Erasmus University Rotterdam

Climate-related physical risk and Dutch property prices: a Synthetic Control Method approach.

Author

J.M. Fennema (580454)

Faculty

Erasmus School of Economics

Master Program

Quantitative Finance

University supervisor

PH. B.F. Franses

Second assessor

B. van Os

Company supervisors

F. van Essen & T. Greeven

Abstract

This paper aims to quantify the impact of climate-related physical risks on average neighborhood property prices in the Netherlands. We use a Hedonic Pricing Model (HPM) to show that the impact of the presence of flood and foundation risk on Dutch neighborhood property prices is limited. On a countrywide level, the results indicate a negative *ceteris paribus* effect on average neighborhood property prices for flood risk and a positive effect for foundation risk. Acknowledging limitations in the dataset and model design of the HPM and understanding the importance of information availability for changes in consumer behavior, we introduce a novel application of the Synthetic Control Method (SCM) to study local risk awareness effects on Dutch property prices. For flood risk, we find that the inundations in Limburg in 2021 had a significant negative effect on property prices in the neighborhood Bunde, ranging from approximately 1.5% to 4.0%. For foundation risk, the results show that the 2021 initiatives by the municipalities of Schiedam, Dordrecht, and Rotterdam were ineffective at raising awareness to have a significant impact on property prices in Walvisbuurt in Schiedam.

8th May 2023

Acknowledgements

I would like to express my sincere gratitude to prof. Franses for his support throughout my thesis journey. I am grateful for his guidance, thought-provoking questions, and considerate attitude that always made me feel comfortable approaching him with any problem I faced.

I would also like to thank my Deloitte supervisors, Femke and Tom, for their knowledge, enthusiasm, and friendship. They managed to keep me calm during stressful times and cheer me up whenever I needed it the most. I always looked forward to the weekly Friday progress meetings and could not have finished my thesis without them.

Last, a special thanks goes out to my family and friends. They all had to put up with my endless stories about floods, foundations, risk, and property prices. Especially Doris; without her patience and care, I would be writing this thesis in gibberish.

Contents

List of Figures

List of Tables

Acronyms

BA	Bucketing Analysis. v, 17, 18, 24, 26, 27, 37
CBS	Centraal Bureau voor de Statistiek. 2, 5, 6, 37
CCA	Complete Case Analysis. 6, 8
HPM	Hedonic Pricing Model. iv, v, 2–9, 14, 17, 18, 23, 24, 27, 28, 36, 37, 43, 45
LIWO	Landelijk Informatiesysteem Water en Overstromingen. 2, 10
LTV	Loan to Value. 1–3
MCAR	Missing Completely At Random. 6, 8, 37
MLR	Multiple Linear Regression. 18
MSPE	Mean Squared Prediction Error. 21–23, 30–32
OLS	Ordinary Least Squares. 6, 17, 37
PC	Principal Component. v, 18, 24, 46
PCA	Principle Component Analysis. 18
PCR	Principle Component Regression. v, 18, 24–27
RMSE	Root Mean Square Error. 25
SCM	Synthetic Control Method. v, 2, 3, 5–8, 12, 13, 15–17, 19–23, 26, 28–30, 32, 33, 35, 36, 38, 39, 43, 45
WOZ	Waardering Onroerende Zaken. 6, 26

1 Introduction

Climate change is causing one of the most significant structural changes in the world today. As first addressed by Meadows et al. (1972), their scenario-based approach showed that our use of Earth's resources was unsustainable and damaging to life on Earth. In 2015, the Paris Agreement finally set goals to limit global warming to below 2 degrees Celsius. A huge energy transition is necessary to reach this goal, which governments try to speed up using different climate policies. These policies, and the effect of global warming in general, are expected to destabilize the financial system. Mark Carney, former governor of the Bank of England and prominent figure in the financial industry, was one of the first to highlight the effect of global warming on the financial system in a speech to the insurance market Lloyd's of London: "Alongside major technological, demographic and political shifts, our very world is changing. Shifts in our climate bring potentially profound implications for insurers, financial stability, and the economy." (Carney, 2015)

Nowadays, these climate-related risks are split into so-called physical risks and transition risks. With transition risks, we refer to risks arising from transitioning towards a lower-carbon economy (Greeven, 2021). For example, one can think of the impact of a carbon tax on the financial situation of shipping companies, which are renowned for emitting a lot of CO_2 (van Essen, 2022). Physical risks are related to physical climatic events. For example, damages caused by extreme droughts and flooding and indirect effects like foundation risk, accelerated by prolonged droughts due to climate change (Kok & Angelova, 2020).

For many financial institutions, a mortgage portfolio makes up a substantial portion of the balance sheet. As such, the mortgage industry is an important channel through which climate-related risks impact financial stability. A perfect example of how the mortgage industry can destabilize the financial system can be found in the global financial crisis in 2008. Ahead of this crisis, near the end of the US housing bubble, people were buying overpriced houses with subprime mortgages (Barberis, 2013). Unfortunately, this meant that people with lower credit scores took out loans with very high Loan to Value (LTV) ratios. The LTV ratio is a financial risk driver used by banks and other financial institutions to assess lending risk. In the case of mortgages, we express it as the ratio of the outstanding loan amount and the appraised property value. If a house is being purchased, the property appraisal determines its value. This value can be higher, lower, or match the house's purchase price. As described by Bian et al. (2018), this means that transaction prices can greatly exceed collateral values in a bullish housing market, where sellers have a stronger bargaining position. When this happens, a mortgage is "under water", greatly increasing its credit risk. This happened in 2008, causing many people to default on their mortgages when the housing bubble burst. Consequently, banks failed, which in turn caused the aforementioned global crisis.

Although an extreme example, it shows that overpriced houses can eventually pose a great risk to the stability of the financial system. This can be caused by an overheated housing market but also by unobserved physical risks as a result of climate change. For example, when foundation risk comes to light, it intuitively hurts a property's value. Put in numbers, damages associated with foundation risk are expected to cost the Netherlands between 8 billion and 54 billion[1] euros by 2050 (Kok & Angelova, 2020). The possible effects these risks eventually pose to businesses and society in the Netherlands have sparked interest in academics (Bosker et al., 2019), politics (Bosker et al., 2016), and financial institutions (Hommes et al., 2023; Reeken & Phlippen, 2022).

This paper builds on these previous works by investigating climate-related physical risks in the Netherlands. However, to narrow the scope of this research, we set the focus on two climate-related physical risk types. As almost a third of the Netherlands lies below (a rising) sea level, and subsidence in cities like Rotterdam is worsening rapidly (Kok & Angelova, 2020), we choose flood and foundation as the research focus. More specifically, the impact of flood and foundation risk on Dutch neighborhood property prices. As noted by Wachinger et al. (2013) in his paper about risk perception, "risks cannot be 'perceived' in the sense of being taken up by the human senses, as are images of real phenomena." This raises the question of whether risk alone will be reflected in property prices. As such, this paper also studies how the awareness of flood and foundation risk impacts Dutch neighborhood property prices.

This leads to the following two research questions:

1. How do flood and foundation risk impact Dutch neighborhood property prices?

2. How does the awareness of flood and foundation risk impact Dutch neighborhood property prices?

First, Section 2 presents previous studies related to these research questions. Similar to the setup we use to explain the relevance of this paper, it starts with research on climate change and financial stability in Section 2.1. Next, we identify the LTV risk driver in credit risk models for mortgage portfolios as an important channel through which climate change impacts financial stability. Consequently, Section 2.2 presents previous works that look into the impact of climate change on the value component of the Loan to Value. Adding to the research on flood and foundation risk in Sections 2.2.1 and 2.2.2 respectively, Section 2.3 discusses papers on the Synthetic Control Method (SCM), in which we find a novel approach to studying the impact of climate-related physical risks on property prices.

Then, Section 3 describes the different datasets we use to answer the two research questions. For this research, we use three public datasets from two different sources. The first source, the Centraal Bureau voor de Statistiek (CBS), collects spatial neighborhood data on various topics. Section 3.1 presents the datasets that include these neighborhood figures. As the two models take advantage of these datasets in different ways, Section 3.1 is split in Section 3.1.1 and 3.1.2. The second source, *klimaateffectatlas.nl*, collects spatial data about climate-related risks in the Netherlands from government-supported organizations like Deltares and Landelijk Informatiesysteem Water en Overstromingen (LIWO). Accordingly, Sections 3.2 and 3.3 cover the flood and foundation risk datasets, respectively.

In Section 4, we outline the econometric techniques that help us find answers to the research questions. We use a different model for each research question, as discussed in Sections 4.1 and 4.2. In Section 4.1, we present a common way to study the impact of housing characteristics on its value in the form of the Hedonic Pricing Model (HPM). Similar to Daniel et al. (2009) and Bosker et al. (2019), we utilize a Hedonic Pricing Model to study the impact of flood risk on property prices in the Netherlands.

[1] Of which 3 to 15 billion are additional costs due to climate change.

Additionally, we use the HPM to investigate the impact of foundation risk. Section 4.1 covers the details of this implementation in more depth. Section 4.2 presents the mathematical theory behind the Synthetic Control Method. Using the SCM, we approach the second research question from a new angle. This method was introduced by Abadie et al. (2010) for inference on the effectiveness of policy changes but finds a new application in this context.

We discuss the results of implementing these methods in Section 5. We use the same split, with Section 5.1 covering the HPM and Section 5.2 presenting the findings to the SCM.

Section 6 presents answers to the research questions. Sections 6.1 and 6.2 include a discussion and pose ideas for further research on climate-related physical risks and property prices using the Hedonic Pricing Model and Synthetic Control Method, respectively.

2 Literature Review

Research about climate change and financial stability is often categorized by the climate risk type it studies. Out of physical and transition risk, we focus on the former. First, we look at papers that discuss climate change and financial stability on a macroeconomic level in Section 2.1. Then, we consider the LTV risk driver, which creates a link between mortgages and credit risk. Moreover, the dependence of the LTV value component on property prices connects climate risk and property prices. This paper considers two climate-related physical risks: flood and foundation risk. In Section 2.2, we present the existing research on flood and foundation risk separately. In Section 2.3, we cover the original paper on the Synthetic Control Method by Abadie et al. (2010). Due to our novel interpretation of this method, there is no previous research relating it to climate risk yet.

2.1 Climate Change and Financial Stability

Research papers about climate change and financial stability often point to the role of regulators (Campiglio et al., 2018). The Basel Committee on Banking Supervision is one of these regulators. In June of last year, the Basel Committee presented the "Principles for the effective management and supervision of climate-related financial risks" report[2]. In this report, the committee presents 18 principles to deal with climate-related financial risks. One of these principles concerns credit risk. A paper by Monnin (2018) identifies three channels through which climate change impacts a borrower's credit risk: cash flow, financial wealth, and collateral value. In their book on credit risk models, Bluhm et al. (2016) mention the Loan to Value risk driver is widely used in credit risk models. Since the LTV ratio incorporates the collateral value, which Monnin (2018) labels as an impact channel for climate risk, it captures the impact of climate change on a borrower's credit risk. As such, the LTV links property prices to credit risk.

2.2 Climate Risk and Property Prices

Not all climate-related physical risks impact property prices in the same way. As such, most research focuses on a single physical risk type. In Section 2.2.1 and 2.2.2, we present previous works on flood and

[2]The report can be found using the following link: https://www.bis.org/bcbs/publ/d532.pdf

foundation risk separately.

2.2.1 Flood Risk

When papers study the impact of flood risk on property prices based on a flooding event, it binds them to specific flooded regions. For example, Atreya and Ferreira (2015) and Atreya et al. (2013) studied the city of Albany in Georgia, United States. Combining a difference-in-differences model with a HPM to control for spatial variation effects, they find that properties in flooded areas received a larger price discount than similar properties in a nearby region that did not flood, but that the effect fades over time. Approximately four to nine years after the flooding event, the difference in property prices between the different areas had vanished.

As we focus on the Netherlands in this paper, we are interested in previous works on flood risk and Dutch property prices. Two papers, by Daniel et al. (2009) and Bosker et al. (2019), use HPMs to investigate the impact of flood risk on property prices in the Netherlands.

The first half of the paper by Daniel et al. (2009) compares the results of previous studies on flood risk and property prices in the United States. They argue that the conflicting results of previous works are due to omitted variable bias. To circumvent omitted variable bias, Daniel et al. (2009) include a river proximity and water surface variable. Additionally, they try to control for spatial variation effects by using highly granular data of six-digit postal codes. For flood-prone municipalities in Limburg around the 1993 and 1995 floods, they find a significant negative effect of flood risk on property prices, accumulating with each flood to 9.1% after the second flood in 1995. Moreover, they showed a positive effect of river proximity on property prices which attenuated the flood risk effects and validated the omitted variable bias concerns.

A recent study by Bosker et al. (2019) takes a different approach to the omitted variable bias in a HPM. Instead of adding the variables that comprise the positive water proximity effect, it controls for them. By using a Border Discontinuity type Design (BDD), they isolated areas with the same distance to a river but different flood risk values. Such a control group that is safe from floods tries to control for any omitted variable bias, creating a fair comparison between similar areas that only differ in their flood risk. By executing this approach, Bosker et al. (2019) find a significant negative 1.0% effect of flood risk on property prices throughout the Netherlands.

2.2.2 Foundation Risk

The number of previous studies that look into foundation risk is relatively small. Still, Foundation risk is a climate-related physical risk type that is particularly relevant in the Netherlands. Section 3.3 explains this in more depth. Unsurprisingly, the research we discuss is done by the Dutch bank ABN AMRO. In their paper, Hommes et al. (2023) examine foundation risk and its impact on property prices in the Netherlands. Their approach is based on a comparison between realized property transaction prices and the property's value according to a machine learning based automated valuation model by Brainbay. In their comparison, they identified transactions in which the house sale advertisement either listed foundation damage, a repaired foundation, or did not mention the foundation conditions. This allows them to isolate the effect of reported foundation damage on the transaction price by comparing it to the automated valuation model value. Hommes et al. (2023) find that properties mentioning foundation damage in their sale advertisement were sold at a price that was on average 12% lower than their automated valuation model value. Reversely, properties that listed a repaired foundation were sold at a price that was on average 2% higher than their model value. In absolute numbers, this comes down to a discount of 47,500 euros and a premium of 13,500 euros, respectively. This shows that, using average repair costs of 50,000 euros[3], foundation issues are priced in - if they are known during the home buying process (Koopman, 2020). However, the work and/or stress related to the foundation repair process is not priced in, as the roughly 60,000 euros shows no premium on top of the repair costs listed above. Furthermore, the fact that for many properties these damages are not mentioned or unknown causes this risk to remain unpriced. According to Bunni (2003), underestimation of risk by the insurance sector can ultimately destabilize the financial sector, further demonstrating the relevance of research into foundation risk.

2.3 Synthetic Control Method

This paper proposes a novel approach to investigating the impact of climate-related physical risks on property prices in the Netherlands. The idea is similar to the difference-in-differences approach by Atreya and Ferreira (2015) and Atreya et al. (2013), which studies the development of property prices over time. The difference in execution is based on our new implementation of the Synthetic Control Method. As such, we present the fundamental SCM paper by Abadie et al. (2010). In their research, they find that it is possible to study the impact of a local event on a region by constructing a Synthetic Control group that acts as a counterfactual comparison group. Abadie et al. (2010) uses the passage of Proposition 99, a major anti-smoking law in California that was enacted by California voters in November 1988, as such a local event. The results of this method show that, as the effect accumulated over the years, the average cigarette sales per capita in the year 2000 in California was 26 packs lower than the counterfactual case without Proposition 99. These significant results display the power of the Synthetic Control Method. In a later paper, Abadie (2021) identifies the versatility of his method, with applications by other authors ranging from analyzing the passage of gun laws, migration and asylum policy, and changes to the tax system.

3 Data

For this research, we use publicly available data from two different sources. That is, the Centraal Bureau voor de Statistiek for the property prices and other neighborhood figures, and *klimaateffectatlas.nl* for the climate-related physical risk data. As the execution of the HPM and SCM both require both data sources, we discuss each data source separately. As such, Sections 3.1, 3.2 and 3.3 cover the neighborhood figures, flood and foundation risk data, respectively. As the HPM and SCM take advantage of the neighborhood figures in a different manner, we split Section 3.1 in Sections 3.1.1 and 3.1.2. Then, in Sections 3.2 and

[3]Repair costs range from 1.000 euros to 1,750 euros per m^2, multiplied by the average ground surface of 50 m^2

3.3, we briefly describe the dataset construction process, after which we provide descriptive statistics on each climate-related physical risk type. Even though we discuss the methodology of the Synthetic Control Method in Section 4.2, Sections 3.2 and 3.3 conclude with the neighborhood selection procedure for the SCM.

3.1 Neighborhood Figures

The CBS neighborhood figures come from the yearly published *Kerncijfers Wijken en Buurten* datasets. For the HPM in Section 3.1.1, we utilize the *Kerncijfers Wijken en Buurten* in cross-section for the year 2021. For the SCM in Section 3.1.2, we need time series data. As such, we gather *Kerncijfers Wijken en Buurten* data from 2013 - 2022. In Section 3.1.1 and Section 3.1.2, we discuss the pre-processing procedure for the HPM and SCM respectively, after which we put forward general descriptive statistics of the data.

