DISCUSSION PAPER SERIES

No. 10307

NETHER LANDS

Maarten Bosker, Harry Garretsen, Gerard Marlet
and Clemens van Woerkens

INTERNATIONAL TRADE AND REGIONAL ECONOMICS

Centre for Economic Policy Research

ISSN 0265-8003

NETHER LANDS: EVIDENCE ON THE PRICE AND PERCEPTION OF RARE FLOOD DISASTERS

Maarten Bosker, Harry Garretsen, Gerard Marlet and Clemens van Woerkens

Discussion Paper No. 10307
December 2014
Submitted 11 December 2014

Centre for Economic Policy Research
77 Bastwick Street, London EC1V 3PZ, UK
Tel: (44 20) 7183 8801
www.cepr.org

NETHER LANDS: EVIDENCE ON THE PRICE AND PERCEPTION OF RARE FLOOD DISASTERS†

Abstract

The Netherlands is one of the most flood prone countries in the world. It also has the best flood defenses in the world that are built to make floods an extremely rare event. This paper uses the unique Dutch setting to provide evidence on the price and perception of rare environmental disasters. In particular, it allows us to precisely identify whether houses in areas that would flood, in case the Dutch flood defenses fail, cost less than otherwise equal homes that do not run any risk. We do this using a border-discontinuity design that heavily relies on the (spatial) detail of our unique dataset covering every housing transaction over the period 1999-2011. Despite €1.1 billion publicly spent on flood protection each year, we find a 1% flood risk price discount on houses in flood prone areas. The median household living in a flood prone area would be willing to pay an additional €69 per year to be fully insured against flood risk. Our estimate implies that people's perceived flood risk is substantially higher than the official protection levels at which the government claims to uphold the country's flood defenses.

JEL Classification: D8, Q54 and R21
Keywords: house prices, implicit insurance and rare flood disasters

Maarten Bosker bosker@ese.eur.nl
Erasmus University Rotterdam, Tinbergen Institute and CEPR

Harry Garretsen j.h.garretsen@rug.nl
University of Groningen, Cambridge University and CESifo

Gerard Marlet marlet@atlasvoorgemeenten.nl
University of Groningen and Atlas voor Gemeenten

Clemens van Woerkens woerkens@atlasvoorgemeenten.nl
University of Groningen and Atlas voor Gemeenten

---

† Please send all correspondence to: bosker@ese.eur.nl. We thank Sacha Kapoor and Dinand Webbink, as well as seminar participants at the London School of Economics, the University of Venice, Oxford University, the University of Sheffield, the University of Barcelona, Wageningen University, VU University Amsterdam, and the Netherlands Bureau for Economic Policy Analysis (CPB), for useful comments and suggestions that have improved our paper.

Nether Lands

Maarten Bosker, Harry Garretsen, Gerard Marlet, Clemens van Woerkens[1]

December 2014

Abstract

The Netherlands is one of the most flood prone countries in the world. It also has the best flood defenses in the world that are built to make floods an extremely rare event. This paper uses the unique Dutch setting to provide evidence on the price and perception of rare environmental disasters. In particular, it allows us to precisely identify whether houses in areas that would flood, in case the Dutch flood defenses fail, cost less than otherwise equal homes that do not run any risk. We do this using a border-discontinuity design that heavily relies on the (spatial) detail of our unique dataset covering every housing transaction over the period 1999-2011. Despite €1.1 billion publicly spent on flood protection each year, we find a 1% flood risk price discount on houses in flood prone areas. The median household living in a flood prone area would be willing to pay an additional €69 per year to be fully insured against flood risk. Our estimate implies that people's perceived flood risk is substantially higher than the official protection levels at which the government claims to uphold the country's flood defenses.

[1]Bosker (1.8 m above sea level) is affiliated with the Erasmus University Rotterdam, the Tinbergen Institute and CEPR; Garretsen (0.9 m above sea level) to the University of Groningen, Cambridge University and CESifo; and Marlet (2.2 m above sea level) and van Woerkens (0.3 m above sea level) to the University of Groningen, and Atlas voor Gemeenten. Please send all correspondence to: bosker@ese.eur.nl. We thank Sacha Kapoor and Dinand Webbink, as well as seminar participants at the London School of Economics, the University of Venice, Oxford University, the University of Sheffield, the University of Barcelona, Wageningen University, VU University Amsterdam, and the Netherlands Bureau for Economic Policy Analysis (CPB), for useful comments and suggestions that have improved our paper.

