Water label under credit restrictions

Understanding the relation between banking and the water label on flood risk price integration

K.C. van den Berg

TU Delft

Water label under credit restrictions

Understanding the relation between banking and the water label on flood risk price integration

By

Koen van den berg

4967530

Chair	Tatiana Filatova
First supervisor	Theodoros Chatzivasileiadis
Second supervisor	Omar Kammouh
Advisor	Asli Mutlu
Project duration	25th of March 2025 – 27th of April 2026
Faculty	Technology policy and Management

Preface

During my master and my bachelor~~,~~ I was always exploring my broad interests. Whenever I had time left, I would throw myself into books, articles and videos about these subjects, anything I could find really. During parties I would walk up to a random person and ask them anything about whatever they were interest in. I would learn more about the functioning of electricity grids to movie genres. Apparently, knowing how to cook garlic is a science in itself, who knew! Personally, I would spend my time on anything from microarchitecture of computer chips to monetary theory or I would start reading medical papers of subjects I would find interesting. In the past I questioned myself whether all these things were a bit too differing to be of any use, struggling to find a good reason as to why I should read them. To be honest, I just kept going because I found it fun. But lately, once I started my Master, Engineering and Policy Analysis, these interests suddenly became relevant.

I had developed strategies without really knowing it at the time. I could read basic medical papers, which helped me during certain projects. Other times I would remember helpful facts given to me from the partying students, providing me the solution to math problems (for an example, see appendix F2.2.4). In most cases it helped me quickly develop ideas, knowing the broad discussion allowed me to understand where I needed to search to find the answers I required.

When I first started this thesis, I wanted to see how these interests would help me. I knowingly started off quite ambitious, giving me plenty of space to grow while understanding I would not be able to do everything I promised. My supervisors recognized this and were willing to help me in that ambition. Tatiana and Theo stopped me when I would go off on a wrong trail and helped keep me bounded to the relevant subjects. Other times them and Asli would give me hints about certain problems, while still allowing me to figure things out on my own. Asli would help me out whenever I would get lost in the model, giving me the knobs I would need to dial in to make the model function properly. Finally, Omar, he mostly helped me in the beginning as before this I was unfamiliar with MESA, the Agent-Based Modelling Python package used for this thesis. He was also incredibly useful in keeping a broader overview of what was necessary. Each of them was instrumental in this thesis and I could not have done this without them. I especially want to thank them for allowing me great freedom and only stepping in whenever they felt necessary.

After I finished most the model and I still needed to write large parts of this thesis, I suddenly fell ill, making it difficult to work. When this started happening, I underestimated the impact this would have. Despite most of it now having passed while writing this, I am reminded of the great people in my life that have helped me deal with this. From my family and friends that have listened to my incessant rants, trying to understand what was happening, to my doctor, nurse practitioner and dietician that helped me figure it out. But also, the people of my graduation committee that have been incredibly understanding and even giving me some guidance while dealing with this. Most importantly, Theo, who, with his infectious personality, has surprisingly become one of my biggest supporters.

While this thesis has taken far longer than I initially anticipated, I can only look back with pride about what I was able to achieve with the help of my village. They have helped me when I was confused, frustrated, or lost, but they have also helped me understand my strengths and how to develop them. Supporting me to become the best version of myself.

Abstract

Netherlands enjoys the best flood protection system in the world. Although probabilities are small, still 33% of population lives in flood-prone areas and up to 80% of new housing is planned in climate-sensitive areas. Surveys repeatedly confirm that homeowners are unaware of these risks. Following the 2021 flood, the government set to ensure that citizens are flood-aware and resilient, instructing the Ministry of Infrastructure and Water to create a 'water labels' for homes as one of the instruments. At the same time, the European central bank and the Nederlandsche Bank are increasing their regulatory pressure towards banks to increase the interest rate in flood prone areas, due to the increased credit risk created by floods.

The goal of this thesis is to investigate the relationship between banking and the 'water label' through the usage of an ABM model. Through this modelling exercise the system is explored to understand the interactions between these two policies and how the outcomes of these policies might affect households. Its main results concern the Loan-To-Value (LTV) ratio valley, where due to strong competition on the market, LTV-ratio of lower-income households could be driven up, combined with the increased interest rate increases for flood-prone houses. The combination of these two policies might thus increase overall credit risk but will require other policies to reduce this effect. For the Ministry of Infrastructure and Water, these could come in the form of subsidies for lower-income households to floodproof their homes, or for the Central banks by introducing LTV-ratio restrictions for flood-prone homes, which can be reduced when structural adaptations to homes are introduced. However, this model does not take into account the difference in demand and flood risk. Flood risk in the Netherlands has large spatial heterogeneity, mainly affecting the areas in the Randstad, which also coincide with high demand. This causes the effect to likely be reduced if this is taken into account.

Table of Contents

1 Introduction

With the increasing risks of climate change, many different types of effects are expected to impact the general liveability through increased chance heat waves, floods, or wildfires. In the Netherlands specifically, flooding remains one of the key climate-induced hazards that might impact people's homes. Even in 2025, housing in a neighbourhood in Enschede has become unliveable following a heavy rainfall. The floods in these streets have become so severe that it is difficult to even pay for the damage, as the next flood will quickly undo their repairs (NOS Nieuws, 2025), with these floods expected to only become increasingly severe as climate change worsens (Schiermeier, 2011). However, in the past, informing households of flood risk has been difficult, due to much of it being up to the individual's own perception of risk (Bamberg et al., 2017; Bubeck et al., 2017; Van Valkengoed & Steg, 2019) New solutions are therefore not only necessary to help decrease the impact of these floods, but also to inform people of the potential flood risk a property holds.

To achieve this goal, households need to be made aware of the flood risk their homes face. However, people tend to experience biases in perceiving risks rather than understanding probabilities objectively. This also has implications for how flood risks are captured in housing markets. Empirical evidence has identified that prices decrease only after flood events, which are consequently forgotten after a few years, leading to prices returning to normal despite no actual new adaptations being introduced to these homes (Atreya et al., 2013; Bin & Landry, 2012; Mutlu et al., 2023). In response to this, information campaigns have been set up but have only been successful in reducing this premium in local buyers after a flood event occurred (Niu et al., 2025).

New solutions were therefore required to inform all buyers and households about these flood risks, especially as they are expected to become a larger problem in the future. In response, the Ministry of Infrastructure and Water has started developing a water label – "de Waterwijzer" – to fill this information gap, helping households understand the climate risk to their homes. Through creating property level flood prediction data, each home can get an accurate flood risk assessment (Tieman, 2025), helping potential buyers understand the flood risk associated with each particular home. For current homeowners, this, in turn, can help understanding the need to prepare for a flood, as even simple measures such as sandbags can decrease damage by up to 50% (Endendijk et al., 2022). Structural adaptations, like dry-proofing or wet-proofing homes, can also play a role, as they can help decrease vulnerability of homes to flooding, reduce the potential damage and eventually improve this label (De Lange et al., 2025; Slager, 2022).

At the same time this improved information availability will help banks in their credit assessment strategies, increasing mortgage rates in flood-prone areas (Phlippen et al., 2023), caused by the unexpected floods putting households in financial difficulties (Deelen et al., 2025) by not only creating large damages but also making locals aware of the flood risk in the area, thereby increasing their default rates (Kousky et al., 2020; Thomson et al., 2023). If these homes more accurately reflect their financial risk, the home's value will more closely reflect the real price. At the same time, this decrease creates a fiscal space for banks to deploy new financial products to help households floodproof their homes. Then the owner is awarded not only a more valuable home but also a lower interest rate (Phlippen et al., 2023). Similar to what is commonly referred to as a sustainable mortgage (ABN Amro, z.d.), usually referring to mortgages including loans for home energy efficiency improvements, it can be argued that expanding this concept to include physical climate risks can improve the adoption of climate adaptations, such as flood-proofing (Phlippen et al., 2023). From here on out, sustainable mortgages will only refer to the flood adaptation side. Adoption of these practices has been rather slow, with few banks actually increasing the interest rates in these areas (Fontana et al., 2025). However, the question remains how these two policies would interact. To help explore this issue, an Agent-Based Model (ABM) will be created to mechanisms of these changes on the housing prices, which can help assess the effectiveness of this policy.

Existing research has used different ABM models to explore the relationship between household flood risk and housing choices, with the focus primarily on how they perceive the risk. For example, Dubbelboer et al. (2017) & Abebe et al. (2020) strongly focus on flood adaptation, which can help improve a home’s resilience in the event of a flood. Both focus on how information might spread, with Dubbelboer et al. (2017) focusing on insurance premiums as a means of relaying information, while Abebe et al. (2020) interpret risk as spreading through social systems. Han & Peng (2019) introduced prospect theory, which posits that households prefer to avoid risk whenever possible. Tierolf et al. (2023) highlight the importance of experience on an individual's flood risk perception. With a focus on the consequences of climate change worsening, making households choose to migrate to safer areas. Finally, De Koning & Filatova (2019) highlight people’s prior experiences with flooding and how these experiences affect their risk assessment. They focused on previous flood experiences and their local knowledge, which allows them to assess a home’s risk more accurately.

None so far have focused on the impacts of a water label and how this might affect a household’s choices, while also accounting for the effects of sustainable mortgages on those choices. This thesis aims to fill this gap by exploring the relationship between credit availability, mortgage risk, flood risk, and the flood label. This research also aims to fill the gap between the previously discussed ABM models and climate-related financial risk in mortgage lending, as current research primarily focuses on understanding whether mortgage lenders are adopting these new credit management practices (Blickle et al., 2024a; 2025; Fontana et al., 2025; Keenan & Bradt, 2020).

This research aims to take a multi actor approach, as due to the complexity of this issue, no single problem owner can be identified. Instead, the system will get looked at from the perspective of the different problem owners, these have been identified as the Nederlandsche Bank and European Central bank, which both supervise the national banks in the Netherlands through the single supervisory mechanism (European Comission, n.d.), their primary goal would be to understand how the introduction of a water label will interact within the current regulatory system and possibly how these could be adapted. The ministry of Infrastructure and Water is responsible for designing the water label, and how its design might affect the outcome for mortgage lenders.

Furthermore, this thesis aims to bring different fields together, offering a novel perspective on how flood risk can be interpreted, not only in conjunction with the water label but also through the lens of interest rates. Combining the views of different fields, such as economic modelling and behavioural economics, can reveal new relationships between these ways of thinking and how they might interact. To achieve this, a previously used model in behavioural economics research will be expanded upon, drawing on practices and results from economic modelling. This model, previously used in Koning & Filatova (2019), Mutlu et al. (2024) and Mutlu & Filatova (2026), will serve as a basis for exploring new relations by introducing credit constraints and adding sustainable mortgages practices. This can not only help test the effectiveness of a water label but also highlight potential difficulties in adopting these policies.

This thesis fits into the EPA curriculum by tackling a wicked problem (or a Grand Challenge), which can be defined as a problem involving multiple actors with competing interests, with uncertainty within the system (bounds), there being considerable consequences on the actors when decisions are being made (Pryshlakivsky & Searcy, 2012). This thesis fits perfectly within these criteria, given the significant consequences if no action is taken, with competing interests of not only households, banks, and central banks, but also many public institutions, such as the Ministry of Finance and municipalities, all affected by decisions made within the system. Moreover, the consequences of climate change are somewhat uncertain, creating a large spread of possible scenarios. But so are the system boundaries, as many other decisions taken outside this system, such as flood defence measures or insurance companies, also have large effects on it. On top of this, this thesis builds on many different courses within the EPA curriculum, with the main course being “Macro Economics for Policy Analysis” for its economic thinking, and “Advanced Simulation” for its ABM modelling practices. These practices have helped build the simulation model used in this thesis. It therefore combines systems thinking with simulation techniques from the EPA curriculum.

As AI was used during the thesis, its usage has been disclosed in appendix G considering the societal discussion about its correct usage. In summary, during coding these tools have been extensively to reduce certain repetitive tasks or to quickly iterate upon ideas. For text creation, none of it was used, with Large Language Models only being used to on occasion summarize papers and to reduce spelling and grammatical errors.

2 Literature Analysis

In this section the general concepts that will be used in the thesis will be explored. These mechanisms will create a system overview that will be used to explore these topics in more detail during their respective sections. This section will mainly explore the actor behaviour within the system, set system boundaries, explore the interactions between the actors in the system while also showing some of the basic conclusions from the various sources explored. The conclusion of which gives us a range of topics that will need to be explored that will deepen the complexity of the model. This section builds on the introduction that discussed the actors, while also providing the problem statement. Each subsection will explore the actors, their behaviour and their interactions in the system, with the households, Ministry of Infrastructure and Water, Insurance companies and banks being discussed in that order.

2.1 The Increasing Effects of Climate Change on the Housing Market

The effects of flooding on home prices depend on subjective risk perceptions. Therefore, there is a strong component of learning the risk through experience rather than calculated risk. This means that housing prices decrease only when floods occur, due to damage and the increased perception of flood risk in the area. This is despite there being no improvements to the flood adaptation measures within the area. Instead, home prices increase to levels similar to those before (Atreya et al., 2013; Bin & Landry, 2012). Unprotected areas are often more expensive due to the water-side view. However, information campaigns have been successful in reducing this price premium (Niu et al., 2025).

The risk of flooding, therefore, does not seem to be adequately internalised in home prices. However, this was not the case for every area. Take, for example, the flooding in Limburg during 93-95. After 9 to 12 years, housing prices had returned to normal levels, coinciding with the introduction of nature-based flood prevention measures. However, this change was not observed with conventional flood prevention measures, such as dykes. This suggests that green areas are perceived as increasing prices, despite their similar effects on flood prevention (Mutlu et al., 2023).

Conversely, local buyers tend to bid lower than non-local buyers in flood-prone areas, suggesting that local buyers are more aware of the risks associated with the area. At the same time, local buyers will take flood risk information campaigns far more into consideration. Similar information campaigns, however, do not work in dike-protected areas without recent flooding. Experience seems to be a major factor in the success of information campaigns. This information asymmetry causes local buyers to bid lower in flood-prone risk areas, as they experience flooding, while those outside of the area do not internalise these risks (Niu et al., 2025).

Thus, it can be concluded that home prices do not properly internalize the area's flood risk, with only the experience of a flood having a significant impact on home prices. At the same time, information campaigns only work on people in high-risk areas or those who have experienced flooding.

2.2 The Water Label

One trial has already been completed to test the flood label in the US, as explored in Fairweather et al. (2024). With over 17.5 million users participating in the year-long trial, they introduce a flood risk label to the platform Redfin, an online brokerage responsible for around 20% of the US internet property market, which is around 8% of the total market. Using FSF's flood maps they were able to construct a label at the individual property level, which consisted of two different numbers, a number ranging from 1 to 10 representing the risk and a percentage chance such a flood was likely to occur every 30 years.

Users are split into two groups, a treated and a control group, thus giving them differing experiences, making it an excellent way to understand the behavioural changes when being confronted with a flood label. Due to the comprehensive data collection from the app, they are able to go even deeper and construct a thorough understanding of a users' micro decisions, allowing them to observe even relatively minor changes in distinct groups. They show how a user starts to learn from their previous experiences and develop strategies in using the label.

While this does not say anything about the mechanisms that lead to this behavioural change, i.e. insurance costs and potential damages, most of the description is instead focused on other potential reasonings for their changes. For example, treated users that were browsing high risk homes instead changed their tactics, resulting in them bidding on homes with an average of 57% reduction in flood risk scores compared to normal users. This result was achieved through a combination of changes in search area, touring decisions, and changes in bidding strategy, all of which resulted in this reduction. At the same time, they reduced their overall search time, highlighting a decrease in the need for information gathering.

Overall, the trial has coincided with a reduction of prices for high-risk properties of $7.000, mainly caused by reduced competition for these riskier homes. Moreover, certain homes that were previously left out of older maps, either due to their risk increasing caused by climate change or better modelling, have experienced the largest decreases, showing how prices can quickly shift when information is readily available. A surprising finding was that there no party voter bias was found, showing no significant differences in behaviour, highlighting the need for easily accessible climate risk data not being bounded to greener parties.

The expectation that introducing a label would decrease housing prices makes sense, as the introduction of this information will lead to buyers pricing this into their bids causing the prices to eventually decrease. However, banks have noted that high demand might actually negate this, leading to a slower pricing integration. As demand increases the competition for houses, buyers become increasingly desperate to secure a house, leading to difficulties in ensuring a discount for the risky homes (Phlippen et al., 2023). This thesis wishes to further investigate this claim and understand the role in banking on this process and how this can be resolved. But first, the relevance of the water label for banks will need to be established.

2.3 The water label and its effects in general

Now, the water label is a lot more of a general tool than what will be discussed in this thesis, as such, a small part should be dedicated to setting proper system boundaries. The water label is first and foremost a tool that is used to inform households of their flood risk, however, the exact consequences of this are far more widespread than what will be discussed in this thesis, so some boundaries should be set beforehand about what part of the water label will be discussed. In the Netherlands and in Europe such large scale trials as seen in Fairweather et al. (2024) have yet to take place, which, in the Netherlands, can be attributed to the difficulties in constructing the label due to the low quality data, making the construction of the label difficult (Hoogvliet et al., 2023; van Ginkel et al., 2024). Although some have been created by private organisations, municipalities, or a combination of both, many, however, show clear limitations. With these labels already having been discussed extensively in Hoogvliet et al. (2023), their results as deemed relevant will only be summarized. Moreover, as a lot of these labels fit more neatly into the climate label category, only the decisions for flood risk will be discussed to simplify the overall discussion, excluding issues such as foundation and heat, despite their strong interconnection with the overall label design.

Many different types of designs were created that highlight the diverse needs of each label. Each of these label focus on various aspects, usually caused by their goal. For example, the "Staat van je Straat" and "BlueLabel" are mostly building level information meant for municipalities to identify problem areas and to create policy solutions around them. This can lead to municipalities in communicating the risk towards their citizens and opening a dialogue to create solutions, leading to a range of different adaptations.

Homeowners may also use these labels to inform them during for example mortgage decisions and insurance (IkBenWaterProof, Mijn Waterrisicoprofiel, Calcasa Woningcheck) and creating awareness around adaptation choices on water management (Waterlabel.net, IkBenWaterProof). All of them serve the same purpose in helping improve communication between the municipalities and citizens by creating easy to access information, aiding in policy design. The main difference is how they are informed, while with the former this was achieved top-down, letting the label aid in discussions as the municipality sees fit, while in the latter the same goal is achieved bottom-up by allowing citizens to find and use this information as they see fit.

The final group, property investors, the focus is more on creating a specialized package of adaptations through frameworks that help identify possible problem areas for a building (KIM-Tool and DGBC Framework). These labels therefore seem a lot more integrated in their design, combining many different elements from their environment as well as building characteristics to create a plan to reduce their climate risk. While in some ways, this can lead to a similar result as a water label, the information is clearly more for experts.

Overall, three distinct types of labels can be identified, which is caused by the groups of adaptations. These can be grouped as:

1. **Private buffer adaptations**

They usually offer little impact in reducing damages but help increase the spread of stormwater runoff decreasing the inundation depth by changing building characteristics and environmental characteristics on private property. These can include increasing greenery, reducing tile usage, rain barrels and green roofs.

Examples: Waterlabel.net, IkBenWaterProof, KIM-tool, DGBC framework, Staat van je Straat, BlueLabel

2. **Private adaptations**

Reduces the damages caused by floods by changing building characteristics. These can include raising entry height, concrete floors,

Examples: IkBenWaterProof, KIM-tool, DGBC framework, Staat van je Straat, BlueLabel

3. **Public adaptations**

The focus is to reduces the inundation depth, for example through improved sewage systems or dyke improvements. But it can also include for example buffer adaptations, through increased greenery or reduced asphalt. These are more characterized by their public ownership and are thus bound by different rules.

Examples: Staat van je Straat, BlueLabel

Some labels (Mijn Waterrisicoprofiel, Calcasa Woningcheck) are missing adaptations, with them being characterized as simply an indication rather than a full inspection due to their focus being more on investment decisions (i.e. mortgages). These are also usually far cheaper, using publicly available data while also being created by private companies, the others are usually created by a collaboration between municipalities, private companies and other public institutions. It is clear that these interests create significant differences in their design goals. With the latter usually being more aligned with creating a larger public discourse, aiding in the creation of innovative solutions and finding the problem areas.

Finally, each of these labels also assigns a score, giving indications of their design goals. These labels once again differ significantly due to their goals, with complexity increasing if the label's target audience is more knowledgeable, therefore those for experts will be left out from now on, with households becoming the only party of interest. One label makes use of an ordinal scale, using descriptions for their labels of how much water nuisance they experience, such as pools of water near their home. While others include some more measurable type of scale when severity increases (BlueLabel), clearly being designed for a more communicative label. All other labels only use a ratio scale, with most of them differing in which data they present. Most of them employ a standard amount of rain, ranging from the amount and duration of rain, with some following EU standards.

Concluding remarks

The water label sits in a complex field of different actor interactions, most of which will not be discussed in this thesis. Setting proper boundaries is therefore particularly important for later. With mortgage lenders mainly being interested in the home itself, the damages that will be faced due to the tragedy of the common, i.e. those labels that focus on the buffer adaptations and public adaptations, despite their substantial impact will not be further discussed in too much detail. This was chosen due to this thesis' focus on buyer-bank interactions, which are more relevant to a bank's decisions. This can be ascertained from banks only mentioning private adaptations (Phlippen et al., 2024). In this relationship, banks are more interested in private adaptations, due to the effect of damages directly increasing the credit risk (Caloia et al., 2022; Holtermans et al., 2024; Thomson et al., 2023; Kousky et al., 2020).

Another important detail also jumps out, with most of these labels meant for mortgages being simple and quick, focusing on being easy to understand and widely available, rather than accuracy, creating low search and information costs. Mortgage lenders also question whether this will happen explicitly, by informing the buyer of the physical climate risks, or rather implicitly "because financial institutions take this risks into account during their assessment of funding applications" (Hoogvliet et al., 2023).

2.4 The water label and banking

As climate change worsens, the expectation is that these flood risks will only become worse, making it even more challenging to live in these areas. To solve these issues, two distinct effects will need to occur. Firstly, buyers and homeowners will need to be informed of their (increasing) flood risk so they can decide how best to address it. This issue can therefore be addressed through the water label, which informs them of the risks they are taking when purchasing a home. Secondly, people need to be able to afford adaptations that can reduce the damage caused by climate change. Here, sustainable mortgages comes into play. As flood-prone houses carry a higher credit risk, the interest rate should rise to reflect this (Kousky et al., 2020). At the same time, rates in safer areas will decrease, as the credit risk caused by floods will properly get assessed by banks. This price difference then creates space for new financial solutions, allowing refinancing to pay for these adaptations and decreasing risk (Phlippen et al., 2024). Very few banks in Europe have already undertaken measures in response to rising interest rates in flood-prone areas, leading to higher premiums (Fontana et al., 2025).

Thomson et al. (2023) explores how the financial situation of households change after a flood event. Two distinct effects occur that affect the home's value, first the home decreases in value due to the knowledge of flood risk now being spread, at the same time many homes are damaged, requiring prohibitive costs to be repaired, sometimes leading to damages homes being sold. Afterwards they can choose to sell their homes, despite some of these homes not even having been repaired, due to them being unable to pay for these damages. Others might be able to pay for their damages, but either rely on insurance or their own savings. This is simply why a water label is so important, as it allows them to prepare, leading to not only reduced damages (Endendijk et al., 2023), but also reducing the sudden flood discount, meaning fewer financial issues afterwards. The following few sections will focus on expanding on Thomson et al. (2023) and how this framework relates to the Netherlands, starting with insurance.

2.5 The poor coverage of flood insurance

Insurance can be used to inform households of a home's flood risk, causing buyers to potentially choose a safer home due to the increase in monthly spending (Dubbelboer et al., 2017). In the US, these have helped improve pricing integration, but due to the differences in the Dutch insurance system, such a price shift is not observed here, as it is impossible to take out an insurance for larger floods, leading to homeowners needing to pay out their own damages (Phlippen et al., 2024). Despite this, many still believe they are insured against these events (Verbond van Verzekeraars, 2018), possibly leading them to be unaware of the potential risks they are taking on.

The reasons for this are quite complex due to spotty insurance policies depending on flood types, with some being insured while others are uninsured. Mainly floods caused by groundwater seepage and failure of primary water defences (from the sea and rivers) are uninsurable. While the "Wet Tegemoetkoming Schade bij rampen" (WTS) does offer partial compensation, it only applies to when a primary flood defence along a river fails (Wettenbank, 2021). While other types of floods, such as by heavy rainfall and secondary flood defence failure are insurable (Bani et al., 2024). Although it is likely that these also will become more difficult to insure as damages by heavy rainfall is predicted to increase due to climate change (Verbond van verzekeraars, 2015).

Because of this complex and obtuse web of insurance types, many homeowners are unaware of the state of their insurance, leading to this type of information not reaching buyers (Phlippen et al., 2024). Now, to bring this back to banking, if a flood were to occur and the homeowner was uninsured, they would need to pay for the damages, creating sudden extra costs. At the same time, their home's value would decrease, making it more difficult to sell their home and recoup their losses. Both of these will lead to a higher level of credit risk, increasing the interest rate (Caloia et al., 2022; Holtermans et al., 2024; Thomson et al., 2023; Kousky et al., 2020). But, as was stated before, this effect is rather limited in Europe (Fontana et al., 2024). However, before this topic is further explored, the basics of banking should be explained.

2.6 The basics of banking

Taking out a loan in its basics is a process of information gathering, wherein the bank will need to assess the creditworthiness of a loan, basing the interest rate on its risk. To set these values, the bank will need to assess the risk of the asset, in this case, the house, the customer and their own certainty of their risk assessment. Then by estimating the chance the loan would default, meaning the chance that the lender will be unable to pay for their loan, times the percentage that will be lost, also called the loss given default (LGD), the interest rate can be calculated (Campbell & Cocco, 2015). The final point is especially important, given that when a loan is originated, banks with the best understanding of the risk can offer lower rates, giving them a competitive edge in the market. Higher information quality enables banks to screen out riskier loans more effectively, leading to better outcomes for both parties: the customer pays less, and the bank holds less risk (Blickle et al., 2024b). To bring this conversation back to floods, given this competition requiring the banks to gather as much data as possible to reduce their risks, the eventual outcome would logically be for banks to take flood risk into account, as this will lead towards a lower risk loan (Holtermans et al., 2024; Thomson et al., 2023; Kousky et al., 2020). However, this has yet to be observed in the Netherlands (Fontana et al., 2024), with data quality being seen as the main issue (De Nederlandsche Bank, 2021).

The consequences of the low data quality led to an overvaluation of homes with unmapped flood risk, while safe homes get undervalued. In other terms, the real value is never realised due to the missing information, while this unrealised flood discount is only expected to increase as climate change worsens. This will increase the credit risk as the real price of the asset is lower than the perceived value (Phlippen et al., 2023), with this credit risk only increasing after flood events due to the sudden increase in monthly expenses and loss in value, which increases the loan-to-value (LTV) ratio (Caloia et al., 2022).

Before explaining this relation, the basics of the LTV-ratio should be explained, as even during normal lending decisions, the LTV ratio is incredibly important due to the "Dual trigger" effect, significantly raising the loan risk. Firstly, the home is seen as collateral, meaning it will be sold in case of a default to recoup the costs, some of which will be lost due to costs related to the sales process. Since the borrower has low or negative house equity (ergo the high LTV-ratio), the lender is recouping less of their loan. At the same time, if the borrower already has low equity overall, in other words meaning they have lower savings, they are more likely to have difficulty paying their loan in times of financial issues (due to for example a drop in income), as they cannot rely on their savings to make up their loan repayment (Campbell & Cocco, 2015).

Flood events add another mechanic to this, after a flood were to occur, two events happen that decrease the house's (and owner's) equity. Firstly, the house experiences a loss in value due to the flood risk knowledge now becoming more widespread, causing the resale value to decrease. Secondly, the house experiences a decrease in value that is equal to the damages. In response, the homeowner can use their savings or their insurance payout to repair the damages, while some even choose to completely abandon their property. In turn, this puts them in a more difficult position due to the loss in home equity and savings, increasing the chance of defaulting. This also creates larger losses in the local areas, creating a systemic financial risk created by the flood event (Thomson et al., 2023; Caloia et al., 2022).

Repayment can also become more difficult due to the sudden drop in income a lot of people experience up until a year after flood events (Xiao, 2011), increasing the potential risk of defaulting. However, there is much difficulty in actually showing the strength of this relationship due to more integrative approaches currently being missing, caused by the difficulty in collecting the necessary data ex-post (Fontana et al., 2025; Kousky et al., 2020). Moreover, these relationships are heavily dependent on the country's institutional context, due to the insurance policies set up to alleviate some of these financial losses, creating major differences between each country (Tran & Wilson, 2024). Ex-ante research in the Netherlands suggests that around a third of the households affected by floods could experience financial issues, with most homeowners needing to pay around 12.000 to 65.000 euros out of pocket, and around a third of these are expected to face financial difficulties in paying for these damages (Deelen et al., 2025).

In response to these losses, the expectation would be that banks would instead increase their premiums in these areas, reducing their potential losses.

2.7 The Poor Integration of Flood Risk into Mortgage Lending

Through exploring the differences in behaviour between the banks in the Europe Union (EU) and the United States (US), some context will be provided for the validity of whether flood maps will improve the credit assessment process of banks in the Netherlands. In short, the US is further in this process, with flood maps being an integral part of many banks' credit assessment processes (Bickle et al., 2024a), while in Europe these lags behind (Fontana et al., 2025). Furthermore, the moral hazard this information causes will also be explored.

First, securitization should be explained to make sense of the later discussions. Simply put, mortgages can be packaged together and sold by banks towards other groups; this "securitization" creates packages of mortgages that are commonly referred to as Mortgage-Backed Securities (MBS). These MBSs will then be used to pass the interest and principal payments through to the new owner, making them now hold onto the risk of these loans. Now, because of the grouping of multiple loans, these are commonly seen as less risky. However, as was evident in the 2008 mortgage crisis, securitization also creates problems caused by the inherent information asymmetry within this practice, leading to the potential moral hazard wherein banks can make use of this information asymmetry to push risky loans towards an unknowing secondary bonds market. This problem is created due to the simple rules set up for securitization, making it possible to push in riskier loans as long as they are allowed by these rules (Ashcraft et al., 2010). This same tactic can be used in flood-mapped areas, due to Fanny Mae's disaster modifications. It can be argued that due to these rules, it disincentivizes banks to account for flood risk, creating the need for banks to account for their liquidity issues during such events (Kousky et al., 2020; Ouazed et al., 2019). The next discussion will use securitization to understand the credit assessment practices within banks and how they use flood risk information in these processes.

2.7.1 The differences between the US and the EU

In the US, FEMA flood maps have been far more commonplace compared to similar maps in the Netherlands, even being used regularly in assessing the insurance premiums (Hoogvliet et al., 2023), while mortgage lenders made less use of them. Although, before this was commonplace, the overall institutional environment did need to support this. This can be seen in Ouazad et al. (2019) and Keenan & Bradt (2020), wherein US mortgage lenders reassess their loans after flood events, changing their expected risk given the added information. In turn, this can be used as proof that mortgage lenders also have a learning process, requiring them to reassess their previous loans after new information becomes available. On top of this, they increase their securitization due to the sudden need of increased liquidity, showing that banks were unaware of the higher risks from these loans during origination. Simply put, if the lenders were aware of the risk posed by these loans, they would not have a higher chance to securitize them, as there would be no incentive to do so.

Ouazad et al. (2019) proposes two different theories that might explain this behaviour. Firstly, mortgage lenders have more local knowledge than the secondary bonds market, thus they can exploit the difference in knowledge when such an information event were to occur. Fanny Mae and

Freddie Mac have rules set up for which a loan has to conform to; as long as these are met, they can securitize their mortgages, creating the incentive for mortgage lenders to reduce their overall risk by securitizing loans from flooded areas. This can be corroborated with the increases in conforming loan origination after flood events. Secondly, the explanation given is that neither party takes into account the flood risk during their risk assessment, meaning that they underestimate the loan risk. In turn, this creates the problem that if they do not securitize these mortgages, they will face liquidity problems. Thus, they will reassess their loan, while also choosing to securitize to improve their liquidity.

A much later study from Blickle et al. (2024a) explores both these dynamics. It also shows that behaviour from the previous study has now significantly changed, with mortgage lenders now normally using flood maps in their credit risk assessment practices. They show this by comparing the FEMA maps with other flood zone maps, meaning they can identify potential unmapped flood-zones. In these areas, they highlight how fewer loans get originated, while also setting higher premiums and lowering LTV ratios. However, when looking at local banks (and non-bank lenders), they still tend to securitize more than national banks, meaning that this practice is still used, just not as largely.

In the EU such a large shift in credit risk assessment practices has yet to take place, with only some changes being seen in the past few years as the regulatory pressure increases from the EU central bank. The expectation is that, with the increase in regulatory pressure, these differences will shift. However, major differences in behaviour between the banks have been observed, with some already pricing the risk properly, while most do not. To fix this, the single supervisory mechanism (SSM) is responsible for regulating banks (Fontana et al., 2025). With the ECB taking largest banks, the Nederlandsche Bank is responsible for the smaller sized banks (European Comission, n.d.). According to the Nederlandsche Bank, Dutch mortgage lenders have cited issues in transitioning towards more climate-conscious banking due to poor data quality, caused by limited availability of data on climate risks and the home's adaptations, making it difficult for banks to accurately assess credit risk (De Nederlandsche Bank, 2021). In short, this seems to be confirmed by the previous discussion on American banks.

Overall, it can be concluded that while some banks do integrate climate risk measures, these changes tend to be relatively small. However, due to regulatory pressure from the ECB, this is likely to shift as European mortgage rates increase in areas facing climate risks. Moreover, there is some evidence that some banks are aware of climate risk in the US, but with the data quality of these flood maps being inconsistent, some banks might choose to securitise these mortgages, pushing the risk into the unknowing secondary bond market. The most important conclusion is that flood map availability does lead to improved credit assessment, showing the need for publicly available and accurate flood maps.

2.8 The data quality of flood maps

As previously stated, the data quality for flood prediction, according to the banks, is currently quite poor, resulting in difficulty extracting the credit risk related to flood risk of an individual home (De Nederlandsche Bank, 2021). These can be attributed to a multitude of different stacking issues, which will be discussed within this section. Overall, they can be summarized in two broad categories: the local environment surrounding the home and the building characteristics. A more in-depth overview of all the complications faced in gathering the necessary data for the water label can be found in Van Ginkel et al. (2024) and Hoogvliet et al. (2023). The next section summarises both these reports to gain an accurate problem demarcation.

2.8.1 The environmental characteristics

Flood maps in urban areas can be difficult to construct due to the many complexities that are necessary to accurately predict the risk of an individual home. At the moment, these maps have far too low a resolution, between twenty-five meters and one hundred meters. This aggregation can lead to situations where elevated homes have the same risk as lower homes, despite these making significant differences in actual risk. This can become an even larger issue once the map grid stops aligning with the street grid, causing homes on dykes and on polders to be grouped together (van Ginkel et al., 2024).

Moreover, these maps are constructed by the provinces and use many different methods to construct them. This leads to varying quality and little verifiability, as they require an accurate assessment of the quality of modelling of the sewage system, as well as interaction effects with the country and city-wide water systems. The accuracy of this is important because of the difficulty in predicting water flows beforehand (i.e. floods due to high groundwater levels, sewage overflow, etc.), as multiple of these threats at the same time are usually the cause of floods.

Pluvial floods are even more difficult to predict. The Netherlands is especially prone to these due to its flat lands, but these predictions are more reliant on knowing the maintenance as well as the layout of local sewage systems, conversely this data can quite often be unreliable. While these systems have been designed for more extreme scenarios, most have not been tested for them, making it difficult to predict their actual effectiveness. Moreover, the category of pluvial flood also contains floods from the sewage system itself, making the issue stacking.

If these issues are looked at from the fluvial flood types, the main concern is the accuracy of the predicted chance for the primary water defences, such as dykes. This difference can be between a factor of 10 and 100, making it difficult to predict the risk from floods caused by rivers.

Finally, some adaptations are used as sponges which decrease the speed at which the stormwater gets put into the sewage system. These can be incredibly minute details, such as tile usage or the amount of greenery (Hoogvliet et al., 2023).

2.8.2 The building characteristics

Minor changes in building characteristics can have large consequences for the flood risk. For example, a heightened doorstep will make a massive difference during a small pluvial flood or a sewage cover if the sewage system becomes overburdened (van Ginkel et al., 2024). These characteristics do become less important as inundation depth increases, wherein other, larger adaptations start to become more prominent (Hoogvliet et al., 2023). In these types of scenarios, more expensive investments, such as concrete floors, are required to reduce damage (Endendijk et al., 2022). Other adaptations can act as a buffer, slowly releasing rainwater, which in turn reduces loads on sewage systems. Overall, there are a lot of small building characteristics that dictate the number of damages suffered by a flood, but they are also strongly dependent on the type of flood, which is mostly decided by the environmental characteristics. In turn, there is a large necessity in choosing the right adaptations for each type of flood.

On the other hand, some adaptations can be used to spread out stormwater runoff, reducing the chance of sewage systems overflowing, highlighting the need for buildings and ground to also act as buffers (Hoogvliet et al., 2023). The complexity of interactions between flood type and the lack of data on these small details cause high uncertainty ranges in estimating the damages caused by pluvial floods (Hoogvliet et al., 2023; van Ginkel et al., 2024).

2.8.3 The difficulties of adopting the water label

These issues combined lead to a large margin of errors, making it difficult to create an individual home's risk profile. While many of these problems are solved by improving data quality, there is still a stronger link for the flood risk between a home's characteristics and its environment, leading to less power towards households. This creates more difficulty in creating the water label, as depending on the flood types, their uncertainty range and severity, other types of adaptations will be necessary, while also requiring the municipalities and other government organisations to step in in cases of poor water management (Hoogvliet et al., 2023).

These results can be corroborated with expert interviews from the Dutch real estate sector, who have raised similar problems for developing the water label, citing the need for clear adaptation measures, scale differences in "building-level and area-level risk" of floods, while also ensuring that the label stays transparent and of high standards (Oerlemans et al., 2025). To achieve this, data quality will need to improve across the board, to ensure that a clear adaptation path for households as well as municipalities can be created. While this thesis does stay out of this discussion as to how this should be done, these will need to be accounted for when designing the flood label for the model.

Conclusion

The ECB has already taken steps to resolve these issues by increasing regulatory pressure on banks to integrate climate change into their risk management practices, with a primary focus on raising mortgage rates in at-risk areas. The Single Supervisory Mechanism (SSM) has successfully improved the risk integration of banks under its oversight, although it does not include all banks within the European Union. More steps are being taken to ensure that banks change their credit risk management policies, such as introducing fines, as the ECB has increased its focus on reducing the effects of climate change. However, current data on its effectiveness is unknown, as mortgage data from 2023 onwards is still lacking, which coincides with the timeframe when the ECB began increasing pressure on banks (Fontana et al., 2025).

De Nederlandsche Bank shares these ambitions, wanting to help improve the adoption of climate conscious policies into banking, with them being responsible for the extension of sustainable mortgages in the Netherlands (De Nederlandsche Bank, n.d.), the focus of this thesis will be on understanding the overall impact of the water label and how accounting for the physical effects of climate change will change the overall housing market. With many of the current issues in the banking sector, as well as for households, being caused by the unavailability of (high-quality) data, the water label forms an essential part in reducing the flood risk within the Dutch housing sector. This will lead to not only the banks competing on a more level playing ground with each other, as the creation and improvement of such a label will improve the data quality, it also helps improve the relationship between the banks and their customer, easing the difficulty of communicating the flood risk of a home.

This thesis aims to highlight the urgency of adapting to these changes quickly through the use of an ABM model that incorporates the flood risk into mortgage lending practices, by building on the Risks & Hedonics Empirical Agent-based land market (RHEA) from Filatova (2015), which uses hedonic regression analysis to predict home prices in the American housing market. This model was later adapted for the Dutch housing market by Mutlu et al. (2023), which also incorporates households' flood risk perceptions into the framework. This thesis will build on these two works by incorporating credit risk into the model. Its novelty mainly comes from the combination of these two fields, for which, as of yet, no analysis has been done of the possible impacts of a water label combined with the effects of the expansion of sustainable mortgages, despite the strong connection between the two ideas. This thesis aims to fill this gap by further expanding the RHEA model with credit risk assessment, which can help predict the potential impact of the water label on top of sustainable mortgages.