Intermezzo: WOZ valuation

For the analysis of property prices in the Netherlands, we are bound to the average neighborhood property price included in the *Kerncijfers Wijken en Buurten* datasets. In *Kerncijfers Wijken en Buurten*, the average neighborhood property prices are denoted by *38_g_woz*. In *38_g_woz*, Waardering Onroerende Zaken (WOZ) refers to the WOZ value of a property. The WOZ value is the legally binding value of a property. The WOZ legislation is a law that sets up municipalities with the responsibility to appraise all real estate within its borders on a yearly basis. Due to the infeasibility of manually appraising all real estate in a neighborhood, estimation models are being used. Along with several house characteristics, such as living space, number of bed- and bathrooms, and construction year, these estimation models use recent sales of comparable properties as input to attach a price tag to a property in a neighborhood. A time frame of one year before and half a year after the valuation date (January 1^{st} of each year) is taken into account to come up with an accurate estimate of the property's true value. As such, the property values on which *38_g_woz* is based, are assumed to be a good approximation of the market value of a property on the first day of each year.

3.1.1 Hedonic Pricing Model

The 2021 *Kerncijfers Wijken en Buurten* dataset contains 17,681 rows, with information on neighborhoods, municipalities, and the Netherlands as a whole. As becomes clear in Sections 4.1 and 4.2, the execution of the HPM and SCM is dependent on variable *38_g_woz*. For both the dependent variable *38_g_woz* and the independent variables in *Kerncijfers Wijken en Buurten*, we find missing entries. As it is the responsibility of each neighborhood individually to acquire and deliver these figures to the CBS, we assume missing entries to be Missing Completely At Random (MCAR)[4]. According to Soley-Bori et al. (2013), performing Ordinary Least Squares (OLS) using a Complete Case Analysis (CCA) remains consistent under this assumption. In a CCA, observations with missing entries are ignored (Little, 1992).

[4]Based on a conversation with a CBS employee during a phone call on the 12^{th} of March 2023

As such, we exclude the 1,989 rows that do not contain a value for *38_g_woz*, which reduces the number of neighborhoods in the dataset to 15,692. The *Kerncijfers Wijken en Buurten* dataset has several variables related to the location specification of a neighborhood. For example, *1_gwb_code_10*, which is a unique and neighborhood-specific code. The purpose of these neighborhood codes is threefold:

1. It links the values from the climate-related physical risk datasets from *klimaateffectatlas.nl*, as discussed in Section 3.2 and 3.3, to the 2021 *Kerncijfers Wijken en Buurten* dataset
2. It acts as a label to the observations in each *Kerncijfers Wijken en Buurten* dataset, such that different years of dependent and independent variables can be linked to create the time series necessary to perform the SCM, as described in Section 3.1.2
3. To make inferences on peculiarities and results from (a group of) data entries

None of the data entries have missing values for *1_gwb_code_10*. However, some neighborhood codes change compared to the previous year, as indicated by the *7_ind_wbi* variable. Such neighborhoods lose their value for the HPM, as follows from the first neighborhood code functionality. As such, we exclude these neighborhoods from the dataset. This leaves us with 14,123 neighborhoods. The variable *5_recs* specifies the region type of an area. As mentioned by Bosker et al. (2019), the level of data granularity is preferred to be as high as possible for the HPM. Consequently, we remove the data entries not of the highest granularity level (i.e., neighborhood level), leaving us with 11,278 rows in the dataset.

Looking at the 120 columns, we identify multiple variables that are either uninformative or contain very high levels of 'missingness'. Examples of variables that we classify as uninformative for performing the HPM are *118_pst_dekp*, *7_ind_wbi* - after it served its purpose as described neighborhood filter above - and *5_recs*. We classify very high levels of missingness as being variables with missingness > 90%. After we exclude all of these variables that contain either very high levels of missingness or are uninformative, we are left with 90 neighborhood attributes in the dataset. Of these 90 attributes, four are spatial label variables like *1_gwb_code_10*, and one is the dependent variable *38_g_woz*. This leaves 85 explanatory variables for the HPM. We provide an overview of all attributes from the *Kerncijfers Wijken en Buurten* dataset and whether they are included in the HPM, in Appendix I.

We do a short analysis of these 85 neighborhood attributes to see how they relate to each other. The correlation matrix shows high degrees of association between variables. Appendix II includes this matrix in full. Section 4 further explains the impact of correlation on the proposed methods and how to prevent multicollinearity from reducing the precision of future estimates.

3.1.2 Synthetic Control Method

As the Synthetic Control Method requires time series data, we gather ten years of *Kerncijfers Wijken en Buurten* datasets. For reasons we discuss in Section 4.2, we choose the 2013 - 2022 period. The composition of these datasets is the same over the years, which makes them suitable for the SCM time series procedure. As a result, we apply a similar data pre-processing procedure as the one described in Section 3.1.1. Nonetheless, the neighborhood attributes in the *Kerncijfers Wijken en Buurten* datasets vary slightly over the years. The absence of neighborhood attributes in one or more years cause a discontinuity that forces us to exclude it from the analysis. As such, the number of neighborhood attributes that span the entire 2013 - 2022 period is lower than the number of neighborhood attributes in *Kerncijfers Wijken en Buurten* for a single year. As mentioned before, we assume missing entries to be MCAR, which turns out to be both a blessing and a curse. A curse, because - as expected - different entries are missing every year, which further decreases the number of neighborhoods in the cross-section. A blessing, because it is in agreement with our assumption and motivates us to continue the CCA we started in the HPM in Section 3.1.1. Subsequently, we end up with 2,750 neighborhoods and 55 neighborhood attributes in the dataset. Four of these 55 columns in the dataset are spatial label variables like *1_gwb_code_10*, and one is the dependent variable *38_g_woz*. As such, we include 50 columns as explanatory variables in the SCM. Appendix I includes an overview of all these variables.

Regarding the descriptive statistics, we include Figure 1 to give an idea of the spatial distribution of property prices in the Netherlands. Figure 1 shows the average property price for all municipalities in the Netherlands.

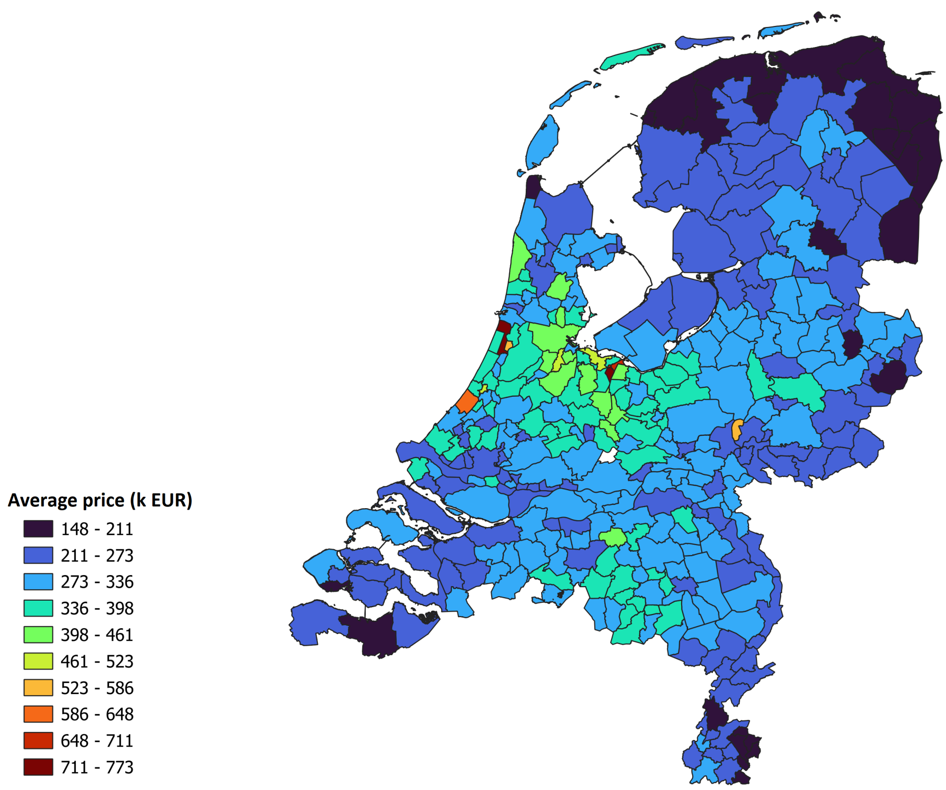

Figure 1. Spatial distribution of average property prices in the Netherlands on a municipality level.

The area in the north- and southwest, known as the Randstad ("Rim City"), has a higher average property value than the countryside. This area includes the major Dutch cities, like Amsterdam, Rotterdam, Utrecht, and the Hague. Also, along the coastline and in the Groene Hart ("Green Heart"), in the middle of the Netherlands, houses are, on average, sold for a higher price than in the rest of the country. In Figure 1, the decision to use a granularity level of municipalities is based on visualization purposes. For the quantitative analysis, as presented in Section 4, we use the *Kerncijfers Wijken en Buurten* dataset on the neighborhood level.

To get an idea of the distribution of the dependent variable *38_g_woz*, Figure 2 shows two distributional histograms. As *38_g_woz* only assumes values bigger than zero, we can take the logarithm to find log(*38_g_woz*). As suggested by Heij et al. (2004), a logarithmic transformation of the dependent variable can reduce the skewness and heteroskedasticity of the data. Since *38_g_woz* is negatively skewed, we include the logarithm in Figure 2b.

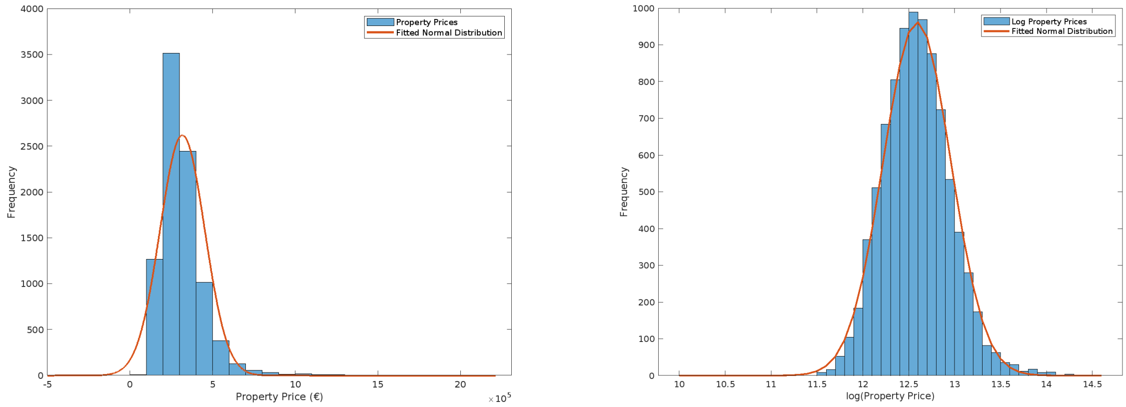

(a) Average neighborhood property prices

(b) Logarithmic transformation

Figure 2. Histogram of average neighborhood property prices in the Netherlands and its logarithmic transformation: a comparison with a fitted normal distribution.

Figure 2 shows that a logarithmic transformation on the *38_g_woz* variable creates a dependent variable that comes close to being normally distributed. Also, it reduces heteroskedasticity and negative skewness. In Section 4, we further discuss functional forms of the HPM, of which the logarithmic transformation is an example. Table 2 contains summary statistics for both the logarithmic transformation and the original variable *38_g_woz*.

Variable	**Maximum**	**Minimum**	**Mean**	**Median**	**Standard Deviation**
38_g_woz	$2.110 \cdot 10^6$	$5.120 \cdot 10^4$	$3.501 \cdot 10^5$	$2.890 \cdot 10^5$	$1.352 \cdot 10^5$
log(*38_g_woz*)	14.562	10.843	12.590	12.572	0.372

Table 2. *Summary statistics for dependent variable 38_g_woz and its logarithmic transformation.*

3.2 Flood Risk

As shown by Pistrika et al. (2014), flood damage costs are positively correlated with flood depth. As such, we assume flood depth to capture flood risk in a way directly relatable to its potential impact on property values. The website *klimaateffectatlas.nl* offers flood depth maps for four different flooding probabilities. Each probability level corresponds to one of these four categories:

- High: on average, an area will be flooded once every 10 years
- Medium: on average, an area will be flooded once every 100 years
- Low: on average, an area will be flooded once every 1000 years
- Extremely low: on average, an area will be flooded once every 100.000 years

Considering the average mortgage term is 30 years, we choose to include maximum flood depth for the high probability level in our analysis. Therefore, the maps provided by *klimaateffectatlas.nl* consist of two elements. First, the determination of the location-specific probability of flooding. Second, conditional on the occurrence of a flooding event, is the maximum local flood depth. Based on the location-specific probabilities of flooding, research institute Deltares created scenarios for each probability category[5]. These scenarios include flood types A-D, as shown in Figure 3 below[6]. Flood types E and F, which are the result of flooding sewage systems and heavy precipitation, respectively, are disregarded in the scenario analysis by Deltares (Slager, 2019).

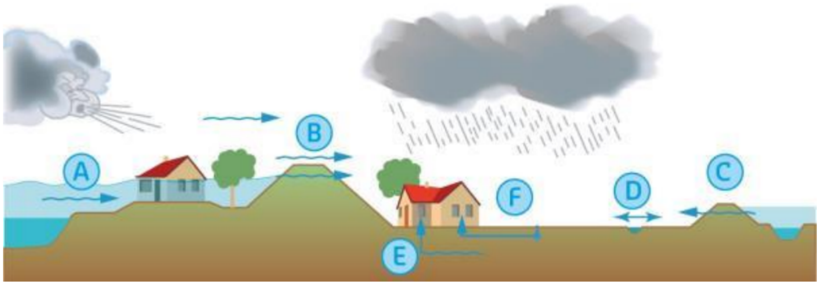

Figure 3. Schematic overview of different flood types in the Netherlands.

Flood types A-D are described in the following way:

(A) Flooding of unprotected areas (e.g., floodplains) along the main water system.

(B) Flooding of protected areas along the main water system, caused by flooding of primary flood defenses.

(C) Flooding of protected areas along the regional water system, caused by flooding of regional flood defenses.

(D) Flooding of unprotected areas from the regional surface water system.

Doing many simulations on these different flooding scenarios allows Deltares to create flood data maps, including minimum and maximum flood depth. In-depth analyses of the methodology and models behind these simulations, as discussed by Slager (2019) from Deltares, are outside the scope of this paper.

[5]Visualisations of these scenarios can be found on the website of LIWO

[6]Schematic overview comes from 2019 Deltares report (Slager, 2019)

To move from the geopackage type files provided by *klimaateffectatlas.nl* to data types suitable for data processing in Matlab, we use the open source Geographic Information System QGIS[7]. Using QGIS, we create an intersection between the *Wijk en Buurtkaart 2021 v2* neighborhood mapping file and the flood depth geopackage file. This intersection links neighborhood codes to the spatial flood depth data. Next, we export the QGIS project to a more common spreadsheet format in Excel, which we can import into Matlab for data analysis. Below, Figure 4 visually displays the intersection in QGIS.

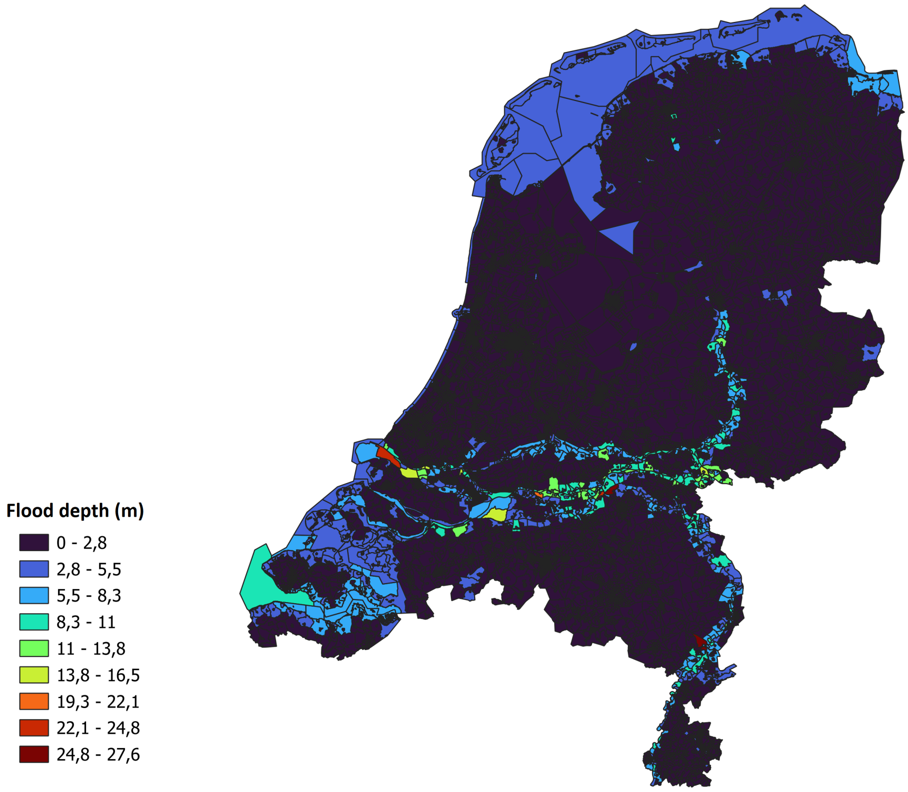

Figure 4. Spatial distribution of maximum flood depths in the Netherlands on a neighborhood level.

