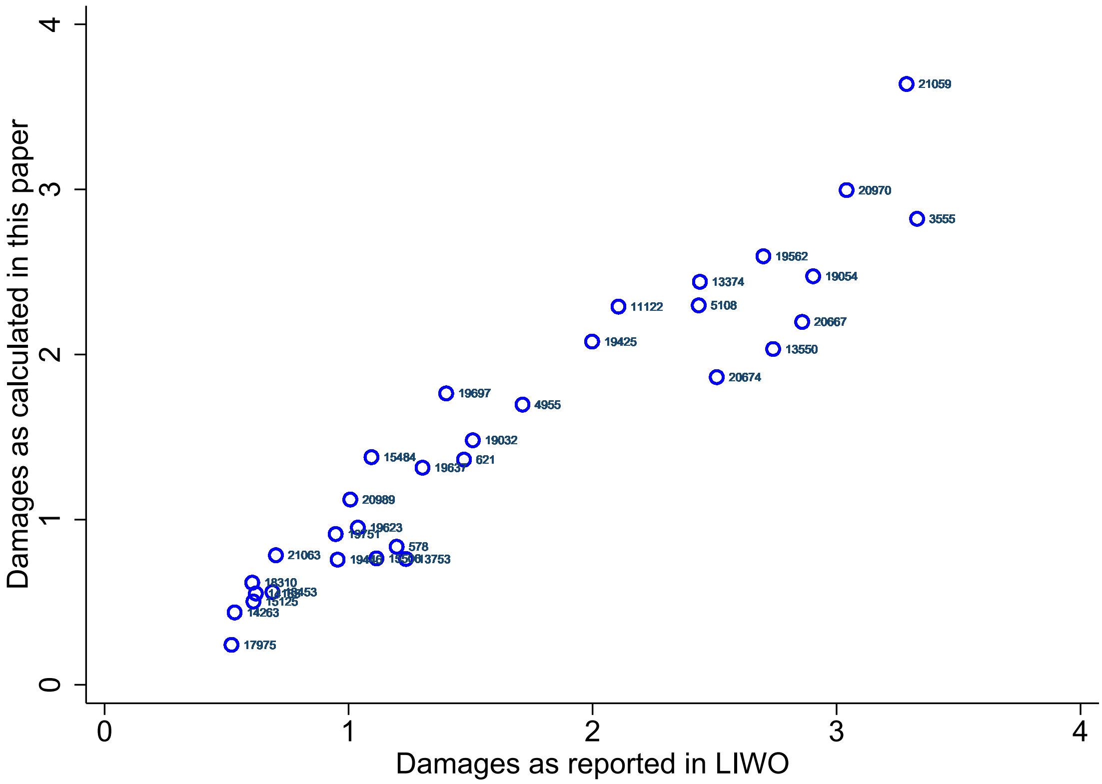

(c)

Note: For 32 single-breach flood scenarios, this figure compares our damage estimates for residential properties (vertical axis) against those available in the LIWO (horizontal axis). All number are in EUR bn. and in price level of 2011. For the numbers on the vertical axis, we apply a rescaling to account for sample differences between our data set and the LIWO. Our data set only has properties which serve as collateral for a loan, while the LIWO contains damage calculations for all properties in a postal-code area. $\rho = 0.95$.

4. Adjusting maximum damages to price level 2020

Our financial exposure data is for end-2020 but the SSM-method computes damages in price level of 2011. To adjust our damage calculations to 2020 prices, we need to calibrate scaling parameters per type of real estate, i.e. the parameter π in the main text. We broadly follow De Grave and Juch (2023), who provide guidance on translating the SSM damage amounts to 2022 prices.

For residential real estate, we use a maximum damage of EUR 1,200 per m2 instead of the EUR 1,000 currently used in the SSM approach. One approach for rescaling is using the consumer price index for housing. This would result in a rescaling by a factor of 1.217. An alternative is rescaling using an index for building costs, which would imply a factor of 1.178. We use a rescaling factor of 1.2 as the middle ground.

For commercial real estate, we start from Table 14 in De Grave and Juch (2023). This table reports maximum damages for different types of properties in 2022 prices. To rescale these into 2020 prices, we use the fact that these numbers are, in fact, based on estimates for 2020 but rescaled by an inflation rate of 12.9%. In the end, we use the following maximum amount (in EUR, per m2):

- For offices, EUR 1,423.
- For retail properties, EUR 1,591.
- For industrial properties, EUR 1,258.

Reference

De Grave, Peter, and Sanne Juch (2023). Update basisinformatie SSM 2022.

5. Inundation depths per postal codes

The two figures below present information on the distribution of inundation depths for the various flood scenarios. Figure A.4 presents results for the single-breach flood scenarios. Figure A.5 provides information for the six extreme multiple-breach flood scenarios. The horizontal axis in each figure represents the postal code (measured at the four-digit level).