3 Research Questions

With no research currently having been conducted on understanding the effects of increasing mortgage rates for flood prone homes in relation to the water label. Despite the effects flood risk has on mortgages (Thomson et al., 2023; Kousky et al., 2020), very few banks in Europe have responded by increasing interest rates in Europe (Fontana et al., 2025). At the same time, home buyers are generally unaware of the increasing risks caused by climate change (Atreya et al., 2013; Bin & Landry, 2012), leading to further credit risks increases (Thomson et al., 2023). Banks can instead integrate flood risk into their portfolios, ensuring they are hedged against the future risk of the physical effects of climate change.

With flood risk labels shown to be effective (Fairweather et al., 2024), a price decrease is expected when the water label is introduced, creating risks and opportunities for banks that they will need to be aware of (Phlippen et al., 2023). This thesis wishes to explore both these policies through a modelling exercise, by understanding their mechanisms, new interaction effects could be discovered that previous literature has not yet discussed.

A significant part of the research questions is therefore focused on how the interaction between these policies. This leads us to the main research question:

Main research question
What effects might the introduction of the water label and sustainable mortgages have on the integration of flood risk into housing prices?

For this, multiple sub-questions will need to be answered. These questions will highlight how these two policies might be designed and how they might interact to integrate the flood risk.

Sub Question 1
What type of flood label characteristics would be most relevant for mortgage lenders and how might this impact its design?

Because climate change is highly region-specific, a rich and complex discussion is currently happening in each country about how to communicate the risks of climate change to homeowners properly. This discussion is equally as complicated in the Netherlands. Initially, this question should be explored, relating it to the needs of mortgage lenders as well as households. This question was already partially answered in the sections 2.2, 2.3 and 2.4. The latter part, about how this might impact design, will be required to construct the water label for the ABM model. More specifically, which requirement would need to be met for the water label to adequately function and thus meet the needs for these actors. Furthermore, its design would also need to be related to the current understanding of biased flood risk perception.

Sub Question 2

To what extent do Dutch and European banks currently integrate flood risk into mortgage pricing and how does this relationship function?

As previously explored, very few banks currently integrate climate risk into their credit assessment strategies in Europe. As such, two different aspects will need to be explored. First, do mortgage lenders raise interest rates significantly in these areas? This have already been answered in section 2.7, with very few mortgage lenders changing their rates. The second part of this question relates to understanding the causal relationship between higher credit risk and flood risk. These have also been answered previously in section 2.6.

Sub Question 3

How do households calculate their utility when buying a home?

Here, multiple questions must be answered. First, an exploration of housing market models needs to be conducted, understanding which aspects of the regulatory system and mortgages they find most important. This will then result in a discussion on the current mortgage regulations. Then, the discussion will focus on how these regulations will affect household budgeting. These can help answer how households set their budgets and how these will then impact the model.

Sub Question 4

How can flood-risk-sensitive interest rate be represented in an agent-based model of the Dutch housing market?

This sub question combines the previous three, taking the causal links discussed in sub questions 1 and 2 and combining it with the understanding of utility from sub question 3. Through combining these topics, an understanding will be gained about the more intricate causal links, creating a framework for the ABM model to function in.

Sub Question 5

How could the interaction between a water label, flood-risk mortgage pricing and credit restrictions affect housing prices and LTV ratios?

This question will consist of two parts. Initially, an exploration of the relevant housing market and ABM literature will be conducted. Highlighting how interest rates are typically incorporated into ABM models and how different ideologies view their impact on the housing market. This then leads us to the second part of this question: how these are integrated into the model. Through this integration, the effects can be explored throughout the system. Thus this question has two parts, part theorizing how the system might be affected by these policies, while the second part relates more to the results of the model and understanding how the interactions in the model create the results. Through answering this final sub-question, the main research question can finally be answered.

4 Methodology

To explore the relationship between the water label and banking on the effects of housing prices, an ABM model will be developed based on the sociotechnical system described in the introduction. Using this model, the previously described research questions can be answered; however, before this can be done, a method should be used to aid in the development of the model. Ideally, a method that allows the complex interactions among actors, such as banks, buyers, sellers, and real estate, to be modelled accurately. For this, Nikolic & Ghorbani (2011) best fits this description.

The methodological framework described in Nikolic & Ghorbani (2011) focuses on modelling Large-scale Socio-Technical Systems, which are interlocking, complex groups of actors centred around a technical system, such as a supply chain. In this system, different actors respond to one another and adjust their behaviour, accordingly, enabling them to adapt to changes in the system dynamically. This allows complex behaviour to emerge naturally from a set of simple components, which can then be used to explain overall system-level behaviour. Most importantly, it allows behaviour to emerge, enabling agents to respond to their changing environment, which can then be used to explain the system's end state.

This perfectly describes the end goal of the ABM model: relatively simple banking behaviour that allows them to set interest rates based on the property's risk, the buyer's risk, and the government's lending rules. Buyers have to respond to these changing bank demands, evaluating which property best fits their needs, while also considering mortgages and potential flood risk. Moreover, both agents will require multiple combinations of possible behaviours, allowing exploration of complex system interactions. Each one of these elements is simple on its own, but together they form a complex web of interactions, allowing for the overall system to be analysed.

Overall, this framework has five components, with each component not necessarily following the previous one. This is due to the process being more iterative: during each step, mistakes are found that require a previous step to be reconsidered, and each step is reconsidered multiple times during the building of this model. The five steps are:

1. *System analysis*
2. *Model design*
3. *Detailed design*
4. *Software implementation*
5. *Model evaluation*

Typically, the forming of this complex web of interactions is an iterative process, but for this model, it was not entirely the case, as this work builds on a previous model by Filatova (2015) and Mutlu et al. (2023). As such, some changes have been made to the framework proposed by Nikolic & Ghorbani (2011), with the focus initially shifting to system analysis and model evaluation to understand better the modelling choices made previously. Moreover, the focus of this thesis will be on the new interactions added to the model, not necessarily on the previous interactions, as these steps have already been completed during the building of their previous iterations. Any additional changes will be discussed in their respective sections.

4.1 System analysis

Before constructing an ABM, the desired system should be explored to identify the interactions among actors. But before this can be done, the relevant actors will need to be identified, of which the problem owners are most important, with them being able to make changes to the system. For this thesis the problem owners decided are the Nederlandsche Bank, the European Central Bank and the Ministry of Infrastructure and Water. After this, multiple questions will need to be answered:

- “What is the problem being addressed? What is the emergent pattern that is of interest to us?
- Why is this a problem? Is this pattern an existing or a desired one?
- Whose problem are we addressing?
- What is the intended model use? Before a working model is ready, it should be decided how the model will be used. Are the stakeholders interpreting the model outcomes themselves or are the modelers involved in the usage? To what end are the model outcomes going to be used (e.g. are we making policy, are we analysing existing policy or are we forming a hypothesis about the functioning of the system?) (Nikolic & Ghorbani, 2011)”

These questions have been answered during the Literature Analysis and Research Questions.

Next follows the System identification, which will lead towards “a description of the environment in which the actors are performing in.” Instead of focusing on the actors responsible for actions within the system, the focus will now be on the interactions between actors, as well as the rules and regulations that influence the system. This widens the focus, allowing a broad overview of the entire system. It should be noted that this was also the starting point after receiving the model, as the focus initially was on understanding all the interactions by listing them; only after that were the problem owners identified. This is due to the model being an expansion of previous models, which require in-depth knowledge of the entire system. Only after this was done was a list of all interactions expanded to include the new modules to be added to the model, all of which were ranked from least to most important.

After this, the findings were discussed with experts in the field, which helped determine which modules were most important and the scope of all actors and their interactions. If, in any way, they determined that previous decisions were wrong, these would be retroactively changed to ensure the system's design was accurate according to their specifications. Based on these specifications, the literature study was further expanded to ensure the system overview was correct and consistent with the overall system overview. These experts were from two different fields, with one mainly focused on the financial systems while the other focused on the effects of flooding on the housing market in relation to people’s perception of flood risk.

Once this was finished, the next step was system conceptualisation, in which the entire system was summarised as a causal diagram, providing a simple overview of what would be modelled. In this case, some parts were left out, with a focus on expanding the model and simplifying the parts that were already provided. Moreover, multiple behavioural diagrams of the different actors, mainly buyers and bankers, were created, each including possible steps for building the model. This was to ensure that each iterative step was feasible and could be communicated to domain experts, enabling better communication. Nikolic & Ghorbani (2011) mention Ontology Formalisation, as different fields can use different words for similar concepts, which was not viewed as an issue during this thesis. Instead, the main problem was that similar concepts were placed at various levels of aggregation, as these modelling choices were logical for that specific model but would not suffice if added to this model. A sizeable portion of this thesis was therefore put into deciding the proper aggregation level and how to translate this to an accurate model.

4.2 Model design

With the system analysis complete, the next step is to translate this into a functional model, starting with the agents' structure and classification. This was a very iterative process, involving much trial and error to determine which combination of classes, agent interactions, and variables would be ideal. To start this process, the environment and its components were identified, with the environment defined as the processes that evolve incredibly slowly. After this, the actors' behaviours were identified. As this required expanding a previous model, the process began by listing all the steps the agents already in the model would take. Afterwards, this list was expanded with the new behaviours identified in the previous step, allowing a broad breakdown of what would happen in the model and its potential effects. Moreover, breaking points were identified for issues that could later arise in the model, such as potential data availability issues, avenues for policy change, and how actions identified as environmental factors can alter the situation, ensuring that this was adequately taken into account. Section 5: Institutional context and detailed agent behaviour will explore this step in most detail, discussing in broad details the decisions behind the modelling decisions.

Desk Research for Model Design

To gain this understanding, desk research will be conducted that combines the institutional context of the Netherlands with the understanding of the previous literature analysis. It will go deeper on the most relevant research for this thesis, discussing some of their limitations and differences in the institutional context, as most of the research about flood risk and residential mortgages was conducted outside of the Netherlands, rather, most originate from the US. When available, however, Dutch research and data was preferably used. These differences resulted in significant limitations of the model, partially due to the limited data available in the Netherlands, but also what was available for this thesis.

Moreover, this section will discuss multiple research fields that were relevant to the ABM model design, and how these then combine within the Dutch context. Most importantly, discussions within these fields will need to be identified, such as mainly the standard housing market models, and how these differences might change outcomes between these other research fields, i.e. the behavioural economics field (Dubois, 2026).

4.3 Detailed Model Design

With this done, the next step was the detailed model design, in which the previous steps are ensured to be turned into computer code. This logical model formalization was mainly achieved by looking at similar models that achieve other goals. These inspirations were then used to derive the necessary components; moreover, the previously imposed limitations were combined with these models to ensure that, while still following the logic of these models, substantial limitations needed to be imposed on the scale of the model so that the result would be feasible, while still being accurate enough to ensure good results. However, as this was relatively novel, not every necessary component could be found for the model. Instead, careful literature analysis was conducted to find ways around this, which sometimes resulted in significant compromises. Other times, the previous experts were consulted to discuss how these goals could be achieved, and to ensure that the system was translated correctly into the detailed design. Afterwards, this design was translated into a class diagram and a swim lane activity diagram to visualise the interactions between components.

Next, the experiments were designed, which were already partially completed, as it was built as an extension of Mutlu et al (2023). These mainly included variables related to selling, buying, biased perceptions of agents, and realtor behaviour. As the focus of this thesis was on the impacts of banking and the water label, these variables were held constant across all experiments, with the length of each run, 30 years, also maintained, as it is the median length of a Dutch mortgage. The number of runs was determined by simply running the model for longer than expected to be required, after which the model's convergence was assessed. However, after 150 runs, the model-generated data exceeded the capacity of the Pandas DataFrame. As such, the runs were limited to 150, and lower numbers led to slower convergence. Afterwards, the scenarios were designed, which required setting the parameters based on the expected behaviour of the actors. With these parameters, two different scenario matrices were set up to ensure that all parameter combinations were met.

With the detailed model design being a large part of this thesis, it was required to split up the section in two parts. The main section, six detailed model design, is in the main text. This will be the main through line of the functioning of the model, while avoiding some of the more intricate model details. In this section a broad overview of the base model is given, with some explanation of the new model. Most of the focus will be on the formalization of the model. The supplemental appendix F will instead discuss the rationale behind the formalization, going into detail of why certain limitations were chosen.

4.4 Software Implementation and Model Evaluation

For each new module added, it was first created outside the model to ensure it would produce the desired results. Afterwards, this module was added to the model, which now required extensive verification to ensure it produced the expected behaviour. This was done by first turning off all previous modules to simplify agent interactions. After this, each module was turned back on individually, with the parameters adjusted to ensure that all components continued to function as they should. If strange behaviour was noticed, some internal parameters were set to extreme values, significantly altering the results, which helped identify these bugs. After these steps were completed, the results were verified through scenario testing and expert review, as no real historical data could be used to confirm that the results would be correct. Instead, a visual inspection was deemed adequate.

5 Institutional context and detailed agent behaviour

In this section, Institutional context and detailed agent behaviour, the actor behaviour described in section 2: Literature Analysis will be used to create the formal rules that the model will use, creating an overall outline that the model should follow, while also adding more environmental details. This section thus is between the Model design and System analysis of the method given by Nikolic & Ghorbani (2011) as there was a strong intersection between both sections within this model given that the agents behaviour based on the environmental factors. This was chosen due to section 2 being a story that explains the problems within the system, while this section will explain more of the rules, thus institutional context, deepening how these actors interact with this system. At the same time, some aspects will need to be simplified to fit into the model, while also still ensuring that it accurately portrays the real world. Through exploring different but similar models, these two discussions can be intertwined, showing a deep understanding of the entire system. Additionally, some new avenues for policies will be discussed, while not all have actually been worked out. The final complication for this is that the model builds on top of previous research that mostly simplifies the relationship between banking and households, requiring some concepts within the model to be rethought. While this section will only discuss some of the broad concepts, the actual modelling design choices will more be discussed in section 6 and its supplemental appendix F.

Moreover, during this step, the detailed model design was also taken into account, ensuring that at each new addition to the model's a functional product was created. This part has been described in appendix E: Model outline. Described in this appendix is the initial system outline being far larger than what has been modelled, due to it being necessary to start with a larger general design, so that this can then be decreased to fit the overall time budget.

Because of the requirements of this section, the overall layout based on the starting question, how do households set their budget, which then creates a storyline through how this question was answered, how these are related to different models this was inspired by and the modelling choices in Mutlu et al. (2023).

5.1 How do households set their budget?

Households prefer to pay a certain percentage of their income towards housing, with that percentage rising only as their income increases. While short-term shocks will change this, in the long term, these percentages seem to remain stable. These short-term shocks seem to be driven more by speculative or psychological factors, whereas, if the fundamentals of a housing market do not shift, these percentages tend to remain stable (De Vries & Boelhouwer, 2008). This relationship can be related to the maximization of mortgages (Van Der Drift et al., 2022; Madsen, 2012; McQuinn & O'Reilly, 2007) of which will be discussed later. However, the main point is that this would disallow any bidding strategy, as households would bid according to their maximum allowed mortgage.

Here, deposits come into play. Buyers can save up a certain amount of their income to put forward as a deposit. This has two effects. It increases the household's maximum bid but also reduces the financial burden of their mortgage. This bidding strategy allows households to, in a way, ignore the interest rates, as a part of their bid now is their own savings. This can, for example, be seen in how the Dutch housing market responded towards a significant interest rate shift, for which housing prices barely budged. This was due to rising incomes and price expectations that set people up for higher bidding (DNB, 2024). It should be noted, however, that this shift is expected to be temporary, as this was in response towards inflation. This shifts price expectations upward, thereby mitigating some of the effects of interest rate changes.

This deposit also causes another change. As people put down a deposit on their homes, their loan-to-value ratio decreases, which is advantageous to banks. This is because the home is used as security for the loan, so that if the loan defaults, the home can be sold for a similar value to pay back the loan. On top of this, during financial issues, the mortgage loan will need to be paid out of savings. Thus, if savings are low, the chance of defaulting is also increased. The loan-to-value (LTV) ratio is thus essential for setting interest rates (Campbell & Cocco, 2015).

The LTV ratio is commonly described through balance sheets, wherein households will need to ensure they have enough liquidity to pay for their mortgage in times of financial difficulties (Campbell & Cocco, 2015). Exactly recreating this would be rather difficult, due to needing to recreate financial shocks, and getting the data for accurate estimations of interest rates. Thus, from this point onwards, a simplification was already necessary to ensure that the model could be simplified. From data it can be seen that within the Netherlands, there is a preference for an 80% LTV ratio (CBS, 2024). This can logically be reasoned from to the previous section due to its lower credit risk, at the same time this 80% ratio is also argued by Campbell & Cocco (2015). Thus, by simplifying this and taking their results as a basis for the interest rate, the other relationships to interest rate can be in some ways estimated.

This is a significant shift from how it is commonly done in other models. However, setting this up correctly would be a significant undertaking. This assumption is therefore also a significant limitation. The consequences of this will most likely affect the lowest and highest income groups, as the former is more likely to rely on credit, while the latter is more likely to rely on deposits (De Vries & Boelhouwer, 2008; van der Drift et al., 2023). This is expected to decrease the actual buying power of higher-income households, while increasing it for the lower-income households. Moreover, this decision has also had significant consequences for the design of the model. The effects of which will be discussed during section 6, detailed model design, as well as its appendix F: detailed model design, which further deepens out the concepts explained its corresponding section.

Moreover, this clearly quite inaccurate due to higher income groups facing fewer financial difficulties due to higher savings, thus they would logically also face lower interest rates (Campbell & Cocco, 2015), even in flood zones (Bickle et al., 2024a). This was still seen as an adequate starting point, simplifying the start of the model and thus ensuring that everything stays simplified for each new addition to the model. This is also seen as the main limitation in the model and will form a vital discussion from this point onwards as each new addition slowly starts to complicate this. Now that deposits have been explained, on what basis mortgages are maximized should be explored in more detail.

5.2 Mortgage lending in the Netherlands

Mortgages are regulated by the Ministry of Finance, advised by the NIBUD, which sets the maximum monthly financing burden a mortgage is allowed to take on each year (Warnaar et al., 2024). These maximums are a percentage of a household's gross income and are set based on each household's current financial situation. These are calculated using the average household budget, which is set to maintain a basic standard of living. This consists of, for example, food, energy, and healthcare costs, for which each household will pay a standard percentage of its gross income, with this percentage decreasing as income increases. As these costs of living differ year over year, these regulations help households balance their budgets, with bouts of inflation reducing the maximum financing burden and thereby decreasing people's mortgage payments. These budgets are then used to calculate the percentage of gross income that can be spent on a mortgage, which sets the maximum monthly financing burden for each income range.

Finally, the Netherlands has a mortgage interest deduction, which allows a household to deduct part of their mortgage interest from their income tax, reducing the amount they have to pay on their mortgage from their net income. As this reduces the financial burden of a mortgage, the maximum financial burden also increases with higher rates.

Within these regulations some interesting details can be noted that may be of use for later discussions. Firstly, due to the decreases in monthly spending on energy, a higher energy label will result in a higher maximum mortgage as it increases disposable income. At the same time, the buyer can increase their mortgage if they want to improve their energy efficiency label, due to the same reasonings. These can be later used as inspiration for similar policies.

Another relation of use that should be noticed here is that, given the interest rate increases, the bid will decrease due to the set monthly spending allowed by these regulations being (mostly) kept stable. Thus, when interest rates increase in the flood prone areas, this relation is expected to be of significant use in estimating the actual decrease in bidding. This observation becomes vital in the next section, which will discuss the logic behind the utility. But before this, some issues surrounding the housing market should be discussed.

5.3 Calculating the Utility of a Home

The unresponsiveness of housing supply to demand is a significant issue, making it difficult to accurately predict a home's actual utility. This problem, also known as the financialization of housing, causes financial credit to act merely as an inflationary force on housing prices, rather than helping ensure they remain at a reasonable level. This issue has been attributed to the scarcity of ground and the behavioural changes it produces for those owning a home, for which a more extensive explanation can be found in Appendix A. In summary, new homes are increasingly challenging to build, due to the competing interests of homeowners wanting to keep their housing prices high, stopping new development, and inflating prices, while those without a home are not adequately taken into account. This then makes it more difficult for developers to increase housing supply, as they are unable to build more, leaving demand for housing unable to influence supply.

The result is housing supply not responding to demand, with more available credit merely causing prices to inflate rather than creating new houses, making the problem worse.

Then the question becomes: how does the utility of a home get calculated? For this, there are two competing visions. Firstly, the home can be viewed purely as a financial asset. Here, the value of a home is based on the difference between expected growth in value and the mortgage, which becomes more expensive as the interest rate increases. However, this faces a significant issue when put into the model as it is difficult to predict the expected growth in the at-risk areas. Secondly, it is unknown how a differing interest rate will affect home value; previous models have focused on region-based expected growth, such as through Ricardian theory (Wright, 2024). Integrating this into the model can be difficult, as flood risk can vary significantly between streets and building layouts (Van Ginkel et al., 2024; Hoogvliet et al., 2023). On top of this, as households are very much constrained by mortgage rates when setting their bids, an alternative approach can also be used. One calculates the household's monthly budget constraints, then differentiates the interest rate.

This assumption aligns better with the model's required specifications, as it assumes people bid based on the expected monthly utility rather than the full price of a home. This monthly payment, however, can quickly increase as housing is a vital necessity. Therefore, other models, such as those from Van Der Drift et al. (2022), Madsen (2012), and McQuinn & O'Reilly (2007), indicate that households are constrained by their maximum monthly lending capacity rather than by expected growth in value. While this theory does not disprove the issue of too much available credit, it does offer an alternative to the view that houses have become purely financial assets. Instead, people need homes and thus want to allocate a sizeable portion of their monthly budget to housing. On top of this, they are restricted by regulations in the Netherlands, as previously discussed, limiting the amount a household can reasonably pay monthly for a home. This also explains why homes have become more expensive: as people's monthly payments as a percentage of their income remain the same, their buying power is determined by the mortgage rate, which increases as the rate decreases (Francke et al., 2014).

Finally, this also helps explain why certain groups get displaced or are left out, for which the previous theory does not explain, as homes become too expensive to fit into their monthly budget.

As such, the decision was made to calculate the effect of the interest rate on utility using the monthly budget. Not only is it a more elegant solution, but it also provides a rationale for why households are displaced, forming a link between climate gentrification and the impacts of differing interest rates. Furthermore, it affects how policy is constructed, as levers such as increasing a household's monthly budget or its maximum lending capacity can affect the options a household considers. Something that is effectively being done through the expansion of sustainable mortgages. But to explain this, the interest rate and more banking basics will need to be explained, which will then be followed up by a deeper explanation of the utility function.

5.4 Setting the Interest Rate

Given that the balance sheets should be simplified, the most common solution it to use the simplified formula of rfree + default probability (PDEF) * Loss Given Default (LGD) to calculate the interest rate. This consists of three parts, with the rfree, it consists of two aspects: the bank's desired profit rate and the interest rates set by the central bank. This profit rate is expected to be relatively homogeneous across banks and, as such, will be mostly left out from now on, while the central bank's interest rate will change over time, thereby changing mortgage rates. The PDEF is the chance that a household will default on their loan, meaning they will declare themselves insolvent and unable to pay off their loan. In turn the loan will need to be paid off by selling off the asset, in this case the home to recoup the costs, of which some will be lost. This last concept is signified by the LGD and as explain previously are strongly interlinked with the LTV ratio, given that the chance of defaulting will increase with the LTV ratio, while also increasing the LGD due to the associated costs (Campbell & Cocco, 2015). Thus, this formula simplifies this relationship further, helping estimate the interest rate.

Other factors also play a significant role in the interest rate. These are commonly used in many banking and housing market models, which include loan type, borrower characteristics (e.g., income and age), and mortgage affordability measures. All of these models have integrated mortgage data in some way. (Baek et al., 2021; Campbell & Cocco, 2015; Lin & Tsai, 2021; Mérő et al., 2023; Van Der Drift et al., 2022, Madsen, 2012; McQuinn & O'Reilly, 2007). This type of data was unavailable, so a workaround was needed that would still allow a similar analysis. Therefore, the case was made stronger to assume that every household would get enough deposit to set its LTV preference at 80%, as this would create the mean mortgage LTV ratio in the Netherlands (CBS, 2024). This would ensure that the average default probability and LGD align, allowing them to be varied to some extent.

To increase the context behind this decision, looking at Mérő et al. (2023) they use their data to create a regression analysis that sets their interest rate. Through sampling from real world data, they can also set their necessary characteristics per agent, which then sets the interest rate through their regression. Given that for this model both are missing, as only income can really be sampled from the housing prices, the assumption of 80% LTV is required to introduce some variances, thereby allowing other parameters to at least be estimated with this assumption as the preference can be treated as a mean from which the rest is varied. This can be viewed as trying to estimate the parameters of the regression analysis. By estimating them through grabbing the parameters from other research or by trying to estimate them based on their results at least some similar type of behaviour can be recreated that would fit the expectations of the regression analysis. However, this does lead to large limitations.

For example, interaction effects between the different variables will be left out, caused by the difficulty in estimating the parameters without standard regression analysis, as simply grabbing the parameters from other research simply will not suffice. At the same time, this requires the simplification of large aspects of banking to ensure that the interest rate can still in some ways be estimated, therefore requiring only the most important characteristics to be included. It thus also requires strong limitations on the number of variables within the interest rate calculation to reduce the chances of large errors, as the further away the interest rate calculation is moved from the mean, the more likely that errors will be introduced that will create behaviour that is not reproducible.

Simply put, the calculation will take the mean and try to estimate what the interest rate would be given the input variables and how much these would change from that mean. However, as this avoids the standard regression analysis, the results can not be verified and thus is expected to become less accurate the further from the mean the measurements are taken. Thus, high- and low-income groups will be less accurate than the middle-income group. This changes the results of this thesis significantly and will need to be taken into account later on.

At the same time, this will need to be accounted for during modelling, requiring the need to think what would model design ideal when using a regression analysis, and how this change might shift certain thinking and trying to create a model that would fit this idealized model. For example, the 80% LTV preference can be seen as a sample, while banks still do prefer this due to the lower risk (Campbell & Cocco, 2015), this relationship is most likely a lot more complicated and is not the case for every single debtor (Van Der Drift et al., 2022). Treating this as a sample does create differences in later model design decisions, as it changes the way the customer is treated by the bank. How all this relates to the utility function will be discussed now how this interacts with the interest rate.

The Utility and Interest Rate

Getting a mortgage can be summarized in three different steps. First the bank wants to ensure that the debtor has sufficient funds and financial stability to reliably pay off their mortgage, requiring them to make changes to increase their savings, to reduce their overall debt by paying off their loans or reduce their overall monthly expenses. During this phase, the bank is interested in decreasing their loan risk and setting a good baseline for later negotiations and ensure that these will run smoothly. Once this is done, the debtor will start their search for homes, for which the bank will give interest rates and the amount they are willing to lend out, for which the debtor then can accept this as an adequate amount. Then follows the third phase, where the debtor will decide to buy a home, together they will go into more detail to set a bidding strategy. In this phase, the bank will try and push the debtor towards less credit risky homes, by informing them of these and try and pursue them to choose the better options, given their financial capabilities and the overall credit risk of the loan. These can also include options to take on several types of insurance policies and building inspections to ensure that potential unknown risks are taken care of. From this point onwards, phase two and three will be repeated until the debtor successfully purchases a home (NVM, n.d.; Campbell & Cocco, 2015).

Some conclusions can be drawn from these processes, mainly that interest rates can be used as a tool to reduce their loan risk due to it being used as a deterrence, it can push debtors towards other better or cheaper homes. At the same time, the banker will inform their client of potential risks, helping them to better search for homes. Now this last point will most likely be solved through the climate label. However, the former perfectly fits into the previous assumption, that households will get fully loaned up (Van Der Drift et al., 2022; Madsen, 2012; McQuinn & O'Reilly, 2007).

Now, given the data, it can be argued that the bank prefers to reduce their LTV ratio towards 80%. However, from looking at this process, it becomes quite clear that due the data only being from after the house was purchased, many details data on these processes are missing, but some conclusion can still be drawn from it. Banks prefer their debtor to be fully loaned up, while at the same time, reducing their loan risk by keeping the LTV ratio at a reasonable level, in this case, around 80%. In turn, if banks would deem the loan to be too risky, they will push their customers towards the safer options, while the opposite would also be true, preferring to increase their loan, as they deem the risk to be more than adequate. If the loan becomes risky, they can decrease their loan amount by increasing the interest rate, the LTV ratio will automatically also decrease, as this decreases the maximum amount of money a debtor can lend through the maximum allowable monthly mortgage payment. This is the fundamental reasoning for the utility function.

Summarized, the debtor comes in with a certain monthly payment in mind, without fully understanding what increases a home's credit risk. However, through the climate label they do ascertain some information, making them possibly reduce their monthly expenditure on this home. They ask the bank what interest rate they are willing to give, not fully understanding what makes a certain interest rate, thus they keep their expected monthly expenditure the same, under the assumption of getting fully loaned up, as this is what they deem as good enough. The bank might instead increase the interest rate, if they deem the home as too risky, thus the debtor will change their mortgage expectations downwards, but keeping the monthly expenditure the same (with the other way also staying true). Now this does treat the relationship between the banker and the debtor as a sort of black box, reducing the impact a banker might have on changing the debtor's behaviour as only the interest rate effect is kept in mind, despite the larger number of tools available. This can be simply seen as a limiting factor on this model, as some aspects of this relationship would naturally need to be simplified, such as the missing of insurance policies, one of the other main tools in this process.

Now, one small point of nuance should be added. In the model, the buyers will increase their bids if their previous bids have failed. But, even in these cases, the 80% preference should still hold. Given that the failed bids are a result of a heated market, the banks prefer the buyer to bid on homes in a lower price range and overbid on those, resulting in a higher chance of obtaining a home. However, even in these circumstances, the banks still prefer to decrease their overall risk, thus holding onto the same preference. Once again this is a limitation of the model, as the preference is treated like a sample, more types of relationships might exist that are currently not being accounted for.

5.5 The Link between Interest Rate and Flood Risk

The reasons for defaulting after floods have already been explained extensively in section 2.6, a summary of this will follow. Overall, three main reasons were found that contribute significantly towards the increasing default rate after flood events. These will be discussed in increasing importance.

Firstly, due to the loss of infrastructure and buildings, the abandonment of buildings significantly increases financial risks (Thomson et al., 2023), at the same time, people lose their income due to job losses, which can take up to a year to decrease (Tran & Wilson, 2024; Xiao, 2011). Property damage and increase in knowledge of flood risk worsens this effect further by decreasing home prices, which is the second factor. As home values decrease, decreasing their home equity, thus increasing default risk. Household can then make use of two different tools, firstly they can take out their savings, further increasing their credit risk. Secondly, they can also use their insurance payout and to pay for their damages, which, finally, may lead to them selling their home, thus repaying their mortgage (Holtermans et al., 2024; Thomson et al., 2023; Kousky et al., 2020). Now, given the poor state of insurance in the Netherlands as discussed in section 2.5, the first cause will be assumed from here on out (Phlippen et al., 2024; Wettenbank, 2021).

The third and final issue is caused by the loan-to-value (LTV) ratio of a mortgage. Due to the home's significant decline in value, the LTV ratio increases, making the loan riskier and making it more difficult to borrow to repair the damage, thereby increasing the likelihood of default (Caloia et al., 2022).

Kousky et al. (2020) gives an excellent explanation of the mortgage risk after a flood event over a two-year period for single family homes, this can be used to extend the interest rate calculation further. In this paper they give the odds ratio of households after hurricane Harvey, by dividing them into groups of whether they have insurance and the severity of their damages, combined with data on different loan characteristics during origination, such as income levels and reserves, they have created a model to estimate the change in default rate compared to the baseline. In this case, this means the households that have not faced any damages. However, due to the data missing on the household's characteristics after the period of the flood, not all previously mentioned causes can be accounted for in this model.

Despite these limitations, this can still easily be used to improve the interest rate calculations. Given the odds ratio, this can be seen as creating a new mean of a different population, meaning that by changing this new mean through the odds ratio and combining it with the LTV ratio, the default probability can be estimated. The effect will mean that the default increase caused by the floods can be extracted from the population, giving an estimation of the default probability for safe homes, assuming all other variables stay equal.

This does leave out some effects that are caused by the damages, as these also increase the Loss Given Default (LGD), caused by the decrease in equity. This effect can be estimated through the discount rate for homes, which is the percentage decrease in price a home should have given the flood risk, given by Beltrán et al. (2018), a meta-analysis of the effects of floods on home prices. Once again, the same extraction process can be used to extract the LGD of the safe homes. The main issue with this data is the missing of inundation depth, making it difficult to estimate the LGD of the lower risk homes.

The combined effects of this will increase the interest rate for unsafe homes, thereby decreasing the amount a borrower can lend, in effect also decreasing the LTV ratio due to the house being overvalued. This is caused by the exact same sequence of events described earlier in section 5.3, whereby the buyer only internalizes the mortgage as a monthly payment, thereby the debtor will only decrease the total mortgage through increasing the interest rate, with households only changing their monthly payment based on new information. At the same time, this will decrease the amount and chances that a buyer will overpay for a house due to the difficulties in receiving the amount of credit necessary, thereby decreasing the value of the home.

5.6 The Functioning of the Bank

To summarize the functioning of the bank, as interest rates will be set higher in flood-prone areas, households will face more difficulty in acquiring credit, decreasing the LTV ratio of these loans, with the bids decreasing due to mortgages being restricted by the regulations up to a certain percentage of income per year. In turn, as these homes are sold at each step, prices should gradually decrease until they are close to the real price, with now the flood risk taken into account. The bank's perspective on mortgage risk has been summarised in Figure 1.

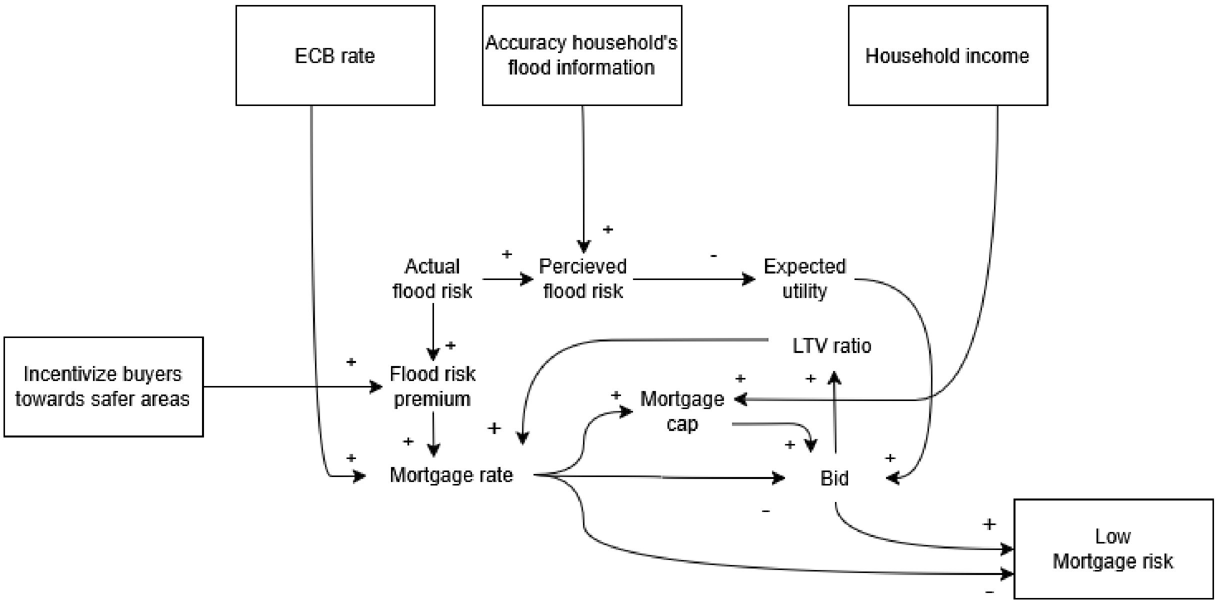

Figure 1: Reduced Bank's perspective on mortgages available in model. This leaves out many other incentives for banks, such as profit incentives, costumer appreciation, but also external factors such as climate change. Adding these would significantly change the requirements of the model.

It should be noted that figure 1 does not represent the entire system, instead it is a reduced representation that fits into the current model design.

5.6 Designing the Flood Label

Given that damage to a home significantly increases its default risk, thereby making the mortgage riskier, a flood label will now need to be constructed around it. To start with, it is best to look at the energy efficiency label to understand the issues these faced when they were introduced. Later, these lessons can be used to construct the flood label while also accounting for current issues with effectively disseminating flood information.

5.6.1 The Lessons from the Energy Efficiency Label

In the Netherlands, the energy efficiency label took a significant amount of time to be readily adopted and used, with initially mainly green voters interested in the concept due to its presentation, as the label was mainly advertised to aid the environment (Brounen & Kok, 2011). It also struggled to gain prominence because the mandatory labels provided limited information to buyers, resulting in little price effect (Stangenberg et al., 2020). Looking at, for example, Belgium, the label only found prominence later, after significant changes. The focus was on achieving a financial benefit, which widened the price gap when savings were significant (Gerassimenko et al., 2024). Similar findings have been reported in France, where energy efficiency significantly increases as energy prices rise (Chareyron, 2024).

This issue highlights a large barrier the water label might face in its adoption, floods, despite their large consequences, simply do not happen frequently enough for the costs to be internalized already. While information campaigns have faced these challenges already, their price effects only work temporarily, while also not being large enough to represent the real price (Niu et al., 2025). Even after flood events, the prices normally go back to normal after 9-12 years (Mutlu et al., 2023). While the label has successfully been adopted in Fairweather et al. (2024), it still raises questions whether prices will decrease enough if no monthly costs are also added to pay for the increased risk.

This leads us to the argument that some form of monthly costs, like the energy payments, will be required to ensure that households are aware of the consequences. Two sides can be argued, by increasing the insurance premiums it creates incentives for households to potentially increasing household adaptation (Hudson et al., 2016), and/or through the increasing the interest rates through the banks (Bani et al., 2024).

5.6.2 Constructing the Flood Label

From this information, a potential flood label can be constructed. Firstly, the label needs to be simplified to avoid building-level risk, as this would make the label needlessly complex for this thesis. As such, only single-family homes will be considered in the label. Then these floods need to be frequent enough so that people remember the events yet damaging enough to ensure people take the labels seriously, so that each step in the label will be significant enough to yield a financial benefit. This is especially important to consider, as the damage follows an S-curve as the inundation depth increases (Endendijk et al., 2023). The label, therefore, needs to ensure that the floods are both high and frequent enough to be considered before households might take them seriously. With flood risk data from the LIWO splits the flood risk level into every 10/100/1000 years (Ministerie van Infrastructuur en Waterstaat, n.d.), every 100 years being found to be ideal as they firstly follow this S-curve closely, with a low label of 1 being at the start of this curve, while 4 and 5 reach the end of this curve, ensuring that each step will be financially interesting enough to invest into.

At the same time, the label will need to account for mainly local buyers taking information seriously, showing the need for buyers to experience or have knowledge of floods to take floods seriously (Niu et al., 2025. It is expected that with a mortgage of around 30 years, an average of 23% of people will experience such a 1 in 100 year flood event, meaning that with a flood label, people will hopefully be able to or have people in their vicinity experience it, avoiding the possible recency bias. Given that the flood event takes 9 to 12 years to be forgotten (Mutlu et al., 2023), it is assumed that the flood label will extend this period.

Banks have offered a different perspective, similar to the question of a climate label. Their concerns have mainly been centred on the extent to which properties get damaged and its overall macroeconomic impacts, highlighting the impact of a 1 in a 1000-year event being too far out of their control, as the damages will be extreme in such a scenario. Rather, they focus on more frequent lower impact floods, mainly flash floods, as they are the single most regularly damaging flood type (Bani et al., 2024). With 1 in a 100-year flash floods reaching upwards of 0.3 meters (Ministerie van Infrastructuur en Waterstaat, n.d.) these can cause damages of up to 5% of total home value (Endendijk et al., 2023), which can result in 1.44 times more defaults (Kousky et al., 2020). These floods can be significant enough for the banks to take interest. As this relationship is exponential (Caloia et al., 2023), it can be assumed that the 1/100-year is ideal, as any lower would decrease overall damage, while any higher would mean other parties should become more involved. It can therefore be concluded that the 1 in 100-year label is ideal given the different interests of banks and households.