Figure 4 shows that the high-risk areas are located primarily around the Rhine-Meuse-Scheldt delta. Table 3 gives an overview of summary statistics for the flood risk variable for all Dutch provinces.

[7]Version 3.30 s'Hertogenbosch

Municipality	Maximum	Minimum	Mean	Median	Standard Deviation
Drenthe	1.580	0	0	0	0.021
Flevoland	1.433	0	0	0	0.022
Friesland	9.881	0	0.061	0	0.283
Gelderland	27.580	0	0.100	0	0.592
Groningen	6.202	0	0.021	0	0.241
Limburg	26.933	0	0.062	0	0.410
Noord-Brabant	24.770	0	0.032	0	0.251
Noord-Holland	5.110	0	0.013	0	0.110
Overijssel	11.502	0	0.031	0	0.212
Utrecht	8.561	0	0.032	0	0.280
Zeeland	10.451	0	0.042	0	0.301
Zuid-Holland	23.501	0	0.041	0	0.251

Table 3. *Summary flood depth statistics for all provinces in the Netherlands.*

The highest values of maximum flood depth belong to provinces that either accommodate the Rhine, Meuse, or Scheldt rivers or their respective river branches. As such, Table 3 shows that Limburg and Gelderland have the highest flood depth values in the Netherlands and Drenthe and Flevoland the least. As higher flood depths lead to more damage in case of a flooding event, Limburg and Gelderland bear a bigger risk than Drenthe and Flevoland. The average flood depth in the Netherlands is close to zero. This shows that even for high-risk municipalities in Limburg and Gelderland, the risk differs substantially between neighborhoods. Next, we discuss what neighborhoods we consider to bear the biggest risks. As we discuss in Section 4.2, the execution of the SCM depends on this.

Neighborhood Selection

Without diving into details, we discuss the neighborhoods we consider for the flood risk analysis of the SCM. We base the selection of a suitable neighborhood on two criteria:

1. The neighborhood contains a high value for the risk type considered
2. The neighborhood is assumed to be aware of being a high-risk neighborhood

In the summer of 2021, rivers flooded throughout Europe as a result of heavy rainfall. This caused flood damage in multiple municipalities in Limburg and was broadly featured in Dutch newspapers. According to Janssen (2023) and Jacobs (2021), Landgraaf, Heerlen, Kerkrade, Meerssen, Valkenburg aan de Geul and Gulpen-Wittem were the most affected. As such, we assume that flood risk awareness in these municipalities increased the most. Table 4 shows summary statistics for each of these municipalities.

Municipality	Maximum	Minimum	Mean	Median	Standard Deviation
Landgraaf	2.73	0	0.002	0	0.071
Heerlen	4.02	0	0.012	0	0.140
Kerkrade	2.38	0	0.009	0	0.090
Meerssen	8.59	0	0.132	0	0.722
Valkenburg aan de Geul	2.640	0	0.018	0	0.140
Gulpen-Wittem	2.682	0	0.021	0	0.121

Table 4. *Summary flood depth statistics for the Dutch province Limburg.*

In terms of maximum and average flood depth, Meerssen stands out as the municipality with the biggest flood risk. Assuming the awareness of flood risk is similar for all municipalities, all that rests is looking for the neighborhood in Meerssen with the highest flood depth value. In descending order, the top three flood depth values belong to Geulle, Bunde, and Weert. Of these three, Geulle is not part of the 2,750 neighborhoods in the SCM dataset. As such, we choose Bunde - with a flood depth of 7.420 m - to be the flood risk neighborhood of choice for the SCM.

3.3 Foundation Risk

To capture foundation risk, we use the foundation risk maps offered by *klimaateffectatlas.nl* as the data source. Leading research institute Deltares is the main contributor to these maps, whose technical methodology has been described by Kok and Angelova (2020) and consists of three parts:

1. Probabilistic assessment of the foundation type per location, given in the form of the likelihood that a property is built on wooden poles.
2. An analysis of the sensitivity to droughts, displayed in the number of days that the wooden pole foundation is exposed to oxygen - resulting from lower groundwater levels during periods of prolonged droughts.
3. An estimate of the damage pole rot does to the foundation condition in 2050.

Wood plays an important role because, as an organic material, it is susceptible to rotting. During periods of prolonged droughts and lower groundwater levels, the availability of oxygen allows fungi to grow and start developing pole rot. As a result, the first two parts are what drive the current foundation risk maps. This relationship defines the risk of pole rot as the product of *the percentage of wooden poles* and *vulnerability.* The percentage of wooden pole foundations in a neighborhood is obtained by analyzing foundation practices of construction periods, following (Deltares) foundation experts. For the vulnerability, Deltares analyzed the groundwater level, the depth of the pile head, and the soil type. However, an in-depth explanation of the methodology behind the foundation risk maps is out of the scope of this paper[8].

[8]For the interested reader: *https://www.deltares.nl/nl/nieuws/nieuwe-kaarten-over-funderingsrisico-panden-op-buurtniveau-in-de-klimaateffectatlas/*

Similar to Section 3.2, we translate the geopackage type data from *klimaateffectatlas.nl* to a spreadsheet file using QGIS' intersection functionality. Then, we import these Excel-type files into Matlab for modeling and data analysis purposes. As shown in Figure 5, the foundation risk data is missing neighborhoods. To end up with a complete dataset - and considering the importance of these risk variables for the analysis - we decided to exclude these neighborhoods. This leaves 8,889 neighborhoods for the Hedonic Pricing Model. We include a visual representation of the aforementioned intersection for the remaining 8,889 neighborhoods in Figure 5 below.

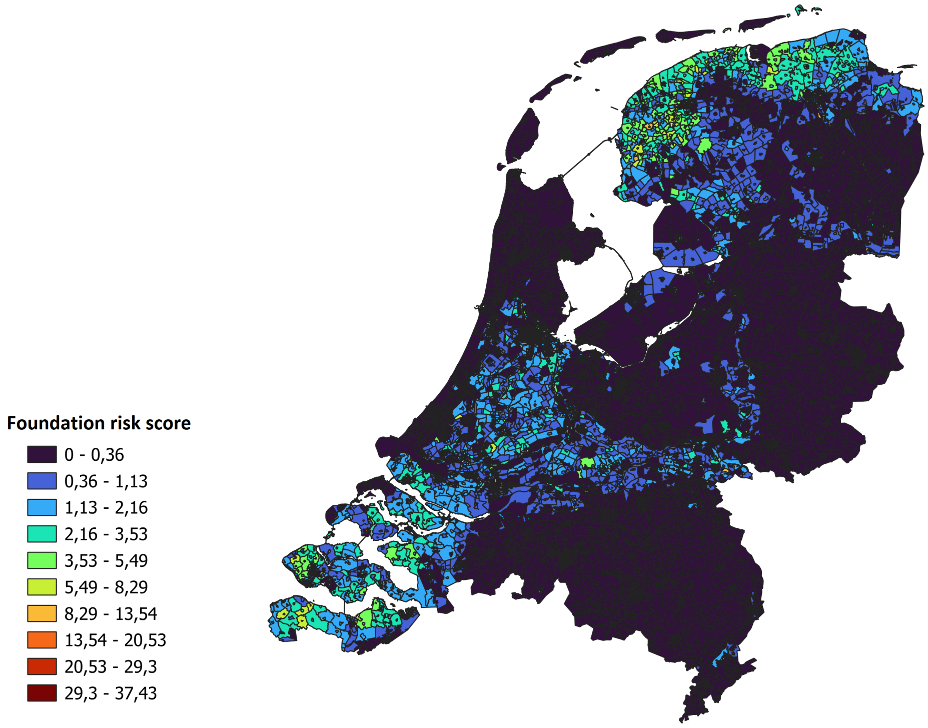

Figure 5. Spatial distribution of foundation risk scores in the Netherlands on a neighborhood level.

Foundation risk is higher in urban areas than in the countryside. Approximately 750,000 houses in the Netherlands are built on wooden pole foundations (Kok & Angelova, 2020). This foundation type was mainly used before 1970, in the post-war rebuilding period. Not surprisingly, this was mainly the case in larger cities and urban areas. As it can be challenging to see in Figure 5, Table 5 gives a global overview of the summary foundation risk statistics for all Dutch provinces.

Municipality	Maximum	Minimum	Mean	Median	Standard Deviation
Drenthe	4.490	0	0.241	0.042	0.460
Flevoland	2.662	0	0.162	0	0.372
Friesland	21.251	0	1.510	0.630	2.253
Gelderland	25.690	0	0.303	0	1.041
Groningen	7.361	0	0.751	0.142	1.151
Limburg	3.012	0	0.052	0	0.240
Noord-Brabant	4.071	0	0.101	0	0.392
Noord-Holland	37.430	0	0.891	0	3.181
Overijssel	7.750	0	0.151	0	0.581
Utrecht	27.860	0	0.821	0.042	2.422
Zeeland	22.731	0	1.932	1.333	2.721
Zuid-Holland	30.961	0	0.661	0.211	1.670

Table 5. *Summary foundation risk statistics for all provinces in the Netherlands.*

In agreement with Figure 5, Table 5 confirms that the highest foundation risk areas are located in provinces that accommodate big cities. As such, Zuid-Holland and Noord-Holland have the highest foundation risk scores in the Netherlands. Table 5 also shows that the mean foundation risk score in Zuid-Holland is lower than that of a more peripheral province like Groningen. This indicates that high-risk areas inside the province are alternated by areas of lower risk, which is confirmed by the high standard deviation values in Noord-Holland and Zuid-Holland. As this potentially averages out local effects, it is important to study the impact of these risks on a more granular neighborhood level. Also, risk can differ substantially from one neighborhood to the next. In the coming section, we examine which neighborhoods are most suitable for the foundation risk execution of the SCM.

Neighborhood Selection

For the foundation risk neighborhood selection of the SCM, we consider the same criteria as before. Considering the first criterium, Figure 5 shows that we find the highest foundation risk scores in larger cities. Table 5 shows that Zuid-Holland and Noord-Holland contain the highest foundation risk scores of all provinces in the Netherlands. A previous study by Hommes et al. (2023) reports that Schiedam, Dordrecht, Rotterdam, Amsterdam, Zaanstad, Haarlem, and Gouda are municipalities where foundation risk is reported more often than in the rest of the Netherlands. Table 6 shows summary statistics of foundation risk in these municipalities.

Municipality	Maximum	Minimum	Mean	Median	Standard Deviation
Amsterdam	37.430	0	2.660	0.111	5.842
Zaanstad	8.803	0	1.181	0.492	1.652
Haarlem	18.150	0	1.680	0.741	2.722
Rotterdam	7.221	0	0.862	0.510	1.010
Dordrecht	25.860	0	1.461	0.590	4.120
Schiedam	30.961	0.041	2.502	0.382	5.841
Gouda	8.732	0	1.210	0.320	1.731

Table 6. *Summary foundation risk statistics for high-risk Dutch municipalities.*

Schiedam is the only municipality with a minimum foundation risk score larger than zero. This means that no neighborhood in Schiedam does not have foundation risk. Only topped by Amsterdam in both maximum and average foundation risk, it is no surprise that the local government started a campaign to raise awareness on this matter in 2021. The municipalities of Schiedam, Dordrecht, and Rotterdam jointly started with multiple initiatives[9] to aid homeowners and homebuyers in solving this issue. As such, we assume the risk awareness to be *high* in these municipalities and *medium* in the others. Table 7 combines the *klimaateffectatlas.nl* risk scores with our risk awareness assessment to facilitate the neighborhood selection process. Finally, a neighborhood must be part of the 2,750 neighborhoods included in the SCM dataset. Therefore, Table 7 provides this information as well.

Neighborhood	Municipality	Risk score	Risk awareness	Included in dataset
1. Marnixbuurt Midden	Amsterdam	37.430	medium	no
2. BG-terrein e.o.	Amsterdam	33.910	medium	no
3. Zaagpoortbuurt	Amsterdam	32.591	medium	no
4. Burgwallen Oost	Amsterdam	31.682	medium	no
5. Stadserf	Schiedam	30.960	high	no
6. Marine-Etablissement	Amsterdam	29.291	medium	no
7. Planciusbuurt Noord	Amsterdam	29.281	medium	no
8. Walvisbuurt	**Schiedam**	**29.180**	**high**	**yes**
9. Oude Kerk e.o.	Amsterdam	28.580	medium	no
10. Zuiderkerkbuurt	Amsterdam	27.970	medium	no

Table 7. *Foundation risk scores for the ten neighborhoods in the Netherlands with the highest risk.*

Of the ten neighborhoods with the highest foundation risk scores in the Netherlands, only one is part of the 2,750 neighborhoods in the SCM dataset. Being located in Schiedam, we assume it to have high

[9]Funderingsloket, Servicepunt Woningverbeteraar, and Fonds Duurzaam Funderingsherstel are aimed at providing detailed information and financial support to those facing foundation risk

risk awareness. Fitting both criteria, we select Walvisbuurt to study the impact of foundation risk on average neighborhood property prices with the SCM.

4 Methodology

To answer the research questions formulated in Section 1, we apply multiple econometric techniques and models to the datasets from Section 3. In this section, we discuss the mathematics and econometric reasoning behind these models and the techniques and analyses that preceded them. Similar to the rest of this paper, we discuss the two models separately. As such, we start by explaining our interpretation of the Hedonic Pricing Model in Section 4.1 and continue with our novel application of the Synthetic Control Method in Section 4.2.

4.1 Hedonic Pricing Model

Before discussing the specifics of the Hedonic Pricing Model, we outline the different parts contributing to this goal. After a global introduction to the HPM, Section 4.1.1 describes the particular model settings that deal with multicollinearity. In Section 4.1.2, we discuss a complementary Bucketing Analysis (BA) that further investigates data peculiarities.

The Hedonic Pricing Model is an econometric method that decomposes an object into its constituent parts or attributes. Taking this idea to object pricing allows us to study the impact of neighborhood risk attributes - flood and foundation risk - on the average neighborhood property price. For the construction of the HPM, we follow guidelines originally proposed by Rosen (1974). We consider average neighborhood property prices P, as a function of $N-J$ general neighborhood attributes, $(x_1, x_2, \ldots, x_{N-J})$, and the J neighborhood risk attributes $(x_{N-(J-1)}, \ldots, x_N)$. As we are working with two risk attributes, flood and foundation risk, $J = 2$. We assume that P can be completely described by the $n = 1, \ldots, N$ neighborhood characteristics x and that the housing market is homogeneous in the sense that all average neighborhood property prices have the same underlying attributes. These attributes are also assumed to have a linear relationship with P. In Section 3, we show that taking the logarithmic transformation of P has desirable properties concerning its distribution. As such, we continue with the following linear relationship between the dependent variable $\log(P)$ and its neighborhood attributes:

$$log(P) = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \ldots + \beta_N x_N + \varepsilon. \tag{1}$$

Under these assumptions, the linear regression results should reflect any effect that flood or foundation risk might have on average neighborhood property prices. According to the Gauss-Markov theorem, using OLS - under six key assumptions - to run this regression provides an estimator that is the Best Linear Unbiased Estimator (BLUE) (Heij et al., 2004). Of these six assumptions, assumption four is called *the absence of serial correlation.* Section 3 shows that the *Kerncijfers Wijken en Buurten* dataset contains strongly correlated variables. A method that can help with highly correlated explanatory variables is the Principle Component Analysis (PCA). Combining PCA with a Multiple Linear Regression (MLR) is often called Principle Component Regression (PCR).