1. Introduction

The Netherlands is one of the most flood prone countries in the world. According to the 2014 World Risk Report, it is the 12th most exposed country to natural disasters in the world. The expected future rise in sea levels and in the frequency of peak water levels in the Rhine and Meuse rivers will only further increase this risk.[2]

Figure 1. Flood risk and elevation levels in the Netherlands.

a. flood risk (blue areas), main agglomerations (yellow) b. elevation (in m above/below sea level)

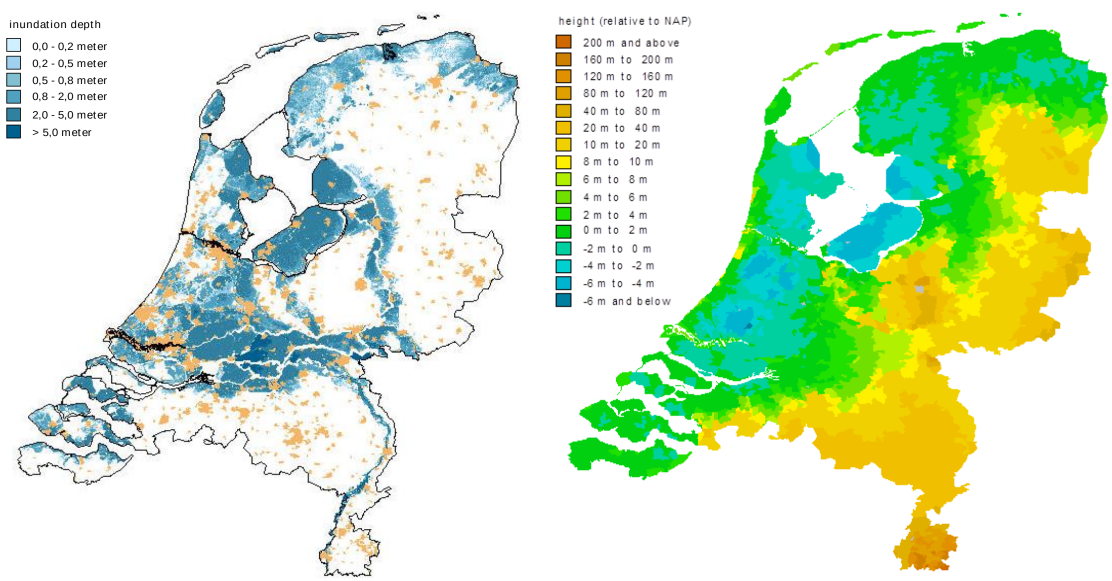

Notes: in Figure 1a, a location faces *flood risk* when it floods in case of a breach in the country's primary sea or river-defenses. Comparing Figures 1a and 1b shows that an exclusive focus on locations below sea level, would miss large parts of the Netherlands facing flood risk from the main rivers.

The Dutch have a long history of dealing with the flood risk posed by the sea and rivers that surround them. Today, the country is protected by arguably the best flood defenses in the world. The Dutch government spends over € 1.1 billion per year on the 3767 km of dikes and (artificial) dunes, and 1458 other primary waterworks (dams, weirs, locks, etc), that protect the country. Without these defenses 36% of the Netherlands floods (see Figure 1a), home to approximately 2.7 million houses/households. They are built to make

[2] A sea level rise of one meter will e.g. put an additional 1.2 million houses/households at risk.

floods an extremely rare event. In the best protected parts of the country they should reduce to likelihood of a flood to once every 10,000 years. [3]

This paper uses the unique Dutch setting to provide evidence on the price and perception of rare environmental disasters. We do this by establishing whether the extremely flood prone, yet also extremely well-protected, Dutch living environment leaves a mark on consumption decisions. Do they still reflect the rare risk of flooding? Or do the Dutch feel so well protected that flood risk plays no role? We look at the Dutch housing market in particular.

We consider the housing market to be one of the best places to look for evidence into how flood risk affects consumption choices. To most people their house is the most expensive purchase of their lifetime. Also, unlike other expensive property, it is hard to move your house to safer ground (or the 2nd floor) when a flood happens. A lost or severely damaged house therefore constitutes the largest monetary risk for people living in flood prone areas.[4]

On top of this, a unique institutional feature of the Netherlands is that people cannot privately insure their house against this flood risk. All private home insurance policies specifically exclude damage caused by flooding from their coverage. Instead, all flood protection in the Netherlands is provided publicly. The Dutch government is responsible for upholding the country's defenses. And, under the Calamaties and Compensation Act (de Vries, 1998), it promises to compensate people in case their house suffers flood related damages. However, how much and/or under which circumstances the government will compensate home owners is completely unclear.[5] It effectively means that the amount that each Dutch citizen pays for the country's flood defenses is

[3] The country's defenses have indeed prevented any major flooding ever since 1953 when the North Sea broke through the dikes and completely flooded the south western part of the country, killing over 1,800 people. In 1995 the Dutch were most prominently reminded of the risk posed by the river-delta they live in. Heavy snow- and rainfall in upstream areas had swelled the main rivers to levels not seen for many decades. 250,000 people were evacuated as some of the country's river defenses were on the verge of giving way. Fortunately, they held up and everyone returned home safely.