6 Detailed Model Design

This section will start with an initial model outline based on the RHEA model by Filatova (2015) and de Koning and Filatova (2019) originally developed using USA data and further contextualized for the Netherlands by Mutlu et al (2026). This outline will help explore how the ABM model's basis functions work, while omitting the parts that have been changed. This section will therefore primarily focus on discussing how the model currently functions in broad strokes, with the differences being left out. The section will, after this, focus on the institutional context, mainly of mortgages. Then, an overview of the steps for agents, households, and banks will be given, with each modelling choice and formula explained afterwards. Finally, the different setups for the scenarios will be discussed, along with how they are expected to impact the outcomes.

6.1 Model outline

The model is based on an expanded version of RHEA, a housing market model that accounts for a house's perceived flood risk by the buyer. The whole model is partially as described in Filatova (2015) and de Koning & Filatova (2019). However, some changes have been made during this period. This section provides a general overview that will later be expanded to include the changes by discussing the logic behind the model from initiation to 1 step. A more detailed explanation can be found in appendix F.2, wherein the formalization of the RHEA model gets discussed. This section will only discuss what is most relevant for the limitations and important modelling choices through natural language, while the appendix will discuss the RHEA model more broadly, and which of its functions needed to be reworked to fit into a more credit-oriented approach.

6.1.1 Model Initialization

The model starts off with a list of single-family homes with certain qualities, such as amount of rooms, size of the lot and most importantly the estimated prices. Each home is assigned to a new household agent, after which the income of the households will need to be estimated. The model follows the basic logic that as households increase their income, they are expected to decrease the fraction of income they spend on housing. Referred to as fraction of income, by ranking the houses based on their expected price the income group and thus the fraction of income can be determined, which then can be used to calculate the income of each respective household. From this, during the initialization, each household is then also given a flood bias, which determines how strongly they take flood risk into account, and a bid margin, which determines how strongly they will either over- or underpay for a home. This can also increase if they fail on the market.

6.1.2 Model Step

Sellers' function is relatively simple: they are randomly chosen to sell their homes on the market after a certain number of steps have passed, creating a consistent trickle of sellers for each step. The risk perception of current homeowners is therefore not taken into consideration when determining when to sell. After their homes are put on the market, the realtor class uses hedonic regression to estimate their prices. The realtor class starts with prices based on data from housing market transactions, using the characteristics of these homes to predict prices. Through this, if the market is hot and prices increase, the asking price of homes naturally responds to this by increasing prices, with the opposite also being true.

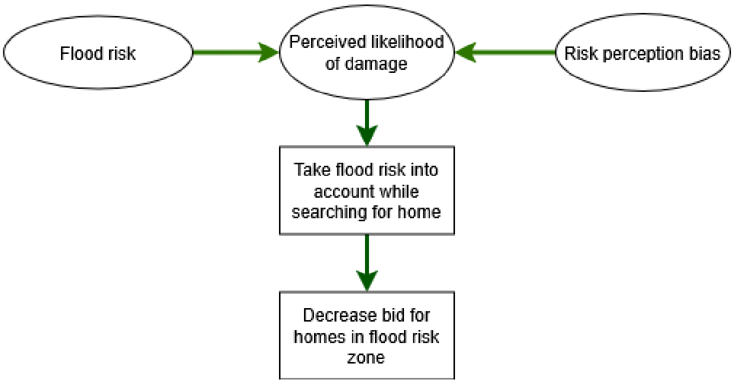

Figure 2: Biased perception of the buyer

After the houses enter the market, every buyer will be given a budget based on their fraction of income, through which they then search for homes. After determining 10 houses they are interested in, the buyers use their own flood risk perception to guide their decisions. This means that instead of properly internalising all the different costs associated with flooding due to the obtuse information gathering, they will use their own experience of the area to guide their flood risk perception. This risk perception allows them to avoid dangerous areas; however, because of their biases, they will underestimate the associated costs, which has been represented in figure 2. After this, they have a short list of 5 homes they are interested in buying. Using their bid margin, they determine the highest utility home for which they will bid the highest, while for the lower utility home they will bid lower.

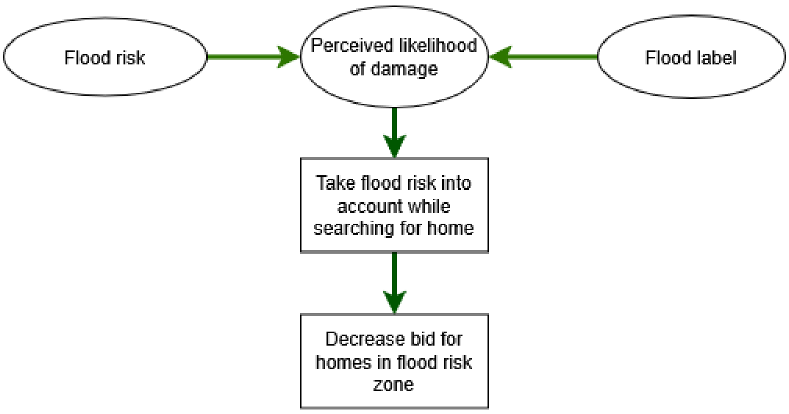

Figure 3: Changed buyer behaviour with the flood label

When the water label is introduced, a shift in buyers' behaviour is expected, represented in figure 3. Instead of using their subjective risk perception, their behaviour will be determined by a more rational approach. Now the buyer will learn about the expected damage they will incur over the life of their mortgage. They have to decide which homes are available, taking into account changes and the risk associated with each home, which will then affect their bidding strategy, lowering the utility of homes in flood-prone areas.

Once the buyer has decided which homes to bid on, the bids reach the seller, who chooses the highest bidder, while the buyer will accept only a bid for their most preferred home. These conflicting interests are resolved by initially accepting the bids, in which both parties get their most preferred homes, while they each slowly move down their preferred lists. For sellers, there is still one caveat: they will only accept a bid that is a certain percentage below the asking price, they may also prefer to wait for another buyer in the next step. Thus both the seller and buyer respond to changes in demand. If demand is low, buyers and sellers expect price decreases, thus they will naturally have lower price expectations. While if the market is hot, buyers will increase their bids through the bid margin to increase their chances of receiving a home.

6.2 General Overview of the New Model

Here, an overview of all the steps the model will go through will be discussed and how the households and banks will affect each other. Starting off with the initialization of the model. Initially, 3000 homes will be sampled from a file containing anonymised single-family home sales. These contain the variables, such as home size and number of rooms, which are needed for the hedonistic regression. These samples come with an initial price, which can be used to reverse the income of the homes based on the assumption that households will almost fully maximize their mortgage. On top of this, it is assumed that households will pay for their home with a 20% deposit on the asking price of their homes. This deposit can then be used to assume the average deposit for a certain income, which will be used for initializing the buyers.

Initially, a certain percentage of homes will be put on sale; these will then be assessed by the realtor class, which will use hedonic regression to predict the home's asking price. The first bid cycle will start; however, interest rates will not yet differ by bank or parcel. Each step starts with new buyers being created, each of whom is randomly assigned one of the available banks, based on their income, and given a deposit. These buyers will then search for homes. However, simply using budgets in this case can be difficult, as interest rates will significantly differ per parcel, instead the bank will set specific upper- and lower-income requirements based on the maximum mortgage allowed, with some space to allow for under and over bidding. This ratio roughly corresponds to the amount of overbidding and underbidding allowed in the scenarios, with additional consideration for the model's stability. Due to the modest number of around 80 parcels sold per step, a few higher-income bidders on cheaper homes can quickly drive prices up. The households then use these income requirements to search for homes within their budget. The buyer then selects 10 or fewer houses based on their income.

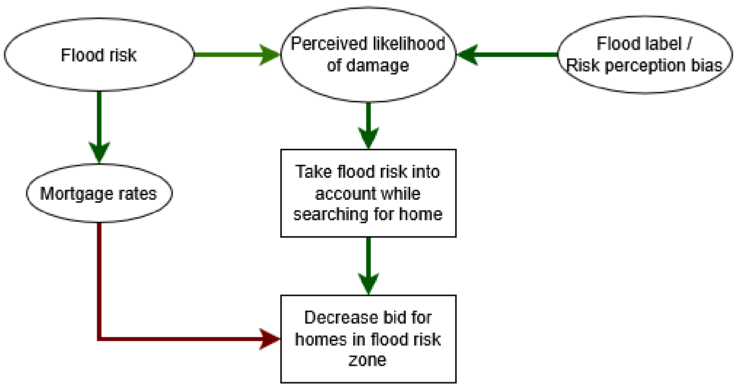

Figure 4: Buyer Behaviour in once banks are added

For these 10 homes, their utility will be calculated based on the expected damage a household will experience from flooding during its mortgage period, with a bias specific to the household. This bias can also be changed to allow them to remove this biased risk perception, in which they instead bid based on the total damages expected during their mortgage period. They will then select the five homes they want to bid on. Then another utility will be calculated in which the interest rate difference from a normal interest rate will be compared. As each bank advertises its interest rates for non-flood-prone properties with a preferred LTV ratio, these expectations will need to be adjusted. Households will bid according to a maximum margin; this margin is a combination of the number of unsuccessful purchase attempts and their preferred bid margin, which is randomly sampled from a uniform distribution.

As discussed earlier, buyers have a specific monthly payment in mind, which is close to the maximum allowed mortgage amount. This maximum mortgage is based on a household's income and the interest rate, with the interest rate now assumed to shift by no more than 50 basis points, which would change the maximum credit. However, buyers also want to pay more for homes when they are desperate, increasing their LTV ratios and interest rates, which is further exacerbated by the increase in default probability for flood-prone homes, thus decreasing their utility even more. Therefore, the utility is first increased by this increase in desperation, also called the bid margin, which defines the upper range of a buyer’s willingness to pay (WTP), which is then converted into a monthly mortgage payment based on an 80% LTV ratio for a safe home. For this upper range, the actual mortgage rate is calculated based on the amount they expect to lend, which is then applied to calculate the actual mortgage. The actual mortgage and expected mortgage are then divided to give the change in utility.

This works due to two assumptions. Firstly, all buyers are assumed to have an 80% preference; therefore, they expect this same rate. While in actuality, there are many different strategies possible. Due to the lack of this data, it was decided to simplify it significantly; secondly, due to the assumption that buyers bid close to or at the maximum allowable credit. As the maximum allowed mortgage is a percentage of a household's income, with a higher interest rate, the actual monthly mortgage payment will remain the same. This will help ensure that households are aware of the decrease in bidding potential they have if they bid on a flood-prone home. Simply put, buyers prefer to bid on a home with the lowest interest rate, as it increases their bid potential.

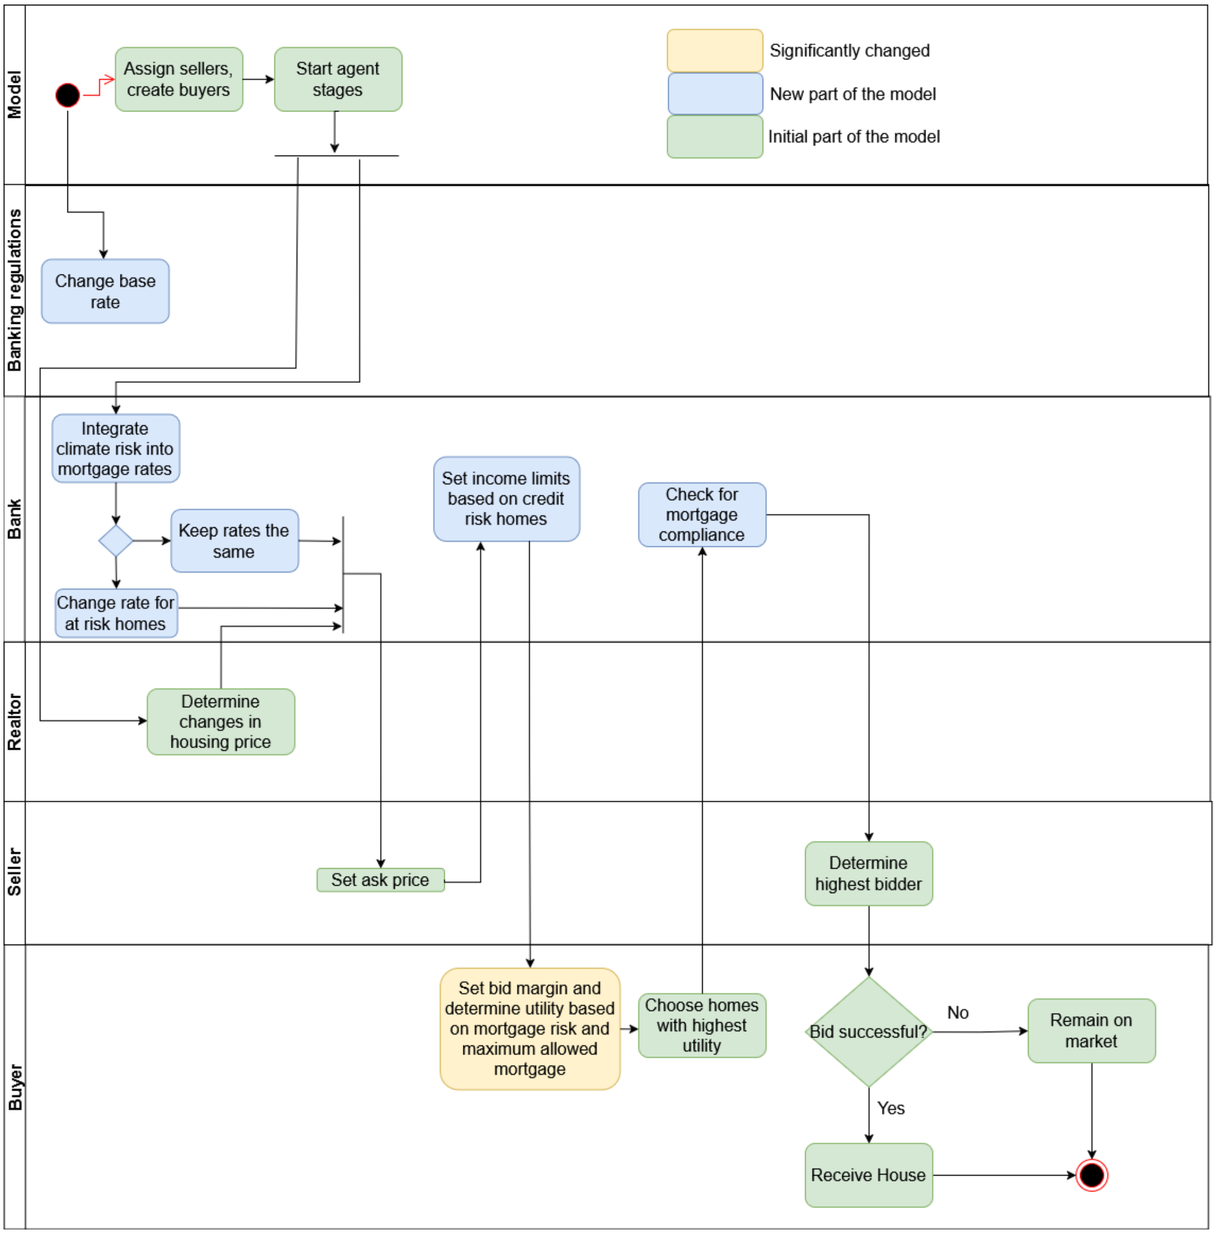

Figure 5: Swimlane activity diagram with changes to the RHEA model.

These homes will then be ranked, with the five highest-utility homes put forward to the bidding round. Each buyer will have a certain amount randomly drawn, and they will either underbid or overbid it. Then, according to their preference for a home, the highest bid goes towards the house most of their liking, with the lowest bid towards the lowest of their liking. However, sometimes homeowners can afford a lower bid but still bid according to a higher asking price. As such, the bid will initially need to be optimized. This is because a higher LTV ratio also increases the interest rate, which then reduces the maximum allowed mortgage. These all differ depending on the bank's loan risk assessment and the buyer's deposit and income. Once the bids are finalised, all the buyers' bids will be resolved according to each buyer's preference, and the seller will receive their maximum possible bid. The seller and their buyer will then transfer their homes, and the sellers will then have a chance to either leave this area or try to find a new home of their own.

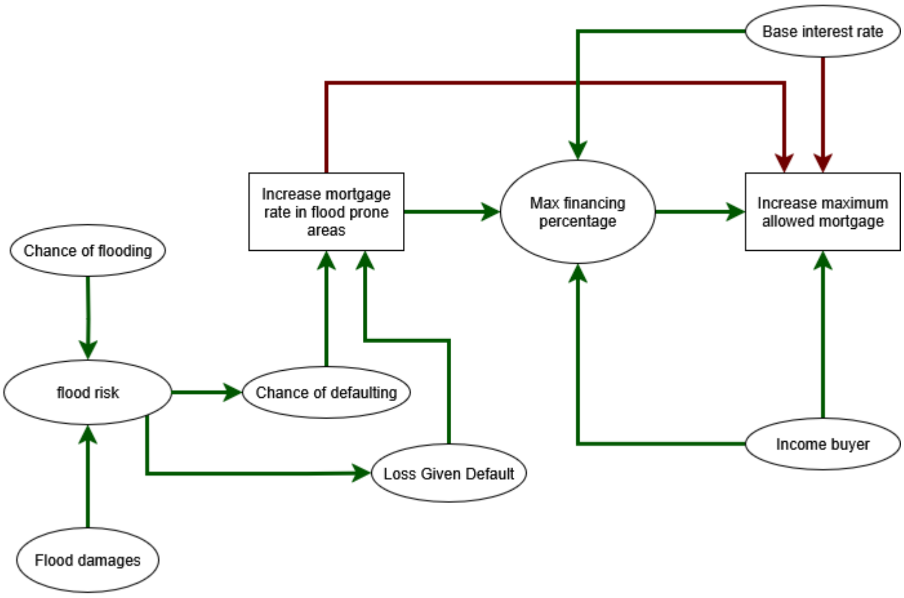

Figure 6: Activity diagram of the bank's perspective.

From the second step onwards, some changes will be added to the model, mainly that banks will slowly integrate climate risk into mortgage rates and that rates will decrease for low-risk homes. Each parcel will be checked by each bank, which will then assign a higher default probability and LGD to that parcel if it is in a flood-prone area. This shift will significantly change the choices a buyer has to make. Firstly, the income requirements for these homes will increase if interest rates rise. Secondly, it is assumed that people will search without fully being aware of the changes in interest rates, as the finding of flood risk is assumed to be obtuse. They will therefore only learn about these changes after finding a home, which will affect how they perceive their utility. As such, households are expected to reduce their LTV ratios for flood-prone properties. This will further shift their preferences towards less flood-prone houses, while ensuring they bid less, as their utility preference dictates how strongly they will over- or underbid.

After calculating their utility, the bid process will also change, as the maximum mortgage is based on the interest rate; thus, an increase will make it harder for the bidder to afford the home. This can be solved by optimizing the maximum mortgage. Ideally, this will be avoided as much as possible through the search function. In summary, buyers will be less interested in buying homes in flood-prone areas and will be unable to afford them as easily due to lower maximum mortgage amounts. This will push them towards safer areas and enable banks to hedge their portfolios more effectively.

Figure 7: Class diagram

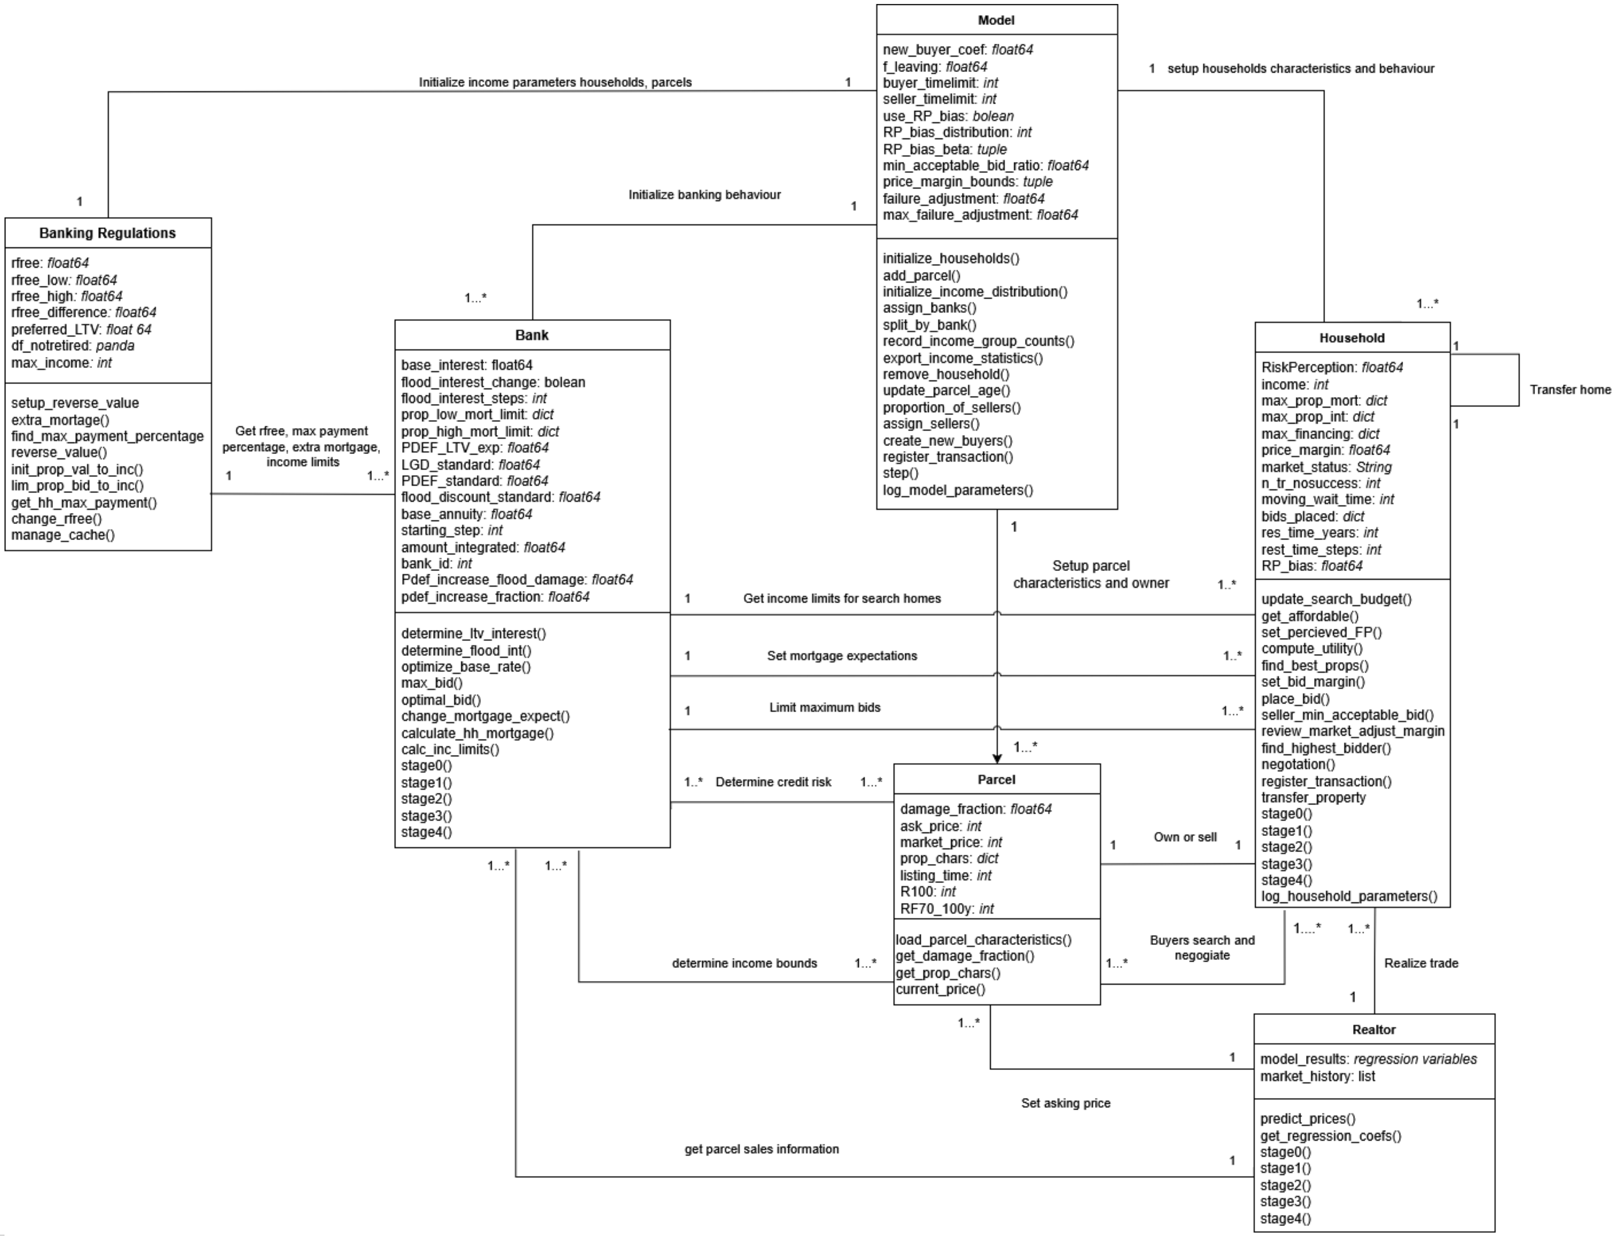

6.3 Formulas

6.3.1 Initializing the Households

As households, on average, have an LTV ratio of around 80% (CBS, 2024) [1], this has been assumed throughout the initialisation process. As such, the current interest rate can be taken. Given the home's initial value, the mortgage can be calculated.

There was a change in how these numbers were recorded in 2018, as they are now measured through the WOZ. However, this model is run on the value of a home. Therefore, the numbers for the year before (2017) were chosen and then estimated for the latest year.

$$M_j^{pref} = P_j^{init} * ltv^{pref} \tag{1}$$

After which, the annuity factor is calculated, which is the percentage of the loan the household has to pay monthly.

$$AF_{base} = \frac{r^{base} * \left(1 + \frac{r^{base}}{12}\right)^{360}}{\left(\frac{r^{base}}{12}\right)^{360} - 1} \tag{2}$$

Then the incomes need to be taken from these mortgages. This can be done by simply reversing the normal maximum mortgage calculations, which will be discussed later. In short, with the average interest rate and mortgage, the maximum financing percentage MF_j^{init} can be calculated based on parcel j's initial value, which in turn can be used to calculate the income. Since only the maximum financing percentage was provided based on income and interest rate, it would differ for each interest rate, as the maximum mortgage amount decreases as the interest rate increases. Then, by calculating the total yearly mortgage payment and dividing it by the maximum financing percentage, the household's income can be determined.

$$inc_i = \frac{12 * AF_{base} * M_j^{pref}}{MF_j^{init}} \tag{3}$$

Finally, due to the preference of 80% LTV, the deposit d for the households can be calculated as the opposite of the LTV. These deposits are later averaged into a percentage of income, so the average deposit for a buyer can be initialised appropriately.

$$d_i = (1 - ltv^{pref}) * P_j^{init} \tag{4}$$

6.3.2 Interest Rate

The interest rate follows the standard formula: it increases as the probability of default (PD) and the loss given default (LGD) increase. These numbers are based on the current average mortgage rates for the first year, as reported by the ECB (2020). Taking the current interest rates r_{base}We can then calculate the amount of interest taken from other sources, such as bonds, and the profit a bank will take from the loan.

$$r_{free} = r_{base} - PD_{init} * LGD_{init} \tag{5}$$

The value of r_{free} is then slowly changed at each step to simulate an exogenous stochastic shock to the risk-free rate. However, this range is quite limited, with there being 17 different steps the interest rate could take, eight up and eight down, each step representing a 1 per cent change of the original r_{free}. Overall, this results in a possible variation of around 55 basis points between the highest and lowest value, which is expected to be sufficient for results. Obviously, this is low compared to possible real world exogenous shocks. However, due to the 150-run limit, it was decided to keep this number low, as the interest rate would be difficult to interpret given the significant variations.

6.3.3 Calculating the Maximum Mortgage

To ensure that the bidding process is done accurately, the current NIBUD financing requirements were added to the model, from which the variable Max payment percentage ($MF_{i,j}$) is based on Warnaar et al. (2024), it takes income and interest rate to give a maximum percentage of gross income a person is allowed to pay for their mortgage. A simple search function gives these. $MF_{i,j}$ times the income divided by the annuity factor of the base interest rate, which is the current average mortgage interest rate, gives us the maximum mortgage a person is allowed to get.

$$M_{i,j}^{base} = \frac{inc_i * MF_{i,j}}{AF_{base}} \tag{6}$$

Then, based on the qualities of a home, which in this case would be flood defence measures of parcel j, an extra mortgage M_j^{extra} can be granted. Which then gives us the total possible mortgage household I can have for parcel j.

$$M_{i,j} = M_{i,j}^{base} + M_{i,j}^{extra} \tag{7}$$

6.3.4 Damage Fraction

All parcels j are given a flood label k, which represents the flood depth each house will experience every 100 years based on the Rijkswaterstaat flood maps. By combining this with the results of Endendijk et al. (2023), the expected damages can be calculated. It should be noted that this does not follow the exact formulas given by Endendijk et al. (2023), as they also include other variables that cannot be accounted for. Therefore, it was chosen to approximate the results, with FD_k representing the flood depth for the label k.

$$DF_k = min[0.05;\ \beta_1 * e^{FD_k * \beta_2} - 1] \tag{8}$$

6.3.5 Setting Interest Rates

From the second step onwards, the banking agent can gradually introduce the new interest rates, which can help them in two significant ways. Firstly, they can supply lower interest rates for less risky homes, while also pushing their lenders towards safer homes. The bank b will linearly increase the difference between the areas, represented by $Pinc_b$. The parcels have six different labels k, which represents the flood label, which increases the worse the average flood event a home will experience every 100 years. Then this is multiplied by the odds ratio for the PD increasing for a home with 5 per cent damage. $PDinc_{d=0.05}$, given by Kousky et al. (2020). This is then multiplied by the total number of homes with this label n_k , divided by the total number of homes n_{all}, then summed over all the different labels. Plus 1, this then gives us the isolated means odds ratio of all the flood-prone properties have in the sample. This can be given as true due to the simple observation being that the mean odds ratio of all homes times the default rate of the home given the odds ratio of 1 gives us the mean default rate. The next step was to then isolate the means odds ratio, which is what formula 9 below does through calculating the means odds ratio.

$$\overline{PDodds_b} = 1 + Pinc_b * \sum_{k=0}^{5} \frac{\frac{DF_k}{0.05} * (PDodds_{d=0.05}) * n_k}{n_{all}} \tag{9}$$

It should be noted that in here the relation to the damage fraction and the PD increase is considered to be non-linear (Caloia et al., 2023), but the relationship given here is clearly linear (the 0.05). This was chosen because the data from Kousky et al. (2020) seem to show a linear relationship. However, due to the lack of additional data, it was decided to keep the relationship linear. Moreover, it is expected that this will have little effect on the results, due to the damage fraction DF_k following an S-curve related to the flood label, with mainly the labels four and five having a high damage fraction. In contrast, one to three have relatively low damage fractions. The discount percentage can then be divided by PD_{init} to give us the $PD_{standard,b}$ for which the bank gives to the non-flood prone properties.

$$PD_{standard,b} = \frac{\overline{PD}}{\overline{PDodds_b}} \tag{10}$$

For LGD, the bank here has a different strategy: it will fully integrate LGD into climate risk areas rather than the safer areas. This is due to the lack of literature on the LGD in at-risk areas. This can be calculated more accurately, as in Caloia et al. (2023). However, due to time constraints, it was instead chosen to simplify this. Instead, the numbers DR_{flood} were taken from the meta-analysis of Beltrán et al. (2018), which gives a discount factor of 4.8% for 100-year floodplains.

$$\overline{DR_b} = Pint_b * \sum_{k=1}^{5} \frac{DR_{flood} * n_k}{n_{all}} \tag{11}$$

Then, minus the initial LGD is the isolated LGD for bank b.

$$LGD_b^{standard} = LGD_{init} - \overline{DR}_b \tag{12}$$

With the assumption that, with an LTV ratio of 80%, the recovery rate is 81% of that, the LGD is 19%. With this, we can calculate the increase in LGD by subtracting the assumed LTV of 80%, which gives us the change in LGD by the LTV. This is, of course, a simplification of the fire sale process, as the total sales costs stay the same for all homes, thereby decreasing the LGD of more expensive homes.

$$\Delta LGD_{i.j}^{LTV} = max\left[-LGD_b^{sandard}; \frac{(LTV_{i.j} - 80)}{100}\right] \tag{13}$$

For banks that do not integrate their flood risk into their interest rate, the values are equal to the initial value. These values can then be used for calculating the interest rates of the parcel j for buyer I, with the β being given by Campbell & Cocco (2015).

$$r_{i,j} = r_{free} + \left(PD_b^{standard} * \left(\frac{DF_j}{0.05} * PDinc_{DF=0.05}\right) * \beta^{ltv\ i,j-\ 80}\right) * \left(LGD_b^{standard} + \overline{DR}_b * FP_j + \Delta LGD_{i.j}^{LTV}\right) \tag{14}$$

With banks, without flood interest integration, simplifying this by not taking into account the differences between areas.

$$r_{i,j} = r_{free} + \left(PD_b^{init} * \beta^{ltv\ i,j-\ 80}\right) * \left(LGD_b^{int} + \Delta LGD_{i.j}^{LTV}\right) \tag{15}$$

This should result in an average rate increase of around eight basis points.

Once again, due to the simplification of LGD, the influence of the LTV ratio on the interest rate has not yet been adequately accounted for. This choice was made for two different reasons. Firstly, no actual data was available to create the regression variables for the interest rate, as it would need to be used with the RHEA model. A combination of mortgage data and housing characteristics would be ideal for this. Secondly, this would require significant rework of the RHEA model, needing the time that was not available. Income initialisation of the households would require significant work, as the income needs to be sampled from the LTV ratio of the new data.

On top of this, the requirement of getting fully loaned up would need to be dropped as higher incomes tend to have lower LTV ratios on top of paying a lower percentage of their income towards their mortgage (De Vries & Boelhouwer, 2008; Van Der Drift et al., 2022). The actual LTV ratio influences both the PD and the LGD. As the LTV ratio increases, the loan's exposure also increases, further increasing the LGD, even with the influence of the damage ratio. Furthermore, as interest and monthly payments increase, the loan will become more challenging to pay off, thereby increasing the PD. It is known that the LTV ratio will most likely play a significant role in sustainable mortgages rates. An example of the influence of the LTV ratio on LGD and PD is shown in Caloia et al. (2023).

6.3.6 The Search Strategy

For households to afford the mortgage, they need to finance the remaining value of the home. At the same time, households tend to allocate a similar share of their income to housing (De Vries & Boelhouwer, 2008; Van Der Drift et al., 2022; Damen et al., 2016). Exactly how this relationship functions is still being debated, with, for example, Damen et al. (2016) arguing that households prefer to spend a fixed amount of their income on their home. However, Madsen (2012) and McQuinn & O'Reilly (2007) argue that these budgets are set solely by credit availability. Van der Drift et al. (2022) combine both points; they argue that higher-income households, due to their greater wealth and age, can spend less on housing. To allow the households to apply a bidding strategy, the last option was chosen. If, for example, credit restrictions become too burdensome, it would be difficult to explain the effects of flood information on housing prices under the previous assumptions. Other factors, mainly households' deposits, would require a significant expansion, which is currently not yet fully explored in the model. Instead, it was decided to recreate this relationship through the households' search function, setting their budgets according to the previously discussed expectations.

To recreate these budget constraints, households have a preference for an 80% LTV ratio, which can then be used to dictate the way the households bid, as they have been initialised to be able to pay for their home with a 20% down payment d_i. The final issue that needs to be solved is that the model will work with multiple banks. Therefore, an average budget would not work, as the household can afford to pay more as the interest rate increases. This is due to the Netherlands having an interest rate deduction, on top of the model using differing interest rates to internalize the increased PD.

While this is not similar to how people search, it would be far more computationally intensive to calculate the household budget for all the different interest rates. This means the same calculations used in the initialisation of households are used to determine the income requirements for a household to afford a mortgage. By taking a margin of around the preferred LTV of around 80%, with the margin being chosen based on model speed and the internal effects on the model. If the top margin was set too high, many high-income earners would be able to pay for a cheaper home, which would decrease the maximum mortgage, thereby negating the previous assumption. If the bottom margin were set too low, lower-income earners would struggle to bid according to their preferences, and sellers would be reluctant to accept their bids because they were too low. Finally, if the bounds were set too wide, it would cause instability in the model, as buyers can bid on only five parcels, and their search range is much wider, leading them to bid on random properties.

For a deeper exploration of this part of the model, see appendix F2.1.2. This will discuss the deeper reasons behind this decision, mostly being caused by modelling limitations.

6.3.7 Utility

According to Mutlu et al (2023), the household will integrate the possible risk into their expected utility. Then, in the standard model, the buyer uses this utility to rank the homes, all of which are then normalized from 0 to 1, which is represented by $NU_{i,j}$, with 0 being the lowest utility. Then the buyer uses their bid margin bm_i, which is based on how unsuccessful the buyer has been in the market, plus a randomly set price margin bound, to set their bids. These values are then multiplied to give the adjusted margin $AM_{i,j}$.

$$AM_{i,j} = \begin{cases} NU_{i,j} * bm_i & if\ bm_i > 0 \\ (2 - NU_{i,j}) * bm_i & if\ bm_i \leq 0 \end{cases} \tag{16}$$

With the bid then being given by multiplying the utility times this adjusted margin to give the bid.

$$b_{i,j} = U_{i,j} * (1 + AM_{i,j})$$

This has been significantly changed to include the mortgage rate. Before the utilities get normalized, the utilities need to change to account for the interest rate changes properly. This is the time when the buyer talks with the bank to negotiate their mortgage. For this, the willingness to pay (WTP) for the home is based on the monthly mortgage. This WTP is assumed to be a range around the monthly payment rate, defined by the bid margin. First, the upper range of the bid margin is taken, which ensures that when the different homes are ranked, the home with the highest utility, but also with the best interest rate, will get ranked higher. Minus the deposit, this gives us the mortgage the buyer wants to pay for the home.

$$M_{i,j}^{expected} = EU_{i,j}^{flood} * bm_i - d_i \tag{17}$$

As the buyer is unaware of the impacts their lending decisions have on the interest rate and banks advertise their interest rates as being based on an LTV of 80 in safe areas, the buyer has a certain monthly payment $bc_{i,j}$ in mind.

$$bc_{i,j} = M_{i,j}^{expected} * AF_{80,k=0} \tag{19}$$

This WTP will be consistent from now on, as it is assumed to be the price the household wants to pay for this specific parcel j. First, the interest rate of this loan will be calculated, as previously discussed. This will then be used to calculate the new mortgage, while keeping the expected monthly payment the same.

$$M_{i,j}^{actual} = \frac{bc_{i,j}}{AF_{i,j}^{actual}} \tag{20}$$

The different mortgages, plus the deposits, are then divided to give the expected decrease in utility due to the reduced potential mortgage. Because the interest rates on flood-prone properties are higher, they will be disproportionately affected. At the same time, higher bidding is also disincentivised, as the bidder is now aware of the risk they are taking on when they increase their bid.

$$U_{i,j}^{mortgage} = \frac{M_{i,j}^{actual} + d_i}{M_{i,j}^{expected} + d_i} * U_{i,j}^{flood} \tag{21}$$

While at first this would seem to have a significant effect on expected utility, it actually has only a marginal effect, around 1-2%. This is because the search function already accounts for the model's assumed LTV ratio. Furthermore, as the deposit increases with income, the actual LTV ratio for the home decreases slightly, while the bid margin each household has is also increased.

It should be noted that the utility here is not fully calculated as previously discussed, with the primary assumption missing: the maximum lending capacity dictates what households can bid, and the search function approximates this assumption. Instead, the utility here merely functions as a way to ensure that the households will bid according to what will bring the most utility, with the lowest monthly payment, while also ensuring they order their bids according to where the interest rate will most strongly increase due to the flood risk premia as the interest rates in these areas will strongly increase when the LTV ratio increases, thereby decreasing their utility more strongly.