4.1.1 Principal Component Regression

Both PCA and PCR are based on the idea of Principal Components (PCs), which stems from Linear Algebra. Essentially, PCs are a new coordinate system with fewer dimensions than the original dataset. A linear transformation projects the information from these $(N-2)$ highly correlated neighborhood attributes onto $k = 1, \ldots, K$ orthogonal - and thus uncorrelated - PCs. To get the desired result of a dimensionality reduction, K must be smaller than $(N-2)$. As such, it is often called a dimensionality reduction technique. The PCs are determined in order, with PC_1 explaining the largest amount of variance from the dataset and every other PC accounting for a decreasingly smaller portion of variation in the data. When the neighborhood attributes are not the same order of magnitude, their variances are not comparable. Therefore, a variance maximizing technique like PCA assumes homoskedasticity of the data. Section 3 showed that the neighborhood's attributes $(x_1, x_2, \ldots, x_{N-2})$ are not of the same scale. As a result, we normalize all neighborhood attributes to ensure proper PCA results. Below, we include the mathematical expression for the PCA relationship in Equation 2.

$$(\mathrm{N}-2)\mathrm{PC}_i = (a_{i_1}x_1) + (a_{i_2}x_2) + \ldots + \left(a_{i_{N-2}}x_{N-2}\right) \tag{2}$$

Here, we represent the number of components by $(N-2)PC_i$. Equation 2 shows each PC is a linear combination of all neighborhood attributes $X = (x_1, x_2, \ldots, x_{N-2})$, with component weights $\{a_{i_1}, a_{i_2}, \ldots, a_{i_{N-2}}\}$. With these PCs, we can transform the regression from Equation 1 into a PCR. With the research question in mind, we augment K PCs with the flood and foundation risk variables (x_{N-1}, x_N) to form the MLR displayed in Equation 3 below.

$$log(P) = \beta_0 + \beta_1 PC_1 + \beta_2 PC_2 + \ldots + \beta_K PC_K + \beta_{K+1}x_{N-1} + \beta_{K+2}x_N + \varepsilon \tag{3}$$

4.1.2 Bucketing Analysis

The homogeneity of attributes is one of the assumptions behind the HPM. Doing a Bucketing Analysis, in which we perform the PCR for different data ranges of a variable, is a way to test this assumption. Also, it explores the dataset and model outcomes in the case of structural breaks. For example, with a BA on the average neighborhood property prices, we split the dependent variable *38_g_woz* into groups with a bucket size of 100,000 euros. As such, the dataset is split into smaller datasets, each containing only the rows belonging to dependent variables in that bucket range.

We can do a similar analysis for explanatory variables. Then, the BA is based on ranges of a specific independent variable, like *80_p_hh_hi*. The variable *80_p_hh_hi* represents the percentage of households in a neighborhood that belong to the group of 20% highest income earners. Out of the 85 explanatory variables, we choose *80_p_hh_hi* for its structural break properties. Random Forest regression results on the same *Kerncijfers Wijken en Buurten* dataset by Kars (2021) depict the *80_p_hh_hi* variable to be an important feature to split observations. This means that a BA for this variable compares buckets designed to be as different from each other as possible. In contrast, observations within a bucket are as similar to each other as possible[10].

4.2 Synthetic Control Method

We discuss the Synthetic Control Method in three parts. First, Section 4.2.1 introduces the general theory behind the SCM. Then, Section 4.2.2 dives into the selection procedure of the neighborhood attributes that precede the construction of the synthetic version of Walvisbuurt and Bunde. In Section 4.2.3, we aim to explain the methodology behind the model significance assessment.

4.2.1 General Model

Although the number of applications of the Synthetic Control Method is numerous (Abadie, 2021), we propose a novel approach here. Contrary to common use in policy assessments, like California's tobacco control program in Abadie et al. (2010), we study the effect of assumed risk awareness on consumer behavior in the residential real estate market in the Netherlands. More specifically, we seek to isolate the effect of the awareness (and presence) of climate-related physical risks on property prices in the Netherlands. We do this for both flood and foundation risk. For each risk type, we study a different neighborhood. For flood risk, we study Bunde in Meerssen, a Dutch municipality in Limburg. For flood risk, we focus on Walvisbuurt in Schiedam. Apart from their specific risk values, we select these neighborhoods, as described in Section 3.1.2, based on events that are assumed to increase risk awareness locally. These events are:

1. The 2021 floods in Limburg, which caused an estimated 350 to 600 million euros in damages (Jonkman, 2021).
2. The 2021 foundation risk awareness initiatives in Rotterdam, Schiedam, and Dordrecht.

Studying the evolution of the average property price in Walvisbuurt and Bunde after these events took place does not provide enough information to infer their effects. That is because it is missing the counterfactual information on what would have happened to the average property price in Walvisbuurt and Bunde if these events had not happened. The SCM aims to fill this information gap by constructing counterfactual "synthetic versions" of Walvisbuurt and Bunde designed to provide this information.

Before we dive into the procedure, we set up a mathematical framework. Based on the notation by Abadie (2021), we consider $J + 1$ total neighborhoods in the dataset: $j = 1, 2, \ldots, J + 1$. As we study Walvisbuurt and Bunde separately, this means that for each analysis, the dataset contains J other neighborhoods: $j = 2, \ldots, J + 1$. Being affected by the previous events, we denote Walvisbuurt and Bunde as "treated units" with $j = 1$. The J remaining neighborhoods from the "donor pool", out of which the Synthetic Walvisbuurt and Synthetic Bunde, or "untreated units", are constructed. For the analysis, we distinguish between two time periods: a pre- and post-event period, or "pre-intervention"

[10] Of course, this only holds if the bucket breaks are chosen according to the Random Forest regression results

and "post-intervention". We describe the number of pre-intervention periods, or "pre-treatment" periods, by T_0. As such, the pre-treatment period runs from $t = 1, \ldots, T_0$, and the "post-treatment" period from $t = T_0 + 1, \ldots, T$.

Let Y_{jt} be the average neighborhood property price of neighborhood j at time t. We can then distinguish between Y_{jt}^I and Y_{jt}^N. We use the superscript I to define Y_{jt}^I as the variable of interest for a unit j affected by the intervention in the post-intervention period $t = T_0 + 1, \ldots, T$. Similarly, we use Y_{jt}^N to denote the variable of interest for the same unit j in the period $t = 1, \ldots, T$, in case of no intervention. We are interested in the comparison between Y_{jt}^I and Y_{jt}^N for the treated unit $j = 1$ and the period $t > T_0$. The difference between Y_{1t}^I and Y_{1t}^N for $t = T_0 + 1, \ldots, T$ is assumed to be attributable to the intervention. The SCM aims to find this difference, which we denote by τ_{1t}. Equation 4 defines τ_{1t} in the following way:

$$\tau_{1t} = Y_{1t}^I - Y_{1t}^N. \tag{4}$$

In Equation 4, τ_{1t} captures the impact of foundation or flood risk on the average neighborhood property price. It represents the average neighborhood property price difference for the treated unit $j = 1$, which for $t > T_0$ is assumed to be attributable to an increased awareness of flood or foundation risk. For the treated unit, Y_{1t}^I is observed. However, inherent to a counterfactual, the property price of the synthetic version Y_{1t}^N is not observed. As a result, the estimation of τ_{1t} depends on the estimate of Y_{1t}^N. This leads to the following equation for the estimation of τ_{1t}:

$$\hat{\tau}_{1t} = Y_{1t} - \hat{Y}_{1t}^N. \tag{5}$$

The estimate of the synthetic version of the treated unit in Equation 5 is given by $\hat{Y}_{1t}^N$. We construct this estimate as a linear combination of neighborhoods from the donor pool. Consequently, we write $\hat{Y}_{1t}^N$ as follows:

$$\hat{Y}_{1t}^N = \sum_{j=2}^{J+1} w_j Y_{jt}. \tag{6}$$

We can combine all weights for the J units in the donor pool in a single column vector, $\boldsymbol{W} = (w_2, \ldots, w_{J+1})'$. To find proper estimates of the effect of the intervention, the SCM requires some key assumptions to be met.

First, the donor pool is supposed to be free of units affected by similar interventions. In this research, the interventions are represented by events that presumably increase awareness among homebuyers. In an ideal scenario, the treated unit went from a population awareness of 0% to 100%, while all units in the donor pool remained at 0%. Being based on assumptions, we understand this comes with a particular sensitivity to errors. As such, Section 5.2 shows results for different donor pool requirements.

Second, the units in the donor pool are assumed to be unaffected by the treatment unit. In their paper on Proposition 99, Abadie et al. (2010) call this the absence of the "spillover effect". It assumes that California's increased cigarette tax does not affect cigarette sales in neighboring states. Translating this to a property price setting, we assume that the population whose risk awareness increased as a result of the intervention does not use this knowledge to influence property prices elsewhere. Similar to the first assumption, Section 5.2 includes the rationale behind excluding some units in the donor pool.

Third, the treatment effect can not precede the moment of treatment $t = T_0$. As time is the dimension over which the effect accumulates, a wrong calibration can lead to incorrect conclusions. For example, if the 2021 floods in Limburg have a significant effect on average neighborhood property prices that wears off within a year, setting the post-intervention period one year later would lead to the conclusion of no effect.

The last conditions, as mentioned by McClelland and Mucciolo (2022), are related to the optimization procedure to find the optimal donor pool weights, $\boldsymbol{W}^* = \left(w_2^*, \ldots, w_{J+1}^*\right)'$. We discuss this in Section 4.2.2.

4.2.2 Predictor Selection

For the construction of the untreated units, we optimize their similarity to the treated units in the pre-intervention period over a small subset of K predictors, which we denote by $s = 1, \ldots, S$, and P pre-treatment values of the variable of interest Y_{jt}. We use two procedures to select S out of K predictors. Figure 6 shows a schematic overview of the entire Predictor Selection process, in which these two procedures are denoted by steps 2 and 3.

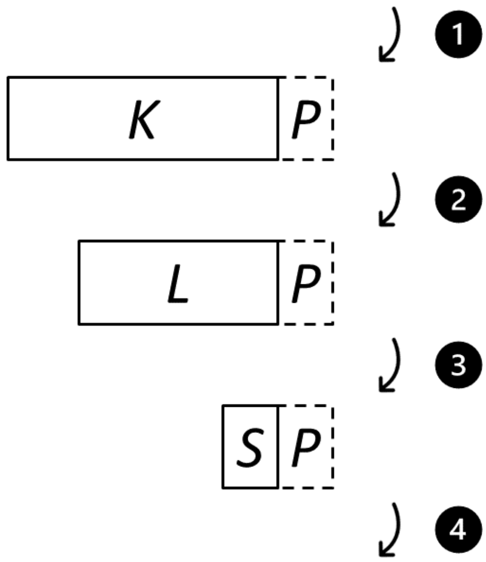

Figure 6. Schematic overview of Predictor Selection procedure.

As such, Figure 6 identifies the following four steps:

1. Data pre-processing, discussed in Section 3.1.2.
2. McClelland and Mucciolo (2022) Convexity Condition procedure, discussed in Section 4.2.2.
3. MSPE optimization procedure, discussed in Section 4.2.3.
4. SCM results, discussed in Section 5.2.

For each neighborhood j, the dataset contains a set of $k = 1, \dots, K$ predictors: $X_{1j}, \dots, X_{kj}$. Additionally, we add P pre-treatment values of the variable of interest Y_{jt} as predictors. To quantify similarity, we follow Abadie et al. (2010) in using $\|\boldsymbol{X}_1 - \boldsymbol{X}_j\|$ as a measure of dissimilarity between treated unit $j = 1$ and some unit $j \neq 1$ from the donor pool. Let z be the variables included in $\boldsymbol{X}_1$ and $\boldsymbol{X}_j$, over which the similarity is optimized in the SCM, with $z = 1, \dots, S, S+1, \dots, S+P$. As shown in Figure 6, S is a subset of L. Let $l = 1, \dots, L$, be the set of predictors that fulfills the convexity condition appearing in Abadie et al. (2010) and Abadie et al. (2015), and a subset of K. To move from K to L, we use a procedure described by McClelland and Mucciolo (2022). In the convexity condition, the weights $\boldsymbol{W} = (w_2, \dots, w_{J+1})'$ are non-negative and sum to one. Consequently, the predictor values for the treated unit must reside inside the convex hull of the donor pool predictor values. Imagine a predictor k for which holds: $X_{k1} \ll X_{k2}, \dots, X_{kJ}$, then no linear combination - using positive weights that sum to one - of the J predictor values from the donor pool $X_{k2}, \dots, X_{kJ}$ can approximate the value of the treated unit for the same predictor. To find predictors that fulfill the convexity condition, McClelland and Mucciolo (2022) propose a method that compares all K predictor values for the treated unit against those from the donor pool. Following this method, we start by averaging all predictors for the treated unit over the pre-treatment period $t = 1, \dots, T_0$. Then, we perform the same operation for the predictors of all units in the donor pool and take the difference with the corresponding values of the treated unit. To compare these differences, we adjust them to have a unit standard deviation. Predictors that fulfill the convexity condition appear in the middle, with an approximately equal number of donor pool units above and below.

With predictors L we can perform the SCM. As such, S can be selected with an optimality analysis. As all predictors L are suitable candidates, we use the Mean Squared Prediction Error (MSPE) to decide which predictors to include in S. As this analysis is also part of the significance assessment, we discuss it in more depth in Section 4.2.3. Section 5.2.1 presents which pre-treatment years of Y_{jt} we include in P and the rationale behind this decision. When the elements that make up S and P are known, we combine them into Z to write the difference between the treated unit and its synthetic version more explicitly in Equation 7.

$$\|\mathbf{X}_1 - \mathbf{X}_0\boldsymbol{W}\| = \left(\sum_{h=1}^{Z} v_h \left(X_{h1} - w_2 X_{h2} - \ldots - w_{J+1} X_{hJ+1}\right)^2\right)^{1/2} \tag{7}$$

To construct synthetic versions that are as similar as possible to the treated units, we find the weights that minimize the expression in Equation 7, $\boldsymbol{W}^* = \left(w_2^*, \dots, w_{J+1}^*\right)'$. Here, $v_1, \dots, v_Z$ represents the set of weights given to each of the Z predictors. On top of the method discussed in Section 4.2.2, weights are allocated to each predictor. For now, we take $v_1, \dots, v_Z$ as given. Following Abadie et al. (2010), we classify the expression in Equation 7 as a constrained quadratic optimization problem. We solve this with the programming platform Matlab[11]. As such, weights $\boldsymbol{W} = (w_2, \dots, w_{J+1})'$ are a function of $v_1, \dots, v_Z$, which leads to the following expression: $\boldsymbol{W}(\boldsymbol{V}) = (w_2(\boldsymbol{V}), \dots, w_{J+1}(\boldsymbol{V}))'$.

Finally, we can take different approaches to decide on $\boldsymbol{V}$. Based on computational simplicity[11], we continue to follow Abadie et al. (2010) with their approach based on the MSPE. We determine $\boldsymbol{V}$ as being the weights for which $\boldsymbol{W}(\boldsymbol{V})$ minimizes the MSPE in Equation 8 below.

$$\sum_{t \in \mathcal{T}_0} (Y_{1t} - w_2(\boldsymbol{V})Y_{2t} - \ldots - w_{J+1}(\boldsymbol{V})Y_{J+1t})^2 \tag{8}$$

For $t \in \mathcal{T}_0$, Abadie et al. (2010) denotes $\mathcal{T}_0 \subseteq \{1, 2, \ldots, T_0\}$ as being the (sub)set of pre-treatment periods.

4.2.3 Confidence Intervals

Referring back to Figure 6, the construction of confidence intervals coincides with step 3 of the Predictor Selection procedure. Consequently, we discuss these together. Depending on the size of the set of predictors that fulfill the convexity condition, denoted by L, we can put together multiple unique combinations of S. Mathematically, we express the number of unique combinations using Equation 9.

$$\begin{pmatrix} L \\ S \end{pmatrix} = \frac{L!}{S!(L-S)!} \tag{9}$$

In Equation 9, we treat the predictors as "distinguishable objects". This means we "do not allow for repetition and order does not matter", as per Zwillinger (2018). The number of unique combinations, as found by Equation 9, does raise the question of which combination we should incorporate as the model outcome in Section 5.2.2. In Equation 8, we find $\boldsymbol{W}(\boldsymbol{V})$ that minimizes the objective function. Following the same reasoning, we choose S based on its MSPE for the pre-intervention period. Computationally, we do this by iteratively examining all unique combinations of S. As only one combination yields optimal results in terms of pre-intervention MSPE, we propose to use all nonoptimal results for a sensitivity analysis. Combining all post-intervention variables of interest for the synthetic version at $t = T$, we can construct a confidence interval for the optimal post-intervention outcome. This idea is based on the percentile method, as discussed by Efron and Tibshirani (1994). Let $N_{NO} = \binom{l}{s} - 1$ be the number of nonoptimal estimates, then for all $n = 1, \ldots, N_{NO}$ we order $\hat{Y}^N_{1T}$. A $100(1-\alpha)\%$ confidence interval for $\hat{Y}^N_{1T}, \ldots, \hat{Y}^N_{N_{NO}T}$ then includes the $100(1-\alpha)\%$ center values of the order transformation for $\hat{Y}^N_{1T}, \ldots, \hat{Y}^N_{N_{NO}T}$ (Puth et al., 2015).

5 Results

In this section, we dissect the outcomes that execution of the econometric models and techniques from Section 4 yield. Doing so in a clear and structured way helps us interpret these results in light of the research questions. Continuing with the same setup as before, we start with our HPM findings in Section 5.1 and discuss the SCM output in Section 5.2.

[11] Using a Synthetic Control software package provided by Jens Hainmueller, a co-author in the famous paper by Abadie et al. (2010): https://web.stanford.edu/ jhain/synthpage.html

5.1 Hedonic Pricing Model

We split up the HPM analysis into two parts. Section 5.1.1 comprises the PCR setup and main model findings. Section 5.1.2 includes the Bucketing Analysis, complementary to the first part.

5.1.1 Principal Components Regression

Before executing the PCR, we need to decide on the number of PCs we want to include. We use Figure 7 as a decision-making tool.