Figure A.4 Inundation depths per scenario (single-breach floods)

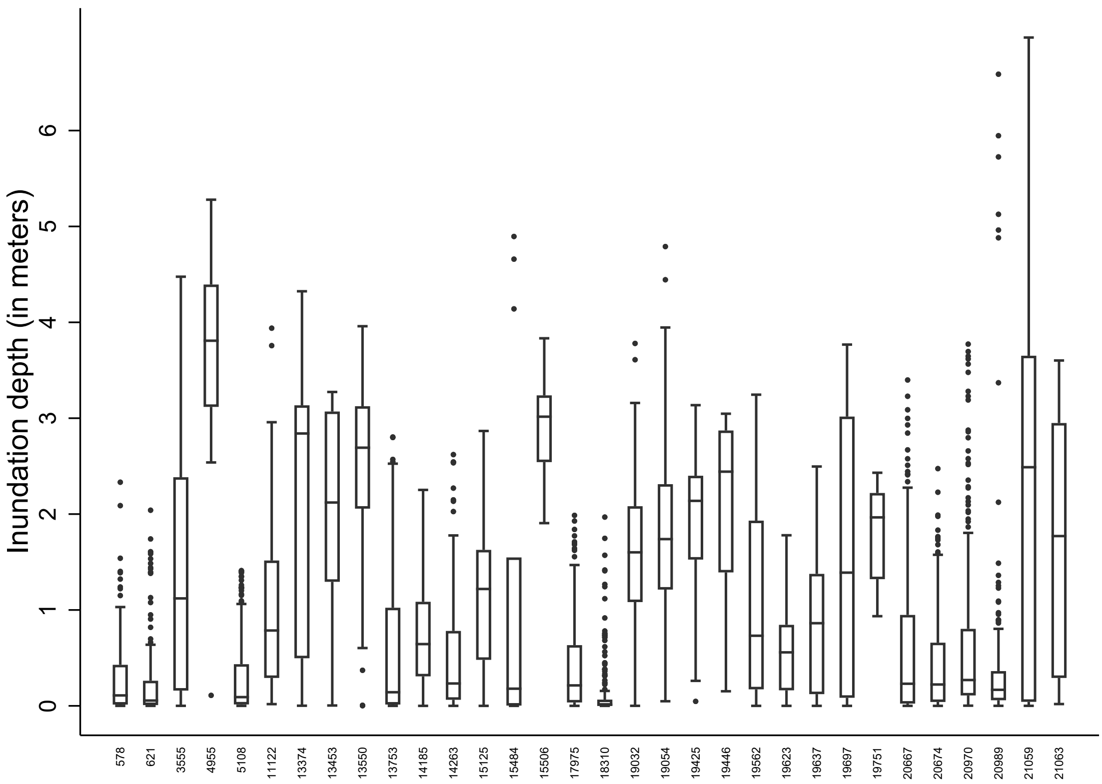

Note: This figure shows, for 32 single-breach flood scenarios, box plots for inundation depth (in meters, vertical axis) per scenario (using the labels available in the LIWO, horizontal axis).

Figure A.5 Inundation depths per scenario (extreme multiple-breach floods)

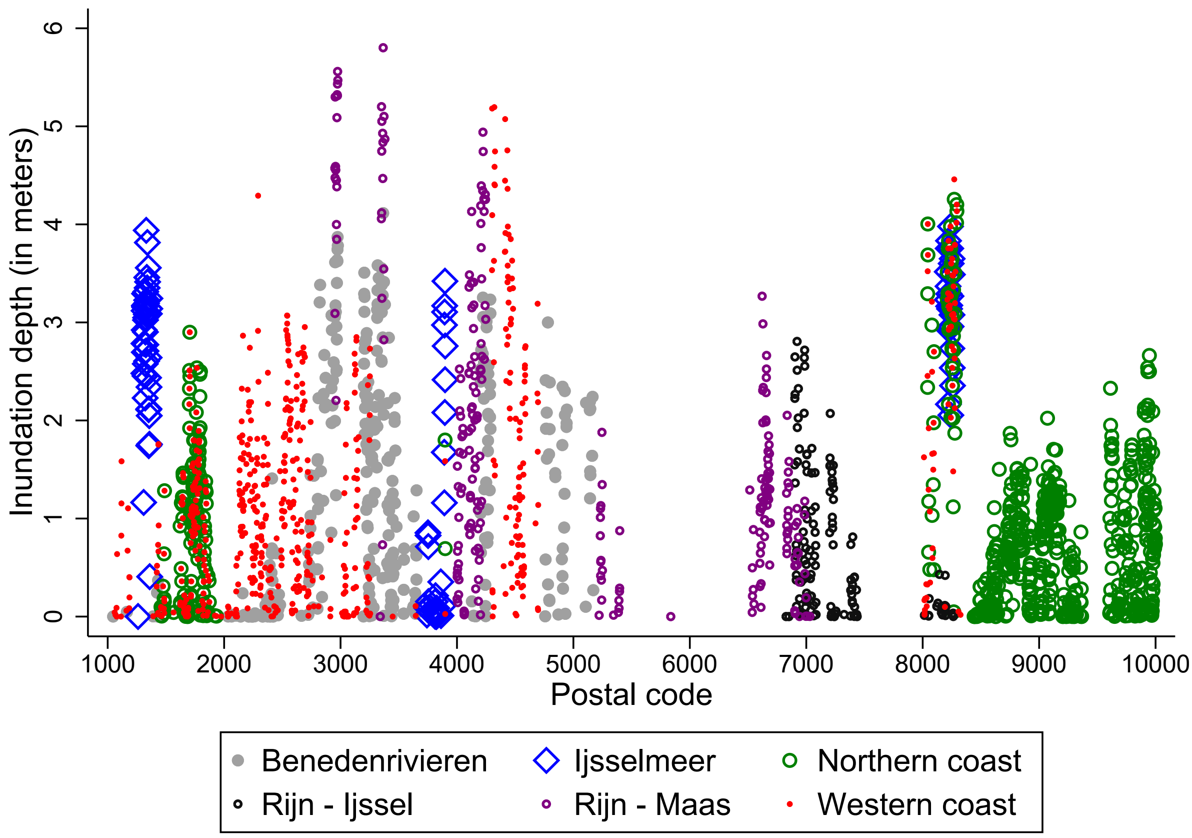

Note: This figure shows, for six extreme multiple-breach flood scenarios, the inundation depth (in meters, vertical axis) per four-digit postal code (horizontal axis) in the Netherlands.