Furthermore, this should suffice under the assumption that the financial constraints will not shift up or down, as set by the NIBUD's interest rate thresholds. If these thresholds are met, the household's ability to lend marginally increases or decreases. Once again, this is handled by the search function, which takes it into account.

Finally, the assumption of maximizing the credit usage is incompatible with the bid margin, as it would require the bidder to always bid according to their maximum lending capacity, thereby negating the bid ordering and bid margin, limiting the effects the bidding process (and thereby the flood perception) will have on the housing price, and negating any behavioural differences between the agents. This is also taken into account while setting up the search function by allowing for some bid margin.

In conclusion, this will still allow the bidder to bid up to their maximum mortgage but will instead punish them for bidding higher on more flood-prone houses, without them truly understanding why. A sort of black box that limits their bidding. This can be explained by looking at two similar homes, one with and the other without any flood risk, while all have the same qualities. Both can be bid on, but the risky home has a higher interest rate, which decreases its value. Both will (at least closely) max out their mortgage due to the search function.

From this point onwards, the model will function as usual, with only the bids checked to ensure they comply with the 100% maximum LTV ratio, and that the bid is not higher than allowed by the NIBUD. If the bid is higher than what is allowed, it creates an optimization issue where the maximum mortgage increases if the bid decreases due to the LTV ratio decreasing. As the difference is generally marginal, usually around 3% above the maximum allowed mortgage due to the search function, the relationship can be assumed as linear. Simply dividing the bid plus the maximum mortgage by two usually suffices. This new bid then gets rechecked to ensure that it complies with the regulations, and the previous steps are repeated until the mortgage is accepted. For a more thorough explanation of this process, please see appendix F2.2.5.

7 Data

For the model, the house transaction data from Mutlu et al. (2023) has been used. This dataset has been gathered by the Dutch Association of Real Estate Agents (NVM) and includes a wide range of variables, such as transaction price, lot size, and garden size, with all variables available in Appendix B. The dataset also includes coordinates and 6-digit zip codes. Instead, the variables were used to estimate the prices of 10,000 homes using randomly sampled values from the dataset. This ensures that the data is fully anonymized by removing any of the possible identifiers, while still ensuring the data is accurate. For each model run, 3,000 of the 10,000 homes were randomly sampled to provide a representative sample of the Dutch Housing stock.

The dataset comprises 68,247 transactions for single-family homes in Limburg between the start of 1990 and the end of 2020. To ensure these prices remain consistent, all transactions have been adjusted to reflect 2020 inflation rates. These houses have then been combined with the LIWO dataset (Ministerie van Infrastructuur en Waterstaat, n.d.) to obtain the current potential water depth every 100 years from flood events. They are represented by a number from 0 to 5, with 0 being safe and 1 indicating a potential flood depth. By further combining this data with that of Endendijk et al. (2023), the damage fractions can be calculated. The damage fraction is the percentage of a home's value it will experience every 100 years. This damage fraction will later be combined with data from Kousky et al. (2020) to give the odds ratio.

In Table 1 these variables are represented for each of the flood labels. Two types of grouping become clear when looking at this table: the flood-prone and non-flood-prone houses, and the group distinguished by the largest increase in damage fraction between labels 3 and 4, which exceeds 2 percentage points. These groups, from here onward, will be called flood-prone and not flood-prone for those representing labels 1 through 5 and 0, while the other group will be called none to low flood-prone and high flood-prone, which represent labels 0 through 3 and labels 4 and 5.

Table 1: The different single family housing options with the water depths from Ministerie van Infrastructuur en Waterstaat (n.d.), combined with the damage fraction of Endendijk et al. (2023) and the default risk odds ratio from Kousky et al (2020)

Classification	***Water depth***	***Percentage of houses***	***Damage fraction***	***Odds ratio***
0	Safe	67.1%	0.00%	1.00
1	0.075 meters	8.9%	0.69%	1.06
2	0.125 meters	7.6%	1.49%	1.13
3	0.175 meters	6%	2.76%	1.24
4	0.225 meters	5%	4.78%	1.42
5	>0.3 meters	5.2%	5.00%	1.44

In these groups some significant differences can be found, as there is a large price increase in the high flood prone area. This becomes especially important as these houses will potentially experience the largest price shifts due to their large damage fraction increase compared to the previous labels, on top of these homes also having the highest average price, it is expected that these will also experience the heaviest losses.

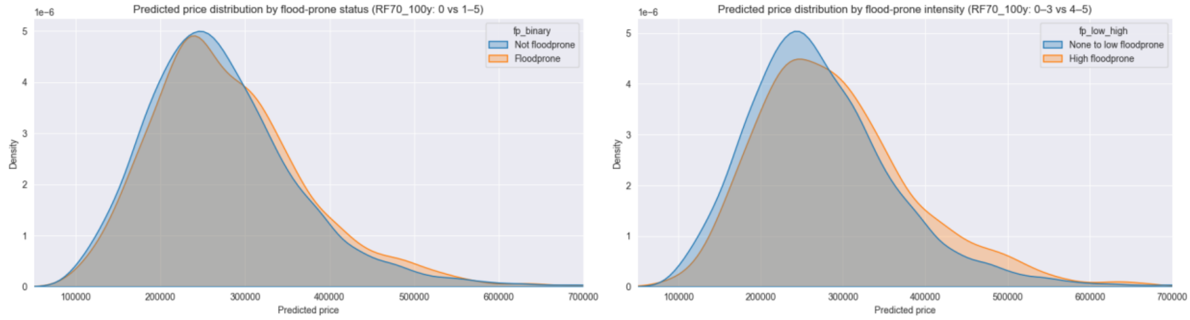

Figure 8: The two distributions: Flood prone and High flood prone are different from each other (Kolmogroff p-value: 1.1e-2 and Mann-Whitney test p-value: 1.6e-03). Means are: Not flood prone (0): €274.000, flood prone (1-5): €282.000, high flood prone (4-5): €293.000.

This difference of around 3% can be explained with previous research from Bosker et al. (2018), as the most expensive single-family homes are found in the Randstad, the urban core of the Netherlands. This highlights one small issue with the anonymization of houses, as this loses the impact of the surrounding utilities, which are more prevalent in the Randstad, explaining the price difference. If these are taken into account, a discount of 1% for flood prone homes would be needed. It is therefore important to note that with the risk perception biases used in the model, the results should be accounting for this difference.

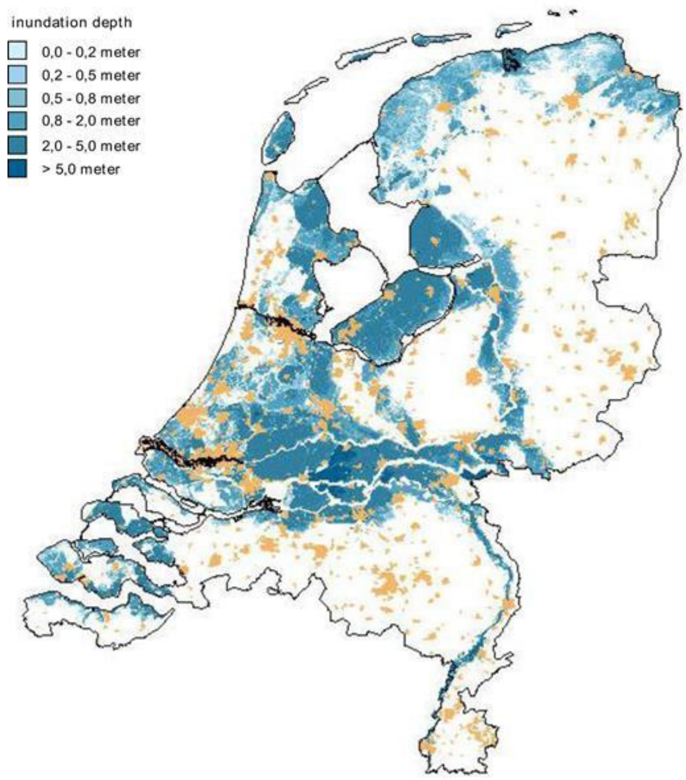

Figure 9: Inundation depth in the Netherlands. The yellow areas represent the urban areas, while the black represents the flood defences. This highlights the overall flood risk in especially the Randstad. Originally from Bosker et al. (2018)

8 Experimental Design

The primary focus of the base scenario is to examine the effects of pluvial (flash and surface floods) on buyers' risk perception. As the damages of pluvial floods are relatively low, around 0.05 of home value, but with a high chance of 1/100 per year for at-risk single-family homes, as discussed earlier. Finally, all parcels will have a uniform flood risk mode, meaning that distance from a river will not be considered.

Buyers will start with a sort of “base bid margin,” called the price margin of under and over 5%. They will also be able to assess flood risk based on their biases. This is to account for any awareness buyers may already have of the flood risk in the area. The model will be run 150 times for each scenario, to ensure that the base state will not significantly impact the results, while also giving a good range of outcomes. While this number is in the lower range of what is typically done, limitations in the amount of data a Panda DataFrame can handle limited later data processing. At the same time, a full run of each scenario could take up to 3 days of compute time, restricting the amount of runs further. Finally, sellers will discount their homes by up to 0.05 for each successive failure to sell.

All other variables are presented in Appendix B, which also provides additional context on how they influence the model. Here, only the most important ones were mentioned.

8.1 Single and multiple bank scenario

Two different scenarios will be run, called the single bank and multi bank scenarios. The first group of scenarios will focus on either 100% uptake or no uptake of integrating flood risk under the different demand scenarios, forming the best and worst case of the scenarios. These banks will either not integrate the flood risk or integrate it in a relatively quick pace of 2 years, seen in table 2. Then, in the multi-bank scenarios in table 3, some banks will slowly integrate the flood risk over 10 years, while others will follow the speed of the previous banks. These will, from now on, be referred to as slow and quick integrators. These scenarios will either have partial quick uptake for 33% and 66% of banks, meaning only a total of 3 banks are simulated in each run. For all scenarios, each buyer will be randomly assigned to one of these banks with a specific strategy, meaning there is no actual competition between the banks beyond the buyers’ bidding competitiveness.

The chosen scenarios can highlight the strategic issues a slow integrator might face, meaning that while they could indeed integrate these rate changes slowly, they could face other issues that will make this strategy (dis)advantageous. Instead, because their customers cannot afford safe homes, they will disproportionately affect poorer communities further by offering favourable rates for unsafe homes. Furthermore, possibly due to decreased competition for unsafe homes, it may lead slow integrators to buy them disproportionately, thereby keeping prices high. Because prices are kept high, these banks might not be as interested in quickly changing their rates.

Next, all scenarios will be run with 1x, 2x, or 3x ***demand ratio***, meaning that for each seller introduced at a step, there will be one, two or three buyers generated. These will from now on be referred to as 1 for low demand, 2 for high demand and 3 for very high demand (or also low, high, vhigh in the scenario names). These buyers will remain in the model for a certain period, even if they are unsuccessful in buying, thereby increasing demand for houses in the model. These can produce two distinct effects regarding the flood label. Firstly, demand for houses can constrict credit availability, reducing price integration as the need for a home outweighs the potential flood risk. The second effect can cause the opposite to happen. Instead, houses are sold more often and more quickly, thereby increasing the number of transactions, which in turn speeds up the full integration of price changes into the market. These two effects are expected to create significant differences between these scenarios. These scenarios have specifically been decided based on this, as during the preliminary runs, very little changes in behaviour was found with higher than three times demand. Thus it was chosen to cap demand at three.

The final variable changed for each scenario is whether the buyer has ***rational risk perception*** or not. If turned off, it means that the buyer will have a bias towards underestimating the number of floods the home will experience, simulating how the market currently functions. Turned on, it will simulate how a water label might impact the housing market and how much it would change prices. It should be noted that both a water label and sustainable mortgages are expected to be introduced and expanded, making it unnecessary to run scenarios of this kind. However, it was decided to include these scenarios, as they improve the interpretation of the results, helping ensure that each policy's effectiveness is properly accounted for and that interaction effects between the two policies are appropriately evaluated.

Given these variable variations, the standard scenario matrices were used to ensure that each variable interaction was accounted for. The given tables below can be read as follows: Table 2 represents the scenario's with a single bank, where if the ***change speed*** is set to *None*, it means that the scenario is run without the bank changing the interest rate according to flood risk, with an integer, it is set to increase the rate difference over that time period slowly. For Table 3, all three banks have different speeds: some take 10 years, while others take a faster 2 years. This is represented by the list of *integration periods* for each bank.

*Table 2: Scenarios for the single bank runs. The naming conventions are: (credit assessment strategy)_(Demand ratio)_(Flood risk perception). **Demand ratio** can be read as Low = 1, high = 2, very high = 3, as it relates to the name of the scenario. For **change speed and Credit assessment strategy** NoFLINT (= None) in the scenario name means no banks take flood risk into account, while FlInt (= 2 years) means all banks do. The final adjective, **Rational,** or nothing at all means the rational integration or subjective risk perception, based on the household's own experiences and biases.*

Scenario name	***Change speed (Years)***	***Demand ratio***	***Rational risk integration***
NoFlInt_low	None	1	False
NoFlInt_low_rational	None	1	True
NoFlInt_high	None	2	False
NoFlInt_high_rational	None	2	True
NoFlInt_vhigh	None	3	False
NoFlInt_vhigh_rational	None	3	True
FlInt_low	2	1	False
FlInt_low_rational	2	1	True
FlInt_high	2	2	False

FlInt_high_rational	2	2	True
FlInt_vhigh	2	3	False
FlInt_vhigh_rational	2	3	True

*Table 3: Scenarios for the multi bank runs. The naming conventions are: (credit assessment change speed, **F**ast**S**low)_(**x** out of **3** banks are fast **Fl**ood interest change **Int**egrators)_(Demand ratio)_(Flood risk perception). **Demand ratio** can be read as: Low = 1, high = 2, very high = 3 as it relates to the name of the scenario. The **change speed** is a list of 3 banks, and the time in years after the first step of the model at which they are fully done changing their credit risk assessment practices, so from the scenario name 1/ 3 = [2, 10, 10] and 2/3 = [2, 2, 10]. The **rational risk integration** is related to whether the buyers are using the flood label to assess flood risk (**True**), instead of using their own experiences (**False**).*

Scenario name	***Change speed (Years)***	***Demand ratio***	***Rational risk integration***
FS_1/3FlInt_Low	[2, 10, 10]	1	False
FS_1/3FlInt_Low_rational	[2, 10, 10]	1	True
FS_1/3FlInt_High	[2, 10, 10]	2	False
FS_1/3FlInt_High_rational	[2, 10, 10]	2	True
FS_1/3FlInt_vHigh	[2, 10, 10]	3	False
FS_1/3FlInt_vHigh_rational	[2, 10, 10]	3	True
FS_2/3FlInt_Low	[2, 2, 10]	1	False
FS_2/3FlInt_Low_rational	[2, 2, 10]	1	True
FS_2_3FlInt_High	[2, 2, 10]	2	False
FS_2/3FlInt_High_rational	[2, 2, 10]	2	True
FS_2/3FlInt_vHigh	[2, 2, 10]	3	False
FS_2/3FlInt_vHigh_rational	[2, 2, 10]	3	True

8.2 Output Metrics

For each run, two different groups of output metrics were saved, one was saved per transaction, as seen in table 4, while the other registered certain characteristics of single family home owners per step, as seen in table 5. Through these, trends could be noticed that indicate which parts of the system impact the results the most. These metrics thus mostly relate to sub question 5, *"How could the interaction between a water label, flood-risk mortgage pricing and credit restrictions affect housing prices and LTV ratios?"*, giving a broad system overview of all the metrics that determine the housing price. One seemingly important metric in table 4 is missing, the bid margin. This variable is based on the price margin that is randomly sampled from a range of -5% and 5%, plus an increase based on successive market failures. Due to this also being estimated through and restricted by other metrics, such as the percentage of maximum mortgage, interest rate and flood label, it was deemed not necessary to track this as it would become increasingly inaccurate as the other metrics increased. These other metrics therefore were more important to understand the effects of credit restrictions, thereby answering the sub-question, since their increases would cause the changes in housing prices and LTV-ratios.

Table 4: The output metrics per transaction recorded per step

Output metric	Dimension	Description
Asking price	€	The price set by the seller, based on the realtors expected value
Transaction price	€	The price after negotiation, thus based on a combination of bid margin, the flood information availability (flood label or personal bias) and interest rate changes
Percentage of maximum mortgage	[Mortgage / Maximum mortgage]	Based on the mortgage negotiated before the transaction divided by the maximum mortgage with the maximum mortgage being based an *estimate* of the actual value, not taking into account any changes in LTV ratio. Thus this output metric is always an underestimate.
Interest rate	Dimensionless	The interest rate set for the mortgage at the time of the transaction.
LTV ratio	[Mortgage / Asking price]	The negotiated mortgage amount as a percentage of the asking price
Income buyer	€	The income of the buyer at the time of the transaction
Bank ID	Dimensionless	The ID of the bank
Flood interest change	years	How many years before the bank fully changes their credit risk assessment practices to account for flood risk
Damage Fraction	[Expected damage after flood / asking price]	The expected percentage of damage after a flood
Flood label	Dimensionless	The flood labels as described in section 7, giving an estimate for
Step	Half a year	The step at which the transaction was completed on

For each step, output metrics for the home-owners were also saved. These metrics were considered less important, but could track overall trends within the market and how the transactions could affect the community and banks.

Table 5: The output metrics for homeowners that were saved for each step

Output metric	Dimension	Description
Owner income	€	The income of the owners at step x
Flood label	Dimensionless	The flood label of the owner's single family home

Last Transaction Price	€	The last recorded transaction price of their home
Bank ID	Dimensionless	The ID of the bank they recorded their last transaction with
Step	Half a year	The step at which these metrics were saved

9 Results

In this section, the model's results will be discussed. These results will be explained in the same order as the scenarios, as each layer on top of the other. The scenarios have therefore been split into three different groups, the scenario where no banks perform any integration of flood risk. Then follow the scenarios in which all banks follow a sped-up integration process over 2 years. As the final group, a mix of strategies will be examined, with some banks following the quick strategy and others slowly integrating it over 10 years. These groupings were chosen as the main scenarios because this thesis focuses on the impact of banking in relation to a water label and on the relative complexity of the interactions between these groupings and the other variables.

Each of these groups, therefore, will follow the same sequence: initially, the overall effects of the scenario are explored, followed by the effects of demand on this grouping. Finally, the interaction between the water label and all the other effects will be investigated, due to the high complexity of the interactions, especially between the different bank scenarios, which require this ordering.

9.1 What happens without a bank that changes interest rates?

Initially, the model was run without the banking changes to see what would happen under different demand and risk-perception scenarios—first, the effect of demand on prices. As expected, demand plays a significant role in home prices. As demand increases, the speed at which these rise increases as well. As prices rise, so do the income groups able to buy a house, until eventually the average buyer is unable to secure a home, forcing them out of the market and increasing the average buyer's income. This results in a more aggressive market, in which higher loan amounts and, therefore, higher LTV ratios are required to secure a home.

Contrary to expectations, both lower- and higher-income groups do take on higher LTV ratios, but for different reasons. While the former can easily be explained by the extra credit and deposit required to secure the home due to intense competition, the latter requires deeper exploration. As higher-income individuals do not need to outbid others as much, they are required to lend less than they are capable of, forcing them to rely more on their deposit. This suggests that there is an oversupply of homes in the higher price segment, and that it takes longer for these homes to sell, which is an essential part of how income groups get displaced. But as the model runs, prices will rise, leading to less competition, fewer transactions, and a lower LTV ratio, thereby decreasing interest rates. For all figures, such as Figure 10, a bubble represents a decile of their respective group.

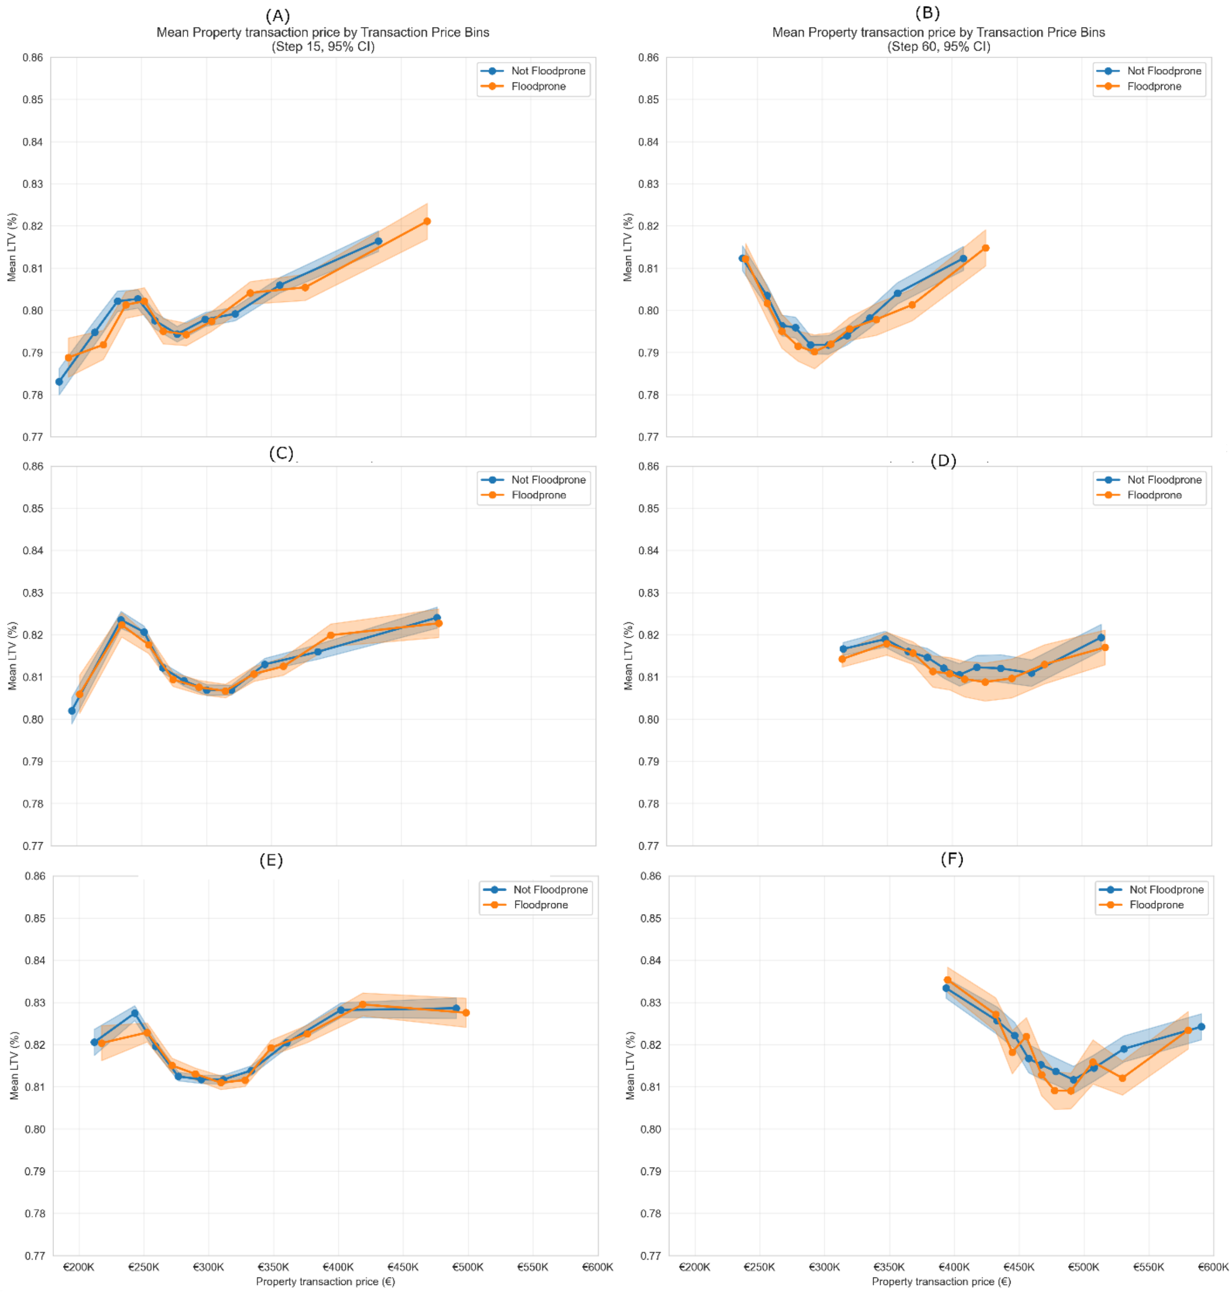

Figure 10: The left figures are at step 15, while the right figures are at step 60, the final step. All figures are without any sustainable mortgages and biased risk perception, forming the baseline. Figure A and B are low (1x) demand, figure C and D are high demand (2x) demand and figure E and F are very high (3x) demand. This shows that as demand increases, the transaction price increases and the LTV-valley also increases. On the y axis, the LTV is represented.

9.1.1 Credit restrictions

Looking at the scenarios from the perspective of credit, at the start, as demand consists of households of all incomes, it causes top income earners to comfortably be able to outbid the lower incomes, thereby displacing them and, due to the rising prices, ensuring that it becomes difficult for them to find another home in the area. Eventually, the prices level off and the houses are distributed purely by the amount of credit that is available, with the higher income earners being able to pay for these homes with less credit than what they are able to borrow, while lower income earners are far more credit restricted.

Seemingly, the ordering of the credit restrictions from Van Der Drift et al. (2022) and Damen et al. (2016) is a natural outcome of the bidding process. Higher-income earners can outbid lower-income earners. This allows for a sort of rank ordering, where each lower-income group must outbid higher-income groups, forcing them to lend more to purchase a home. Each higher-income group, therefore, puts pressure on the group below it, culminating in those in the lower-income group having to rely on credit the most.

Especially early in the scenario, credit is essential, as competition is incredibly high, making it difficult to secure a home. Once prices start rising, fewer households will bid for more homes, decreasing the need for credit, but with a broader spread. This can be explained by the central bank-set interest rates in the model. As they lower the interest rate, the maximum lending capacity will also increase. However, it will take some time for the realtor to set prices accordingly, while households will need to adjust their bidding strategies, thereby widening the spread. Intense competition will instead negate this spread, as all buyers are bidding aggressively, requiring maximum credit to secure a home.

This order, however, does not come back in the low-demand scenario. Instead, both lower- and higher-income earners are not limited by credit. This also highlights one of the limitations of the model, as the reverse of the decreases in credit reliance for higher incomes can quite easily explain it. Because income is sampled from housing prices, the sample therefore has a bias, as it does not account for the entire range of possible income groups below this cutoff that are also bidding. Instead, it has a hard cutoff, rather than the demand slowly trailing off. This cutoff also explains why this issue is exclusive to low-demand scenarios: when the model is run in higher-demand scenarios, it naturally solves itself. Over time, housing prices will increase, but demand will remain the same across price groups, thereby increasing competition for the now lower-priced section of the market. This means lower-income groups in low-demand scenarios experience less competition than in high-demand scenarios.

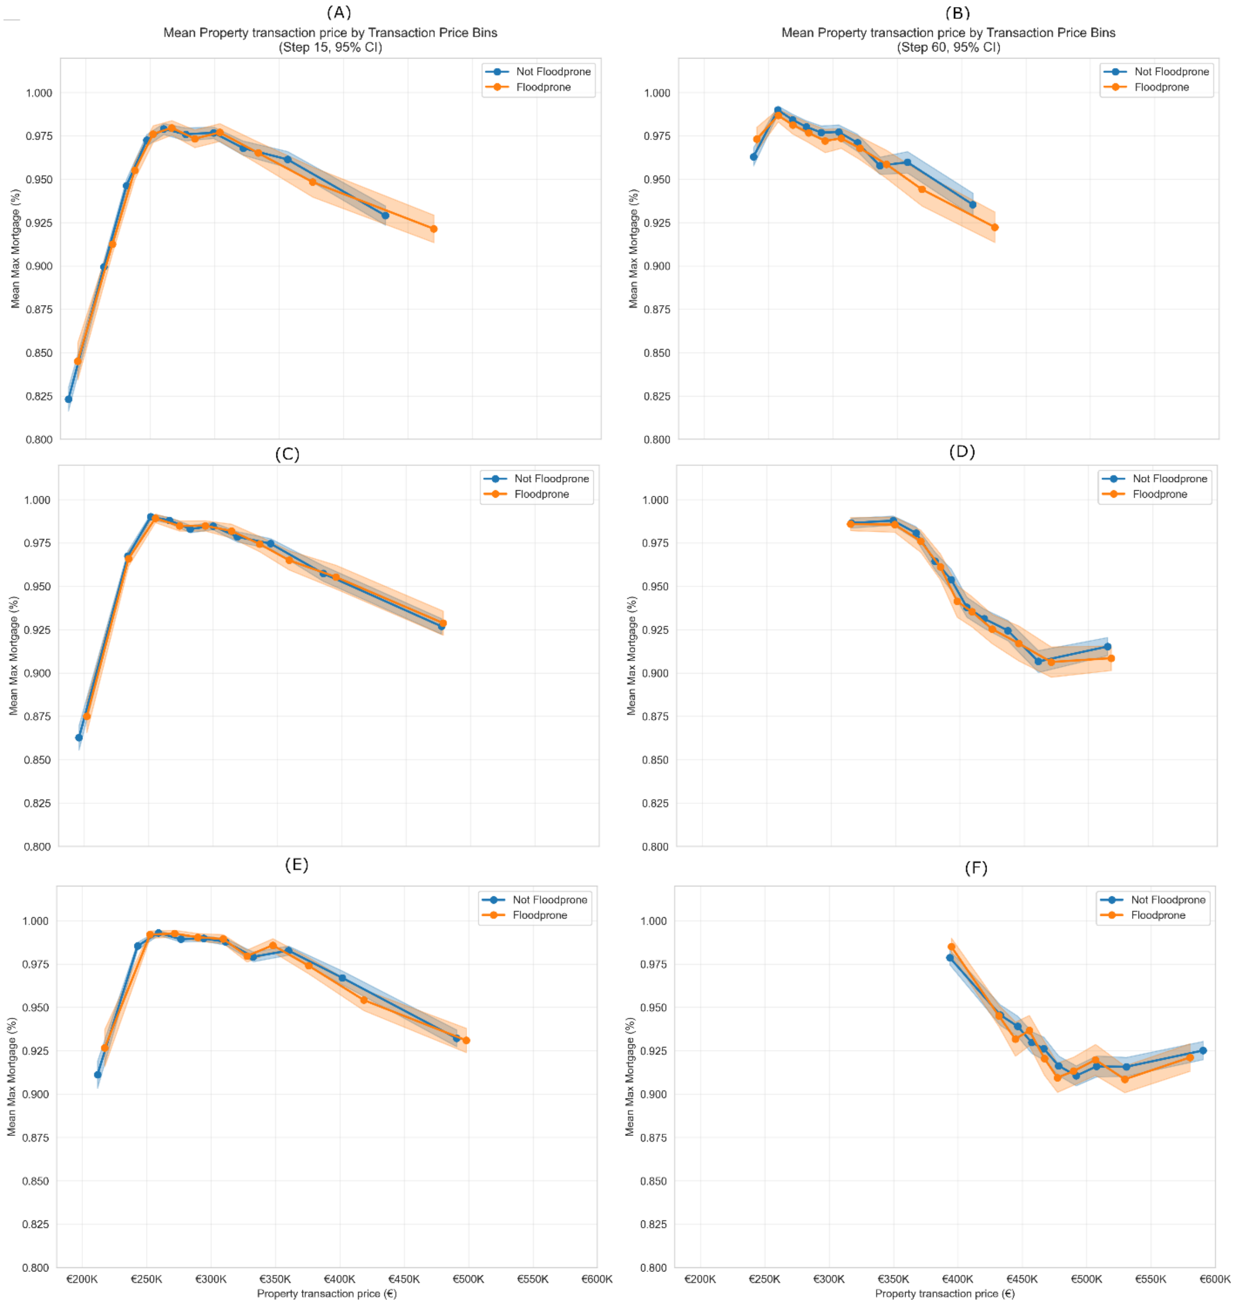

Figure 11: The left figures are at step 15, while the right figures are at step 60. All on the y axis are the maximum percentage of mortgage. Figure A and B are low (1x) demand, figure C and D are high demand (2x) demand and figure E and F are very high (3x) demand. This shows that as demand increases, the credit restrictions are more quickly lifted, due to the income groups being able to more easily afford their home through the deposit.

9.1.2 Unbiased perception

Furthermore, no actual displacement was found, with homes in the risky areas priced slightly, but insignificantly, higher in all scenarios, even without the water label. There are some significant effects, such as the lowest 10% of flood-prone houses now being lower in price than the lowest 10% of safe homes. Nevertheless, this effect is not significant enough to cause any major shifts in the overall housing market. This, however, shifts significantly when the buyer rationally accounts for the potential damages they might experience over the lifetime of their loan, especially in the high (2x) demand scenario. In this scenario, the buyers start to actually reduce their bidding for the risky homes, while increasing their bids for the safer homes, resulting in decreasing prices of around 6%

for the flood-prone properties and 8% for the highly flood-prone homes. This behaviour, however, does not result in the opposite, where the prices increase in the safer areas.

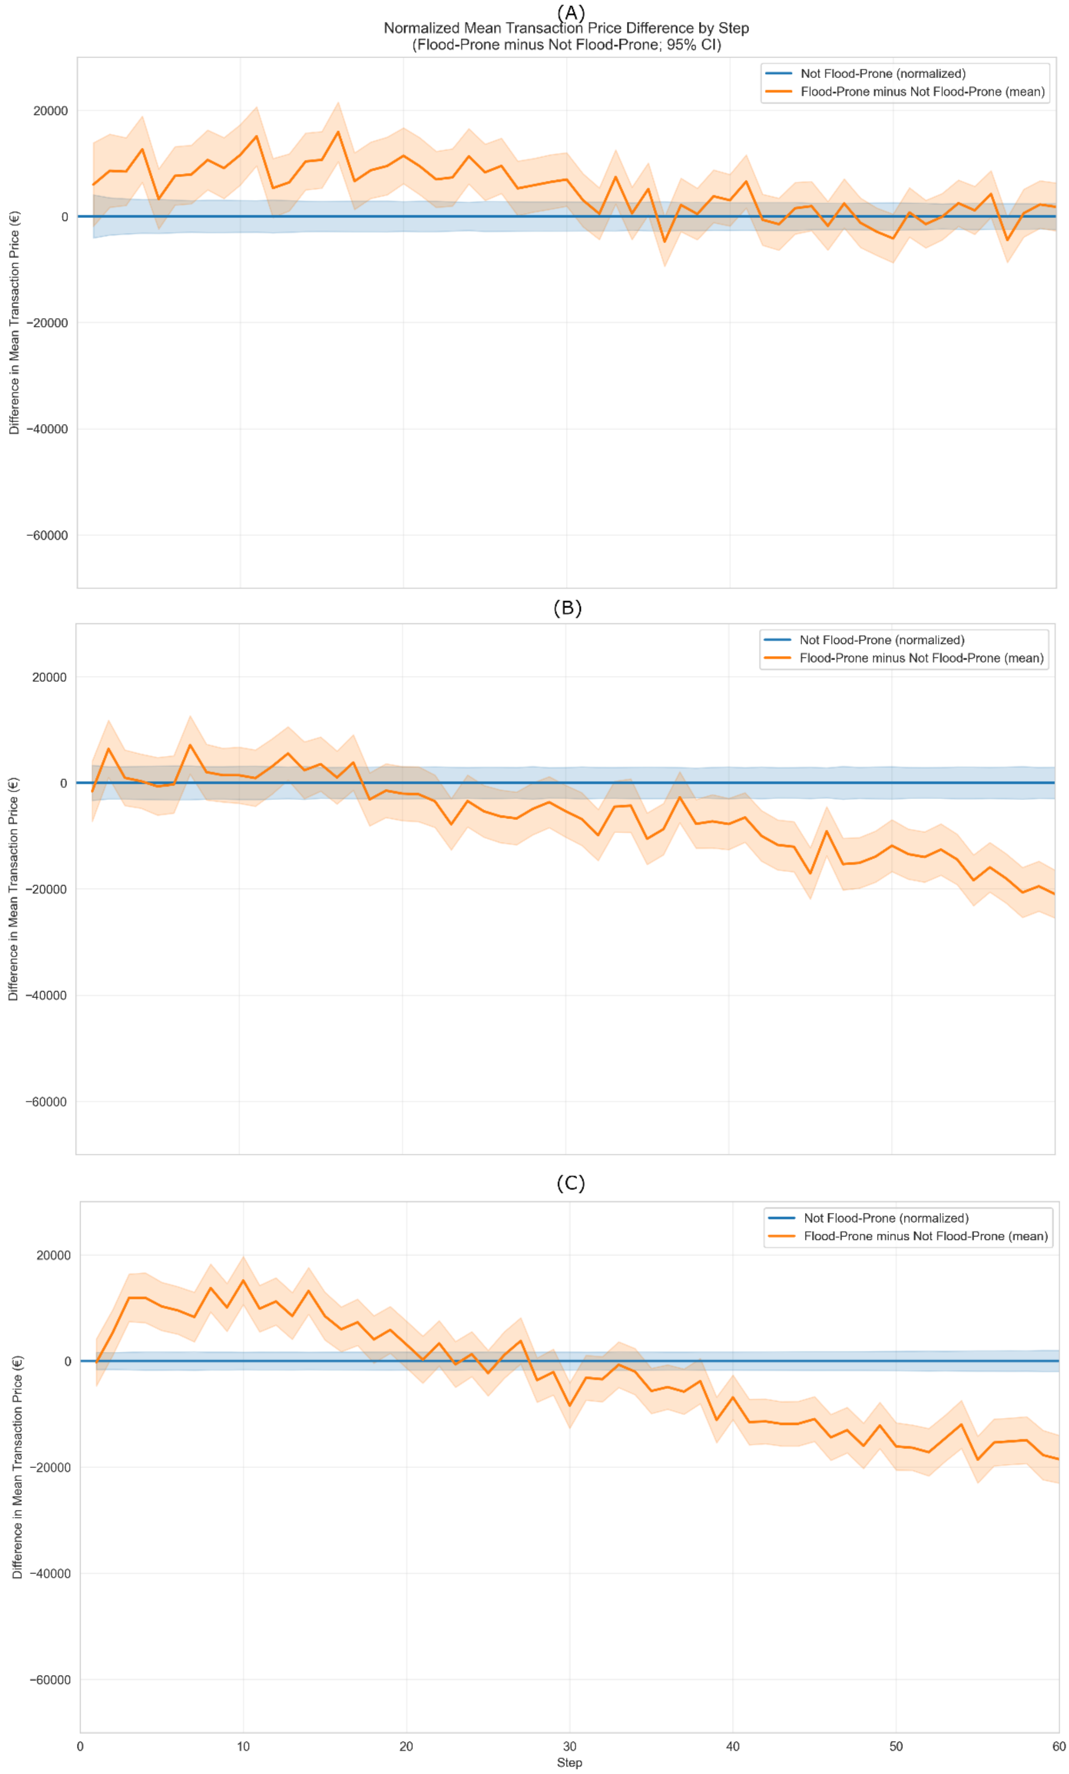

Figure 12: All scenarios are without sustainable mortgages and with the water label, figure A is low demand (1x), figure b is high demand (2x), figure C is very high demand (3x). As demand increases, the water label becomes more effective, but hits a limit at very high demand

The behavioural changes cannot be seen in the credit restrictions, as especially buyers of lower income still rely heavily on credit to secure a home, but rather through the LTV ratio. As higher-income households can put down larger deposits on homes and lend more, they are able to outbid on non-flood-prone homes, increasing their LTV ratio. Lower-income groups, however, are not able to bid as much, leading them to bid less for flood-prone homes. Nevertheless, it simply means they can snatch up a home at a price higher than what would generally be expected, with the downside that it has a higher flood risk.

In the low-demand scenario, prices decrease only slightly, and the noticeable market effect is too slow to cause many changes. While still faster than the very high-demand scenario, with the LTV ratio difference the highest across all scenarios, it becomes easier for people to lower their bids in the flood-prone housing sector. As credit restrictions play a lesser role in this scenario, this should actually help increase the effectiveness of the flood label. But due to low transaction volume, these price changes are slow to take effect.