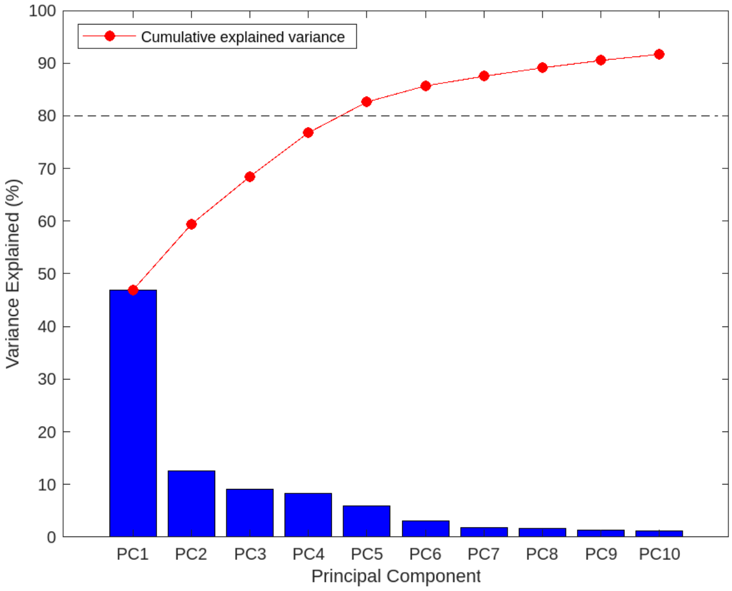

Figure 7. Explained variance and the cumulative explained variance for the first ten Principal Components.

In Figure 7, we display the variance and cumulative variance explained by the first ten PCs. As the number of PCs increases, each PC accounts for an increasingly smaller portion of variation in the data. This causes a trade-off between model sparsity and the variation in the data we seek to explain. As suggested by Oladunni and Sharma (2016), we choose the number of PCs based on where the slope of the cumulative explained variance intersects with the horizontal threshold for the first time. In this case, that occurs for PC_5, which coincides with a variance explained $> 80\%$. We include a qualitative interpretation of each PC in Appendix III. We can now include the five PCs in the PCR and run the regression. For the implementation of the PCR, we focus on the following three questions:

1. Can flood and foundation risk help explain average neighborhood property prices?
2. Does a higher risk variable, *ceteris paribus*, lead to higher or lower values in the dependent variable?
3. Can we label this effect as statistically significant?

We try to answer the first question by means of comparing the (Adjusted) R^2 of the PCR without risk variables against the PCR with risk variables. An increase in the reported (Adjusted) R^2, as a result of adding the risk variables, indicates that the risk variables have an effect on the average neighborhood property prices. Table 10 shows this comparison. For completeness, we add the Root Mean Square Error (RMSE) and other significance measures (i.e., the F-statistic and p-value).

Model	R^2	**Adjusted** R^2	**RMSE**	F**-statistic**	p**-value**
Excluding risk variables	0.625	0.625	0.226	$2.961 \cdot 10^3$	0
Including flood risk[1]	0.626	0.626	0.225	$2.480 \cdot 10^3$	0
Including foundation risk[2]	0.626	0.625	0.226	$2.474 \cdot 10^3$	0
Including both risk variables	0.627	0.627	0.225	$2.131 \cdot 10^3$	0

[1] Flood risk variable has a negative sign and is significant ($p < 0.05$);

[2] Foundation risk variable has a positive sign and is significant ($p < 0.05$);

Table 8. *PCR regression results: explanatory power of flood and foundation risk.*

Table 8 shows the difference in both R^2 and Adjusted R^2 is limited. This shows that the effect of climate-related physical risks, like flood and foundation risk, on the dependent variable is limited. To get an idea of the exact PCR results, Table 9 gives an overview of the regression estimates and their statistical significance. The model in Table 9 is the same as the model from the fourth line of Table 8.

Variable	**Estimate**	**SE**	t**-statistic**	p**-value**
(Intercept)	5.601	0.007	732.220	0.000
PC_1	0.220	0.003	65.113	0.000
PC_2	0.433	0.006	66.723	0.000
PC_3	-0.041	0.008	-5.150	$2.651 \cdot 10^{-7}$
PC_4	-0.501	0.008	-62.820	0.002
PC_5	0.420	0.011	43.711	0.000
Foundation risk	**0.004**	**0.001**	**3.230**	**0.001**
Flood risk	**-0.009**	**0.002**	**-5.550**	$\mathbf{2.880 \cdot 10^{-8}}$

Table 9. *PCR regression results: model estimates of flood and foundation risk.*

Using Table 9, we focus on questions 2 and 3. It shows that a higher foundation risk, *ceteris paribus*, increases average neighborhood property prices. The size of the coefficient shows that, compared to the other variables included in the PCR, the impact that foundation risk has on the dependent variable is approximately an order of 10x smaller[12]. This is not in line with previous findings by Hommes et al. (2023). For flood risk, the sign of the effect is in line with previous findings by Bosker et al. (2019), Daniel et al. (2009) and Reeken and Phlippen (2022) and its size is around double the size of the foundation risk effect. The effects of both risk variables are statistically significant.

Combining all three questions, we conclude that flood and foundation risk have a statistically signi-

[12] A fast calculator would report an order difference of 100x, but the size difference between the risk variables and the general neighborhood attributes corrects for this

ficant but limited impact on average neighborhood property prices. The sign of the regression coefficient of foundation risk is not in line with previous results by Hommes et al. (2023). The difference in research design could explain this. The paper by Hommes et al. (2023) investigates transactions of property listings that mention either foundation damage or foundation repair. As such, we assume home buyers are aware of a property's (no) risk situation. The difference in research findings hints at information gaps in a lot of high-risk neighborhoods[13]. A change in consumer behavior impacts the WOZ value through property transactions but depends on this information. The second research question, to which we present findings in Section 5.2, takes the importance of risk awareness into account. It is also possible that stronger effects locally might be averaged out on a country-wide level. Again, the SCM, which is devoted to analyzing these effects specifically, takes this into consideration. Below, Section 5.1.2 reports the additional PCR data examination results using a Bucketing Analysis.

5.1.2 Bucket Analysis

The results from the BA for the dependent variable *38_g_woz* are shown in Table 10. Each row contains the outcome of a separately run PCR.

Bucket	**Adjusted R^2**		**Variable sign**	
	excluding risk	including risk	foundation risk	flood risk
min - 100[1]	0.612	0.990	-	-
100 - 200[1]	0.310	0.311	-	+
200 - 300[1]	0.209	0.211	-	+
300 - 400[1]	0.130	0.133	+	-
400 - 500[2]	0.011	0.026	+	-
500 - 600[1]	0.014	0.011	-	+
600 - 700[1]	0.026	0.014	+	-
700 - 800[1]	0.088	0.058	-	+
800 - 900[1]	0.112	0.054	+	-
900 - 1000[1]	0.271	0.201	+	0
1000 - 1100[1]	-0.210	-0.161	+	+
1100 - 1200[1]	0.522	0.432	+	0
1200 - 1300[1]	0.018	0.410	+	0
1300 - max[1]	-0.442	-0.550	+	0

[1] Both results are not significant ($p > 0.05$);

[2] Flood risk is not significant ($p > 0.05$)

Table 10. *Summary statistics for the Bucketing Analysis of the dependent variable 38_g_woz.*

With the three questions in mind, we equipped Table 10 with all the necessary elements to answer

[13]Deltares foundation expert Marco Hoogvliet also mentioned this during a phone call on the 7^{th} of February 2023

them. Regarding question one, Table 10 shows that the general explanatory power of the HPM has diminished. More importantly, the explanatory power of the risk variables in most buckets completely disappeared. Adding flood and foundation risk to the PCR in some cases even decreases the Adjusted R^2 value. Signs of both risk variables are flipping, moving from one bucket to the next. Finally, apart from the 400-500 bucket, both risk variables' effects are insignificant in every bucket. A smaller variation in the values of the dependent variable loses a lot of information. This causes the performance to decrease drastically. To understand how Adjusted R^2 values can be negative and decrease with the addition of flood and foundation risk variables, we look at the following Adjusted R^2 formula:

$$\text{Adjusted } R^2 = 1 - \frac{N-1}{N-K}\left(1 - R^2\right). \tag{10}$$

The Adjusted R^2 is designed such that increasing the number of explanatory variables k is penalized. This means that when the explanatory power of an additional variable is low, the penalization factor is bigger than the gain in the R^2 value, causing the Adjusted R^2 to decrease. Negative values of the Adjusted R^2 can appear if the value of R^2 is small relative to the ratio of the number of observations N and the number of explanatory variables K. We included a visual representation of this result in Appendix IV.

Regarding the PCR output for the BA of the explanatory variable *80_p_hh_hi*, Table 11 shows the necessary statistics for the aforementioned questions.

Bucket	**Adjusted R^2**		**Variable sign**	
	excluding risk	including risk	foundation risk	flood risk
0 - 10[2]	0.444	0.450	-	-
10 - 20	0.511	0.513	+	-
20 - 30	0.531	0.540	+	-
30 - 40[1]	0.640	0.641	+	-
40 - 50	0.691	0.703	+	-
50 - max	0.822	0.834	+	-

[1] Both results are not significant ($p > 0.05$);

[2] Flood risk is not significant ($p > 0.05$)

Table 11. *Summary statistics for the Bucketing Analysis of explanatory variable 80_p_hh_hi.*

Although the general explanatory power of most buckets is lower than the original PCR outcome from Tables 8 and 9, it is still much higher than the BA of the dependent variable, as shown in Table 10. It shows that the variability in the dependent variable is lower, but not to the extent as before. That is because we do not use the exact decision node values for *80_p_hh_hi*, which Kars (2021) finds with his Random Forest regression, as bucket breaks. The inconsistency of a significant impact of the risk variables on the dependent variable indicates that structural breaks exist in the dataset. As such, the homogeneity assumption, which states that the contribution of each attribute to the average neighborhood property price is the same for all values of the dependent variable, seems unrealistic. This suggests that consumers' property preferences and the value they attach to these (neighborhood) property attributes differs across price ranges. Put differently; different price ranges attract different consumers with different preferences.

All in all, the results from the HPM seem to question the presence of a consistent effect of climate-related physical risks on the average neighborhood property price. Considering the necessary, but seemingly unrealistic, homogeneity assumption of the HPM, it remains unclear if flood and foundation risk have no impact on average neighborhood property prices or that the HPM falls short of proving it. Section 5.2 discusses the results of a different approach, in the form of the SCM, that might answer this question.

5.2 Synthetic Control Method

Here we build on the findings from Section 5.1 by investigating question two. We divide the findings of the SCM into two parts. First, we present the outcome of the Predictor Selection procedure for each risk type and donor pool selection in Section 5.2.1. Section 5.2.2 consists of the major results for each climate-related physical risk type. It includes confidence intervals to put the significance of the results in perspective.

5.2.1 Predictor Selection

As introduced by McClelland and Mucciolo (2022), we compare all $K = 50$ predictors on the convexity condition. As P represents pre-treatment outcome values Y_{jt} for some $t = 1, \ldots, T_0$, we include it as item 51 in the analysis. Following the steps described in Section 4.2.2, we find figures for both risk types. We discuss each risk type separately.

Flood Risk

In Figure 8, we compare all predictor values for the treated unit against those from the donor pool. We describe the neighborhood selection procedure in Section 3.2. For flood risk, the treated unit is Bunde in Meerssen, Limburg. Using Figure 8, we select L suitable predictors out of K.

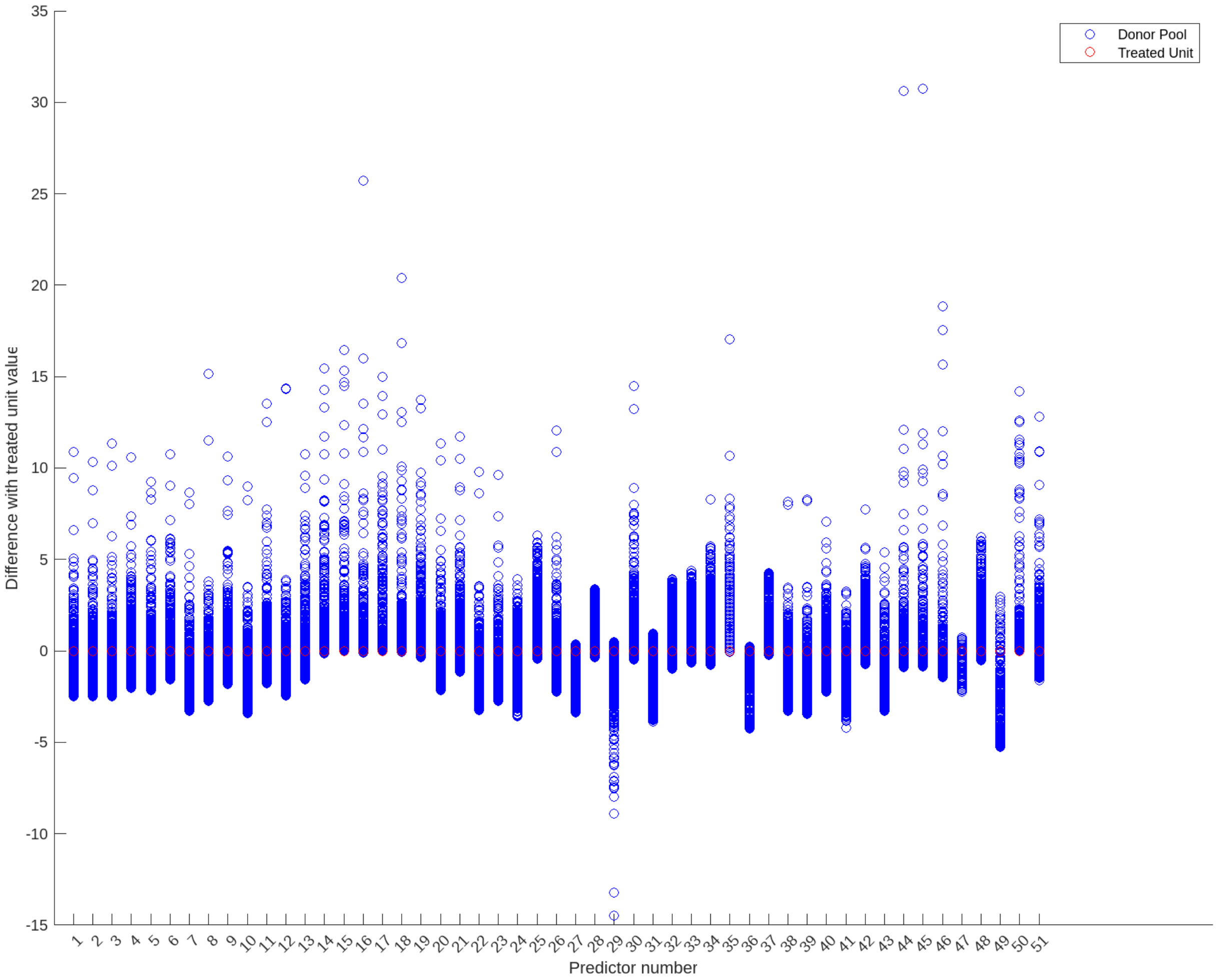

Figure 8. Comparison between all predictor values for treated unit Bunde and the Synthetic Control donor pool.

Ideally, the number of donor pool units with values above and below the treated unit value should be approximately equal. Figure 8 shows that this holds for predictor subset L, {1, 2, 3, 7, 8, 10, 12, 22, 23, 24, 26, 38, 39, 40, 41, 43}. Abadie et al. (2010) include six predictors in their SCM, divided among P pre-intervention outcomes Y_{jt} and a subset S of the neighborhood attributes in L. Following Abadie et al. (2010), we take three pre-intervention outcomes and three neighborhood attributes, such that their ratio is 50/50. Including too many predictors in the optimization procedure increases the risk of overfitting. Including more pre-intervention outcomes and fewer neighborhood attributes similarly increases this risk. In that case, the Synthetic Control is constructed to focus on following the average neighborhood property price of the treated unit, without being structurally similar[14] to it.

Continuing the Predictor Selection procedure, we present the elements p included in Z. To choose subset P, we follow Abadie et al. (2010) in spreading out the pre-intervention outcomes Y_{jt} over the entire pre-treatment period. As such, elements p correspond to $t = \{1, 5, T_0\}$, or 2013, 2017 and 2021. We choose subset S out of the L suitable predictors that follow from the procedure by McClelland and Mucciolo (2022). With $S = 3$ and $T = 16$, we find $\binom{L}{S} = 560$. Of these 560 unique combinations, we

[14] By structurally similar, we mean that the synthetic version closely resembles the treated unit in terms of neighborhood attributes.

include only one in S. Since we fix elements p, Z also has 560 unique combinations. By executing the SCM for all 560 combinations, we can find the optimal combination in terms of MSPE.

As mentioned in Section 4.2.1, the selection of elements j that make up the donor pool is subject to awareness assumptions. Mainly, we exclude neighborhoods of which we assume awareness to be impacted by the intervention. Hence, we exclude Schiedam, Dordrecht, Rotterdam, Amsterdam, Zaanstad, Haarlem, and Gouda from the donor pool for the flood risk Synthetic Control (Hommes et al., 2023). Then, we exclude all neighborhoods with flood variable values larger than zero. This creates another composition of the donor pool, with fewer neighborhoods included. A comparison between the results for both donor pool compositions allows us to make an inference about this assumption. Table 12 displays the optimal combination of elements s in Z.