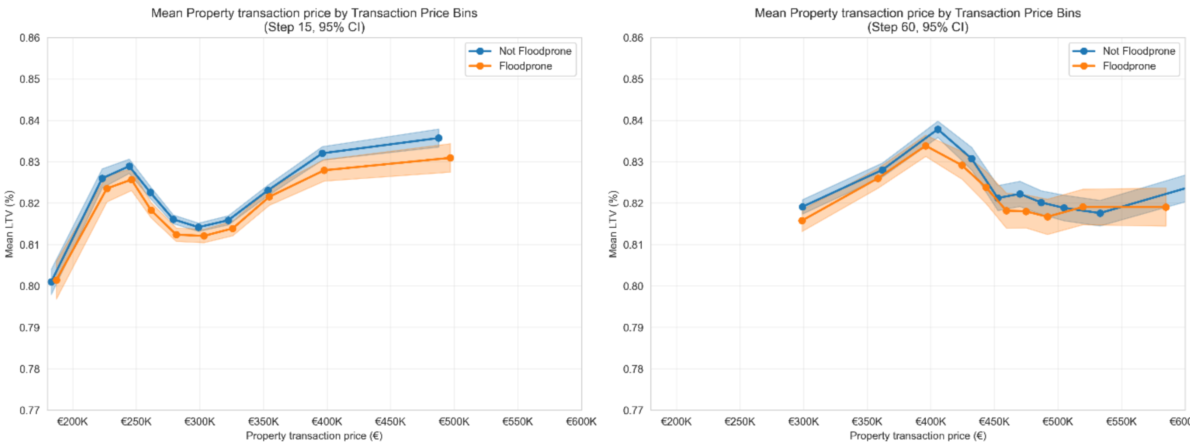

Figure 13: very high demand scenario, unbiased, Y mean LTV, X property transaction price

It can be concluded that, up to a point, demand for housing actually increases the speed at which the water label is integrated into pricing, rather than slowing it. This can be explained through a multitude of reasons. Firstly, due to the aforementioned behavioural change, the higher-income groups have the means to change their bidding strategy, which affects the bidding strategy of the lower-income groups. The high-income groups, therefore, have significant market power to set prices. A typical example of the musical chairs effect in the housing market. The second explanation was theorized earlier, which is simply due to the increase in transactions, resulting in more buyers integrating the water label into their bidding strategy, speeding up the price shift. Both of these effects are strengthened by increases in demand, as there are not only more buyers per parcel but also more high-income buyers, which further amplify the effect.

The effectiveness of the water label appears to decrease as demand further increases in the very high (3x) scenario, with the LTV ratio difference now decreasing. While this difference is still significant enough to reduce prices, the actual difference between flood-prone and non-flood-prone properties (so not high-flood prone) is now far lower than in the previous scenario, at around €12.000, rather than €30.000. Now, credit restrictions are more important, making it harder for prices to reflect increases in flood risk. Even the highly flood-prone properties only slightly decrease by around €15.000.

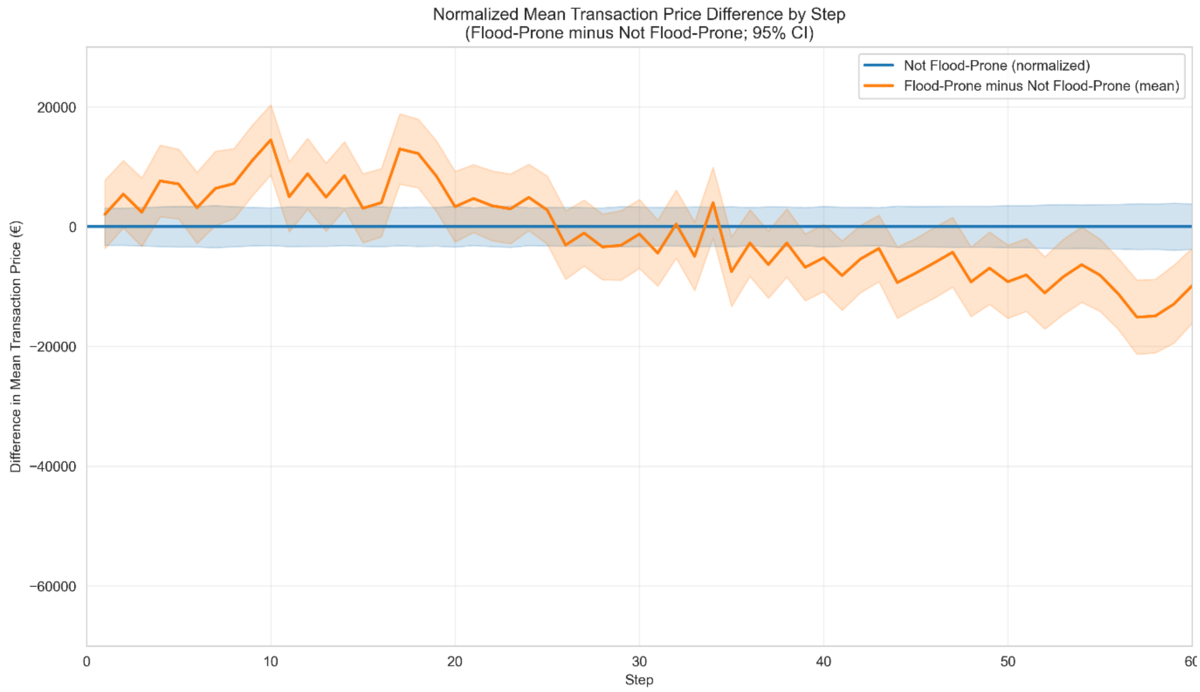

Figure 14: Price difference between flood-prone and non-flood-prone areas, unbiased perception, very high demand

Although this issue can also be viewed from the perspective of highly flood prone properties, which does change the story up quite a bit. Instead, integration still happens at a steady pace, even matching the results of the high demand scenario. This is still contrary to what should be expected, as demand should in theory increase this difference. It can therefore be concluded that does definitely play a role. However, it is not significant enough to impact the results of the most dangerous properties. Instead, the safer homes do have a tendency to stick in their prices, while the other group will have their prices adjusted.

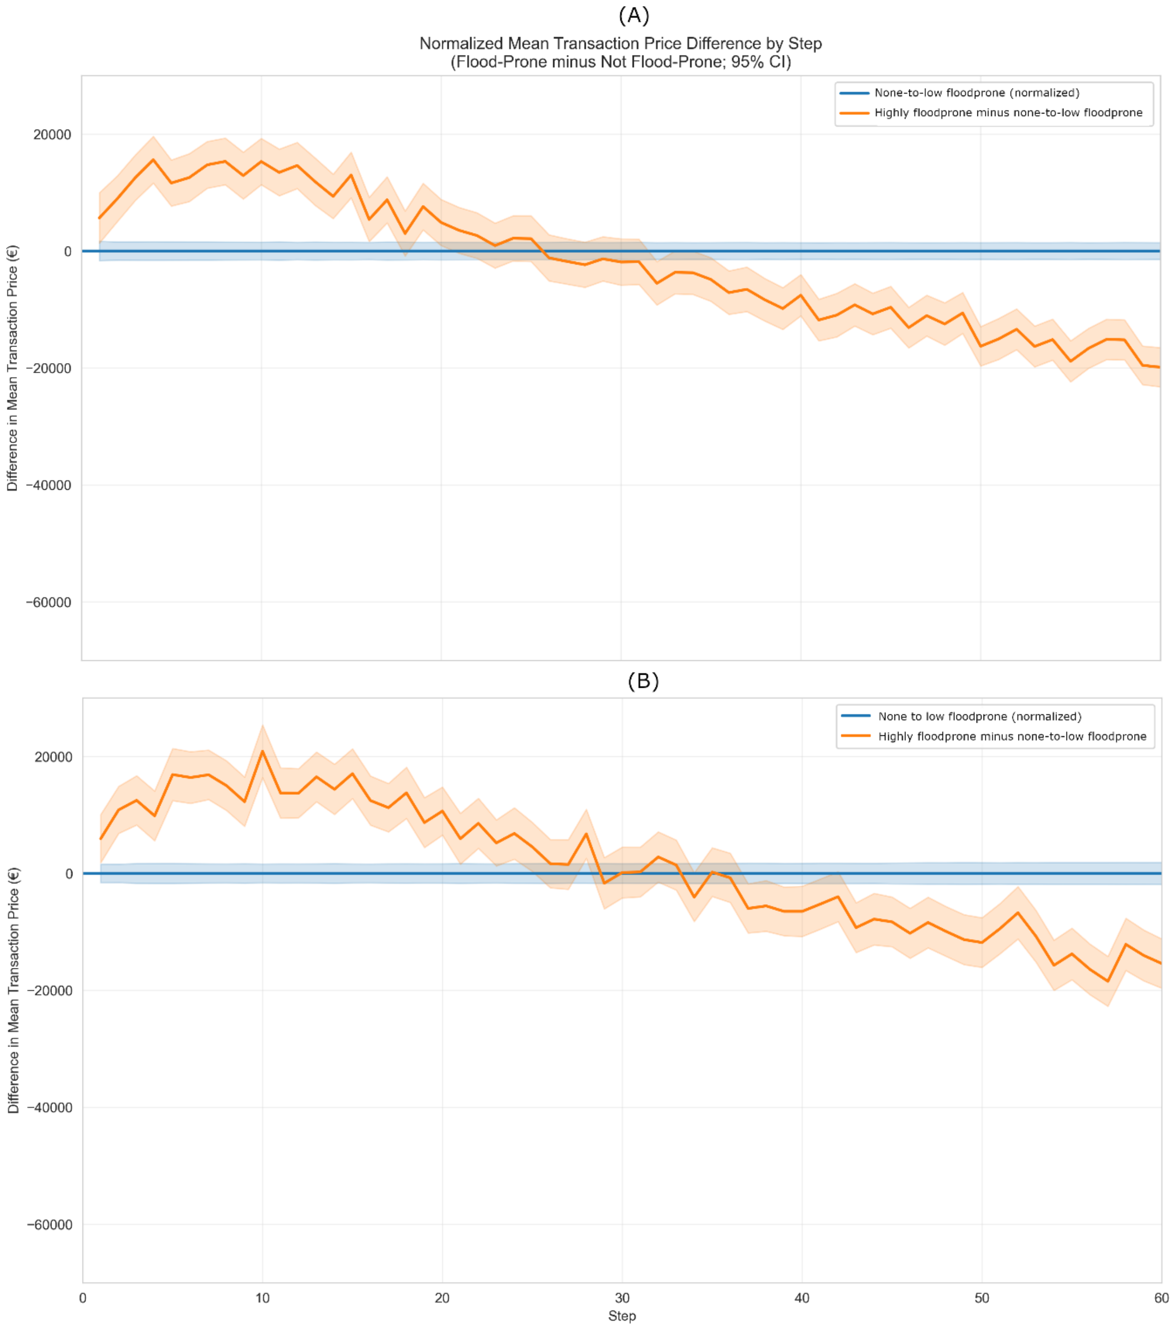

Figure 15: Figure A in low demand, while figure B in very high demand, the main difference between this figure and figure 16 is that these are none to low flood-prone and highly flood-prone, showing that while demand might be low, highly flood-prone do show quicker integration speed.

9.2 Introduction of the flood risk premium

Once all banks start integrating flood risk, buyers significantly change their bidding strategy. Now, the LTV ratio is playing a significant role in the bidding strategy. This is due to the exponential increase in interest rates as the LTV ratio rises, which is further amplified by the PD and LGD for homes in flood-prone areas, making these differences quite significant. Buyers can therefore lower their monthly payment by a significant amount by bidding less for a home. This strategy is clearly seen in scenarios where all banks raise interest rates, thereby mitigating some of the effects of interest rate increases on homes. However, this effect is not long-lived. As the prices start to stabilize and demand naturally starts to decrease, the mean LTV ratio starts to decrease for the safe homes, reducing the difference in interest rate between the areas. This causes demand to play a significant role in the speed of price declines, as it widens the interest rate differential due to higher LTV ratios. Shortly afterwards, the LTV ratios for the unsafe homes will start to increase, widening the interest rate gap again until the LTV ratios of both start to converge. Now, the unsafe homes will once again become interesting as the price decrease has been fully integrated.

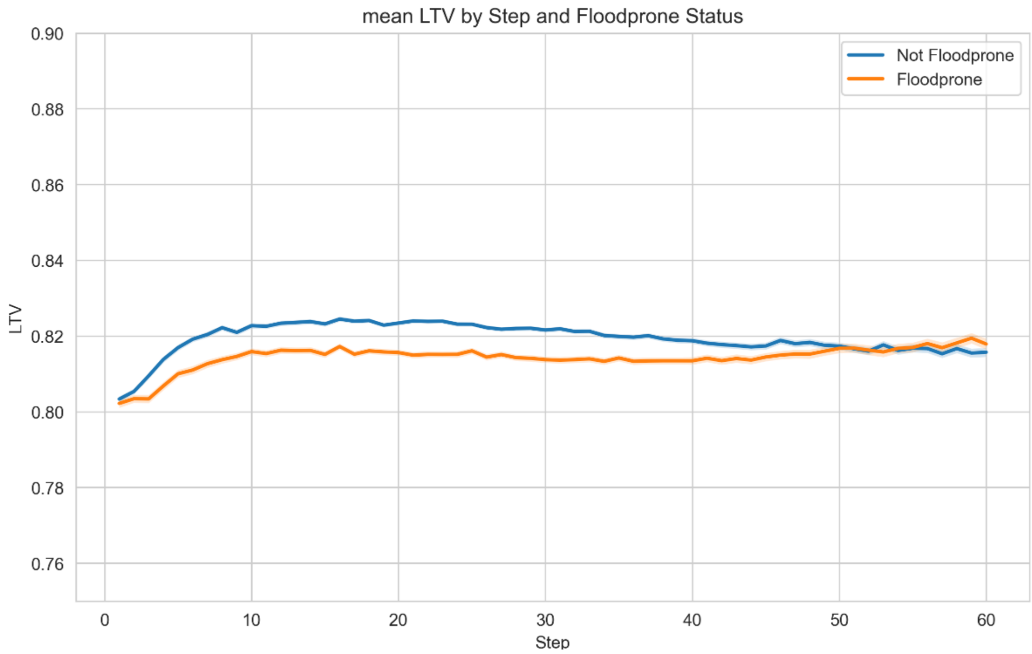

Figure 16: X LTV Y Timestep, very high demand, with sustainable mortgages

However, at this point, a major issue arises. In scenarios with large numbers of bidders, the result is an eventual mismatch between demand and supply. As the homes were built to match the demand of their time, they do not yet account for the flood risk in that area. This is caused by an upward price shift in the safe homes, while the unsafe homes shift downwards, making the highest- and lowest-priced homes the most sought after. After all, the distribution of buyers' incomes will not change, thereby increasing the LTV ratio in both market segments. This valley has already been spotted to exist even in scenarios without the risk integration. However, once interest rates shift, this effect is amplified by the splitting of supply and demand. The LTV ratio, therefore, seems to form a more literal, numerical valley, with buyers at either end being punished more harshly. Yet the price difference seems to have a different cause: even with lower incomes, buyers still face harsher credit restrictions, especially later on, and need to lend significantly more to pay for the cheaper, but unsafe, homes. This valley is amplified when demand is higher; in the lower-demand scenario, it is not as significant because the difference between the top and bottom is not as large.

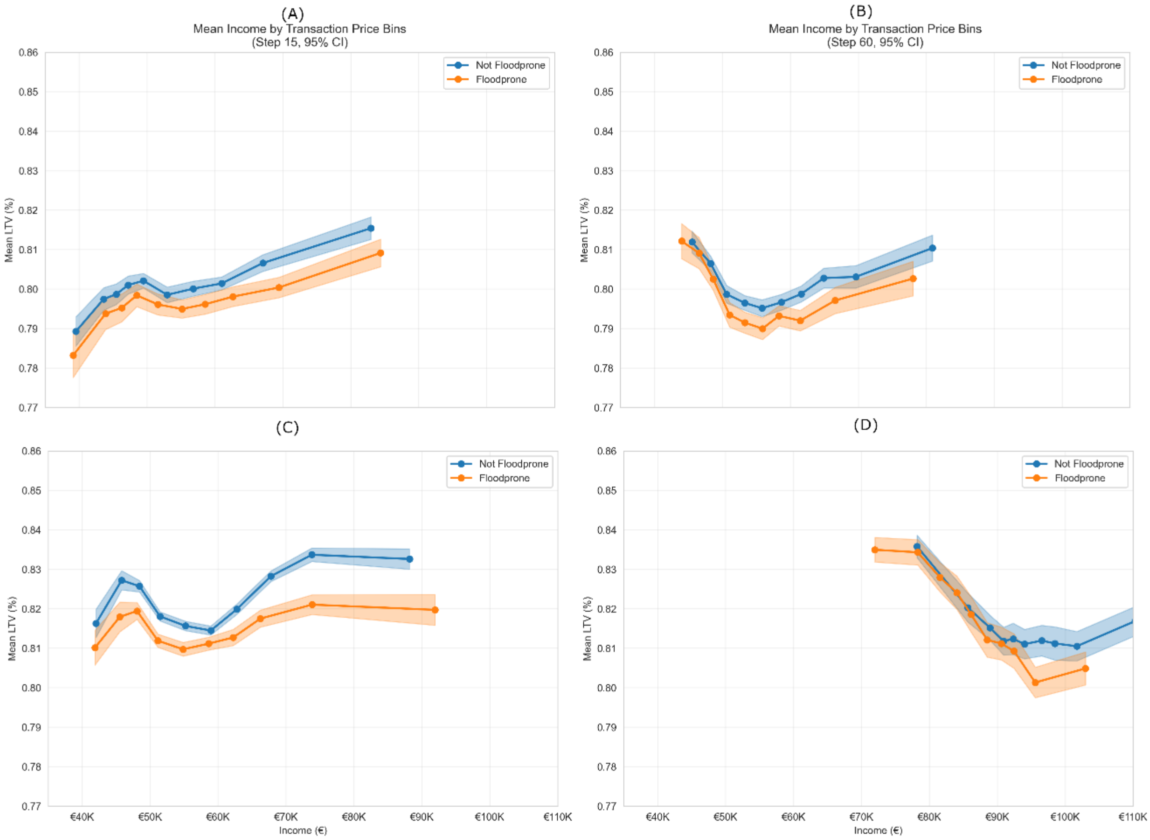

Figure 17: A and B are low demand while C and D are very high demand. All scenarios are with sustainable mortgages and biased risk perception. This shows that as demand increases, even with sustainable mortgages, more risky loans are pushed towards higher income groups, showing that mainly low income groups will be affected by sustainable mortgages with riskier loans.

This valley likely plays a significant role in the functioning of credit restrictions, as these groups must pay higher interest rates than those below, where houses have lower interest rates. Still, this will disproportionately affect lower-income groups, as credit restrictions tighten due to higher interest rates, forcing them to pay higher mortgage payments. In contrast, those in the higher-income group will experience little effect, as they require less outbidding to secure a home, resulting in only a slight increase in credit needed due to increased competition in their section. While this difference is notable, it is still relatively small. However, it mainly signifies a larger potential issue in the market.

Another significant contributor to changes in housing prices is credit restrictions, which are tightened when interest rates increase, thereby decreasing housing prices in flood-prone areas. This results in the ordering as previously discussed, with the lower-income households being more credit-restricted, eventually forcing them towards these homes. It can therefore be argued that these regulations will significantly affect how this ordering process takes place. Firstly, interest rates will significantly reduce the credit available for these homes, thereby lowering their prices. Here, demand mainly acts as a catalyst via the LTV ratio, accelerating the change. Afterwards, prices stabilise, and ordering will follow the credit restrictions set by NIBUD, as higher-income groups are better able to obtain credit.

Overall, this will result in no actual change in the maximum monthly mortgage payment, with similar income groups in safe and unsafe housing having comparable monthly payments. As such, across all scenarios, the actual payment difference between the sections is small to non-existent; if it differs, the maximum is only around a 1% decrease in monthly payment for a flood-prone property. Instead, the flood risk premium only affects the amount buyers can bid, thereby decreasing the property's value. Therefore, under the regulations, the flood risk premium does not seem to affect the monthly mortgage payment. Only the interest rate a household has to pay and the value of the home.

9.2.1 sustainable mortgages with the water label

By combining sustainable mortgages with the water label, both become significantly more effective. Even under the very high-demand scenario, it now takes around 8 years for prices in the flood-prone area to fall below those in the safer area, far quicker than the 27 years in the same scenario without sustainable mortgages, but at the same pace as the same scenario with only sustainable mortgages. This can be seen in the decreases in the LTV ratio and the amount of credit taken on by buyers, which are greater for the flood-prone houses than in the same scenario without the water label. Once again, the higher-income groups are the most important drivers, decreasing their bids, while the lower-income groups tend to take on higher credit, still mostly able to bid less due to the flood risk premium. It can therefore be argued that the increase in the effectiveness of the water label mainly comes in the form of higher-income groups now being able to better integrate their risk perception, rather than the overall market being able to make use of this change.

Across all other demand scenarios, a similar story emerges: the combination of these two policies is more effective than either alone. Even in the low-demand scenario, the price integration of flood risk is larger, with the change occurring far more quickly than in scenarios with only one of these policies active, taking 12 years rather than 18. Demand, therefore, seems to play a smaller role in increasing the speed at which risk is integrated into prices compared to the one-policy scenarios. The only major difference in bidding strategy is that, unlike when only the water label is introduced, the LTV ratio difference now increases between flood-prone and non-flood-prone housing mortgages as demand increases, rather than decreasing when there is only a water label. The only other result of note is that no actual large changes in the prices for non-flood-prone homes were seen, compared to only sustainable mortgages.

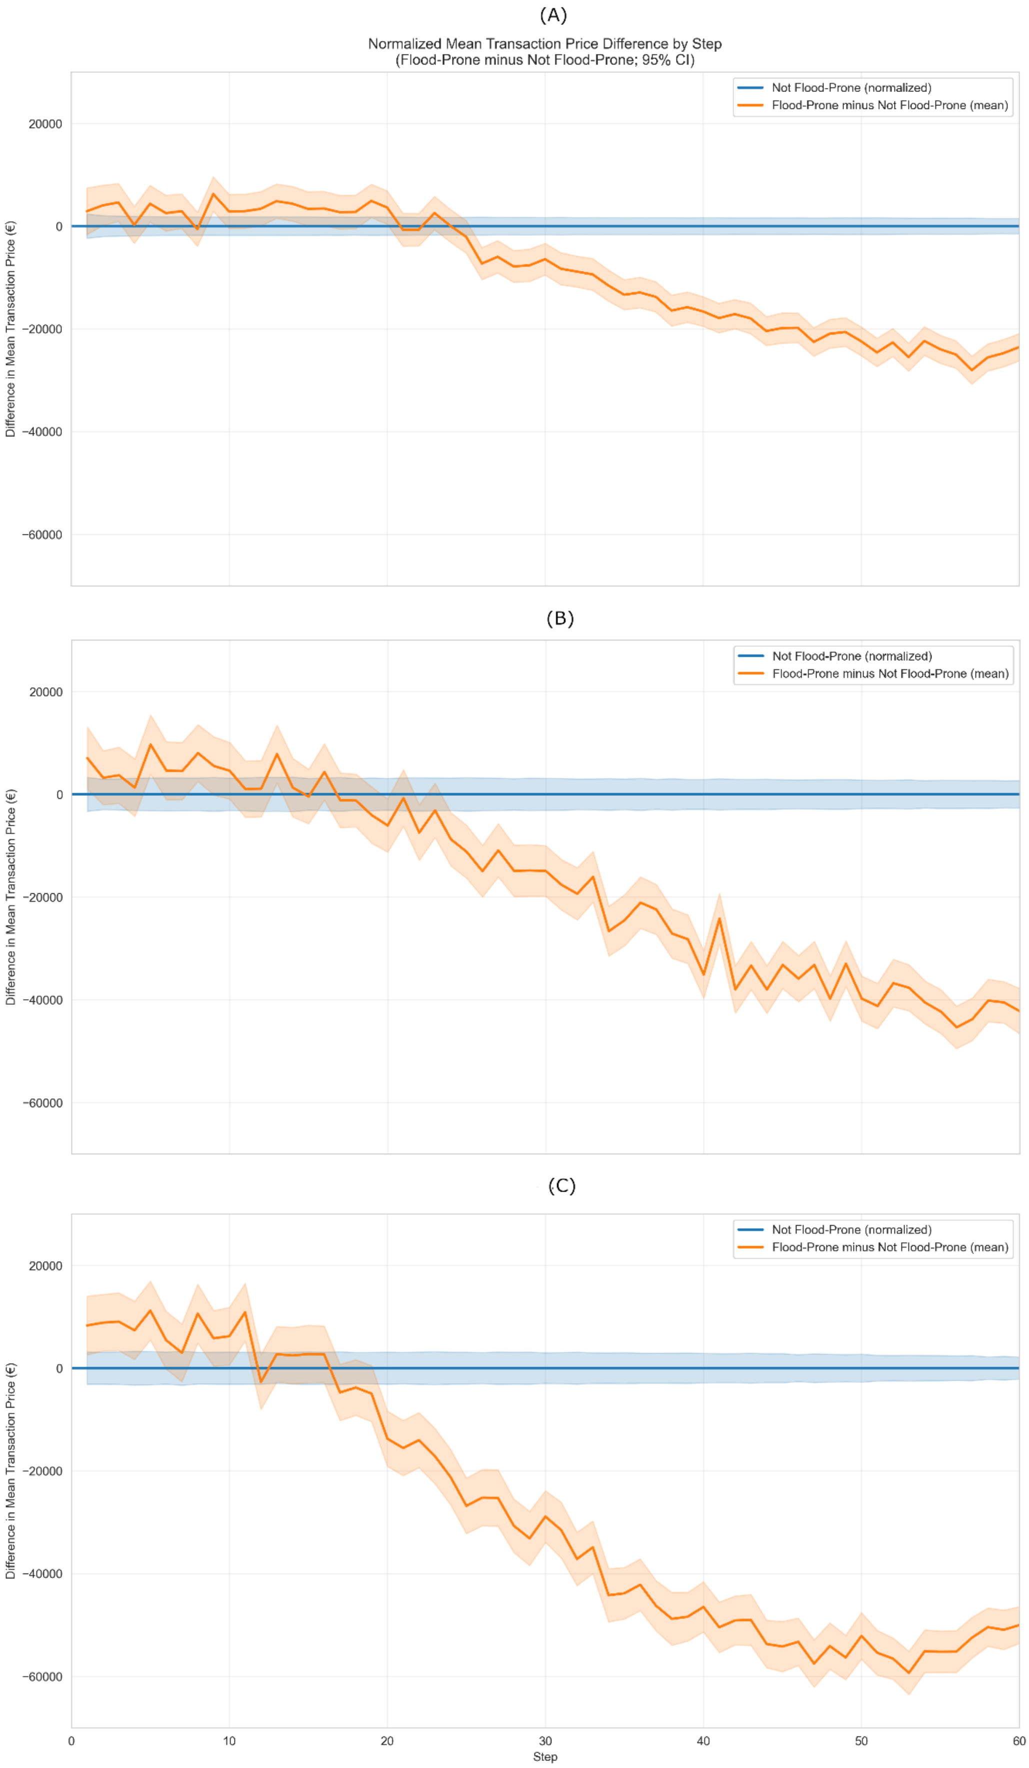

Figure 18: Pricing integration of sustainable mortgages with biased risk perception. A is with low demand (1x), B is with high demand (2x) and C is with very high demand (3x). The speed of integration increases as demand increases, but compared to the label, the amount of pricing integration also increases.

9.2.2 Effect on income

All these effects cause significant changes in housing prices: where previously there was no difference between the sections, the flood-prone areas have decreased by around €15.000, while those in the safe areas have increased by €22.000 compared to the standard high-demand scenario, a difference of around €37.000. This difference decreases significantly when demand is also lower. Within the low-demand scenario (€13.000) and the very high-demand scenario (€50.000). Incomes also shift by a significant amount, as the houses in the standard scenario were slightly higher in the dangerous areas; as a result, the incomes there were also higher. With the introduction of the flood risk premium, incomes change places across all scenarios, with incomes mainly increasing in the safer areas. Lower-income groups will be forced into unsafe areas, as prices will decrease to reflect the interest rate change across all demand scenarios, with demand only changing the speed at which this shift occurs.

9.3 Introducing the slow integrators

Sustainable mortgages can significantly help reduce the rates a bank can offer. By focusing on houses in safer areas, they avoid homes with a higher risk of default. This gives them a competitive advantage compared to other banks. However, this advantage needs to be significant enough before it can be interesting for banks to adopt this strategy.

Suppose only one out of three bank accounts quickly integrates the flood risk. In that case, it is expected that the banks will gain an advantage if demand increases, as the rate advantage over other slow integrators will grow as the LTV ratio rises. Despite this, the quick integrators are still able to effectively capture the safer houses early on, during low demand, when this advantage is at its lowest, while also avoiding the riskier homes. Despite this, it does not significantly impact the overall transaction volume of the quick integrators or the average total transaction value during the introductory period for all demand scenarios. This can be explained by the fact that only a small portion (around a third) of these more valuable homes are at risk.

Early adopters in these scenarios find no competitive disadvantage compared to the slow integrators. Early on, they can capture mostly safer homes and therefore enjoy a double advantage: no risk of losing actual transaction value while also ensuring their portfolio decreases in risk. With the introduction of the flood label in this scenario, minimal shifts occur, making it especially advantageous in the very high-demand scenario. There is no significant difference with a water label compared to the other scenarios.

A slower integration speed, however, does not significantly decrease the rate at which the risk is integrated into housing prices in the high and very high demand scenarios. With the price shift taking only slightly longer, taking around 15 years, rather than the previous 10 years in the very high demand scenario, with similar numbers being seen in the high demand scenario. This issue will shift when low demand comes into play, with the price shift taking increasingly longer as integration speed decreases, making it 15 years longer when there is only one quick integrator. But once the water label is introduced, this difference drops significantly, towards only 5 years.

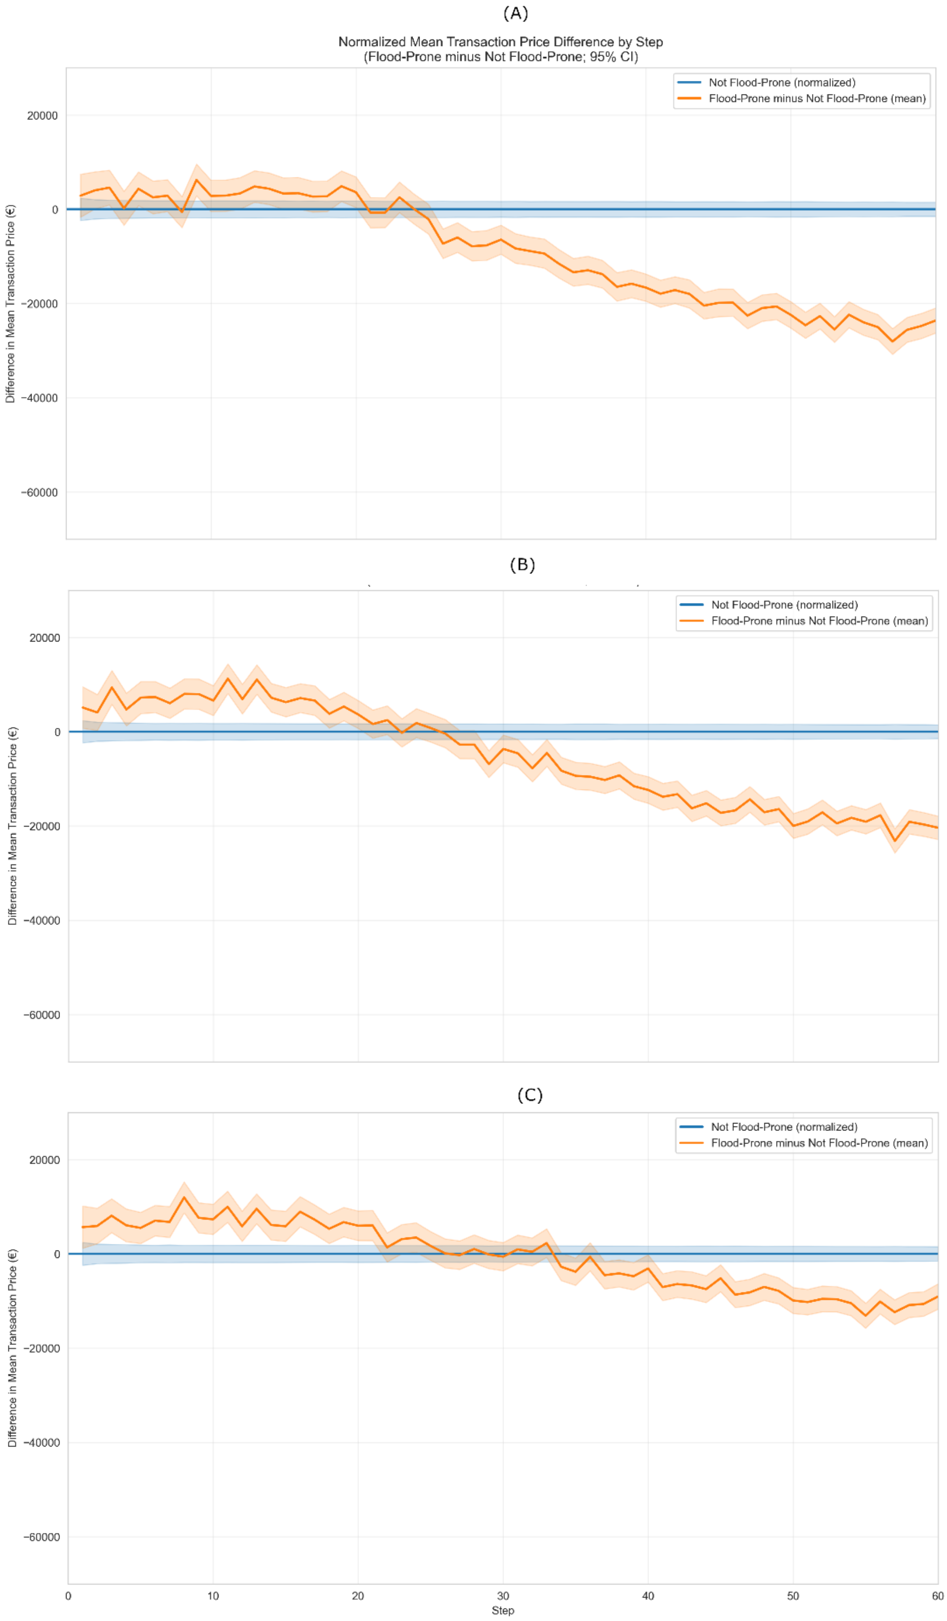

Figure 19: All scenarios are with sustainable mortgages, but differs per amount of fast integrators (2 vs 10 years to change credit risk assessment practices) under high demand. A is all banks, B is 2 out of 3 banks, while C is with 1 out of 3 banks taking flood risk into account during credit risk assessment. The speed has an especially quick drop off after scenario C, showing that at least 1 out of 3 banks would need to take flood risk into account before integration speed decreases.

9.4 Sensitivity analysis

To ensure that the results hold given the uncertainty in the model, a sensitivity analysis has been performed, for which the larger explanation can be viewed in appendix D. Here the appendix will be summarized. To analyse the sensitivity of the model, One-factor at a time analysis was performed on the base interest rate, with a range of 10 percent around the initial value of 3.5%. This resulted in very little changes, which can possibly be attributed to the range chosen. The results mainly pertain to the water label, which under reduced rates consistently face a decrease in the risk integration, which could be the result of more lower income groups being introduced into the model. As these groups face stronger credit restrictions it is possible they will face more difficulties in purchasing a home under very high demand scenarios. This indicates that some non-linearity might not yet be explored in the model. Especially combined with the results that in highly flood prone areas the risk is getting integrated, just less than expected, with for example lower than expected risk, the integration might further suffer.

10 Discussion, Conclusions and Future Work

10.1 Broad implications of sustainable mortgages

Demand plays a significant role in the effectiveness of both sustainable mortgages and the climate label. With sustainable mortgages, demand significantly accelerates the integration of flood risk prices, eventually leading to a price decrease of around 7% for current flood risk. This price difference helps increase the use of financial tools by refinancing the mortgage to allow for flood adaptations that reduce flood risk, thereby decreasing this loss in value. High- and very-high-demand areas will be able to achieve this pricing difference more quickly. Once this difference is significant enough, these tools can make use of a double effect. Not only can the interest rate be reduced due to lower loan risk, but the lower rate also allows the property's value to increase, improving its resale value.

However, given the angle of lower demand, this price integration is significantly reduced, thereby making it more challenging to create space for these financial tools to become relevant. This is due to two reasons: first, low-demand areas will experience lower price growth, making it more challenging to incorporate these flood-adaptation measures into the increase in value.

Secondly, it takes far longer for these differences to become significant, thereby delaying the possible uptake of these financial tools. It is therefore important that other policy tools address this gap, as sustainable mortgages may not be sufficient to achieve this shift quickly enough. These could take the form of subsidies targeted at low-demand areas to increase the uptake of flood adaptations.

Credit restrictions also play a significant role in helping households purchase a home; in short, they ensure households pay only a certain percentage of their income towards their mortgage. These restrictions, therefore, lower maximum bids, thereby reducing the LTV ratio of a mortgage, especially in high-demand scenarios. This LTV ratio becomes more important, as Caloia et al. (2023) found that these ratios play a significant role in increased credit risk during floods, for which sustainable mortgages could serve as at least a temporary solution. As credit restrictions tighten amid rising demand, the LTV difference widens with sustainable mortgages when comparing safe with unsafe homes, as buyers rely more on credit to secure a home. This can create a natural mechanism that reduces the LTV ratio difference between areas, as they will experience greater flood risk premium increases due to their greater reliance on credit.

This could reduce overall credit risk during floods and the volume of defaults during these events. But this does come with a significant caveat: this difference occurs only once, during the integration phase, allowing for a grace period before flood adaptations should be adopted more readily. This issue is flipped when looking at the water label. As the demand increases, the LTV ratio difference starts to decrease. From this, it is suggested that when the climate label is introduced, it should be combined with some uptake of sustainable mortgages, as it could unnecessarily increase overall credit risk. Otherwise, due to the low integration speed, high-demand areas could face unnecessarily high credit risk.

10.1.1 Slow versus fast integration of flood risk

Overall, little difference was found between housing prices and the speed at which risk is integrated in high- and very-high-demand scenarios. It should therefore be less important to focus on the speed at which sustainable mortgages practices are adopted in the high- and very-high-demand areas, whereas in the low-demand areas, speed becomes far more important. Due to the lower demand in certain areas, the risk integration into housing prices takes significantly longer, delaying the potential usage of adaptations, and creating unnecessary credit risk. It is therefore recommended that the speed of climate risk integration be taken into consideration, as it will affect the lower-demand sector of the housing market, ensuring that these areas are not left behind.

10.1.2 The effect of a climate label

Because the value of a home is initially calculated with little to no integration of flood risk, housing prices are overinflated relative to their true value. Even when introducing a biased risk perception, due to the low damage fractions and significant biases, the bids decrease only slightly across all scenarios. However, once the climate label is introduced, the story becomes significantly more complicated, with demand feeding into two mechanisms that will eventually dictate the speed at which flood risk is integrated into housing prices. Firstly, the increase in demand will help speed up pricing integration, as more transactions ensure more bids, thereby accelerating the integration of risk. Buyers can still bid successfully on houses without being overly restricted by their credit, allowing them to buy a house while lowering their bid. Eventually, around 3x demand, these credit restrictions become much more significant, with risk pricing integration only starting later, once demand for the area settles, reducing the flood label's effectiveness. This does come with a significant nuance. If the damage increases are large, the price integration still happens, just more slowly. It is therefore recommended that, when introducing a flood label of some kind, especially in higher-demand sectors, it may be combined with sustainable mortgages. Moreover, much consideration should be put into homes that are most at risk in the future, that still have a climate label, as these homes will face the largest decrease in integration speed.

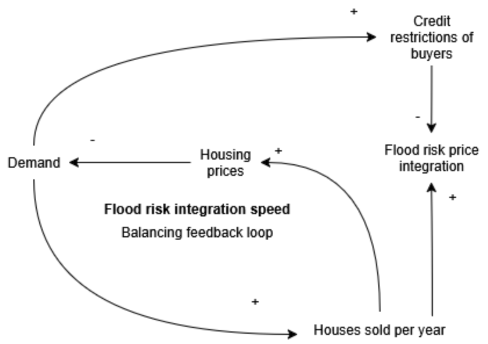

Figure 20: Causal diagram of flood risk integration speed

10.1.3 The impacts of the flood label and sustainable mortgages on wealth inequality

Due to price decreases, flood labels, and differences in interest rates, flood-prone areas within the model are attracting lower-income groups, worsening wealth inequality. This is significant for two different reasons. Firstly, as the most vulnerable groups will be pushed towards these homes, it can make it more difficult for adaptations to take place, as these groups have less funds available to pay for them. This can be put in opposition to the claim that these adaptations also help solve the wealth inequality problem, requiring pricing integration before this will be feasible. The second primary reason is that the pricing integration is indeed significant enough to support adaptations; however, it has important caveats that need to be considered. Firstly, these shifts occur in tandem with rising wealth inequality, as the so-called "game of musical chairs" shifts from higher- to lower-income groups. In turn, it can be argued that for the sustainable mortgages practices to be effective, wealth inequality will inevitably increase. Secondly, as climate change worsens, homes will become more vulnerable to floods. These policies could therefore have the opposite effect, pushing the most vulnerable groups into the most vulnerable areas. Moreover, not all areas can be saved, as discussed in the introduction, potentially leaving them with a home that is now worth far less, making it more difficult for them to move out. Thirdly, there is the issue of the LTV valley. As these groups are more reliant on credit, they are prone to take on riskier loans, thereby incurring large financial risks. In combination with LTV being a primary driver of financial losses according to Caloia et al. (2023), floods can hit these incredibly vulnerable communities, resulting in greater financial losses and further damaging them. Finally, because these price shifts take a long time to become effective, assuming both sustainable mortgages and the water label will be used, low-demand areas will have the most difficulty achieving the real price. This will leave low-income groups in these low-demand areas especially vulnerable. Therefore, they will require the most assistance and focus, as these policies are the least effective for them.

For these reasons, it is necessary to focus on other policies to ensure that these groups are not left out and can receive aid to pay for these adaptations before large price shifts occur. On top of this, assisting lower-income groups in moving out of the most dangerous areas will be essential, as they will most likely not be able to pay for the move. Finally, as banks are a large part of the information-gathering process, it is recommended that banks also work pre-emptively, primarily informing them of the risks, while also assisting these groups to develop potential solutions and help them avoid the most dangerous areas. All these policies are required to work before significant climate damage occurs, as the issues with sustainable mortgages and the water label are that they take time to cause change. Moreover, a strong focus should be placed on low-demand areas, as they are also the most vulnerable. But because the pricing integration time is the longest, these policies also have the largest window for use. Pushing vulnerable buyers away from these areas is therefore recommended.

Sub question 1:
What type of flood label characteristics would be most relevant for mortgage lenders and how might this impact its design?

From understanding the design of current iterations of flood labels and discussions from the Dutch banking industry it became clear what qualities they prioritized. Most importantly, the label needed to be based on property-level data, ensuring that each home would get an accurate assessment. At the same time, the information would need to be simple and easy to understand, ensuring low information costs.

With inspiration from previous iterations of the energy efficiency label, expert opinion on the water label, on top of previous issues of spreading flood information, an iteration of the water label was constructed. The most important fact learned from the energy efficiency label is that households understand the financial implications of buying a flood-prone property. Each label will therefore require a significant increase in damages. As these damages follow an S-curve, the labels would ideally follow this S-curve. Moreover, flood events are often forgotten, leading to prices returning to normal levels around a decade after a flood. Therefore, these needed to be frequent enough not to be forgotten. If a household buys a home with a 1/100-year chance of flooding, around 23% will experience such an event, which will hopefully lead to enough people being aware of the consequences. Coincidentally, these labels also align with the S-curve, making the 1/100-year flood label ideal for this model.

Sub Question 2:
To what extent do Dutch and European banks currently integrate flood risk into mortgage pricing and how does this relationship function?

First, how interest rates are set has to be explained. Banks use interest rates to ensure mortgages are profitable. To do this, they take into account the chance that the loan will fail (the chance of defaulting), the amount of money they will lose in case of a default (the loss given default), the ECB rates, and their profit margins. Furthermore, the banks treat the house as a liability in the event of a default to ensure they can recoup some of their costs. Therefore, if the loan is high relative to the home's value, it will increase the risk of default.

Floods directly influence this credit risk, as homes sustain significant damage after flood events, decreasing their value and increasing the risk of the loan. Furthermore, as the surrounding areas are also affected by this disaster, economic activity decreases, leading, for example, to job losses, further increasing this credit risk.

Sub Question 3:
How do households calculate their utility when buying a home?

Due to a weak link between demand and supply, the housing market operates very differently from other markets: when prices rise, it does not respond by building more homes, due to a weak supply side response. What this results in, as housing is an essential good, credit becomes an important factor as buyers are required to maximize their bid to secure a home. With credit restrictions set by NIBUD, these can be used to set a household's budget. These credit restrictions force buyers to be able to lend only a specific amount based on their yearly income and the interest rate. Other factors also come into play, such as the deposit, further increasing the bid. With this deposit, the LTV ratio can be decreased, lowering the interest rates.

Households experience this quite differently, as utility from a home is treated like a monthly expense. Thus if the interest rate increases, the household will respond by keeping their monthly payment the same, but their bid will decrease. This results in buyers only being able to pay a certain monthly payment, which does not react to any changes in interest rate, only the bid responds to this.

Sub Question 4:
How can flood-risk-sensitive interest rate be represented in an agent-based model of the Dutch housing market?

To isolate this credit risk from the compounding variables after floods, the fractional increase in default chance was taken from the uninsured population that experienced damages. This ensures that only the default chance increase related to these damages is taken into account, but leaves out other factors. In turn, this simplifies the interest rate calculation, giving a lower bound estimate.

When interest rates increase, through the credit restrictions, the maximum loan is decreased, thus lowering the bids. Buyers respond by keeping their monthly payments the same, as this is seen as their utility, with the interest rate payment being treated like a black box. This relationship automatically decreases the LTV-ratio, thus also decreasing the credit risk.

The monthly payment can only be decreased through the flood label, directly decreasing the amount of monthly expenses a household is willing to pay.

Sub Question 5:
How could the interaction between a water label and flood-risk mortgage pricing affect housing prices and LTV ratios?

This question can now be answered by combining the previous two sub-questions. As the current price was set with a normal interest rate in mind, prices in flood-prone areas are too high for their credit risk. Once interest rates rise, the bid falls, as buyers can obtain less credit than before, leading to a lower LTV ratio. Moreover, as flood risk is correctly integrated, it will also lead to lower interest rates for safe properties. This will lead to price increases as credit becomes more available to them.

What effects might the introduction of the water label and sustainable mortgages have on the integration of flood risk into housing prices?

Starting with an ABM housing market model that focuses on households' flood risk perceptions, credit restrictions were added to reflect how banks will adjust their credit management strategies. This expansion was based on the fact that the credit restrictions, interest rates, demand, and income seem to be the largest drivers in predicting the prices of housing (Damen et al., 2012), which have been used in housing market models to model the housing prices accurately. While this expansion does not fit neatly into the models, they have exerted a key influence on the design, particularly regarding how these credit restrictions should function, with the results from Van der Drift et al. (2022) serving as a basis. These have then been simplified and recreated in such a way that the behaviour of buyers should be similar and sufficiently recreate the intended behaviour.

The model was also expanded to include the flood labels, as discussed in sub-question 1. These have been based on discussions with experts in the field to ensure they accurately portray a possible flood label, while remaining simplified enough to fit within the model's constructs. As such, it was decided to use only one timespan, despite a home being vulnerable to different possible floods over varying timespans. Finally, both sections were linked using default rate data and discount rates to estimate the potential interest rate increases these homes would face. This data was then used to identify the at-risk population from the rest, enabling more accurate credit risk estimates. The final result is a model that links the flood label to a home's credit risk, allowing it to predict the home's real price.

Using this model, different demand scenarios were tested, combining the flood label and sustainable mortgages practices to understand the interactive effects of these two policies. It was found that the flood label would not achieve its intended effect, making pricing integration difficult during very high demand. This is due to the credit restrictions being the primary factor in determining a buyer’s bid, rather than their risk perception, as the desperation to buy a home outweighs the flood risk. While in low demand, the pricing integration was slowed by lower sales, reducing the risk integrated into the pricing. These two factors make introducing such a label more difficult, necessitating a combination of policies to ensure the label has its intended effects in these scenarios, as it creates a U-curve in demand rather than demand constantly increasing the difference. Sustainable mortgages, however, seem more successful because it uses these restrictions to its advantage. This allows the policy to be more successful across all demand scenarios, significantly increasing the price difference and integration speed as demand increases. If uptake of sustainable mortgages is slow, the effects are expected to be relatively low, with mainly the lower demand experiencing the most considerable slowdowns.

10.2 Scientific relevance and future research

Agent-based models have been commonly used to study the relationship between an individual's risk perception and their behaviour in the housing market. In these models, the households use their own risk perception to decide to migrate towards safer areas, with a large focus on the agent’s risk perception of floods through different types of risk perception (De Koning & Filatova, 2019; Tierolf et al., 2023; Tonn & Guikema, 2017) or the effects public adaptation on housing choices (Mutlu et al., 2023). The previous literature, therefore, mainly focuses on the effects of information on the buyer’s decision-making. However, none of these models have taken the possible ramifications of financial constraints on flood risk perception into account, nor have they integrated the increased credit risk that flood risk poses to lenders. This study aims to fill this gap, aiding policymakers in their decision-making regarding the integration of climate change risk into housing prices, ensuring these prices more closely reflect the real price and that houses' risks are adequately integrated into their value. Its main finding is an exploration of the complex interplay between the introduction of flood risk information into the housing market for banks and buyers, and demand and financial constraints.

The effects of financial constraints are especially relevant for the Dutch case, given the incredibly overheated market (Langen, 2025), which may make it challenging to introduce a flood label. According to the model's findings, this heightened demand increases buyers' desperation to secure a home, making it more difficult for risk to be adequately integrated into the housing market. However, initially, the opposite is the case: as houses sell more quickly as demand increases, the integration speed actually increases. Low and very high demand could therefore significantly slow adoption speed. This, in turn, brings the discussion back to the energy label, which has had significant difficulties in being adopted, with the literature’s current primary focus on information dissemination and the buyers' requirements of the label being the most significant contributors in aiding its adoption (Brounen & Kok, 2011; Chareyron, 2024; Gerassimenko et al., 2024; Oerlemans et al., 2025; Stangenberg et al., 2020). It therefore follows that financial constraints could become a significant issue later on, with the caveat that no prior literature has yet found a link between financial constraints and the integration of energy labels into pricing. These findings, therefore, need to be taken with a pinch of salt; instead, they may aid potential new avenues of research. Even if these new changes fail to materialise in proper pricing integration, heightened demand could still play a role. As the focus of this study was to understand and explore the issues of financial constraints in conjunction with the housing market and flood labels, empirical evidence is lacking to support these findings.

This result was achieved by combining the results from other housing models that focus on these credit restrictions (Van Der Drift et al., 2022; Madsen, 2012; McQuinn & O’Reilly, 2007), and combining these with the model from Mutlu et al (2023) and De Koning & Filatova (2019), which instead used the assumption that households pay a fixed percentage of their income towards housing that Damen et al. (2016) also used. As these models fundamentally varied in their assumptions, great consideration was given to ensure that these two functioned as naturally as possible together, with mainly the assumption from Van der Drift et al. (2022) being used to build on top of, which stated that while lower income groups are credit bound, higher income groups rather have a preference to pay a percentage of their income, combining both ideas.

While this assumption was not fully met, it was decided to reduce credit restrictions for higher-income households, ensuring that their effects on buyer behaviour could be understood. This novel approach shows that these assumptions can be easily introduced into similar models, enabling new approaches to understand the effects of credit availability on the housing market. As the goals of these models are significantly different from each other, with the models focusing on credit restrictions mainly being used to predict the housing prices, while the climate risk models focus on the effectiveness of policies, the combination of these allows for new types of policies to be explored within the literature. This research also examines the effects of a home’s characteristics on mortgage credibility, an area that has not yet been explored. Moreover, it shows how these restrictions and bidding processes interact. However, as this research was more exploratory in nature, the effectiveness of these policies is unknown, requiring more empirical evidence to validate these results.

Previously, the effects of the water label and sustainable mortgages on LTV were discussed: the former increases LTV, while the latter decreases it as demand increases. As was discussed in Caloia et al. (2023), these effects have been linked to increased credit risk during flood events. Because of the slow uptake of the energy efficiency label, it could take time before the water label becomes widely accepted, creating scenarios in which uptake of the water label is significantly slower than that of sustainable mortgages, and vice versa due to the potential moral hazard which was touched upon in section 2.3.

10.3 Societal relevance

The findings also suggest that sustainable mortgages can help alleviate these financial constraints, as it instead uses them to their advantage, realizing the housing price changes required, even if banks are slow to adopt these changes. This can ensure that even if banks struggle to adopt these new credit management strategies, which require significant time investments due to the current poor availability of information of an individual house's flood risk, actually taking the time to ensure that the data is accurate will not lead to significant slowdowns of pricing integration, except if demand is low as it decreases the integration speed. Furthermore, it will allow banks to invest in housing more safely. Not only because the price better reflects the real price, but also because of a more accurate risk assessment, which later will create more financial space for investment in flood adaptation measures. This will also have broader consequences, ensuring that the overall housing stock adapts to these changing conditions and decreasing overall risk in the mortgage market. Moreover, it ensures that their clients are informed about the home's risks, helping them adjust their bidding and search strategies, allowing banks to manage their overall portfolio naturally.

Even insurers might adapt to these policies, as the overall decrease in flood risk, driven by potential increases in adaptation and improved pricing integration, means fewer payouts are necessary. It will allow insurance companies to more easily inform households of their overall flood risk, while also ensuring that other measures are available to reduce this risk, lowering the likelihood that households will be dropped by them and decreasing overall unexpected flood-related expenses.

As the water label is currently being developed by the Ministry of Infrastructure and Water (Tieman, 2025), the results can also influence their decision-making. Based on the results of this thesis, it is recommended that their highest priority be to ensure that the water label uses high-quality data. This will, in turn, help banks properly isolate these homes in their datasets, ensuring they can set their rates appropriately. While this strategy will be slower, no significant delay was observed in the pricing integration when banks were slow to adopt different credit management practices. The opposite was, however, true for the water label. If banks struggle to isolate flood risk from their data because the water label is not supported by the underlying data, they may avoid setting appropriate risk premia, leading to unnecessary credit risk in very high-demand areas. It is therefore advised to introduce the water label gradually and expand the areas as the data improves.

Lastly, prospective homeowners will be better informed about their flood risk and may be able to take further steps through adaptation measures. This will help them be better informed of the risks they are taking on, helping them to change their choice of homes. Moreover, the majority of homeowners and buyers will also receive lower interest rates, increase their value and ensure that riskier homes get the correct rates. While outside the scope of this research, this could lead to lower damage to a household after a disaster, allowing them to more easily bounce back, decreasing the default rate, and increasing the overall value of their home.

11 Limitations

This model treats the Dutch market as a single, unified market, rather than as a collection of separate regional markets with different demand pressures, housing stocks, and flood risk profiles. In practice, the Dutch housing market is highly segmented: the Randstad operates under vastly different conditions than, for example, the shrinking regions in the periphery. By abstracting away from this spatial heterogeneity, the model cannot capture how local market tightness interacts with flood risk integration. This interaction as the results suggest, matters considerably. This simplification, along with several others, should be kept in mind when interpreting the findings. In this section, the limitations of the model will be discussed, as some of these have already been discussed previously within their relevant sections, mainly in the formalization of the model. This will provide a short overview of the different issues, primarily related to the data, and highlight their impacts on the results.

11.1 Model Limitations

11.1.1 Economic modelling

Firstly, many effects have been strongly simplified. This includes the way interest rates have been calculated. In reality, the bank considers many factors in assessing a loan's risk. This includes, for example, the loan-to-income ratio, the age(s) of the borrower(s) or the security of income. These can significantly change the mortgage rate. For example, lower-income households pay more for their loans because they pose higher risks, reducing the maximum loan they can get (Campbell & Cocco, 2015). On top of this, the greater difficulty of saving for a deposit results in a significant reduction of their buying power (Van Der Drift et al., 2022) As a result, the actual buying power of lower-income households is expected to be significantly overestimated in this model. This is because the default probability has a significant effect on the interest rate; if it increases with the odds ratios for flood risk, it could further decrease their buying power. Future research should focus on creating effective policy solutions to reduce the potential ramifications of this effect.

This issue becomes larger when considering how the pricing integration works, as it reduces overall prices by a percentage of home value, potentially making it more difficult for these policies to work for lower-income groups living in lower-valued homes. Due to flood adaptation's reduced impact on wealth inequality, focusing on groups that could be left behind becomes more important.

This is also reflected in the substantial simplification of the LTV. Usually, these are calculated through a so-called fire sale, in which a certain percentage of a home's value is lost in the sale, plus a fee. However, the simplification of the deposits leads to the Loss-Given-Default being overestimated for higher-income earners. Overall, the model overestimates the buying power of lower-income groups, potentially leading to underestimated price decreases.

These issues can be attributed to the simplification of the saving process, which is currently missing in the model’s iteration. Ideally, the model should draw more inspiration from Van Der Drift et al. (2022) and their conclusions, by expanding on the saving process that allows household agents to accrue wealth more naturally, it will not only allow the results from Van Der Drift et al. (2022) to be more naturally integrated into the model, but will also allow the interest rates to more accurately be estimated.

Moreover, improving the modelling of savings could have a large impact on the results, potentially causing a larger reliance on deposits, which would in return reduce the effects of credit restrictions. This will cause an increase in richer households' potential to buy themselves into a safer home through overbidding, resulting in a faster adoption of the real price.

This, in turn, can also help fix the awkward spot the current version of the model is in, as it resides between the economic modelling of Van Der Drift et al. (2022) and the behavioural economics from Mutlu et al. (2023) and de Koning & Filatova (2019), from which this model is originally from. Many of these issues have already been discussed in 6.3, where a great deal of the workarounds was mentioned, stemming from a single fundamental issue. The actual end state of this model would be closer to Dubbelboer et al. (2017), in which the increase in monthly payments informs households of their flood risk.

11.1.1.1 Reduced competition

Due to the absence of certain lower-income groups, the buyers in the model experience less competition in the lower market segments. This reduces the overall need for credit for this group, overstating the effects of credit limits. The effect is assumed to be low, as credit limitations are less significant in low demand. Moreover, the effects of these were also verified in scenarios where this forms less of an issue.

11.1.2 Behavioural economics

In combination with the water label, the current risk perception is entirely rational, even if buyers do not perceive it that way. Instead, they could have a strong bias to avoid flood-prone areas, which can be better explained through prospect theory. Here, buyers might choose to avoid taking on risk (Osberghaus, 2016), thereby increasing the effect of a water label. In combination with the changes to the bidding strategy, in which households adjust their bids to maintain a certain deposit to later pay for adaptations, based on the information provided, much of which can be informed by this behavioural theory. This would fix one of the identity crises this model has.

11.1.2.1 Changing of market conditions

Because flood adaptation is missing from this model, the effect of risk integration is expected to be significantly overstated. As prices decline, they create more substantial incentives for homeowners and buyers to invest in these solutions, thereby increasing home prices. As this thesis was more of an exploration of the effectiveness of policy solutions on risk integration on prices and whether these solutions are sufficient to introduce these adaptations, further research should focus on whether homeowners will accept these adaptations. At the same time, it is necessary to consider whether these adaptations will sufficiently reduce credit risk in the event of a flood.

Currently, the model also lacks a capacity market. This can significantly affect the overall timescale of the pricing integration. The effects of this market are commonly explained by the game of musical chairs, in which adding homes to the housing market lowers prices across the market. This is due to the market's overall demand ratio decreasing. However, this part has already been taken into account and can be explored by combining the different demand scenarios. Moreover, this anecdote assumes that housing characteristics have already been adequately incorporated into market pricing and that this will not hold if a new housing characteristic is added. Instead, it should be viewed through the lens of its effects on pricing integration speed. These conclusions largely depend on the overall regulations governing flood safety and the share of safe homes in the overall market. If this share stays the same, nothing will change. If the share increases, buyers will be able to bid on more safe homes, making it easier for them to find homes in safer areas and decreasing demand for unsafe homes. Because of these reasons, under the assumption that the Dutch government will enact regulations to create safer homes, the integration speed of risk should increase.

11.2 Data limitations

Another major difference between this thesis and Mutlu et al. (2023) is that the data used there is non-anonymized, this is important due to the very chain through which the credit restrictions act. With location being one of the most important predictors of housing prices, while also being incredibly difficult to get right due to its many compounding factors (Bourassa et al., 2010), missing this makes the actual wealth inequality significantly down stated within the used dataset. With the most expensive homes usually being located in the Randstad, they are also commonly the most at risk of floods (Bosker et al., 2018). Removing this important boundary changes a lot of the market dynamics, as these geographical boundaries creates separate markets, instead of united as in this model. This may result in higher difficulty in finding safer homes, reducing the actual amount of competition between safe and unsafe homes, slowing down the pricing integration.

Currently the dataset only includes single-family homes, which does not cover the entire housing market. If these were added, the effects will be reducing the overall flood risk in the data, as single-family homes are also the most flood-prone due to their lower height (Slager & Wagenaar, 2017). Leaving the flats out of the market, therefore, means that reaching the real price should happen quickly. These effects can be significant as the overall Dutch housing market consists of 2/3 single-family homes (CBS, 2013), giving the safer homes a larger advantage in the market.

With the calculations of the Loss-Given-Default (LGD) some simplifications were made that limit the accuracy of this. Looking at equation 13, it does not take into account the differences in discount factor the labels can have, instead it is the same number for each of them, resulting in inaccurate estimations of the interest rate. These can have large consequences, creating a floor for the flood premium, rather than a smooth line, resulting in issues with the hedonic regression. The result of which means that the hedonistic regression takes a while before it finds a fit, on top of this it decreases the accuracy of the pricing estimation of the damage fraction. Within earlier versions of the model this did not take place, due to the simplification of the interest rate. This resulted in a decrease of focus on the integration speed, as it was deemed not accurate to estimate, which was taken into account within the results.

11.3 Differing modelling choices

The other one being the type of work Mutlu et al (2023) is compared to the other models, such as Van Der Drift et al. (2022), which this thesis is mainly inspired by, and what both are trying to achieve, with the former focusing on extracting buyer behaviour while the latter focuses on price prediction (ex post vs ex ante). This makes the validity of these results questionable, as many more techniques have been used in these models (e.g., Van Der Drift et al., 2022) to make these predictions accurate, which are lacking in the current model. This conclusion was taken into account while writing this thesis. As such, greater focus was placed on the mechanisms underlying the model and on how its components interact, rather than on the price predictions. Many of these reasons have been further explored in 10.2.

11.4 Changes in agent's behaviour

As these banks can advertise lower interest rates, around six basis points, they could hold a competitive advantage over other banks, thereby attracting more customers. On top of this, the bank could advertise itself as climate-conscious, thereby potentially further increasing customers' interest. But it still misses one vital feedback loop, which the model does not account for. Essentially, as banks can poach customers, they can also lend more due to increased savings, creating a positive feedback loop. This could potentially incentivize banks to invest in their own data, helping them have an edge on the competition.

Then there is the issue of information availability for the agents. Currently, the model does not account for homeowners' incentives to sell their homes. As prices shift, it could incentivise homeowners to sell and move to less flood-prone areas, thereby decreasing the value of the surrounding area. The owner could also choose to invest in flood adaptations, negating this price decrease. More research is needed on how this climate label will affect home values in these areas and the speed at which it could change them.

Furthermore, the search for homes with a water label has not yet been fully accounted for, despite being described as playing a significant role in Fairweather et al. (2024). Rather than search within their price range, buyers could also focus their search on safer areas. More research should focus on how exactly buyers adjust their search when confronted with a water label, while also accounting for what else drives their incentives to change their search.

Finally, as credit restrictions play a significant role in housing prices, the mortgage rate deduction poses a major issue for two reasons. Firstly, suppose the flood risk premium is widely used. In that case, the deduction will effectively subsidise the movement of lower-income groups into these areas, as they will be able to deduct a larger share of their mortgage payments due to the higher interest rate. Instead, the flood risk will get pushed towards the government budget. Secondly, as it increases housing prices in these areas, it reduces the monetary space created by this flood rate policy, thereby making the creation of new financial solutions less attractive. Future research should, however, repeat these experiments under different credit regulations to verify the extent of the deduction on housing prices above the flood rate, as it could significantly affect risk integration.

References

Abebe, Y. A., Ghorbani, A., Nikolic, I., Manojlovic, N., Gruhn, A., & Vojinovic, Z. (2020). The role of household adaptation measures in reducing vulnerability to flooding: a coupled agent-based and flood modelling approach. *Hydrology And Earth System Sciences*, *24*(11), 5329–5354. https://doi.org/10.5194/hess-24-5329-2020

ABN Amro. (z.d.). Sustainable Home Mortgage - Duurzaam Wonen Hypotheek. ABN AMRO Bank. https://www.abnamro.nl/en/personal/mortgages/sustainable-living/sustainable-home-mortgage.html

Atreya, A., Ferreira, S., & Kriesel, W. (2013). Forgetting the Flood? An Analysis of the Flood Risk Discount over Time. *Land Economics*, *89*(4), 577–596. https://doi.org/10.3368/le.89.4.577

Baek, I., Lee, S., Lee, J., & Kim, J. (2021). Analysis of Housing Market Dynamics Considering the Structural Characteristics of Mortgage Interest. Sustainability, 13(19), 10523. https://doi.org/10.3390/su131910523

Bani, M., Barendregt, E., Blom, M., Burgers, S., De Groot, C., Hordijk, R., Nobel, A., Phlippen, S., & Vendel, B. (2024, 21 februari). *Climate change and the Dutch housing market: Insights and policy guidance based on a comprehensive literature review*. ABN AMRO Bank. https://www.abnamro.com/research/en/our-research/climate-change-and-the-dutch-housing-market-insights-and-policy-guidance

Beltrán, A., Maddison, D., & Elliott, R. J. R. (2017b). Is flood risk capitalised into property values? Ecological Economics, 146, 668–685. https://doi.org/10.1016/j.ecolecon.2017.12.015

Bernstein, A., Gustafson, M. T., & Lewis, R. (2019). Disaster on the horizon: The price effect of sea level rise. Journal Of Financial Economics, 134(2), 253–272. https://doi.org/10.1016/j.jfineco.2019.03.013

Bin, O., & Landry, C. E. (2012). Changes in implicit flood risk premiums: Empirical evidence from the housing market. *Journal Of Environmental Economics And Management*, *65*(3), 361–376. https://doi.org/10.1016/j.jeem.2012.12.002

Blickle, K. S., Perry, E., & Santos, J. A. (2024a). Do Mortgage Lenders Respond to Flood Risk? In *Staff Reports*. https://doi.org/10.59576/sr.1101\

Blickle, K., He, Z., Huang, J., & Parlatore, C. (2024b). Information-Based pricing in specialized lending. https://doi.org/10.3386/w32155

Bosker, M., Garretsen, H., Marlet, G., & Van Woerkens, C. (2018). 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

Bourassa, S., Cantoni, E., & Hoesli, M. (2010). Predicting House Prices with Spatial Dependence: A Comparison of Alternative Methods. Journal Of Real Estate Research, 32(2), 139–160. https://doi.org/10.1080/10835547.2010.12091276

Brounen, D., & Kok, N. (2011). On the economics of energy labels in the housing market. Journal Of Environmental Economics And Management, 62(2), 166–179. https://doi.org/10.1016/j.jeem.2010.11.006

Bubeck, P., Botzen, W. J. W., Laudan, J., Aerts, J. C., & Thieken, A. H. (2017). Insights into Flood-Coping Appraisals of Protection Motivation Theory: Empirical Evidence from Germany and France. Risk Analysis, 38(6), 1239–1257. https://doi.org/10.1111/risa.12938

Campbell, J. Y., & Cocco, J. F. (2015). A Model of Mortgage Default. The Journal Of Finance, 70(4), 1495–1554. https://doi.org/10.1111/jofi.12252

Caloia, F., Van Ginkel, K., & Jansen, D. (2023). Floods and financial stability: Scenario-based evidence from below sea level. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.4672119

CBS. (2013, 21 januari). *Twee derde van alle woningen eengezinswoning*. https://www.cbs.nl/nl-nl/achtergrond/2013/04/twee-derde-van-alle-woningen-eengezinswoning

Chareyron, S. (2024). The Valuation of Energy Efficiency Labels in the French Housing Market. https://www.tepp.eu/images/pdf/working-paper-2024/valuationenergyefficiencylabelsfrenchhousingmarket.pdf

Damen, S., Vastmans, F., & Buyst, E. (2016). The effect of mortgage interest deduction and mortgage characteristics on house prices. Journal Of Housing Economics, 34, 15–29. https://doi.org/10.1016/j.jhe.2016.06.002

De Koning, K., & Filatova, T. (2019). Repetitive floods intensify out-migration and climate gentrification in coastal cities. *Environmental Research Letters*, *15*(3), 034008. https://doi.org/10.1088/1748-9326/ab6668

De Lange, M., Groot, P., Van Meurs, S., & Koeckhoven, T. (2025). *Water en Bodem Sturend*. https://www.eib.nl/wp-content/uploads/2025/03/RHWB-rapport-v14-FINAL2.pdf

De Nederlandsche Bank. (n.d.). Risico's van klimaatverandering. https://www.dnb.nl/groene-economie/risico-s-van-klimaatverandering/

De Nederlandsche Bank. (2021). Op weg naar een duurzame balans: Integratie van duurzaamheidsrisico's in de kernprocessen van de financiële sector. https://www.dnb.nl/media/shoftigm/web_brochure_op-weg-naar-een-duurzame-balans.pdf

De Vries, P., & Boelhouwer, P. (2008). Equilibrium between interest payments and income in the housing market. *Journal Of Housing And The Built Environment*, *24*(1), 19–29. https://doi.org/10.1007/s10901-008-9131-z

Deelen, A., Tijm, J., Trinks, A., Schippers, V., & CPB. (2025). Overstromingsrisico voor Nederlandse huishoudens: financiële omvang en verdeling. In CPB PUBLICATIE (pp. 2–30). Centraal Planbureau. https://www.cpb.nl/system/files/cpbmedia/CPB_publicatie_overstromingsrisicos-voor-Nederlandse-huishoudens.pdf

DNB. (2024). Verklaringen voor het snelle herstel van huizenprijzen. In *DNB*. https://www.dnb.nl/algemeen-nieuws/achtergrond-2024/verklaringen-voor-het-snelle-herstel-van-huizenprijzen/

Dubbelboer, J., Nikolic, I., Jenkins, K., & Hall, J. (2017). An Agent-Based Model of Flood Risk and Insurance. *Journal Of Artificial Societies And Social Simulation*, *20*(1). https://doi.org/10.18564/jasss.3135

Dubois, C. (2026, April 2). Desk research Done right: How to conduct a literature review that actually adds value | ProjectBist blog. ProjectBist. https://projectbist.com/blog/desk-research-literature-review-guide

ECB. (2020). Trends and risks in credit underwriting standards of significant institutions in the Single Supervisory Mechanism Main findings from the credit underwriting data collection 2019. European Central Bank. https://www.bankingsupervision.europa.eu/ecb/pub/pdf/ssm.creditunderwriting202006~d2a9e3329c.en.pdf

ECB. (2022). 2022 climate Risk stress test. https://www.bankingsupervision.europa.eu/ecb/pub/pdf/ssm.climate_stress_test_report.20220708~2e3cc0999f.en.pdf

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 Comission. (n.d.). Single supervisory mechanism. Finance. https://finance.ec.europa.eu/banking/banking-union/single-supervisory-mechanism_en

Fairweather, D. (2024, 29 October). *Home buyers, take heed: Climate Risk Scores are a Must-Use tool*. Forbes. https://www.forbes.com/sites/darylfairweather/2024/10/29/homebuyers-take-heed-climate-risk-scores-are-a-must-use-tool/

Fairweather, D., Kahn, M., Metcalfe, R., & Olascoaga, S. S. (2024). *Expecting Climate Change: A Nationwide Field Experiment in the Housing Market*. https://doi.org/10.3386/w33119

Filatova, T. (2015). Empirical agent-based land market: Integrating adaptive economic behavior in urban land-use models. Computers Environment And Urban Systems, 54, 397–413. https://doi.org/10.1016/j.compenvurbsys.2014.06.007

Fontana, A., Jarmulska, B., Scheid, B., Scheins, C., & Schwarz, C. (2025). *to fire: is physical climate risk taken into account in banks' residential mortgage rates?* https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5176497

Francke, M., Van de Minne, A., & Verbruggen, J. (2014). The Effect of Credit Conditions on the Dutch Housing Market. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.2526144

Van Ginkel, K., Endendijk, T., Van Veggel, W., & Schasfoort, F. (2024). Kansen en gevaren van klimaatrisico informatie op huisniveau. https://www.deltaprogramma.nl/site/binaries/site-content/collections/documents/2024/11/29/signaal-3.-kansen-en-gevaren-van-klimaatinformatie-op-huisniveau/Signaal+3.+Kansen+en+gevaren+van+klimaatinformatie+op+huisniveau.pdf

Geijtenbeek, L., Mulder, P., & Niessink, R. (2026). Beyond energy labels: estimating housing energy efficiency and energy underconsumption using administrative microdata. Energy Efficiency, 19(1). https://doi.org/10.1007/s12053-025-10405-2

Han, Y., & Peng, Z. (2019). The integration of local government, residents, and insurance in coastal adaptation: An agent-based modeling approach. *Computers Environment And Urban Systems*, *76*, 69–79. https://doi.org/10.1016/j.compenvurbsys.2019.04.001

Hudson, P., Botzen, W. W., Feyen, L., & Aerts, J. C. (2016). Incentivising flood risk adaptation through risk based insurance premiums: Trade-offs between affordability and risk reduction. Ecological Economics, 125, 1–13. https://doi.org/10.1016/j.ecolecon.2016.01.015

Holtermans, R., Kahn, M. E., & Kok, N. (2024). Climate risk and commercial mortgage delinquency. Journal Of Regional Science, 64(4), 994–1037. https://doi.org/10.1111/jors.12681

Hoogvliet, M., Slager, K., & Dolman, N. (2023). Verkenning waterlabel: Stap 1: inventarisatie van (ervaringen met) bestaande labels en voorziene (on)mogelijkheden, met focus op woningen. https://open.overheid.nl/documenten/dc70a8b5-07a6-48fd-8de3-8c3e635748a3/file

Hu, S. (2020, 27 August). What Is Climate Gentrification? *NRDC*. https://www.nrdc.org/stories/what-climate-gentrification

Hudson, P. (2020). The Affordability of Flood Risk Property-Level Adaptation Measures. *Risk Analysis*, *40*(6), 1151–1167. https://doi.org/10.1111/risa.13465

Keenan, J. M., Hill, T., & Gumber, A. (2018). Climate gentrification: from theory to empiricism in Miami-Dade County, Florida. *Environmental Research Letters*, *13*(5), 054001. https://doi.org/10.1088/1748-9326/aabb32

Keenan, J. M., & Bradt, J. T. (2020). Underwaterwriting: from theory to empiricism in regional mortgage markets in the U.S. *Climatic Change*, *162*(4), 2043–2067. https://doi.org/10.1007/s10584-020-02734-1

Kousky, C., Palim, M., & Pan, Y. (2020). Flood Damage and Mortgage Credit Risk: A Case Study of Hurricane Harvey. Journal Of Housing Research, 29(sup1), S86–S120. https://doi.org/10.1080/10527001.2020.1840131

Langen, M. (2025). Housing market monitor: EU housing markets. https://assets.ctfassets.net/1u811bvgvthc/7o7A4FBrdsygsrVyKFdCIe/f7a5c57023d064c68e749f804a87e262/wmm_topic_2025q2_EN.pdf

Lin, C., & Tsai, I. (2021). HOUSE PRICES, RENTAL COSTS, AND MORTGAGE INTEREST RATES. International Journal Of Strategic Property Management, 25(5), 356–368. https://doi.org/10.3846/ijspm.2021.14966

Lloyd, T., Ryan-Collins, J., & Macfarlane, L. (2017). Rethinking the Economics of Land and Housing. Bloomsbury Academic. https://www.bloomsbury.com/au/rethinking-the-economics-of-land-and-housing-9781350374270/

Madsen, J. B. (2012). A behavioral model of house prices. Journal Of Economic Behavior & Organization, 82(1), 21–38. https://doi.org/10.1016/j.jebo.2011.12.010

McQuinn, K., & O'Reilly, G. (2007). Assessing the role of income and interest rates in determining house prices. Economic Modelling, 25(3), 377–390. https://doi.org/10.1016/j.econmod.2007.06.010

Mérő, B., Borsos, A., Hosszú, Z., Oláh, Z., & Vágó, N. (2023). A high-resolution, data-driven agent-based model of the housing market. Journal Of Economic Dynamics And Control, 155, 104738. https://doi.org/10.1016/j.jedc.2023.104738

Ministerie van Infrastructuur en Waterstaat. (n.d.). LIWO [Dataset]. https://basisinformatie-overstromingen.nl/#/maps

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

Mutlu, A., & Filatova, T. (2026). Urban housing markets under flood risk: Modeling demand pressure, risk perception bias, and public interventions. Computers, Environment and Urban Systems, 128, 102440. https://doi.org/10.1016/j.compenvurbsys.2026.102440

Nikolic, I., & Ghorbani, A. (2011). A method for developing agent-based models of socio-technical systems. International Conference On Networking, Sensing And Control, 44–49. https://doi.org/10.1109/icnsc.2011.5874914

Niu, D., Eichholtz, P., & Kok, N. (2025). Asymmetric Information Provision and Flood Risk Salience. *Journal Of Housing Economics*, 102060. https://doi.org/10.1016/j.jhe.2025.102060

NOS Nieuws. (2025, 3 april). Zestig gezinnen raken huis definitief kwijt na hoosbui Enschede. NOS. https://nos.nl/artikel/2562192-zestig-gezinnen-raken-huis-definitief-kwijt-na-hoosbui-enschede

NVM. (n.d.). Kopen. https://www.nvm.nl/wonen/kopen/

Ouazad, A., Kahn, M. E., & National Bureau of Economic Research. (2019). Mortgage Finance in the Face of Rising Climate Risk (Nr. 26322). https://www.nber.org/system/files/working_papers/w26322/revisions/w26322.rev1.pdf

Oerlemans, C., Hoogvliet, M., Bayer, M. L., Niu, D., Nobel, A., & Taylor, Z. (2025). Tagging the Threats: Unpacking Propositions for Real Estate Climate Risk Labels in the Netherlands. Deleted Journal, 4(1), 69–83. https://doi.org/10.3138/jccpe-2025-0010

Phlippen, S., Schreuder, C., & Vendel, B. (2023, 28 november). *Stapeling klimaatrisico's en financiële draagkracht woningmarkt*. ABN AMRO Bank.

https://www.abnamro.com/research/nl/onze-research/stapeling-klimaatrisicos-en-financiele-draagkracht-op-de-woningmarkt

Pryshlakivsky, J., & Searcy, C. (2012). *Sustainable Development as a Wicked Problem. In Topics in safety, risk, reliability and quality* (pp. 109–128). https://doi.org/10.1007/978-94-007-5515-4_6

Saltelli, A., Tarantola, S., Campolongo, F., & Ratto, M. (2004). Sensitivity Analysis in Practice: A Guide to Assessing Scientific Models.