Donor pool	**Number**	**Variable**	**Description**
With risk	38	a_pau	number of personal vehicles
	39	a_bst_b	number of personal gasoline vehicles
	41	g_pau_hh	average number of personal vehicles per household
Without risk	10	a_gehuwd	number of married inhabitants
	23	a_hh_m_k	number of households with children
	40	a_bst_nb	number of personal vehicles driving on other fuel sources

Table 12. *Description of predictor subsets for Synthetic Control of Bunde.*

In Table 12, “With risk” denotes the donor pool that excludes only Landgraaf, Heerlen, Kerkrade, Meerssen, Valkenburg aan de Geul, and Gulpen-Wittem. In addition to Landgraaf, Heerlen, Kerkrade, Meerssen, Valkenburg aan de Geul, and Gulpen-Wittem, “Without risk” excludes all neighborhoods whose flood risk variable has a value larger than zero. It also shows that the predictor subsets S of each donor pool have no overlap. As the SCM optimizes over both the donor pool weights and the relative importance of each predictor s, a difference in the donor pool composition causes a different predictor subset to obtain the lowest MSPE.

Foundation Risk

Similar to the flood risk procedure, we investigate the convexity condition for foundation risk using Figure 8. As described in Section 3.3, the treated unit is Walvisbuurt in Schiedam.

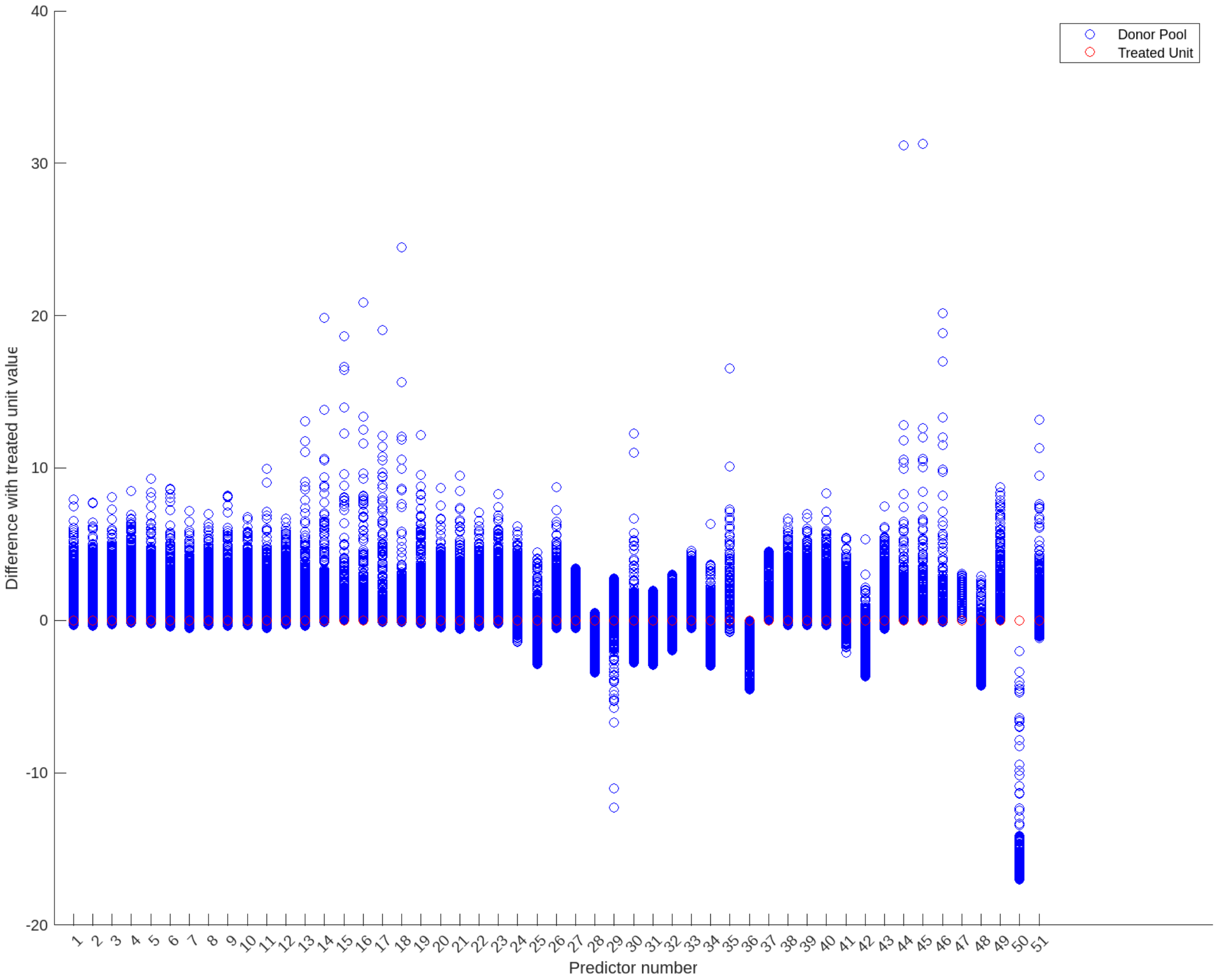

Figure 9. Comparison between all predictor values for treated unit Walvisbuurt and the Synthetic Control donor pool.

Figure 9 shows that the first 23 predictors in K are not recommended to include in the optimization procedure. The values of the treated unit Walvisbuurt for these predictors are among the lowest of the entire donor pool. Selecting predictors L for which the number of donor pool units with values above and below the treated unit value is approximately equal, we find $L = \{25, 29, 30, 31, 32, 34, 41, 42,$ and $48\}$. With $L = 9$ and $S = 3$, we find $\binom{L}{S} = 84$. Testing our awareness assumption, we create two donor pool compositions. Continuing with the same notation, "With risk" denotes the donor pool that excludes Schiedam, Dordrecht, Rotterdam, Amsterdam, Zaanstad, Haarlem, and Gouda, but allows other neighborhoods with foundation risk larger than zero. We use "Without risk" for the donor pool composition that excludes all neighborhoods with foundation risk larger than zero. Out of the 84 unique combinations, Table 13 shows those with the lowest MSPE for both donor pool compositions.

Donor pool	**Number**	**Variable**	**Description**
With risk	31	p_koopw	percentage of rental properties
	32	p_huurw	percentage of rental properties
	34	p_ov_hw	percentage of properties owned by other institutions
Without risk	32	p_huurw	percentage of rental properties
	41	g_pau_hh	average number of personal vehicles per household
	48	ste_oad	number of addresses per km^2

Table 13. *Description of predictor subsets for Synthetic Control of Walvisbuurt.*

Similar to the flood risk predictor subsets in Table 12, Table 13 indicates that the less restrictive donor pool setting obtains the lowest MSPE for a subset of highly similar predictors S. In addition to predictor specification, Appendix II shows high levels of correlation for the elements s in that subset. This suggests that, for a less restrictive donor pool setting in which more neighborhoods lead to higher degrees of freedom, a subset of highly correlated predictors S obtains lower MSPE values more easily. To explain this, we consider an optimization problem with three constraints. When every predictor s is analogous to a constraint, three perfectly correlated predictors lead to three identical constraints. This essentially reduces the number of constraints, thereby easing the optimization procedure. Next to the predictors S, shown in Table 13, we include P pre-intervention outcomes Y_{jt} in Z. Equal to the years we use for flood risk, we include 2013, 2017, and 2021 as the elements corresponding to each p. Next, in Section 5.2.2, we present the SCM findings for the optimal combinations of S displayed by Table 13 for both flood and foundation risk.

5.2.2 General Model

Without the predictor selection method from Section 4.2.2 and the optimization procedure that followed, we could still compare Bunde/Walvisbuurt to the rest of the Netherlands. This way, the donor pool consists of all 2,750 neighborhoods in the SCM dataset, excluding Walvisbuurt and Bunde. Consequently, the donor pool is made up of units with and without risk and with and without awareness of that risk. Figure 10 shows such a comparison.

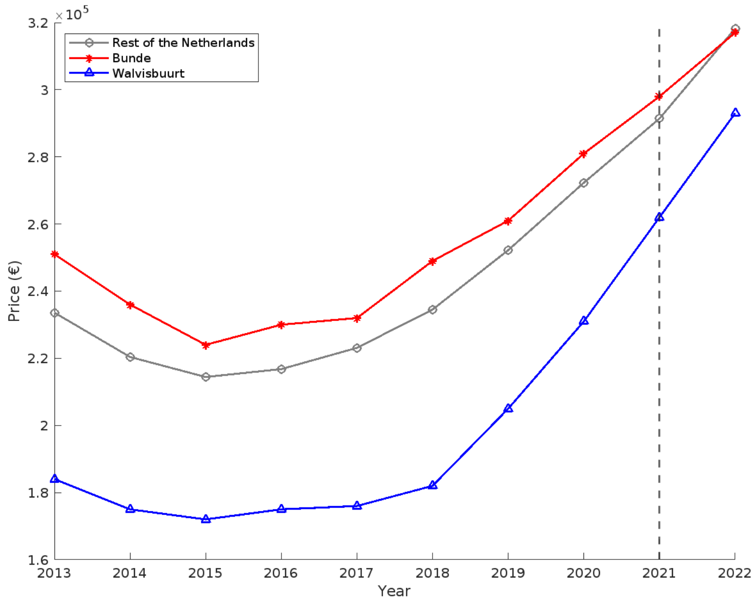

Figure 10. Average property price evolution in euros per year (2013 - 2022): Bunde, Walvisbuurt, and the rest of the Netherlands.

A comparison is difficult. First, due to the difference in average property price ever since 2013. Second, and more importantly, control group *the rest of the Netherlands* contains units with and without risk and with and without awareness of that risk. This is a problem because any differences in the property price evolution from 2021 to 2022 can not be attributed to risk or the awareness of that risk. The careful selection of neighborhoods for each donor pool reflects the importance of the donor pool composition for making inferences about the Synthetic Control. We use these model settings, which each donor pool composition reflects, to check the validity of the awareness assumption for the Synthetic Control. A difference between these model settings indicates one of two things: the assumption of the unaware neighborhoods in the donor pool is wrong, or there is some unexplained baseline effect in 2021 related to the risk variable in question. To find out, we present the SCM for each risk type separately and for both donor pool settings.

Flood Risk

In Figure 11, we include the SCM results for both donor pool settings. It displays the average property price evolution in euros per year over a ten-year period (2013 - 2022) for Bunde and Synthetic Bunde. As we are interested in the effect of the 2021 floods in Limburg on the property prices in Bunde, Table 14 quantifies the difference between the average property price of treated unit Bunde and untreated unit Synthetic Bunde in 2022.

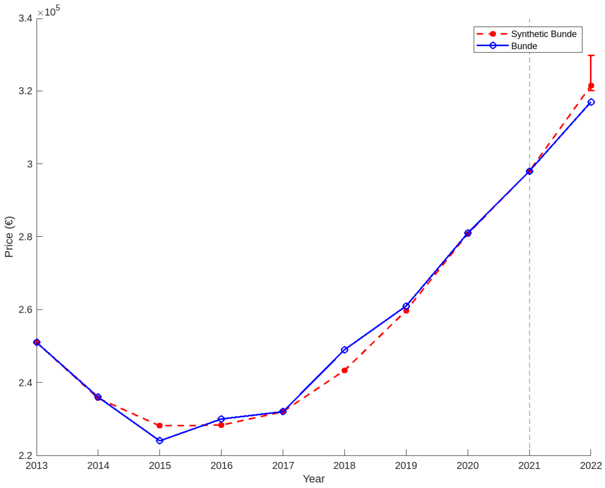

(a) Setting one: including risk larger than zero.

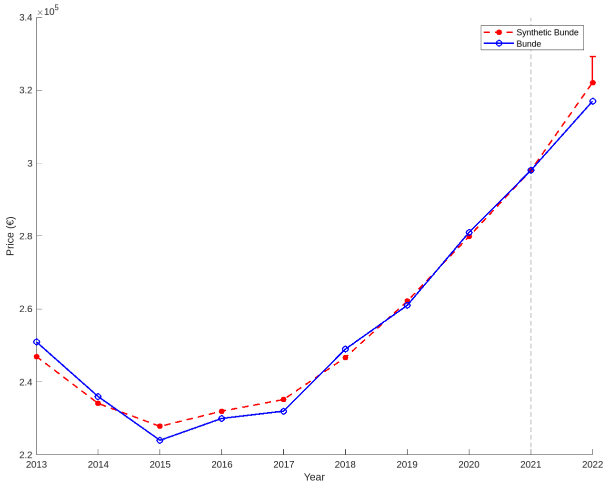

(b) Setting two: excluding risk larger than zero.

Figure 11. Average property price evolution in euros per year (2013 - 2022): Synthetic Control of Walvisbuurt for two donor pool composition settings.

Donor pool	**Variable**	**95% Confidence Interval**					
		lower bound	[%]	estimate	[%]	upper bound	[%]
Including risk	τ_{1T}	$-5.737 \cdot 10^3$	-1.810	$-4.455 \cdot 10^3$	-1.405	$-1.278 \cdot 10^4$	-4.033
Excluding risk	τ_{1T}	$-5.036 \cdot 10^3$	-1.589	$-5.123 \cdot 10^3$	-1.616	$-1.230 \cdot 10^4$	-3.881

Table 14. *Significance of flood risk impact on 2022 property prices in euros for Bunde.*

Figure 11a and 11b contain the results for donor pool setting one and two, respectively. Being less restrictive on the neighborhood requirements by allowing neighborhoods “With risk”, setting one includes 2,653 neighborhoods in the donor pool. In setting two, “Without risk”, 2300 neighborhoods make up the donor pool. Although the donor pool composition and the elements s included in the predictor subsets differ in both settings, the estimates for τ_{1T} are congruent. This suggests that the additional 353 neighborhoods excluded in setting two contain no unexplained effects that impact property prices in 2022. Following the awareness assumption, it indicates that the flood risk awareness in these 353 neighborhoods was not sufficiently impacted to have a significant effect on the 2022 average property price of Synthetic Bunde. In addition to being congruent, both estimates for τ_{1T} find a significant negative impact of the 2021 inundations in Limburg on the average property price in Bunde. In both donor pool settings, the size of the effect ranges from approximately -1.5% to -4.0% of the 2022 average neighborhood property price in Bunde. This result is roughly in line with previous findings by Daniel et al. (2009) and Bosker et al. (2019). Although Daniel et al. (2009) presents a cumulative effect of two floods of -9.1% over a multi-year period, it implies a yearly effect that falls within the 95% confidence interval from Table 14. As suggested by Wachinger et al. (2013), the materialization of flood risk in the form of the 2021 floods effectively raises awareness. Changing consumer behavior then translates into a significant negative effect on property prices through the WOZ estimation procedure.

Foundation Risk

Figure 12 includes the results for each donor pool setting. We add Table 15 to quantify the difference between Walvisbuurt and Synthetic Walvisbuurt in 2022.

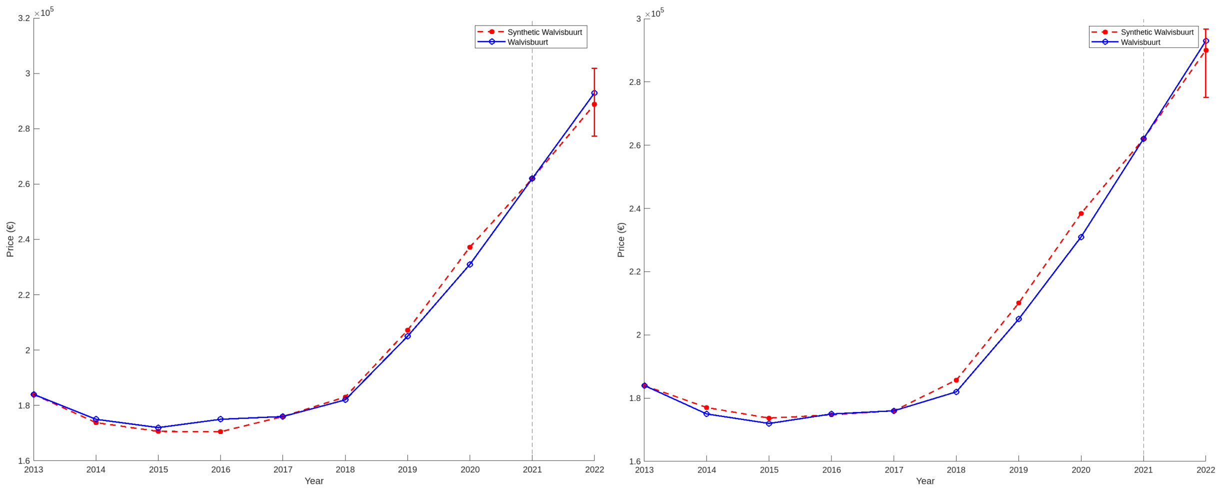

(a) Setting one: including risk larger than zero.