Samuelson, P. A. (1952). Spatial Price Equilibrium and Linear Programming. *The American Economic Review*, *42*(3), 283–303. http://www.jstor.org/stable/1810381

Schiermeier, Q. (2011). Increased flood risk linked to global warming. *Nature*, *470*(7334), 316. https://doi.org/10.1038/470316a

Slager, K., & Wagenaar, D. (2017). *Standaardmethode 2017 Schade en slachtoffers als gevolg van overstromingen*. https://iplo.nl/publish/pages/132789/standaardmethode-2017-schade-en-slachtoffers-als-gevolg-van-overstromingen-definitief.pdf

Slager, K. (2022). Meer voorbereid op overstromingsrisico's door duidelijke informatie. In Deltares. https://cms.deltares.nl/assets/common/downloads/Onderzoek-limburg_2410.pdf

Stangenberg, L., Van Wickeren, S., & Zhang, L. (2020). The information value of energy labels: Evidence from the Dutch residential housing market. https://www.cpb.nl/system/files/cpbmedia/omnidownload/CPB-Discussion-Paper-413-The-information-value-of-energy-labels.pdf

Thomson, H., Zeff, H. B., Kleiman, R., Sebastian, A., & Characklis, G. W. (2023). Systemic financial risk arising from residential flood losses. *Earth S Future*, *11*(4). https://doi.org/10.1029/2022ef003206

Tierolf, L., Haer, T., Botzen, W. J. W., De Bruijn, J. A., Ton, M. J., Reimann, L., & Aerts, J. C. J. H. (2023). A coupled agent-based model for France for simulating adaptation and migration decisions under future coastal flood risk. *Scientific Reports*, *13*(1). https://doi.org/10.1038/s41598-023-31351-y

Tonn, G. L., & Guikema, S. D. (2017). An Agent-Based Model of Evolving Community Flood Risk. Risk Analysis, 38(6), 1258–1278. https://doi.org/10.1111/risa.12939

Tran, B. R., & Wilson, D. J. (2024). The Local Economic Impact of Natural Disasters (Nr. 2020–34). https://www.frbsf.org/wp-content/uploads/wp2020-34.pdf

Van Der Drift, R., De Haan, J., & Boelhouwer, P. (2022). Mortgage credit and house prices: The housing market equilibrium revisited. Economic Modelling, 120, 106136. https://doi.org/10.1016/j.econmod.2022.106136

Van Valkengoed, A. M., & Steg, L. (2019). Meta-analyses of factors motivating climate change adaptation behaviour. Nature Climate Change, 9(2), 158–163. https://doi.org/10.1038/s41558-018-0371-y

Verbond van verzekeraars. (2015). Klimaatverandering & schadelast. https://www.verzekeraars.nl/media/1878/klimaatverandering-en-schadelast.pdf

Verbond van Verzekeraars. (2018). POSITION PAPER OVERSTROMING. In Position Paper Overstroming. https://www.verzekeraars.nl/media/8163/vvv-popa_overstroming_2020.pdf

Warnaar, M., Bos, J., & Van den Enden, G. (2024, september). Rapport Advies hypotheeknormen 2025 (Nibud, 2024) - Nibud. Nibud. https://www.nibud.nl/onderzoeksrapporten/rapport-advies-hypotheeknormen-2025-nibud-2024/

Wright, K. (2024). Financialization of the Housing Market: A Contribution to Modern Urban Rent Theory [PhD Thesis, University of Waterloo]. https://dspacemainprd01.lib.uwaterloo.ca/server/api/core/bitstreams/b749ee4b-4dfe-43a1-bfea-66681f98ba0c/content

van der Straten, Y. (2023). Flooded House or Underwater Mortgage? The Implications of Climate Change and Adaptation on Housing, Income, and Wealth. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.4393731

Xiao, Y. (2011). LOCAL ECONOMIC IMPACTS OF NATURAL DISASTERS*. Journal Of Regional Science, 51(4), 804–820. https://doi.org/10.1111/j.1467-9787.2011.00717.x

Appendix A: A short summary of Rethinking the Economics of Land and Housing

In this appendix, the book “Rethinking the Economics of Land and Housing” from Lloyd et al. (2017) will be summarized. The focus will be on the proof the book uses to show how the financialization how housing has not led to improvements to the affordability of housing. Instead, it has led to inflated home prices. This assumption is essential to the functioning of the model. It first introduces the idea and history of the financialization of the housing market. Then it gives a short overview of where these rents have gone towards and what incentives it creates. Finally, it ends with how these assumptions change the model.

During the liberalization wave in the 80’s a new many different countries have enacted policies that would incentivize homeowner into taking up mortgages, hoping to increase homeownership, and ensuring homes stay affordable. These changes came in tandem with decreases in subsidies towards developers, which used to be introduced during times of low demand in the housing market. The expectation at the time was that by reducing interest rates, creating higher demand in the housing market leading to increases in housing supply. Firstly, due to the higher interest rates incentivizing more homeowners to sell their homes. Secondly, and most importantly, it would encourage developers to build new homes due to the increasing prices. However, these policies have resulted in mixed effects, with countries like the United Kingdom and Australia increasingly struggling to keep housing accessible, while Japan has been successful in improving the affordability of housing. This is proof that these types of policies have had little effect on enabling the accessibility of cheap housing. Instead, in many areas housing prices have increased, while having little effect on affordability. As such, a different theory was needed to explain this discrepancy.

The authors argue that this is due to the inflationary effects of financialization on the ground prices. During the creation of these policies, the perverse incentives of increasing land-prices was left out. Instead, it was assumed that the increases in the ground prices would incentivize capital to flow towards these locations, as these fixed costs of buying the ground allow access towards successful a market. Simply put agglomeration. The authors argue that “this is circular reasoning – it is not that the value of the land has pushed up the fixed costs, but rather that the value of land has absorbed the additional return that was generated via the production.” (P. 56) These increasing prices instead create less room for capital investment. This can for example be seen in the price of the house having not shifted significantly, while ground prices have substantially increased from 1990 to 2008 (P. 8). At the same time, these rents are unfairly being captured by landowners, creating perverse incentives for them to keep these high (such as through NIMBYism). Without fixing these mechanisms, the housing market will function as a stock market. With increasing interest rates only increasing the prices in the market, without then also increasing the supply of housing. Policies should therefore focus on increasing capital investments by either decreasing the perverse incentives through a land value tax or increasing the ease at which housing can be built.

The consequence to the model is that only demand will strongly shift the affordability of housing, as more households will be competing for the same home. The changes in interest rate will only shift the value of the home, as the effects of interest rate changes is purely seen as inflationary.

Thus, the utility of a home is treated solely as a financial asset, wherein the value of the home is the utility. This assumption is thus based on that supply will not strongly increase; however, this assumption can be wrong.

Appendix B: Model setup

Variable name	**Variable**	**Notes**
Runs	150	
Buyer demand	1/2/3	Depended on the scenario
Use_sample	True	Use the anonymized sample
Sample_size	3000	Number of houses sampled randomly from the 10.000 anonymized sample
Flood_probability_scenario	0.01	The chance a flood will happen, without any biases
Flood_risk_mode	Uniform	Uniform, meaning that the flood risk is evenly spread out, if there is any
NBS_mode	Uniform	Will not change anything, mostly for functioning model. Means distance from Nature Based Solution for flood events, more relevant for river floods
Buyer_util_mode	EU_V1_Dutch	To simulate the Dutch market
Use_rp_bias	True/False	Depends on Scenario
RP_bias_distribution	[4, 1.5]	The biases people have, and how they are distributed on a Beta distribution.
Use_market_trends	False	Checks whether the market is hot or cold, and buyers will change their strategy accordingly
Use_uniform_res_time	False	Sets the residence time to random
Use_mortgage_debt	False	Changes bidding strategy to include mortgage, not compatible yet with banking module
Seller_mode	Probabalistic	Random chance at which a homeowner will sell their home
Failure_adjustment	0.05	Increases the buyer's bid if unsuccessful in the market
Max_failure_adjustment	0.1	Maximum 10% decrease of realtor price.
Listing_time_discount	0.02	Two percent decrease for each failure in asking price compared to realtor price
Price_margin_bounds	[0.05, -0.05]	From this, the starting range at how much a buyer will over or underbid from their utility is calculated. This will after unsuccessful attempts then increase.
Dynamic_income_groups	False	The income groups will dynamically increase depending on the prices in the market
Banking	True	Turns on banking in the model
Base_mortgage_rate	0.035	A base mortgage rate of 3.5%, from this the central bank rate will be calculated to get their rate, which will float.
Flood_interest_change	False/True	Depends on scenario, sets the banks behaviour
Flood_interest_change_years	[2/10]	Depends on scenario, this sets the speed at which the banks start to integrate the flood rate.

Appendix C: Variables in the dataset

Variable	**Regression**	**Squared**	Definition
Age	V	V	Age in years
Garden size	V	X	In M^2
Lot size	V	X	In m^2
Quality	V	X	Dummy variable 1, 2, 3, represents higher quality
Rooms	X	X	Number of rooms
Flood label (RF70_100Y)	V	X	Flood label as given by LIWO, 0 to 5, with each starting from one representing increasing flood risk every one-hundred years.
Floodprone	X	X	Whether the property might flood in one hundred years
Predicted price	X	X	Price that has been predicted while sampling

Appendix D: Sensitivity analysis

The sensitivity analysis has been performed with a standard 10% margin around the chosen factors. This is method is also called the one-factor at a time approach, which imposes limitations on the strength of the results, which will further be discussed in the limitations section of this appendix. A form of global sensitivity analysis would fit this type of model far better, as more of the scenario and policy space uncertainty can be explored more effectively. A good sensitivity analysis can ensure that other potential feedback loops will not be overlooked (Saltelli et al., 2004). The short answer for this choice is the necessary computing power required being significant with the number of runs required to complete such an analysis, as each takes around 10 seconds to complete, with tens of thousands usually being required. This shortage was already a significant limitation within this analysis, with only the variable with the largest uncertainty range currently being tested. This is done as the reason for this analysis is to explore the uncertainty space, and it can still verify the same (or at least similar) results given this uncertainty (Saltelli et al., 2004). As such, it was chosen to only vary the interest rate, as this can change quickly, with a range of 1.5 to 4% not being unheard of. Ideally, other variables should be tested, which will also be explored more in the limitations.

The focus of the results will be on the price, price integration speed, and the overall risk integration, as these are the most important results from this thesis. It was chosen to only include the scenarios from the single banks runs, which assumes a high adoption rate. This was done for similar reasons as the creation of this group of scenarios, so that the relations between the different components can properly be extracted and reasoned, making it also ideal for the sensitivity analysis. As the interest rates were chosen to be varied, the range can be viewed in the table below. Afterwards the most important results will be highlighted, with only the final step 60 being taken into consideration for the significance.

Table D1: the variables of interest for the sensitivity analysis

	LOW	BASELINE	HIGH
INTEREST RATES	3.15%	3.5%	3.85%

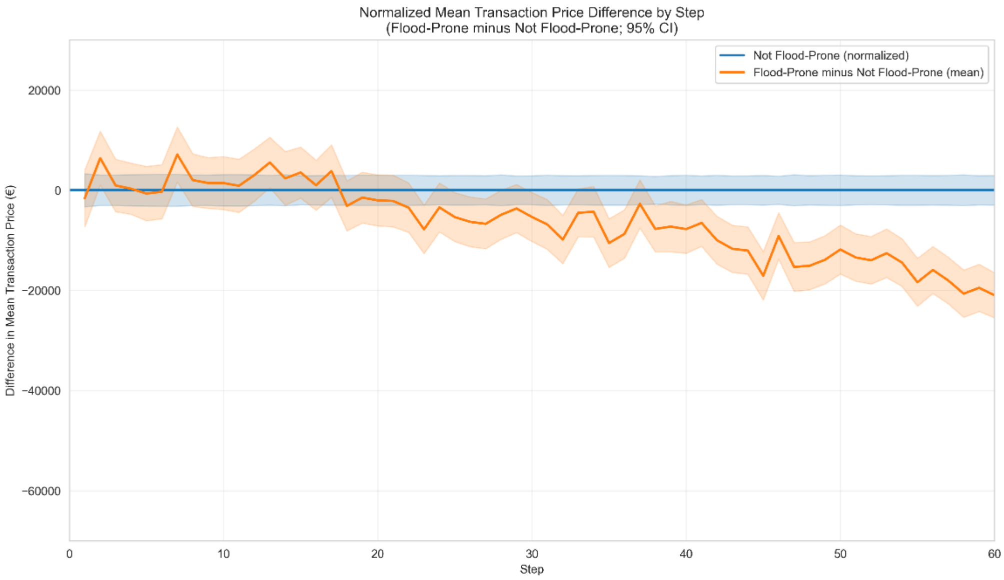

Figure D1: Base scenario, high demand, objective risk perception, no sustainable mortgages

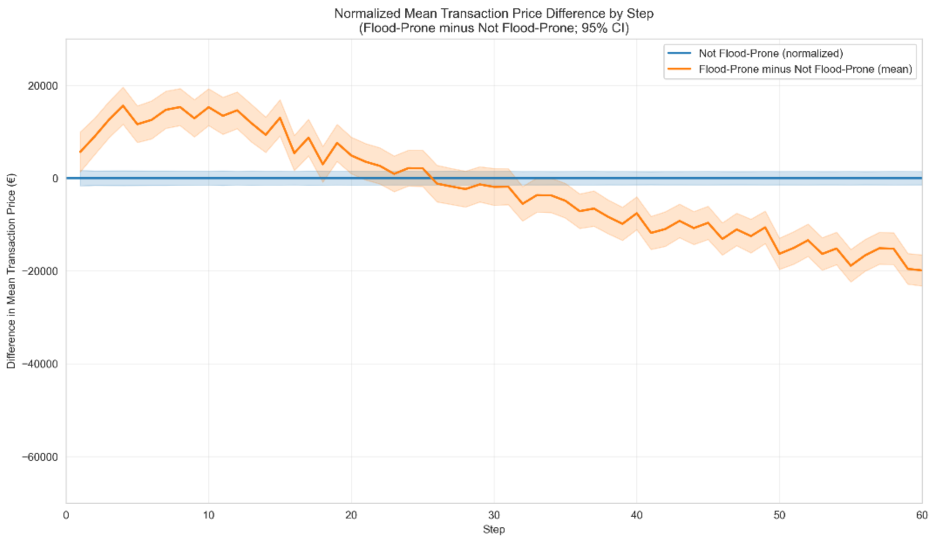

Figure D2: High scenario, objective risk perception, high demand, no sustainable mortgages

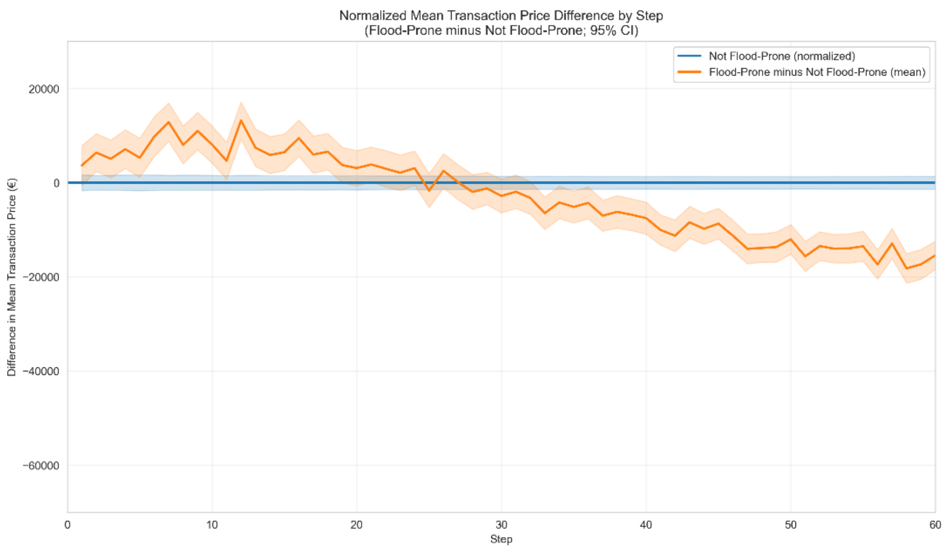

Figure D3: Low scenario, objective risk perception, high demand, no sustainable mortgages

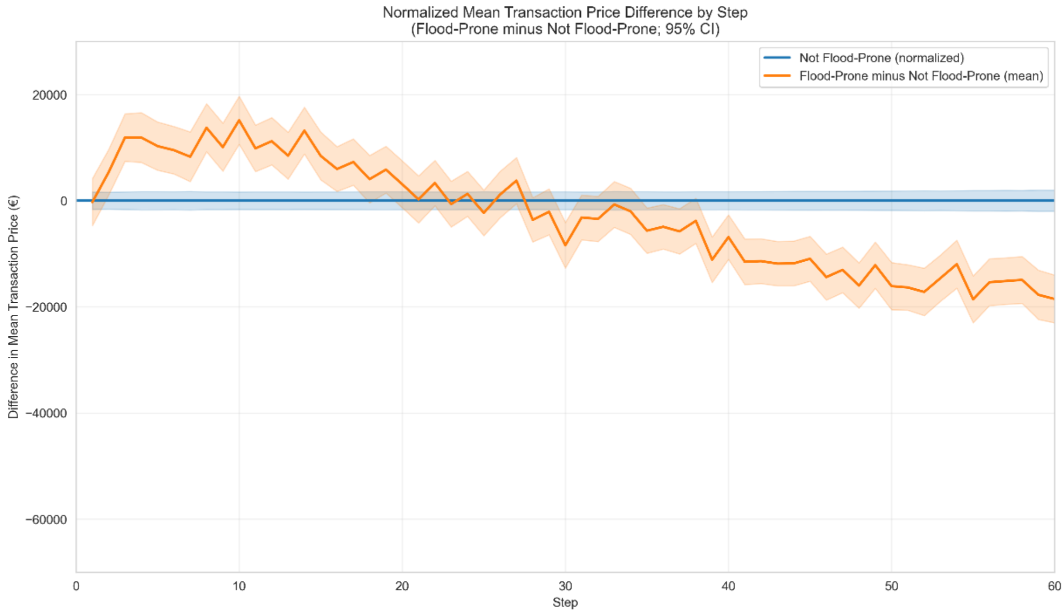

Figure D4: Base scenario, Objective risk perception, very high demand, no sustainable mortgages

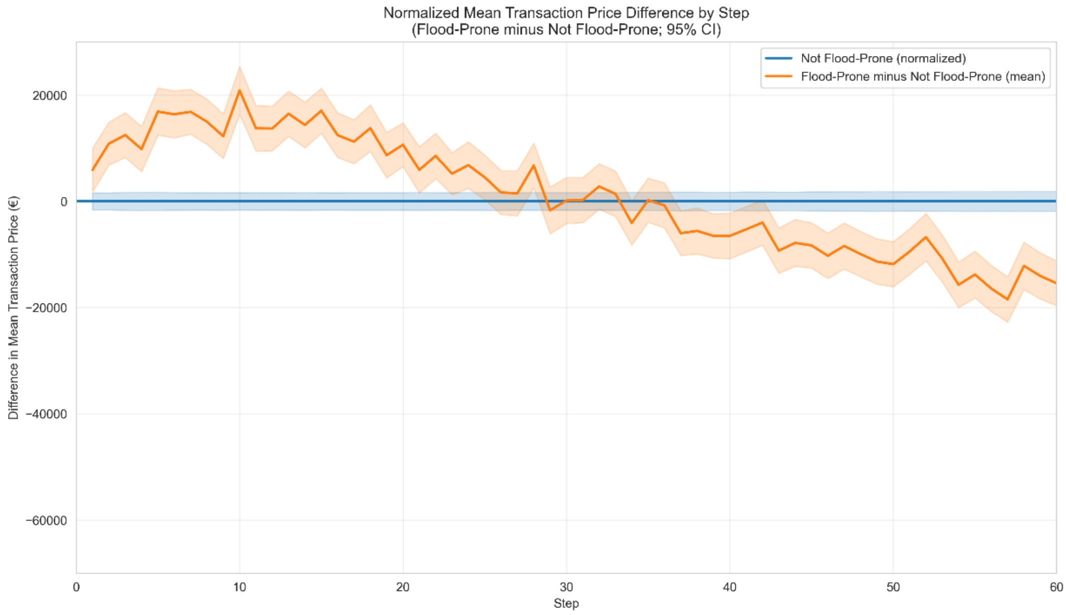

Figure D 5: High scenario, objective risk perception, very high demand, no sustainable mortgages

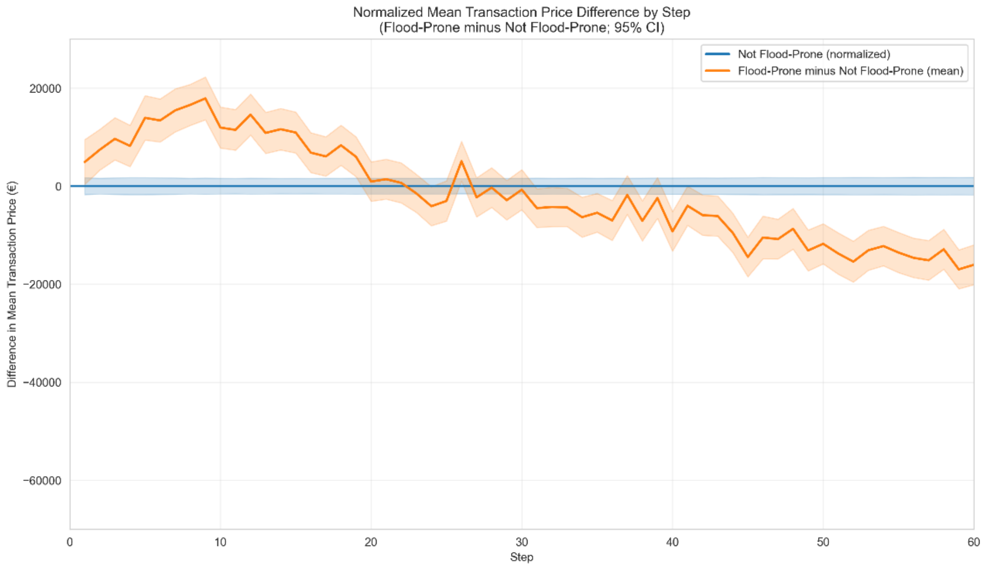

Figure D 6: Low scenario, objective risk perception, very high demand, no sustainable mortgages

Above the most important results concerning the water label can be found. All these results are from the group highly flood prone compared to none to low flood prone. Here it can clearly be viewed that the results even under different interest rates are still holding. In Table D2 below, it can be concluded that most differences, however, are significant when compared to the low sensitivity. Thus, it can be concluded that the strength of the results can be dependent on the interest rates. This can also be seen in the graphs above, as they highlight that the significance of the base scenario is only upheld there, as the difference in price is significant in the other analyses. Although the difference is quite small, only around 3000 euros. Thus, it can be argued that the conclusion still holds as this can also be attributed to differences in income levels, as the lower interest rate also decreases the prices. This is due to the interest rates also setting the income levels of the households. Thus, a lower interest rate also decreases the average income of the households. At the same time, the logic discussed in the results still holds, as increasing demand should also increase the amount of pricing integration. Something which can clearly not be seen in the results. These are also confirmed by Table D3, which give the same results but for the flood prone and non-flood prone group.

Further looking into table D2, it can be concluded that no other group of scenarios is sensitive to the interest rate changes, with mainly the water label being significant. While sustainable mortgages in no regard is, as only a few show, such as scenario FLINT_VHIGH, a significant change. Therefore, it can be concluded that the sustainable mortgages are not sensitive to the changes in variables.

Table D 2: significance of the sensitivity analysis for the Highly flood prone group

SCENARIO	METRIC	T LOW	P LOW	T HIGH	P HIGH
NOFLINT_HIGH	mean price	4.455	0.000	-1.497	0.136
	price difference	2.069	0.039	0.472	0.637
NOFLINT_LOW	mean price	1.771	0.078	1.841	0.067
	price difference	-1.234	0.218	-5.234	0.000
NOFLINT_VHIGH	mean price	1.509	0.132	-8.184	0.000
	price difference	-0.847	0.398	1.271	0.205
FLINT_HIGH	mean price	0.222	0.825	-1.389	0.166
	price difference	0.068	0.946	-0.069	0.945
FLINT_LOW	mean price	-0.003	0.998	-1.775	0.077
	price difference	-1.118	0.264	2.014	0.045
FLINT_VHIGH	mean price	4.021	0.000	-1.457	0.146
	price difference	0.741	0.459	4.792	0.000
FLINT_HIGH_OBJ	mean price	6.706	0.000	0.655	0.513
	price difference	-1.384	0.167	-0.452	0.652
FLINT_LOW_OBJ	mean price	-3.340	0.001	0.364	0.716
	price difference	2.008	0.046	1.070	0.286
FLINT_VHIGH_OBJ	mean price	1.478	0.140	-5.976	0.000
	price difference	5.094	0.000	1.632	0.104
NOFLINT_HIGH_OBJ	mean price	3.493	0.001	-1.394	0.165
	price difference	-2.265	0.024	-1.062	0.289
NOFLINT_LOW_OBJ	mean price	-3.784	0.000	-2.261	0.024
	price difference	0.479	0.633	-2.862	0.005
NO_FLINT_VHIGH_O	mean price	-1.996	0.047	-1.626	0.105
	price difference	-1.070	0.286	-1.012	0.312

Table D3: none to flood prone properties sensitivity significance

SCENARIO	METRIC	T LOW	P LOW	T HIGH	P HIGH
NOFLINT_HIGH	mean price	4.455	0.000	-1.497	0.136
	price difference	3.189	0.002	2.140	0.033
NOFLINT_LOW	mean price	1.771	0.078	1.841	0.067
	price difference	-0.478	0.633	-3.721	0.000
NOFLINT_VHIGH	mean price	1.509	0.132	-8.184	0.000
	price difference	-1.907	0.057	-0.302	0.763
FLINT_HIGH	mean price	0.222	0.825	-1.389	0.166
	price difference	0.316	0.752	0.348	0.728
FLINT_LOW	mean price	-0.003	0.998	-1.775	0.077
	price difference	-0.331	0.741	2.950	0.003
FLINT_VHIGH	mean price	4.021	0.000	-1.457	0.146
	price difference	0.864	0.388	3.989	0.000
FLINT_HIGH_OBJ	mean price	6.706	0.000	0.655	0.513
	price difference	-2.080	0.038	-0.049	0.961
FLINT_LOW_OBJ	mean price	-3.340	0.001	0.364	0.716
	price difference	2.533	0.012	0.549	0.584
FLINT_VHIGH_OBJ	mean price	1.478	0.140	-5.976	0.000
	price difference	3.839	0.000	1.090	0.276
NO_FLINT_HIGH_OBJ	mean price	3.493	0.001	-1.394	0.165
	price difference	-3.285	0.001	-1.506	0.133
NOFLINT_LOW_OBJ	mean price	-3.784	0.000	-2.261	0.024
	price difference	1.028	0.305	-0.713	0.477
NOFLINT_VHIGH_OBJ	mean price	-1.996	0.047	-1.626	0.105
	price difference	0.075	0.940	-0.169	0.866

D.2 Limitations

The application of the one-factor-at-a-time method in its current states holds significant limitations, of which most were choices given the resources at hand. Large compromises had to be made mainly due to the limitations of computing power. It was mainly chosen to limit the variables within the analysis to one, which can mainly be attributed to the one-factor at a time analysis requiring significant computing power to achieve these results, with the high and low run taking around 12 hours each to finish. Global sensitivity analyses are far more effective at these types of loads due more effective spreading of the scenarios, rather seeing it as a space to be discovered instead of different endpoints, resulting in better outcomes. Ideally this would mean other variables could also be included in the analysis, with the most important being the default odds ratio, the damage fractions as well as the discount factor, as these all have the most impact on the results. These ranges for the most part can be taken from their respective papers, as most of them give the uncertainty interval for their results, with the damage fraction Endendijk et al. (2022) and the discount factor from Beltrán et al. (2018) being the most accurate, as they are the result of a literature analysis. Moreover, small changes can impact the outcome. For example, if the damages of floods are on the higher end, the credit restrictions will play a less important role, as was discussed in the results.

Heightened damage fractions might therefore impact the validity of the results and in turn falsify the results of the model. Moreover, this can lead to the discovery of other feedback mechanisms currently not discussed in the results, leading to differing conclusions. A proper sensitivity analysis can therefore also be used to further enhance the policy analysis of this thesis, as the previous variables mainly pertain to the scenario space. By allowing the sensitivity analysis to be used for policy discovery, it can increase the current policy options with the uncertainty, allowing for a more robust analysis and making the sensitivity analysis a far greater part of the overall experimental design (Saltelli et al., 2004).

Another limitation is the chosen range of the interest rate, only being 10% around the original value. This is especially interesting for the water label, due to the possibility of lower interest rates further slowing the price integration, which at the moment is too small for strong conclusions to be drawn. This can become even more important considering how quickly the interest rates might change, changing the amount of credit available, in turn increasing the chances of proper price integration. The previously discussed method should fix this.

Appendix E: Model Design

In this appendix the overall method used during the model design will be more broadly discussed, as especially during this phase, many different techniques were used that differ significantly from the methods described in Nikolic & Ghorbani (2011). The first step was starting off with a broad overview of the initial model. In this document, the basic steps that are described in Nikolic & Ghorbani (2011) are re-treaded by broadly describing each step in the model and explaining why many choices were made. Mainly actor behaviour and how this behaviour was performed in the model was summarized. Afterwards a second document would describe how overall banking models and housing market models with functioned, with some of them being more thoroughly summarized if they were relevant to the model. Much of this has been reworked into section 5: Model Design. Finally, a third document was created that would combine both of the previous ones, combined with the literature analysis that was performed in section 2.

This section will discuss the initially perceived end state of the model, constructed through the activity diagrams given below. These have been helpful in breaking down the steps necessary to build this model and formed the overall outline of the model. This was necessary to ensure that at each step a functioning model could be presented, so that at each point the thesis could be completed. It was essential that the model design would.

E.1 Large outline of the model

After finishing the model design, as is described in section 5: model design, the description was used to create an activity diagram that summarizes the banker behaviour. This was then used to create the smaller steps that were necessary to create the functioning model, essentially making sure that at each step a functioning model was created that could be used for this thesis, while each step would only increase the accuracy at which the model function.

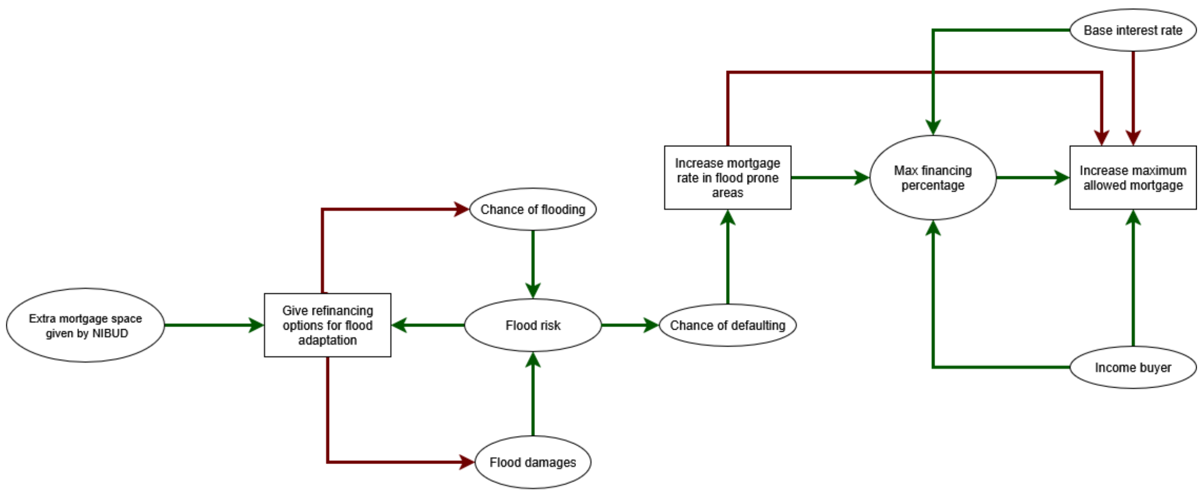

Figure E1: The initial model outline, the maximalist version of the model.

E.2 Minimum viable product

Through this, the minimum viable product was created, essentially the heart of the model, dictating which requirements were most essential for the model to function. In short these were the NIBUD requirements, while also having some basic calculations for the interest rate and the utility function. These are essential in forming the bank's behaviour, creating the basic building blocks from which the model was build.

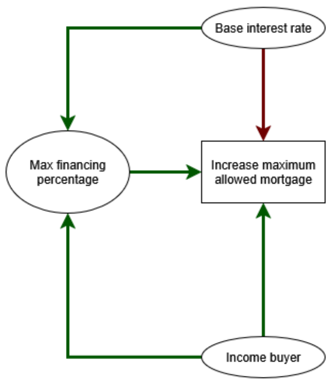

Figure E2: The basic version of the bank agent activity diagram

Through the utility function, some basic calculations could be used to change the expected prices through differing interest rates. In short, giving a simple way to add basis points for risky homes can be used to create simple pricing expectations. This formed an essential part of verifying whether the expected pricing integration given certain basis points would be accurate.

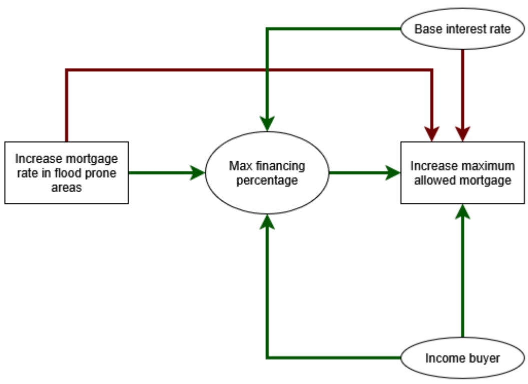

Figure E3: The slightly expanded version of the model, adding the changes in basis points

E.3 Expansion of the model

With the basic model completed, the focus was then put on expanding the model, making it more accurate and adding to its capabilities. In these steps the more complicated sections were added, such as improving the interest rate calculations and then adding the flood risk, as seen in the figures below.

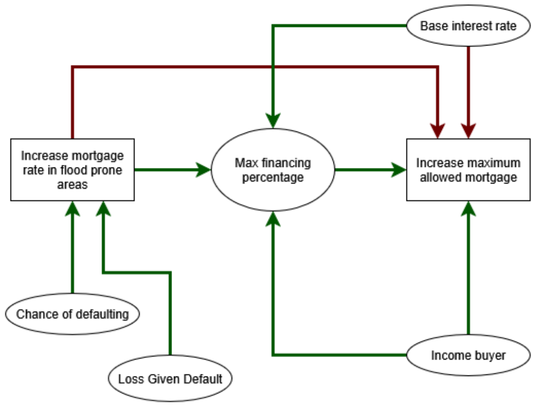

Figure E4: Improvement on the interest rate calculations

From this point onward, more comprehensive flood labels were added that created the necessity to create more accurate flood risk calculations, resulting in this expansion. Many of the more complicated parts of the model was added during this step. While these will not be discussed, this step was most significant, due to the large changes necessary to improve the model. Each new step required the entire process, from system identification to detailed model design, to be redone, while each new step added more time spend during the detailed model design due to the many changes required before a new step could be added. Thus, this was also the final endpoint of the model.

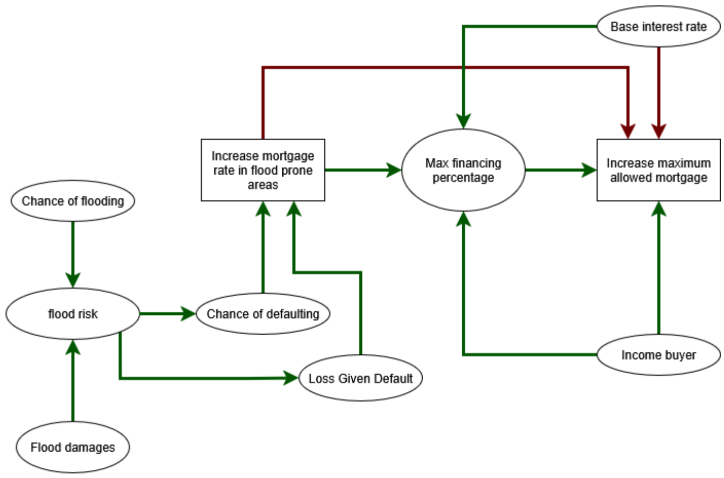

Figure E5: Adding of flood risk and multiple flood labels, improving the credit risk assessment, changing, and improving large parts of the model, requiring significant rewrites

Appendix F: Detailed Model Design

This appendix will explore more of the detailed model design; how certain solutions were found and mainly the reasoning behind these solutions. This differs from the section six: detailed model design in the sense that this explores more deeply as to why certain design choices were made as it relates to modelling limitations, time spend creating the solutions or which areas were more lacking for answering the research question in the starting state of the model. This will be explored more naturally, with different parts of the previous subjects being explored within this appendix through the lens of speed optimizations. This choice was made simply due to the fact many choices and limitations had to be made to reduce the overhead of the model, as the expansion required usage of many repetitious floating-point operations, causing the need for many small improvements needing to be made to improve the speed of the code. Without these optimizations, the model's runtime would easily increase over 5x, while now it is reduced to around 15% over the standard model. These have helped significantly in ensuring that experiments can be ran easily, allowing for faster development of experiments, bug finding and adding new modules.

F.1: Code optimization

In this section, the overall types of optimizations will be discussed, their strengths and weaknesses, and their impact on the overhead of the model. As floating-point operations require significant compute time, a large portion of time was spent on improving these through different means. These will be discussed through 3 different sections, firstly, the most basic type of optimizations. In the second section, through the usage of look up tables and their reasoning behind it. Finally, in the third section, the impact these choices have on the overall class structure of the model. These will then be explained more thoroughly in section G.2, which explores the class structure of the model.

F.1.1 floating point operations optimization

The first and most basic optimization used are the rounding the floating points, which if done, can already save significant time. It only requires some considerations around the level at which the numbers are rounded, too low, the results are fast but inaccurate, too high, and the code becomes slow but accurate. Therefore, a balance should be struck. From this followed the most basic lesson, either avoid, remove or reduce floating point calculations. Through this, some more complex optimizations could be done as for example LTV ratios were not done in percentages but in permilles. As the LTV is an exponent, these are especially slow, as they require large amounts of estimations. Instead, through simple mathematical tricks, you can simply pre-calculate the root 10 of beta and then make the LTV an integer. This simple trick can already save a large portion of time. If the results are not accurate enough, the exponent can simply be increased to the desired accuracy.

$$\sqrt[10]{\beta^{LTV*10-800}} = \beta^{LTV-80}$$

Then through the usage of the previously made assumptions, other optimizations can be found, by reducing the need to calculate the same thing over and over again. Some calculations were done too regularly but would need to always create the same results due to the assumption of 80% LTV, so only calculating variables that used this assumption once per step improved efficiency.

F.1.2 look up table optimization

The second group are look up table optimizations. Once again, due to the need for more repetitious being avoided, simply saving them lead to more time gains. This is caused by taking the solution out of storage being quite a lot quicker than recalculating them at the cost of increasing RAM usage. For this, 2 types of look up table optimizations groups were made: For the floating-point operations and for the regulatory side, starting with floating points. As they were rather slow, storing them became essential, requiring the need for an efficient storage method, while also requiring proper keys to be set. These keys were therefore very dependent on the starting conditions of the formula. These in turn follow the same principles as the rounding, if they are too broad, you create inaccuracies, if they are too tight, the time saved decreases. Plus, it became important to store these tables in the right places, leading to much of the decisions on class design.

For the regulatory side, some other issues started to prop up. As the need to adhere to regulatory standards became important, these also resulted in the need to optimize these tables themselves. As will be explored later, a lot of pre-calculations were necessary to save time, creating large tables that in themselves started to take up unnecessary time. So, these tables then needed to be optimized themselves. The first solution would be to change look up tables commonly used from pandas with the look ups as well as the dataset being put into 3 different numpy array's, which are significantly faster than Panda's. The second solution would be to then store these different tables in a dictionary, so that these optimizations could be stacked, these created a natural tree wherein these tables were quickly found. Now, all these do cause other slowdowns as these dictionaries start to fill, requiring a third optimization, by introducing a basic caching system. Here every table is checked each step to ensure that they are being used. This helps in more high demand scenarios, where these tables could quickly stack up or reduce random slow runs when the random walk of the central bank would increase the overhead. Others are simply emptied out per run as the results differ too much per run.

F.1.3 Class Design optimization

This leads to the final type of optimization, which relates to the class design. Class design and optimization are strongly linked, by cleverly saving certain variables and where to focus most of the optimizations will not only help improve overall performance, but through the usage of the basics of Object-Oriented Programming, the code also stays readable. By assigning commonly used methods towards their appropriate class, while also taking into account the different class interactions. These two different ideas did prove difficult to combine, as the banks, due to needing many repetitious calculations, would be more efficient with one class. This decision will be explored more in depth in the next section.

F.2: Class design

In this section, the class design of RHEA as well as the expansion will be discussed, starting with the class design of RHEA. By starting off with a general overview of functions needed within RHEA, the modelling choices will be explored. While this will not discuss everything available, it will merely just what is relevant for the task, as the model is rather expansive. This will lead to how certain concepts will need to be translated, which will then lead into another section about a more in-depth overview of the expanded class design. Moreover, at certain points the starting points of modelling choices will be explored, how these affect later decisions and their resulting limitations. These were put into their most relevant section, as their starting point was the result of either necessary optimizations, or due to certain concepts needing to be translated into a more credit-oriented model. This section will go into more detail than the short outline of model design given at 6.1 and 6.2, while avoiding the formulas already given at 6.3. Instead, this will more focus on the general coding logic and the implementation of the optimizations discussed previously, while avoiding the formulas already given, as they are assumed to be understood.

F.2.1: Class design of the RHEA and the expansion

This section will discuss the relevant class design given within the RHEA model. This does not include all the functions RHEA has, as much of the model has been designed for a purpose of researching nature-based solutions explored in Mutlu et al. (2023), while later additions are not publicly available. Instead, this will focus on the relevant functions available within the model, by reducing them to their basic concepts and will afterwards be translated to fit into a more credit focused model. This subsection will also explore some of the consequences of the modelling choices and the resulting discrepancy between the two models. The section afterwards will explain the class design of the entire model, with the discussion centred around the choices of class design, which strongly intersect with the optimizations found in the model.

F.2.1.1 Short overview of Class Design in RHEA

Households can have 3 different states, either they are occupants, sellers, or buyers, with them being able to change their state from occupants to sellers after a certain time period or due to pure chance. After their decision to sell their home, they receive their asking price from the realtor and put their house on sale initially for this price. After receiving their bids, they check whether they find them acceptable and accept them if they do. If the bid fails, at the start of the new step, they get the new asking price and reduce it according to a set decay rate. These steps are then repeated until the seller either finds an acceptable bid or if they have entered the market too many times. The model allowed decay rate to be calculated in two different ways, firstly, through a slowly increasing number each step or by calculating their leftover mortgage. This could have been used for buyers to save money in a simplified way, wherein if they want to sell their home, they will only allow the sale if they are fully able to pay off their mortgage with the bid. However, none of those ideas have been materialised within this thesis, instead the deposit was kept simplified. This choice was made due to the complexity required, needing too many changes for it to make sense.

Buyers are randomly sampled with different income ranges or from previous sellers, which according to their income group are assigned what fraction of their income will be designated towards housing. This fraction can then be used to set their search budget according to what they will pay over 30 years. After selecting a certain number of homes, buyers change their expected utility based on their expected flood damages, as well as their failures in the market. With successive failures, buyers will start to increase their bid margin, allowing them to not only bid higher, but also search for more expensive homes. This was noted as this could be expanded later on to introduce the deposit into the model, with the functioning being very similarly to saving. Remember, there are currently no credit limitations in the model, which requires the deposit and the mortgage to set the budget.

To expand on this, it would require significant changes to the model. Most importantly the income fraction, the percentage of income spend on housing, dictates many things from bidding, to initializing the households. Moreover, it is incredibly simple while also being multifunctional to use; as during the initialization, the incomes of the households can be calculated quite trivially. With the model using transaction data, the incomes of all households can easily be calculated by reversing the transaction price, starting off with sorting the value of the houses and relating them to their respective income group. This section would need to be heavily changed for it to fit into the credit context, but this will be discussed after in the next subsection.

Other sections of the model required very little to no changes that would fundamentally change its functioning. For example, the objective flood risk perception was already in the model, decreasing bids if a buyer perceives the home to be risky. But also, a more biased risk perception, allowing buyers to underestimate their overall damages if a flood were to occur. The realtor acts as a price setter, giving the initial asking price of homes, which they seller then uses to set their asking price. This uses hedonic regression, enabling the prices to change dynamically throughout a run by responding to changes in buyer behaviour. These dynamics would shift due to the banks' behaviour, but the hedonic regression would respond to these changes. This leaves this general class diagram as a starting point for the model:

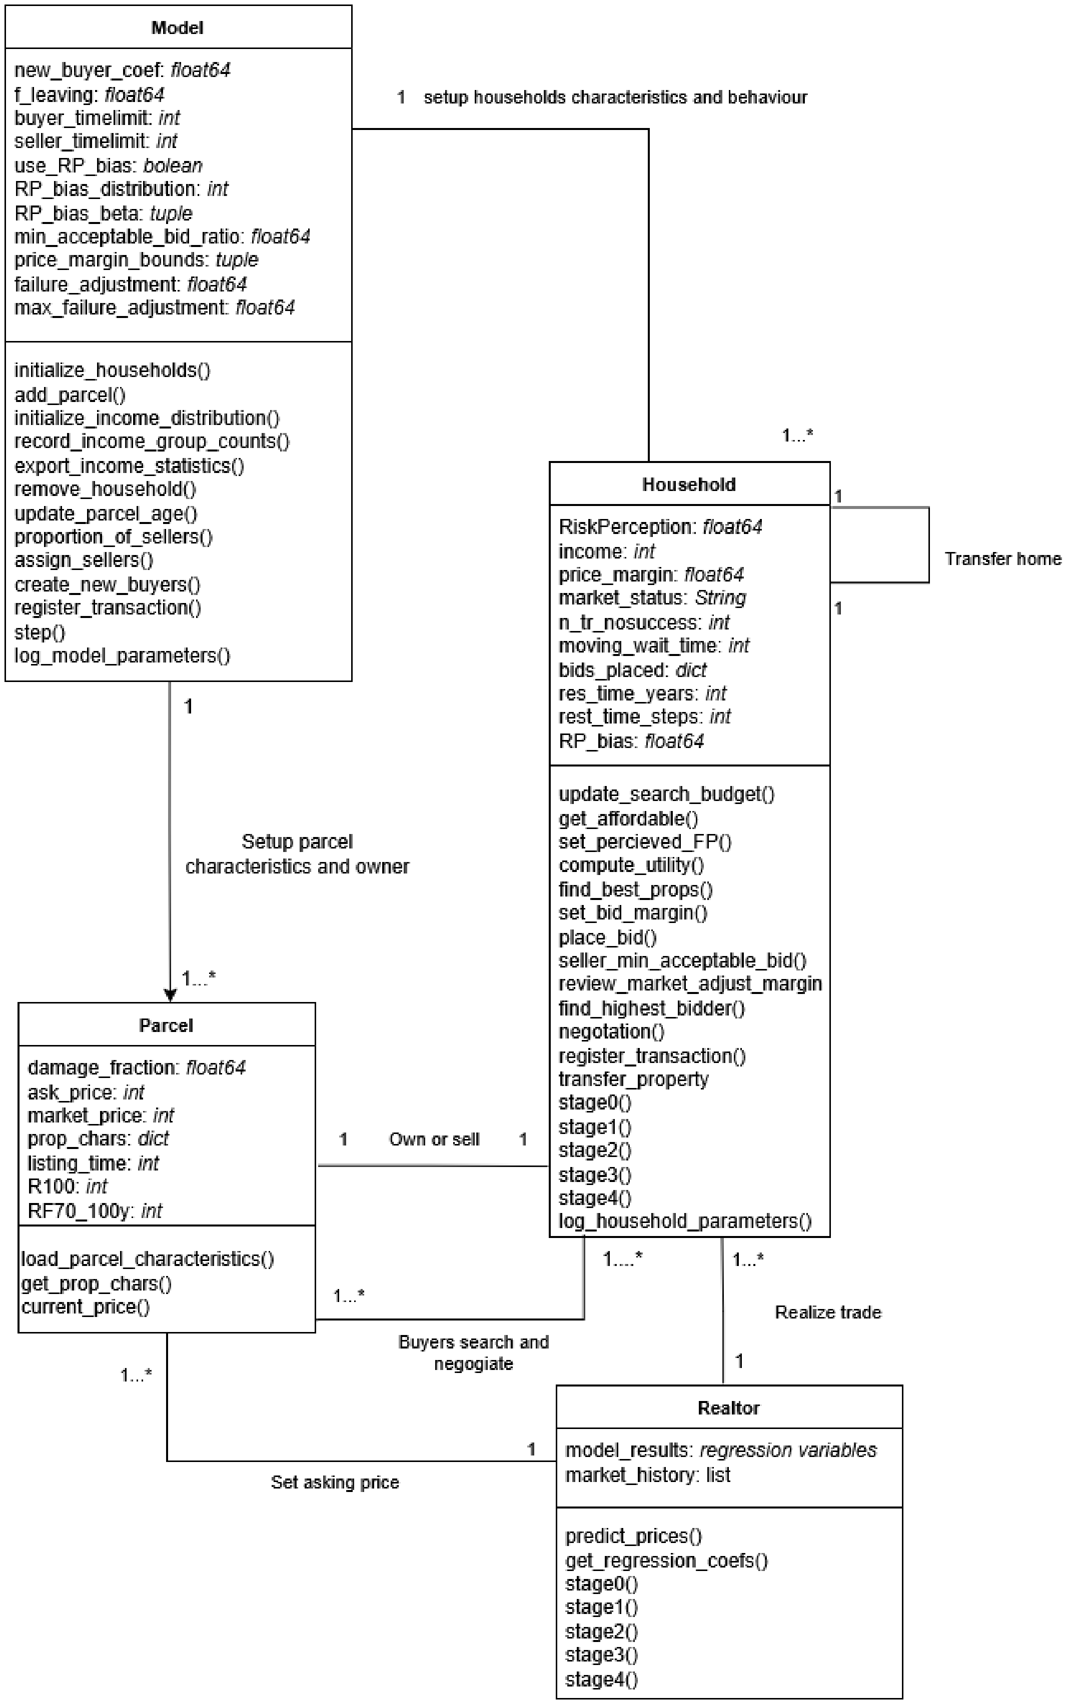

Figure F 1: reduced class design of RHEA

F.2.1.2 Translating towards a credit mindset

According to the previous analysis, the largest changes to the model would be in setting the budget, due to the fraction of income and bid margin inherently being a reinterpretation of the deposits. Thus, many aspects of the bidding process were easily translated to a more credit focused model. For example, the bid margin is obviously another way to simulate saving money, although it does heavily deviate from the deposit and mortgage system discussed previously, with it only being present implicitly. Moreover, the income fraction should also not be confused with the maximum allowed mortgage (see chapter 5.2), which actually increases when the income increases, contrary to what is modelled here (Warnaar et al., 2024). This can be supported by research highlighting that higher income groups tend to pay less of their income towards housing, as they are able to save more for a deposit (Van Der Drift et al., 2022),), resulting in the fraction of income decreasing. Thus, the adding credit under the assumption of maximizing mortgages increases this fraction, which is a debatable fact (Van Der Drift et al., 2022; Madsen, 2012; McQuinn & O'Reilly, 2007). In turn, under this new model design, the fraction of income put towards housing increases, rather than decreases, with income, making a fundamentally different choice within model design.

But the relationship between fraction of income to calculate was the inspiration to use these NIBUD regulations to calculate the household income during the initialization, with them being initialized under the assumption of maximum mortgage. Now at the same time, after some warming up in the model, this relationship will eventually no longer hold as the natural tendency of the model is to maximize the mortgages of lower income groups, while the higher income groups will rely more on deposits, as was discussed in the results (specifically 9.1.1). This discrepancy highlights an issue in modelling choices, with this iteration of the model simply being a choice necessary due to the large undertaking required before deposits could be modelled accurately.

The changing of the income fraction also proved far more difficult later on, as it became rather difficult to create the search function. Adding a changing interest rate, the budget can shift wildly due to the mortgage cap being reliant in two different ways on the interest rate. This became even more complicated once the damage fractions and the central bank behaviour was added, increasing the amount of possibilities even further. Therefore, it required something that was universal, which resulted in the choice to instead reverse the search function to make it about income instead. Simply put, the search can be simplified to:

$$Budget = f_{income} * income * 30 \qquad (F1)$$

Which can easily be changed to set the new range,

$$income = \frac{budget}{f_{income} * 30} \qquad (F2)$$

With the changes to the mortgage and the maximum fraction of income being calculated as:

$$inc_i = \frac{12 * AF_{base} * M_j^{pref}}{MF_j^{init}} \tag{F3}$$

This leaves one issue to be solved, as the maximum financing percentage MF_j^{init} requires an income before it can be calculated. But as these regulations set ranges for income to give the maximum fraction, these can then be used to pre-calculate a range of mortgages, which is calculated beforehand per interest rate so that they can then be used to once again reverse the calculation to get the maximum financing percentage to then give the precise income inc_i.This solution is then also universally applicable, with the interest rate during the initialization of households being calculated based on the assumption that the loan-to-value (LTV) ratio stays at 80%, the other influences on interest rate can be freely be changed while still keeping the calculation simple.

This knowledge was then used to instead set a range of possible LTV-ratios for which households would use to search, allowing them to naturally bid on homes, ensuring that bidding strategy is still possible. This also ensures that strategies between banks given (close to) credit maximization is still possible, as a households' budgets are based on the bank's credit strategy. Although this will limit the overall possibilities of this model later on if deposits were to be expanded, but it suffices given the current progress of the model.

A fix would be relatively simple, requiring the households to have different budget for each different damage fraction, from which then they will search. While this solution was initially proposed, the idea was rejected as it would require a relatively large amount of the compute budget due to the scaling issues. Each buyer for this version of the model would require 5 different budgets, meaning that with each increase in demand, the computing time would scale linearly with demand, while also requiring 5 times as many calculations.

The idea chosen instead, by setting budget by income, was mainly to save compute time, with it only scaling if more banks are added, instead of it scaling with demand. On top of this, due to the repetitive nature of these income calculations, many of them can be saved, reducing the actual computing load compared to the other solution, minimizing the increase when banks would be added. However, the same can be argued with the proposed budget solution, due to the many of the efficiency improvements laid out in the next section reducing the computational load.

It should therefore be noted that this choice was made before many of these improvements, meaning that the previous explanation under current conditions is not correct anymore, moreover making estimating the time saved difficult and most likely wrong. Because of the many savings mainly being related to the budget solution, it can even be that this solution is actually faster. This can be attributed to the many improvements in the calculations of the interest rate and mortgage.

F.2.2: Detailed Overview of Model

In this section, the model will be explained step by step from a coding standpoint, explaining the different variables, optimizations, and steps necessary for a full step to be completed. It will only give a general overview of the model already given but will go into more detail with the decisions related to the additions of the model. This section will use natural language to describe this process, leaving out many other details. For these other explanations, in section 6.3 the mathematics used can be found and in section 6.1 and 6.2 an overall outline of the model is given. These are assumed to be understood for the next section, while this appendix mainly applies to expanding on the class diagram, clarifying its structure, and reasoning. This diagram is given in figure F2.

F.2.2.1: Initialization of the classes

The model starts with the initialization of the banking regulator, a class that contains all the regulations and also does much of the calculations done by the banks. This was necessary to skip a large amount of repetitive computation as the banking classes at times function quite similarly, depending on their chosen strategy. The banking regulation class starts off by loading in the regulations file given by the NIBUD, which sets the maximum allowable financing percentage, translating it into a usable DataFrame. Moreover, it starts off with a base rate, from which the risk-free interest rate can be calculated, which this class from then on floats each step by a maximum of 8% from the start. From then on, the banks will be introduced, which are link with the banking regulation class and their chosen behaviour. After which the realtor will be introduced. Then the parcel data is loaded in, which includes the parcel characteristics, such as its water label, price, and its other qualities. These are then sorted in price, and each assigned a bank, based on its ordering to ensure all banks start off with an equal portfolio.

F.2.2.2: reversing the mortgage

The prices of these homes are then used to estimate the income of their owner. This process has been explained in detail in F2.1.2. These estimates are being done by the banking regulator class under the assumption that each household held a 20% deposit, with the interest rate given by the bank. This can be summarized in that the maximum financing percentage will need to be searched in reverse, rather than requiring income, it requires the mortgage to find the maximum financing percentage, which is then used to calculate the income. Because this process is used many times (later also in the search process), it has many little optimizations in it to ensure this process can be handled quickly. Firstly, a new data frame needs to be set up which holds the income ranges. Secondly, this is then combined with another cache, which saves the income and maximum financing with the key being the rounded mortgage to the thousands. This improves the speed when many banks have similar credit strategies. Both of these will then be saved in a dictionary with the key being the interest rate rounded to the hundredths of a percent, ensuring that it can quickly be found and the data frame does not expand too quickly. Moreover, the step will be saved in a dictionary with the interest rate as a key, which can later at the end of a step be used to reduce the overall size of this cache. The deposit as a percentage of income is also saved for when new buyers are introduced to estimate their deposit and is then used to recalculate homeowners' deposit for consistency. With the income, the household can now be created and given its property. Income distributions are then created, which were already a part of the base model and is thus skipped. These distributions are used to sample incomes when new buyers are introduced. Now all basic parameters are introduced, the model starts with its normal steps.

F.2.2.3: Start of a model step

The step starts off with assigning a random number of sellers, which will set up their parameters, such as minimum allowed bid and creating/assigning new buyers, with some sellers sometimes staying in the market afterwards. The realtor then uses these new homes to set their asking prices using hedonistic regression analysis. Afterwards, the banking regulator changes the interest rate and manages its cache, checking whether the cache has not been used often enough. Banks will then start off by getting the new central bank rate and setting their strategy by computing their Loss-Given-Default and Default rates for the different water labels. Moreover, the base annuity, so the annuity factor given 80% LTV and the safest water label will be saved for later reuse. Now the income targets can be set given an income range, as was previously described in F2.2.2 and F2.1.2. These ranges have been set to ensure that most of the bids will reach or be close to the 100% maximum allowed mortgages.

F.2.2.4: The bidding process

Households will start off with setting their bid margin, based on their success in the market, after which they are then searching for homes based on their income, giving them to a maximum of 10 parcels to compare. They start off by taking their subjective flood perception into account, decreasing its expected utility. As these buyers also have an expected mortgage payment in mind, based on 80% LTV ratio and a safe home (thus the advertised rate), they go to the bank, which prefers they bid on lower value and safer homes. Their utility change is equal to the change in mortgage, due to the changes in interest rate, with the monthly payment staying the same. Thus, their expectations are based on the previously mentioned base annuity, saving much time, as the calculation would be repeated 10 times per buyer. Now the interest rate given the LTV and the water label will need to be computed. Once again, due to many steps being required that need to use this function, it has been optimized. First, the loan risk is saved as the LTV ratio stays relatively consistent throughout the entire run, as such, this was chosen as a key rounded to the permille (1/1000) combined with the damage fraction. Rather than calculating the interest rate each run, this easily safes significant time. These caches will then be cleaned after each change in banking strategy. Moreover, the LTV ratio is in the permille instead of percentage, making it an integer. This avoids the slow calculation as the LTV ratio is used as the exponent, using this as a float would slow this process down. With the utility now also decreased by the changes in mortgage expectations, the buyers will choose 5 homes to bid on, after which these bids will need to be optimized. The bid will initially be checked whether it is 1% higher than what would be allowed, if this is not the case, the bid will need to be optimized, which is the most computationally intensive part of the model. Balancing the income ranges around the bid optimization, as well as the maximum mortgage was therefore vital. Nonetheless, this function required a lot of different optimizations to ensure that it would function efficiently, which will be discussed now.

F.2.2.5: Optimizing bid

This subsection will explore more modelling decisions and its resulting limitations concerning the bid optimization problem. This is caused due to the bid being higher than what is allowed. The bid optimization was rather difficult due to two opposing formulas. Firstly, the bid sets the interest rate, while the interest rate then sets the maximum mortgage. Therefore, if the bid increases, the interest rate and thus the mortgage decreases, resulting in an optimization problem if the bid is higher than the mortgage. In order to discuss this, imagine that the solution is simply 2 lines with on the X-axis the mortgage/bid minus the deposit, and on the y axis is the Interest rate. The intersection of those lines would be the optimal bid. Given these, the lines would be dictated by the following formulas from chapter 6.3:

Line 1:

This represents the bid, which due to the LTV ratio increases, which then increases the interest rate, this is in this example given by this simplified equation:

$$r_{i,j} = r_{free} + \left(PD_b^{init} * \beta^{ltv\,_{i,j} - 80}\right) * \left(LGD_b^{int} + \Delta LGD_{i.j}^{LTV}\right) \qquad (F4)$$

With the LTV ratio being the bid minus the deposit, divided by the house value, giving the relationship between the mortgage and the interest rate. This line can be simplified to a standard exponential curve.

Line 2:

This represents the maximum mortgage, with the interest rate being given by the previous formula, which can then be used to calculate the maximum mortgage through:

$$M_{i,j}^{base} = \frac{inc_i * MF_{i,j}}{AF_{i,j}^{base}} \qquad (F5)$$

With the Annuity factor AF being given by:

$$AF_{i,j}^{base} = \frac{r_{i,j} * \left(1 + \frac{r_{i,j}}{12}\right)^{360}}{\left(\frac{r_{i,j}}{12}\right)^{360} - 1} \qquad (F6)$$

The resulting line 2 will jump each 0.5% of interest due to the changing of the maximum fraction $MF_{i,j}$, causing sudden jumps and increases. Although its relationship can be seen as linear outside of these jumps.

Theoretically these two lines cross, which then gives the optimal mortgage, as then the interest rate given will give the optimal mortgage. This is the explanation in the main text, but it does not tell the entire story. Firstly, this would assume that it can simply be solved through normal equations by combining the two lines and solving the equation or by using a linear optimization algorithm, but there are issues making either solution impractical or far too slow for linear optimization. Firstly, concerning both solutions, the basic concept is rather simple, with the solution being:

$$M_{i,j}^{base} - Bid_{i,j}^{mortgage} = 0 \qquad (F7)$$

This simple formula can then be solved to give the formula with the optimal solution for all mortgage rates. There is, however, a massive flaw in this reasoning as this does assume every mortgage (x) can be solved. Because of the NIBUD regulations, represented by $MF_{i,j}$, the maximum mortgage jumps based on the interest rate thresholds. Meaning that for certain cases no optimal solution could be found if it is between this gap, meaning that it is not possible to always solve the equation, requiring a satisficing solution. Thus, only linear optimization would be possible, as they are created for market clearing problems which regularly have these jumps in their solution space. While they were initially implemented, they would prove to be rather slow, requiring too many iterations to find the optimal solutions. These methods are generally used with multiple input variables, making them inappropriate for these types of tasks. While some optimizations could be made within other methods and by increasing the maximum allowed solution difference, it would prove insufficient.

$$\left|M_{base,i,j} - bid_{mortgage,i,j}\right| \leq 1000 \tag{F5}$$

Thus, the final solution was created, wherein rather than finding the optimal solution and sometimes satisficing, instead the priority was to create a solution that satisfices and is quick. Using the linear optimization algorithms as a starting point, it can be observed that their fundamental reasoning relies on the invisible hand to come towards an equilibrium due to lines and markets working fundamentally similar (Samuelson, 1952). Using this logic, an extremely simplified algorithm can be created that finds a satisficing solution. By simply dividing the bid plus the maximum mortgage by 2 and taking this new number as the bid, if done infinitely the optimal answer will eventually be found, assuming there is one of course. Let us apply this logic to the problem at hand. The bid is 150k, while the maximum given the rate is 100k, assuming perfect opposite lines, 125k should be the answer. If one or both lines is steeper, flatter or exponential, the same natural tendency towards equilibrium still applies, the only difference being how often the solution will need to be divided. This can be done by simply using the new number as the bid, in this case 125k and repeating the same process, this new bid will then decrease the interest rate, thus increasing the maximum mortgage. Eventually this natural tendency will find the equilibrium, even in the case where the mortgage is now higher than the bid, as the same logic still applies. This will only work if the following conditions are met:

a) both lines are either linear or exponential
b) if there is only one solution possible.

From both formulas given above, it can be observed that both lines are either linear or exponential, thus condition a) passes. Condition b) will not pass as stated previously. This can simply be circumnavigated by reducing the maximum iterations to 4. Moreover, this solution also already fits into the current codebase, as the solution given bid (x) gives interest rate (y), which is then used to give the maximum mortgage (x). This circularity reuses already created methods, reducing the time spend creating the solution. This solution thus solves all the problems laid out previously. While the solution is not exact, is accurate enough to be functional, thus this solution satisfices.

F.2.2.6 Optimizing the maximum mortgage calculations

With the bid optimization being a large part of the overhead, other sections of which the previous function calls upon many times would also need to be revised. Some of these have already been discussed, such as the interest rate calculations under F2.2.4, however, most of them are related to calculating the maximum mortgage. As such, it was chosen to calculate these under the banking regulations class, this ensures that repetitive and similar calculations could be optimized for all banks and under higher demand conditions. Households safe their maximum financing percentage, decreasing the necessary search calls dramatically. Finally, the data frames containing the maximum financing percentage have been changed towards an assortment of 3 NumPy arrays, each containing the index, column, and data. While this change has been made redundant by the previous fix, it still saves time, and thus was left in.

After the bid was optimized, the best acceptable bids are selected by the sellers, after which the step is repeated. Including all optimizations to the model, the following class diagram was created:

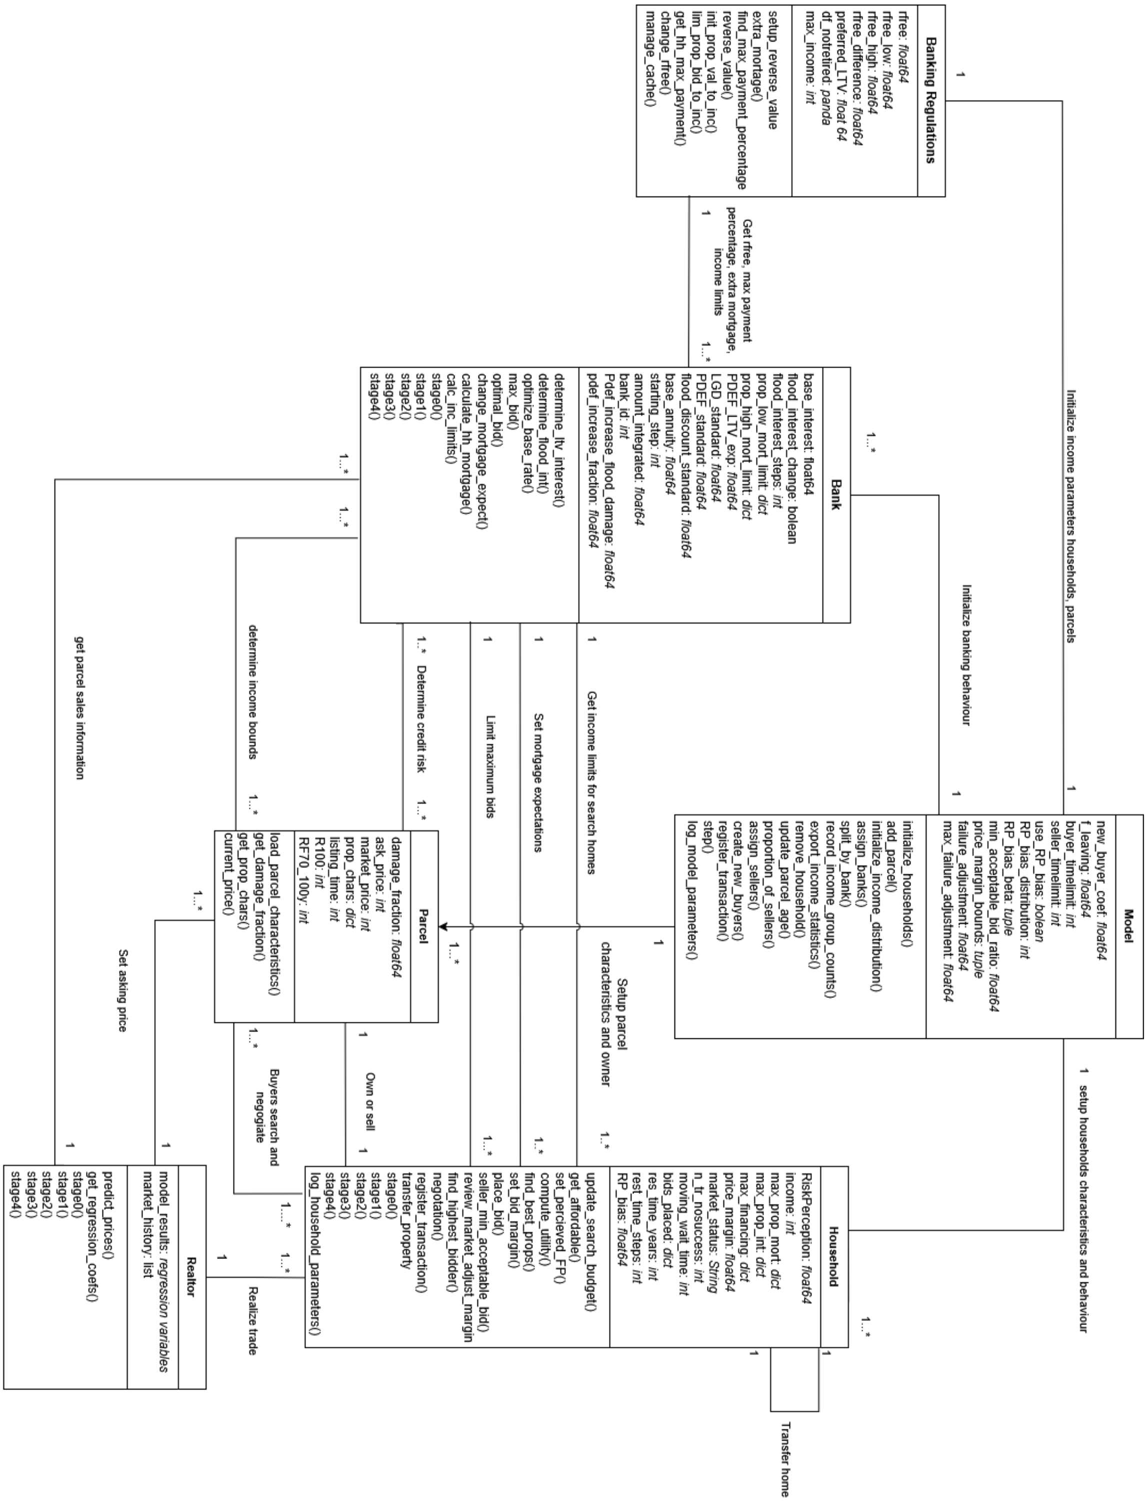

Figure F2: Class diagram of the entire model

F.3 Conclusion

Showing the improvements in runtime of the new model can be rather difficult due to the strong interconnection between the class design and optimization. Removing one aspect can quickly culminate into new bugs, making it difficult to establish a baseline performance. Moreover, during development, the codebase went through many iterations, sometimes removing entire methods entirely in favour of a more streamlined process. Therefore, actually stating with definitive numbers how important these changes are is difficult. This in conjunction with certain optimizations being helpful in an earlier iteration, they may become slower once other processes change or start being optimized further. However, during development it was not rare for the runtime to increase up to 2 to 3 times the total runtime when new features were added, requiring the need for changes. Modelling is an iterative process, and therefore these changes could not be anticipated beforehand.

To highlight the importance of this, the model was run in 4 different configurations over 10 runs for the 3 demand scenarios, with the average runtime being registered. All were given the same scenario, with rational risk integration. These configurations are:

Base model: The model without any changes, as was given.
Current model: The model with all the discussed changes.
Reduced model: The model with some of the optimizations removed. This does not include all of them, only those that were simple to remove. Those that are integral to model functioning have been left as is.
Older model: This version was created for the midterm in which a proof of concept was necessary. This only calculates the interest rate through basic means; by simply adding a specified amount of basis points, it includes the regulations, only has one bank, and contains the utility function as described. It therefore misses many aspects added later on, much of which is incredibly taxing.

Table F1: The mean runtime of the models

	BASE	CURRENT	REDUCED	OLDER
LOW DEMAND	13.2 seconds	17.8s (+34.8%)	18.8s (+42%)	19.1s (+44.7%)
HIGH DEMAND	16.7 seconds	24.0s (+43.7%)	26.9s (+61.1%)	27.9s (+67.1%)
VERY HIGH DEMAND	19.8 seconds	31.1s (+86.2%)	32.5s (+94.6%)	38.0s (+128.5%)

Appendix G: AI Usage Disclosure

During many of the processes of this thesis AI has been used in numerous ways to improve the end result. However, discussion around AI usage have highlighted many pitfalls that are the result of usage of this new technology. As integrating it in these types of learning processes has led to mixed results and doubts about its efficacy due to their links in reducing a students' abilities and their tendency to dream causing wrong conclusions with made up sources. Rather than avoid the topic, it was chosen to disclose its usage and discuss the different rules set up to ensure that this new technology is used properly. AI was chosen to be used as it was believed that, while their current capabilities lead to mixed results, even during the making of this thesis major progress has been made, leading to the expectation that this new technology will, at the very least, become an important new cornerstone during research, writing and coding. By opening up this discussion, it is ensured that oversight is being kept on proper usage and that AI usage will not only lead to a better outcome but will also improve the research process rather than stifle it. Through its proper usage it can lead to a better and more creative thesis, as it can easily be used to critique and improve certain aspects quickly. But, in its current state, AI does require a certain level of oversight before becoming useful.

In this section, the general use cases of AI will be discussed and the rules set up to ensure that their usage will lead to a better outcome. These 3 general usage cases are defined as:

- Supporting the research process
- Improving the software implementation
- Assisting in the writing process of this thesis

Rather than replace human understanding of the subject matter, it was chosen to set up general rules that would bound AI usage during this project. These rules were set up to be rather conservative in their nature caused by scepticism to the current state of this new technology. This conservatism was leading in ensuring that, while there was still room for exploration of its usage, it ensured that if large choices were made using AI, they were made with considerations in mind that would either lead to a better or a faster outcome. These were set up as 3 rules:

- Enhance creativity
- Improve efficiency
- Support understanding

These 3 rules were set up to avoid the AI trap, which is caused by the AI functioning more as a black box for its user, leading to difficulty in verifying the results. As AI has a tendency to dream, it can become difficult to understand why it makes certain decisions (i.e. whether the results were part of bad data or poor training). This fundamental issue with AI can make the usage of this technology quite difficult, especially with to how early this technology still is in its development cycle. Moreover, the complexity of the subjects discussed in this thesis make dreaming more impactful, as the results cannot easily be verified, potentially leading to wrong decisions being made or mistakes being made that would fundamentally alter the understanding of the water label. Bounds were therefore necessary, leading it being used mostly with concepts already understood or for exploring new ways to understand certain concepts. Verifiability thus became the most essential part of its usage with the general idea being it is easy to make sense but difficult to be right. How these rules have been used in the different use cases will now be through the different use cases. This appendix will mostly explore the experiences learned while using AI and how its usage evolved over the course of this thesis, through discussing its positives and limitations as well as some surprising discoveries along the way.

G.1 Supporting Research Process

With the research process being the start of every new addition, it is most sensitive to mistakes caused by dreaming and therefore more caution was taken to avoid this as much as possible, with each small detail being research fields on itself, these tiny decisions could potentially cause more issues down the line. Moreover, with understanding of each new addition during this process being rather low, it can be relatively easy to fall into the AI trap, thus in this section it was chosen to avoid the usage of AI as much as possible, only allowing it in specific use cases. Adhering to tight rules here is especially important, which helps create behavioural patterns helping avoid these traps. For example, if AI were used, everything said would need to be verified by finding the papers discussed. The answers early on needed to explore mostly general concepts of these fields, rather than extremely specific questions, as it was found out that the more specific a question, the more likely the result would be wrong. This could be circumnavigated by feeding the papers and then asking for summaries of said paper to gather information. However, this was also avoided as much as possible as the results would need to be verified before being seen as usable.

During the deeper research process, when the most important papers were found and where a deeper understanding of each paper was necessary, the AI trap now being solved. Due to a deep understanding of the subjects being necessary to understand when dreaming is happening, it can now be helpful to summarize or to use it to explain the concepts that are not understood by feeding it the papers in question. With this, it could help gain a wider understanding of this field if language used was too obtuse or its concepts were difficult to understand. Although this was only used a handful of times, it was helpful when certain math was not well understood, although this was only used on one occasion. Moreover, in this case it was only used example results were given, so that the results could be verified.

Overall, this meant that AI was not an important part of this process, as simple learning is much more important. Only once it is found that certain details were not well understood, AI can be useful to function as a teacher on call. Even then it was rarely used as once the papers were understood, it quickly lost its purpose. Due to the nature of a thesis, a deep understanding of the nuances in the literature is necessary, otherwise at later points of the process the mistakes and misunderstandings have the potential to build up. Although other forms of AI, such as deep research could make this process easier, these were not used in fear of creating the issues mentioned before.

G.2 Improving the Software Implementation

During the coding portion of this thesis AI was used significantly in multiple steps of this process. This was due to the fact that an agent was found to be mostly helpful during certain processes, while at other times could cause a lot of problems. To ensure that the most benefits can be gained from its usage, certain things had to be kept in mind to avoid AI traps. For example, from the resulting code it was clear that it was mostly trained on a data set, mainly containing more standard credit models. This caused the agent to be very heavily biased in a certain direction, causing very consistent errors that could easily be fixed. Many of these were caused by the code requiring the result to be normalized from 0 to 1, for which can then easily be accounted. At the same time, errors would commonly mix certain solutions, which at times could be helpful. As it is trained on a data set wherein the code is written with a certain reasoning, this could often times be reversed, helping in developing new ideas. Although this was rare, in most cases it led to better understanding of coding practices.

Other times it would create something that was too difficult to understand. At times like this, it was important to understand why these choices were made, and if they were seen as useful, it was chosen to use it only if the logic was inferred. Sometimes these could be used to recreate the idea and change it to fit the code base better. For the usage of AI agents, this would be all the use cases in terms of coding it was found to be helpful. It was mostly found that it could help reduce the time spend looking up and choosing the appropriate methods. While at a lot of other tasks, the agents were found to rarely be useful.

With these limitations, work arounds could easily be made, with all of the code structures being man-made, AI could be used to fill in the gaps. Giving it pseudo code with a description of the steps necessary, which methods already available in the code base, general concepts of methods that need to be used and the necessary results the code was often usable, only requiring some changes. This still required much preparation beforehand to ensure that the formulas as well as the general code structure made sense.

For bug testing the agent was found to be rather difficult to work with. It would often randomly change certain aspects and missed obvious bugs. Rather than fix it, it would merely change the code, rarely creating something correct. Other times it would change other parts of the code, creating random bugs. Thus, the agent was only allowed to create code, rather than change it. Instead, it was helpful to develop more personal bug testing skills, asking which features to use to diagnose the issues. Although this was also found to be unhelpful at times. Simply put, the bug testing required for creating a model is different from normal coding, another issue created by poor training data.

With a large part of the thesis requiring incredibly efficient code, at this area, the AI agent was also found to be insufficient. Once again, it could be used to diagnose issues rather well, showing and explaining certain features quite well. It did fail at giving different ways to write more efficient code, using rather simple techniques that would very rarely result in little improvements.

When creating graphs, the AI agents was found to be especially helpful in creating the necessary code. In most cases it was found to be using the correct methods and rarely made mistakes. While these were still present, and at times the result was too complicated, they could easily be fixed through standardizing with some methods. Moreover, mistakes could easily be verified by comparing graphs, since most of them are different ways of presenting similar enough data.

G.3 Assisting in the Writing Process

At no point during the writing of this thesis was generative AI used and was at all points avoided. All the reasoning and structure in the texts presented are my own. An AI writing assistant was, however, used. The most important goal of their usage was to improve the legibility of the text by reducing the repetitious words and sentences and improve sentence structure. This writing tool has significantly helped to improve these, and its criticisms have been used to improve writing later on. These tools do have their own limitations such as their tendency to homogenize language, making it have a style on its own. All people have their own writing style, and these should not get lost through their usage. Most importantly this could also cause more niche words to get missed and replaced by more generally accepted words, meaning that words used in a specific way within a scientific field might get replaced by something that would entirely change its meaning. At all times, these limitations have been avoided, and it has not been used replace the learning of language skills. By choosing which prompts to follow or change the sentence structure based on the feedback these were avoided. The most important part of writing this thesis is to communicate my ideas and my understanding of other people’s work.