(b) Setting two: excluding risk larger than zero.

Figure 12. Average property price evolution in euros per year (2013 - 2022): Synthetic Control of Walvisbuurt for two donor pool composition settings.

Figure 12a shows the SCM for the less restrictive donor pool setting one. Setting two puts more restrictions on the donor pool, leading to Figure 12b. Combining the difference in the composition of S with the difference in donor pool settings leads to the difference between the two SCM results. Setting one has more neighborhoods at its disposal, leading to a slightly better fit. Still, both model settings find similar results. This suggests that the additional neighborhoods excluded in setting two are not clouded by any remaining awareness-raising effects. Put simply, the neighborhoods with a foundation risk score larger than zero, which we exclude in setting two, have not become more aware of foundation risk in 2021, as this would have caused an underestimation of the effect in setting one compared to setting two. We recognize that this explanation assumes the existence of an impact of the awareness of foundation risk on property prices and the validity of our awareness assumption. To put this assumption in perspective, Tversky and Kahneman (1973) state that the “availability heuristic” drives decision-making in humans. This means that home buyers in Walvisbuurt base their perception of foundation risk on substantial events that are fresh in their memory. The substantial event(s) in Walvisbuurt are the 2021 initiatives in Schiedam, Dordrecht, and Rotterdam. It is uncertain if these initiatives can be classified as substantial and effective at increasing awareness among inhabitants of high-risk neighborhoods. Research by Hommes et al. (2023) shows that public knowledge of the presence of foundation damage negatively impacts a property’s price. As such, we are inclined to question the effectiveness of the 2021 municipal initiatives around foundation risk in Schiedam, Dordrecht, and Rotterdam. Mainly because the results for both SCM settings show large variations in the 2022 property price values, as indicated by the confidence intervals in Figure 12 and Table 15.

Donor pool	**Variable**	**95% Confidence Interval**					
		lower bound	[%]	estimate	[%]	upper bound	[%]
Including risk	τ_{1T}	$1.562 \cdot 10^4$	5.332	$4.172 \cdot 10^3$	1.424	$-8.829 \cdot 10^3$	-3.014
Excluding risk	τ_{1T}	$1.792 \cdot 10^4$	6.117	$2.902 \cdot 10^3$	0.999	$-3.849 \cdot 10^3$	-1.292

Table 15. *Significance of foundation risk impact on 2022 property prices in euros for Walvisbuurt.*

Table 15 presents this result more clearly, adding the impact as a percentage of the 2022 average property price in Walvisbuurt. For donor pool setting one, the impact of foundation risk has a 95% probability of falling between −3.014% to 5.332%. For donor pool setting two, the interval is shifted, ranging from −1.292% to 6.117%. The donor pool in setting one consists of 2,589 neighborhoods, compared to the 1,356 neighborhoods that make up the donor pool in setting two. More neighborhoods translate to higher degrees of freedom but do not reduce the size of the confidence interval.

6 Conclusion

To conclude this research, we base the structure on the two research questions:

1. How do flood and foundation risk impact Dutch neighborhood property prices?
2. How does the awareness of flood and foundation risk impact Dutch neighborhood property prices?

For consistency and clarity, we continue discussing each model separately. Using the Hedonic Pricing Model, Section 6.1 tries to answer the first research question. Then, we study local effects related to risk awareness with a novel Synthetic Control Method approach in the second part. Hence, Section 6.2 studies the second research question. Both Section 6.1 and 6.2 discuss model limitations and raise questions on its assumptions. We also cover the significance of the results and propose ways to improve further research.

6.1 Hedonic Pricing Model

To analyze the results of the HPM in a structured manner, we identify three subquestions that help us draw conclusions on their meaning:

1. Can flood and foundation risk help explain average neighborhood property prices?
2. Does a higher risk variable, *ceteris paribus*, lead to higher or lower values in the dependent variable?
3. Can we label this effect as statistically significant?

We answer the first question using the (Adjusted) R^2 value. Including flood risk increased the (Adjusted) R^2 from 0.625 to 0.626. Including foundation risk increased the R^2 value from 0.625 to 0.626 but did not improve the Adjusted R^2. This shows that the explanatory power of both risk variables is, if present, limited. We assess question two using the coefficient sign of each risk type in the regression results. For the risk variables, a negative sign means that more risk, *ceteris paribus*, leads to a lower average neighborhood property price. This would be in line with previous results (Bosker et al., 2019; Daniel et al., 2009; Hommes et al., 2023; Reeken & Phlippen, 2022). We find opposite signs for flood and foundation risk, with foundation risk having a positive sign. The size of the coefficients is related to the *ceteris paribus* effect of each risk variable. As we deal with a logarithmic transformation of the dependent variable, the size does not directly translate to an impact in euros. Also, the measurement unit of each risk type is different, further increasing the difficulty of comparing their price effect. Still, with the value ranges of each risk variable being approximately equal, we note that the flood risk coefficient of -0.009 is roughly twice the size of the foundation risk coefficient of 0.004. Finally, both flood and foundation risk have a p-value < 0.05. As such, we reject the null hypothesis of 'no effect'.

For inference on the significance of these results, we investigate the validity of our HPM assumptions. To start with our assumption on the absence of omitted variables, we recognize the limitations of the dataset we used. During the pre-processing procedure, neighborhoods and neighborhood attributes are lost. As these missing data entries are a Missing Completely At Random, they create no bias in the OLS results. However, the lost neighborhood attributes can cause omitted variable bias. For example, in the absence of a variable that captures property price premia for houses with a direct view on the water, higher flood risk values for houses next to a river could appear to have a positive effect on property prices. That is because flood risk is correlated with the omitted variable that captures a direct view on the water, which is linked to a property price premium. When the latter variable is omitted, the correlation remains undetected, and a bias clouds the regression results.

Then, the Bucketing Analysis shows that the explanatory power of the independent variables differs across buckets. Hence, the impact of flood and foundation risk is not uniform across the different property price ranges. This reveals structural breaks and heterogeneity in the impact of the neighborhood attributes on the average neighborhood property price, which violates the homogeneity assumption.

All taken into account, we find that the impact of flood risk on average neighborhood property prices is statistically significant and in line with previous flood risk results by Daniel et al. (2009), Bosker et al. (2019) and Reeken and Phlippen (2022), yet smaller. On the other hand, the impact of foundation risk is not in line with previous results by Hommes et al. (2023). However, violating multiple HPM assumptions makes it hard to make inferences from these findings. As such, Section 6.2 reinvestigates the first research question on a more local level with a different approach. Also, we study research question two to take a different viewpoint on the subject. For further research, it is key to meet the model assumptions. Regarding omitted variables, we can achieve this by extending the dataset to include all neighborhood attributes that explain average neighborhood property prices. For example, combining datasets from different sources, as opposed to only using *Kerncijfers Wijken en Buurten* from the CBS. We can do this using Geographic Information Systems like QGIS, which projects geopackage-type datasets on top of spatial neighborhood border maps. Then, if supported by the dataset extension, we can obtain even more granular data using spatial border maps of individual homes. To fix the second model assumption, we should do a more thorough analysis to find the exact location of the structural breaks in the data. Breaking up the dataset into different parts, according to these structural breaks, allows for the homogeneity of neighborhood attributes within each group. This distinction also creates the opportunity to discover whether the impact of flood and foundation differs between groups. An example of an econometric method that would be suitable for such an analysis is the Random Forest method.

6.2 Synthetic Control Method

We answer the second research question separately for flood and foundation risk. As such, we cover flood risk in Section 6.2.1 and foundation risk in Section 6.2.2. The discussion jointly covers both risk types.

6.2.1 Flood Risk

Section 5.2 shows that increasing the awareness of flood risk has a significant negative impact of approximately 1.5% to 4.0% on the average property price of the neighborhood Bunde in Limburg. Although we obtain this result using a novel application of the Synthetic Control Method, we find it to be in line with previous results by Daniel et al. (2009) and Bosker et al. (2019). The range corresponds to the boundaries of the 95% confidence interval. Also, the absence of varying outcomes across donor pool settings suggests that no significant awareness effect prevailed in the neighborhoods that we additionally exclude in the more rigorous donor pool setting two. Put simply, excluding all remaining neighborhoods with flood risk larger than zero from the donor pool did not change the evolution of the average property price of Synthetic Bunde. This demonstrates that no unexplained effects that impact the 2022 property prices in a way similar to the 2021 floods were present in these neighborhoods. Moreover, these findings reflect the dominance of the "availability heuristic" in human decision-making, as per Tversky and Kahneman (1973). This illustrates how humans assess risk on (the absence of) observable and recent dramatic events.

6.2.2 Foundation Risk

For foundation risk, execution of SCM shows that the 2022 average neighborhood property price of Walvisbuurt is not significantly different from Synthetic Walvisbuurt for both donor pool settings. In Table 15, the upper and lower bound of the 95% confidence have opposite signs. This result indicates that the presumably awareness-raising initiatives around foundation risk in Schiedam in 2021 do not significantly impact property prices in Walvisbuurt. This result is not in line with research done by Hommes et al. (2023), which reports an 8% average discount on properties with reported foundation damage. The difference in research setup could explain this result. Where Hommes et al. (2023) base their awareness assumption on reported damages in the property listing, we assume the 2021 initiatives around foundation risk in Schiedam to create sufficient awareness that a significant impact on property prices in Walvisbuurt can be observed. Although the effectiveness of the initiatives to inform those interested in buying a house in Walvisbuurt is uncertain, we expect it to be lower than in the Hommes et al. (2023) setup. Restating Wachinger et al. (2013) about risk perception, we reckon that people do not perceive risks until they are materialized. Materialization of foundation risk would ultimately mean the collapse of a house. Hence, without conspicuous materialization of foundation risk and with seemingly ineffective initiatives to raise awareness, property prices in Walvisbuurt fail to reflect a significant impact of risk awareness on consumer behavior. This is in contrast to the dramatic 2021 inundation event in Limburg, which we identify as a successful risk awareness-increasing event, with the corresponding results to show for it.

We jointly discuss the opportunities for further flood and foundation risk research by covering assumptions, data requirements, and the research setup.

Most assumptions relate to the donor pool composition, which must be free of neighborhoods with similar awareness raining events. Although we perform the SCM for two donor pool configurations, we suggest a more careful analysis of neighborhoods and their risk awareness. Any neighborhoods in the donor pool with an increased awareness of the specific risk type weaken the effect, as portrayed by the τ_{1T} variable. We also assume the absence of a spillover effect. Translated to the property transaction case, this means people do not move across neighborhoods. This way, the increased awareness effects only affect the neighborhoods that experienced the event. Even though the impact of a few crossovers might be small, we cannot guarantee that inhabitants of Bunde/Walvisbuurt do not take their risk awareness to other provinces and vice versa. This might cloud any effects we observe.

With regards to data requirements, we understand the limitations of the *Kerncijfers Wijken en Buurten* dataset. Aggregating these datasets over ten years loses many predictor variables in the cross-section. This limits the variety of information that the predictor variables hold. Abadie (2021) also acknowledges the importance of large variability of predictor variables in the dataset. An adequate Synthetic Control needs to be as structurally similar as possible, which requires socioeconomic, demographic, and many more categorical variables that capture a neighborhood's complete identity.

Future research setups can include similar analyses for the other neighborhoods. For example, the neighborhoods excluded from the donor pools would serve as suitable candidates to conduct this research on. Aggregating these results could paint a broader picture of the impact in a bigger area. Furthermore, we can think of the awareness-increasing event decision. As intervention effects accumulate over the years, Abadie (2021) repeats the value of a large post-intervention period. The limited availability of events that increase the awareness of flood and foundation risk bounds us to a post-event period that starts after 2021. As such, the impact we find comes from a single year. Extending this research to include data from years to come has the potential to see how this effect evolves. Will it get stronger or slowly fade as the water flows away?

References

Abadie, A. (2021). Using synthetic controls: Feasibility, data requirements, and methodological aspects. *Journal of Economic Literature*, *59*(2), 391–425.

Abadie, A., Diamond, A., & Hainmueller, J. (2010). Synthetic control methods for comparative case studies: Estimating the effect of california's tobacco control program. *Journal of the American statistical Association*, *105*(490), 493–505.

Abadie, A., Diamond, A., & Hainmueller, J. (2015). Comparative politics and the synthetic control method. *American Journal of Political Science*, *59*(2), 495–510.

Atreya, A., & Ferreira, S. (2015). Seeing is believing? evidence from property prices in inundated areas. *Risk Analysis*, *35*(5), 828–848.

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.

Barberis, N. (2013). Psychology and the financial crisis of 2007-2008. *Financial innovation: too much or too little*, 15–28.

Bian, X., Lin, Z., & Liu, Y. (2018). House price, loan-to-value ratio and credit risk. *Journal of Banking & Finance*, *92*, 1–12.

Bluhm, C., Overbeck, L., & Wagner, C. (2016). *Introduction to credit risk modeling*. Crc Press.

Bosker, M., Garretsen, H., Marlet, G., Ponds, R., Poort, J., van Dooren, R., van Woerkens, C., et al. (2016). Met angst en beven: Verklaringen voor dalende huizenprijzen in het groningse aardbevingsgebied.

Bosker, M., Garretsen, H., Marlet, G., & van Woerkens, C. (2019). Nether lands: Evidence on the price and perception of rare natural disasters. *Journal of the European Economic Association*, *17*(2), 413–453.

Bunni, N. G. (2003). *Risk and insurance in construction*. Routledge.

Campiglio, E., Dafermos, Y., Monnin, P., Ryan-Collins, J., Schotten, G., & Tanaka, M. (2018). Climate change challenges for central banks and financial regulators. *Nature climate change*, *8*(6), 462–468.

Carney, M. (2015). Breaking the tragedy of the horizon - climate change and financial stability - speech by mark carney. https://www.bankofengland.co.uk/speech/2015/breaking-the-tragedy-of-the-horizon-climate-change-and-financial-stability

Daniel, V. E., Florax, R. J., & Rietveld, P. (2009). Floods and residential property values: A hedonic price analysis for the netherlands. *Built environment*, *35*(4), 563–576.

Efron, B., & Tibshirani, R. J. (1994). *An introduction to the bootstrap*. CRC press.

Greeven, T. M. (2021). *The relation of climate-related risk drivers with probability of default models of a dutch mortgage portfolio* (Doctoral dissertation). Rijksuniversiteit Groningen.

Heij, C., Heij, C., de Boer, P., Franses, P. H., Kloek, T., van Dijk, H. K., et al. (2004). *Econometric methods with applications in business and economics*. Oxford University Press.

Hommes, S., Phlippen, S., Reeken, J. v., Schreuder, C., & Ypma, F. (2023). Gemelde funderingsschade leidt tot forse prijskorting bij woningverkoop. https://esb.nu/gemelde-funderingsschade-leidt-tot-forse-prijskorting-bij-woningverkoop/

Jacobs, C. (2021). Overstroming juli 2021. https://www.vrzl.nl/application/files/5116/4301/9488/Rapport_onderzoek_overstroming_VRZL.pdf

Janssen, A. (2023). Gevolgen overstromingen limburg. https://publications.deltares.nl/11207700_007_0011.pdf

Jonkman, B. (2021). High water in limburg during the summer 2021 was more drastic than the river floods in 1993 and 1995. https://www.tudelft.nl/en/2021/tu-delft/high-water-in-limburg-during-the-summer-2021-was-more-drastic-than-the-river-floods-in-1993-and-1995

Kars, J. C. (2021). *Predicting neighborhood prices: Machine learning and hedonic pricing in the dutch housing market* (Doctoral dissertation). Tilburg University.

Kok, S., & Angelova, L. (2020). Impact droogte op funderingen. https://www.verzekeraars.nl/media/7875/20200930-rapport-impact-droogte-op-funderingen.pdf

Koopman, P. (2020). Deltaplan aanpak funderingsschade - verzekeraars. https://www.verzekeraars.nl/media/8532/manifest-funderingen.pdf

Little, R. J. (1992). Regression with missing x's: A review. *Journal of the American statistical association*, *87*(420), 1227–1237.

McClelland, R., & Mucciolo, L. (2022). An update on the synthetic control method as a tool to understand state policy. *Tax Policy Center*.

Meadows, D. H., Meadows, D. L., Randers, J., & Behrens III, W. W. (1972). The limits to growth-club of rome.

Monnin, P. (2018). Integrating climate risks into credit risk assessment-current methodologies and the case of central banks corporate bond purchases. *Council on economic policies, discussion note*, *4*.

Oladunni, T., & Sharma, S. (2016). Hedonic housing theory—a machine learning investigation. *2016 15th IEEE International Conference on Machine Learning and Applications (ICMLA)*, 522–527.

Pistrika, A., Tsakiris, G., & Nalbantis, I. (2014). Flood depth-damage functions for built environment. *Environmental Processes*, *1*, 553–572.

Puth, M.-T., Neuhäuser, M., & Ruxton, G. D. (2015). On the variety of methods for calculating confidence intervals by bootstrapping. *Journal of Animal Ecology*, *84*(4), 892–897.

Reeken, J. v., & Phlippen, S. (2022). Is flood risk already affecting house prices? https://www.abnamro.com/research/en/our-research/is-flood-risk-already-affecting-house-prices

Rosen, S. (1974). Hedonic prices and implicit markets: Product differentiation in pure competition. *Journal of political economy*, *82*(1), 34–55.

Slager, K. (2019). Handboek overstromingsrisico's op de kaart. https://www.helpdeskwater.nl/publish/pages/132175/handboek_overstromingsrisicos_op_de_kaart.pdf

Soley-Bori, M., et al. (2013). Dealing with missing data: Key assumptions and methods for applied analysis. *Boston University*, *4*(1), 1–19.

Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive psychology*, *5*(2), 207–232.

van Essen, F. (2022). *The impact of climate-related risk on the probability of default in a corporate shipping portfolio* (Doctoral dissertation). Rijksuniversiteit Groningen.

Wachinger, G., Renn, O., Begg, C., & Kuhlicke, C. (2013). The risk perception paradox—implications for governance and communication of natural hazards. *Risk analysis*, *33*(6), 1049–1065.

Zwillinger, D. (2018). *Crc standard mathematical tables and formulas*. CRC press.

Appendices

Appendix I

Table 16 gives an overview of all variables from the *Kerncijfers Wijken en Buurten* 2021 dataset, and which variables we include in the HPM and SCM.

Neighborhood Attribute	Description	Hedonic Pricing Model	Synthetic Control Method
1_gwb_code_10	10 figure neighborhood code	x	x
2_gwb_code_8	8 figure neighborhood code	x	x
3_regio	neighborhood name	x	x
4_gm_naam	municipality name	x	x
5_recs	region type	x	x
6_gwb_code	6 figure neighborhood code	x	x
7_ind_wbi	indicator of changing neighborhood code	x	x
8_a_inw	number of inhabitants	x	x
9_a_man	number of male inhabitants	x	x
10_a_vrouw	number of female inhabitants	x	x
11_a_00_14	number of inhabitants, whose age ranges from 0 until 14	x	x
12_a_15_24	number of inhabitants, whose age ranges from 15 until 24	x	x
13_a_25_44	number of inhabitants, whose age ranges from 25 until 44	x	x
14_a_45_64	number of inhabitants, whose age ranges from 45 until 64	x	x
15_a_65_oo	number of inhabitants, whose age is 65 and above	x	x
16_a_ongeh	number of unmarried inhabitants	x	x
17_a_gehuwd	number of married inhabitants	x	x
18_a_gesch	number of divorced inhabitants	x	x
19_a_verwed	number of widowed inhabitants	x	x
20_a_w_all	number of inhabitants with a western (i.e., Europe, North America, Oceania, Indonesia or Japan) migration background	x	x
21_a_nw_all	number of inhabitants with a non-western (i.e., Africa, South America, Asia, or Turkey) migration background	x	x
22_a_marok	number of inhabitants with a Moroccan migration background	x	x
23_a_antaru	number of inhabitants with a Dutch Antilles migration background	x	x
24_a_suri	number of inhabitants with a Surinamese migration background	x	x
25_a_tur	number of inhabitants with a Turkish migration background	x	x
26_a_ov_nw	number of inhabitants with a non-western migration background, excluding Morocco, Netherlands Antilles, Suriname, and Turkey	x	x
27_a_geb	number of births	x	
28_p_geb	number of births per 1000 inhabitants	x	
29_a_ste	number of deaths	x	
30_p_ste	number of deaths per 1000 inhabitants	x	
31_a_hh	number of households	x	x
32_a_1p_hh	number of single-person households	x	x
33_a_hh_z_k	number of households without children	x	x
34_a_hh_m_k	number of households with children	x	x
35_g_hhgro	average household size	x	x
36_bev_dich	number of inhabitants per km^2	x	x
37_a_woning	number of residential properties	x	x
38_g_woz	average residential property value	x	x
39_p_1gezw	percentage of single-family homes	x	x
40_p_mgezw	percentage of multi-family homes	x	x
41_p_bewndw	percentage of inhabited homes	x	x
42_p_leegsw	percentage of uninhabited homes	x	x
43_p_koopw	percentage of owner-occupied properties	x	x
44_p_huurw	percentage of rental properties	x	x
45_p_wcorpw	percentage of properties owned by housing corporations	x	x
46_p_ov_hw	percentage of properties owned by other institutions (i.e., individuals, institutional investors, companies)	x	x
47_p_e_o_w	percentage of properties owned by unknown parties	x	x
48_p_bjj2k	percentage of properties built before year 2000	x	x
49_p_bjo2k	percentage of properties built from year 2000 onwards	x	x
50_g_ele	average electricity consumption [kWh]	x	

Neighborhood Attribute	Description	Hedonic Pricing Model	Synthetic Control Method
51_g_ele_ap	average electricity consumption apartments [kWh]		
52_g_ele_tw	average electricity consumption terraced houses [kWh]		
53_g_ele_hw	average electricity consumption corner houses [kWh]		
54_g_ele_2w	average electricity consumption semidetached houses [kWh]		
55_g_ele_vw	average electricity consumption detached houses [kWh]		
56_g_ele_hu	average electricity consumption rental properties [kWh]	x	
57_g_ele_ko	average electricity consumption owner-occupied properties [kWh]	x	
58_g_gas	average gas consumption [m^3]	x	
59_g_gas_ap	average gas consumption apartments [m^3]		
60_g_gas_tw	average gas consumption terraced houses [m^3]		
61_g_gas_hw	average gas consumption corner houses [m^3]		
62_g_gas_2w	average gas consumption semidetached houses [m^3]		
63_g_gas_vw	average gas consumption detached houses [m^3]		
64_g_gas_hu	average gas consumption rental properties [m^3]	x	
65_g_gas_ko	average gas consumption owner-occupied properties [m^3]	x	
66_p_stadsv	percentage of properties with district heating		
67_a_opl_lg	number of inhabitants with lower educational attainment (i.e., highest achieved level of mbo, vmbo or lower)		
68_a_opl_md	number of inhabitants with only high school education		
69_a_opl_hg	number of inhabitants with higher educational attainment (i.e., highest achieved level of hbo or university education)		
70_p_arb_pp	percentage of inhabitants belonging to the working population		
71_p_arb_wn	percentage of working people belonging to the group of employees		
72_p_arb_zs	percentage of working people belonging to the group of self-employed workers		
73_a_inkont	number of inhabitants that earn income	x	
74_g_ink_po	average income per income earner		
75_g_ink_pi	average income per inhabitant		
76_p_ink_li	percentage of people in a household belonging to the group of 40% lowest income earners	x	
77_p_ink_hi	percentage of people in a household belonging to the group of 20% highest income earners	x	
78_g_hh_sti	average standardized household income		
79_p_hh_li	percentage of households belonging to the group of 40% lowest income earners	x	
80_p_hh_hi	percentage of households belonging to the group of 20% highest income earners	x	
81_p_hh_lkk	percentage of households with a low income (i.e., standardized and deflated income below 9,249 euros)	x	
82_p_hh_osm	percentage of households under or around the social minimum (i.e., required income to provide for livelihood, as determined by law)	x	
83_p_hh_110	percentage of households under 110% of social minimum (i.e., required income to provide for livelihood, as determined by law)	x	
84_p_hh_120	percentage of households under 120% of social minimum (i.e., required income to provide for livelihood, as determined by law)	x	
85_m_hh_ver	median capital of households	x	
86_a_soz_wb	number of inhabitants living of social security 'bijstand'	x	
87_a_soz_ao	number of inhabitants living of social security 'AO'		
88_a_soz_ww	number of inhabitants living of social security 'WW'	x	
89_a_soz_ow	number of inhabitants living of social security 'AOW'	x	
90_a_jz_tn	number of young inhabitants (i.e., until the age of 23) with professional help		
91_p_jz_tn	percentage of young inhabitants (i.e., until the age of 23) with professional help		
92_a_wmo_t	number of inhabitants with tailor-made work arrangement (social care)		
93_p_wmo_t	percentage of inhabitants with tailor-made work arrangement (social care)		
94_a_bedv	number of business locations	x	
95_a_bed_a	number of business locations in agriculture, forestry, and fishing	x	
96_a_bed_bf	number of business locations in industry and energy	x	
97_a_bed_gi	number of business locations in trade and catering	x	
98_a_bed_hj	number of business locations in transport, information, and communication	x	
99_a_bed_kl	number of business locations in financial services and real estate	x	
100_a_bed_mn	number of business locations in business services	x	
101_a_bed_oq	number of business locations in government, education, and health	x	
102_a_bed_ru	number of business locations in culture, recreation, and other services	x	
103_a_pau	number of personal vehicles	x	x
104_a_bst_b	number of personal gasoline vehicles	x	x
105_a_bst_nb	number of personal vehicles driving on other fuel sources (i.e., diesel, LPG, electricity, hydrogen, alcohol, LNG, and CNG)	x	x

Neighborhood Attribute	**Description**	**Hedonic Pricing Model**	**Synthetic Control Method**
106_g_pau_hh	average number of personal vehicles per household	x	x
107_g_pau_km	average number of personal vehicles per area [per km^2]	x	x
108_a_m2w	number of motorcycle vehicles	x	x
109_g_afs_hp	average distance to a general practitioner [km]	x	
110_g_afs_gs	average distance to a large supermarket [km]	x	
111_g_afs_kv	average distance to a daycare [km]	x	
112_g_afs_sc	average distance to a school [km]	x	
113_g_3km_sc	average number of schools located within 3 km	x	
114_a_opp_ha	total surface area [ha]	x	x
115_a_lan_ha	surface area land [ha]	x	x
116_a_wat_ha	surface area water [ha]	x	x
117_pst_mvp	most common postal code	x	x
118_pst_dekp	coverage rate of most common postal code	x	
119_ste_mvs	measure of urbanity (i.e., number of adresses per km^2)	x	
120_ste_oad	density of adresses [per km^2]	x	

Table 16. *Description of all neighborhood attributes included in the Hedonic Pricing Model and Synthetic Control Method.*

Appendix II

Due to the large similarity between some variables in the HPM dataset, high degrees of correlation can occur. In Figure 13, we combine the correlation between all HPM variables into a correlation heatmap.

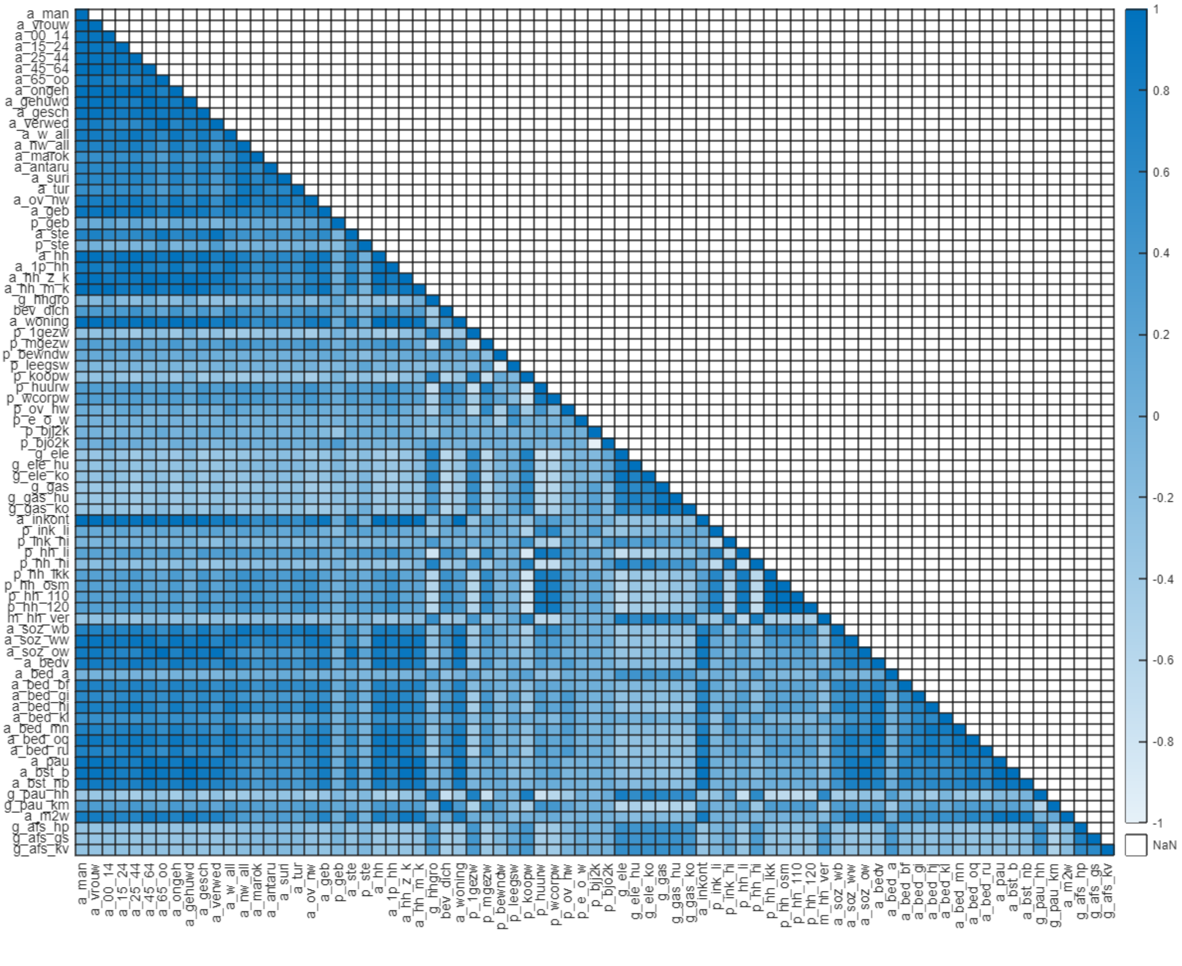

Figure 13. Correlation matrix of all Hedonic Pricing Model variables.

Appendix III

In Table 17, we give an interpretation of the PC loadings. Based on PC weights for each neighborhood attribute, we see clear themes arise for each Principal Component.

Principal Component	**Description**
$PC1$	income and urbanization
$PC2$	building period (before or after the year 2000)
$PC3$	neighborhood demographics
$PC4$	property development
$PC5$	property ownership

Table 17. *Description of the loadings of the first five Principal Components.*

Appendix IV

Restating the formula for the Adjusted R^2 in Equation 11, we investigate when negative values are obtained.

$$\text{Adjusted } R^2 = 1 - \frac{N-1}{N-K}\left(1 - R^2\right) \tag{11}$$

For example, by adding the flood and foundation risk variables, bucket 1300 - max went from an Adjusted $R^2 = -0.442$ to -0.550. In that case, $N = 10$ and adding flood and foundation risk changes K from three to five. Then, Figure 14 shows the variation in Adjusted R^2 for varying K and R^2.

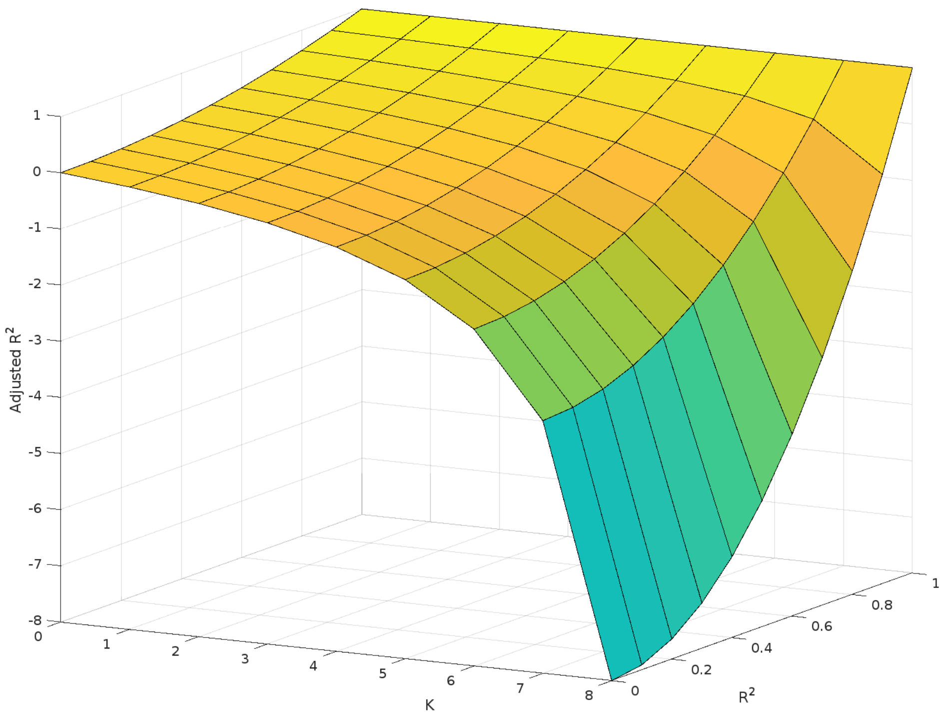

Figure 14. 3D visualization of the Adjusted R^2 for $N = 10$ and varying K and R^2.

As shown by Figure 14, increasing K from three to five, with major improvements in R^2, causes the Adjusted R^2 to fall below 0.