[4] A nice anecdote complementing our analysis comes from Dutch history. The Dutch historian Van der Woude (1972) calculated that houses in the Dutch village of Egmond were worth between 10 and 36 Dutch guilders in 1733. In 1755 they were worthless as a recent flood had moved up the coastline to such an extent that they were now located directly at the sea side. Despite the fact that these houses had not yet suffered from any actual flooding, people were apparently anticipating on what they expected to happen with these houses in case of the next flood.

[5] In 2003 there was a small-scale flood event in Wilnis, causing about €16 million of damage. Only about half of it was compensated. See Aerts and Botzen (2011) for more on this.

independent of the actual flood risk he/she and his/her property faces. The only way one can privately insure oneself against (the negative consequences of) flooding is to actually move to a house that, even if the country's defenses were to fail, would not be flooded.

The absence of private flood insurance helps the interpretation of the results from our hedonic house price analysis. In particular, no price discount for houses at risk of flooding would imply that the Dutch are fully satisfied with the existing public flood protection measures in place. By contrast a significant price discount would reveal that they do not believe that these existing flood protection measures fully protect them from (the negative consequences of) flooding. Moreover, the size of the price discount would give us an indication of people's flood risk perception, as well as their marginal willingness to pay (MWTP) for additional private or public insurance (over and above the insurance already provided publicly).

The difficulty with identifying people's MWTP to avoid flood risk using a hedonic house price approach is that it is in principle hard to compare otherwise equal houses, one at risk and one not at risk of flooding, using non-experimental data. Our identification of the effect of flood risk on house prices crucially relies on the plausibility of the specific exogeneity assumption(s) that we make on our flood risk indicator. The most difficult problem to tackle is that houses in flood and non-flood risk areas may differ for many other reasons. For example, different types of houses might be built in flood-prone areas, or people with different characteristics may sort into living in flood-prone areas. Also, flood prone areas may offer better recreation possibilities or nicer views related to the very same water that threatens to flood them. If these differences are not adequately controlled for, they could easily, and wrongly so, be picked up by our flood-risk indicator.

To make our identification assumptions as plausible as possible, we heavily rely on our unique, spatially very detailed dataset. It covers all Dutch 459,279 six digit postal code (6PPC) areas over the period 1999-2011 (a median 6PPC area is only 60 by 60 meters, containing 20 houses). It contains information on house prices, flood risk, and a wide range of other variables possibly influencing house prices. Besides controlling for all these variables in our regressions, the spatial detail of our data allows us to control for almost all other unobserved house prices determinants. We do this in two different ways.

First, we identify our flood risk effect using only variation in house prices and flood risk within five digit postal code (5PPC) areas. Hereby we control for any unobserved, possibly time-varying, house price determinants that are specific to areas defined by the first five digits of their postal code. These areas have a median size of only 294 by 294 meters, containing on average 17 different 6PPC areas. Given this extreme spatial detail, our focus on within 5PPC variation controls for almost all unobserved neighbour(hood) characteristics possibly related to house prices (e.g. accessibility of parks, restaurants, jobs, or schools, air quality, or type of neighbors).

Second, following e.g. Black (1999), Bayer et al. (2007), or Gibbons et al. (2013), we adopt a border discontinuity design (BDD) and restrict our sample to only those 6PPC areas that are located within a pre-specified distance to what we call “the flood line”. That is, we only consider 6PPC areas that do not flood and are located within 100 meters from a 6PPC area that does flood, and vice versa[6]. By only considering houses within the same, very small, distance band around the flood line, we further increase the likelihood of comparing homes that are, apart from their flood risk, equal in all other respects.

In our preferred specification, using both the above-described strategies to control for unobserved heterogeneity, we find a one percent price-discount on houses facing flood risk. It implies that the average Dutch household is not fully confident in the current publicly provided flood protection measures, and still prices in flood risk when deciding what to pay for their house. In fact, our estimate implies that people’s perceived likelihood of a flood happening, conditional on the country’s defenses, is much higher than what the government claims these defenses are built for. The median household living at risk of flooding would be willing to pay an additional €69 per year to be fully insured (publicly or privately) against future floods; 18% more than the current €400 per year per house at risk of flooding that the Dutch government spends on flood protection.

In extensions to these baseline findings we further explore whether our estimated *average* MWTP to avoid flood risk hides substantial heterogeneity related to people’s flood awareness. We find that people’s willingness to pay to avoid flood risk only depends systematically on the actual extent of the flood risk facing their property: the

[6] We use 100m distance band in all our baseline results, but also show results using a narrower or wider band around the flood line.

higher the expected water level in case of a flood, the more people are willing to pay to avoid this flood risk. We find no evidence that MWTP to avoid flood risk depends on the visibility of water in the neigborhood, nor on the flood protection levels at which the Dutch government claims to uphold the local flood defenses. Besides looking at these awareness measures, we also provide some evidence of income-based sorting into flood prone areas, and probe into the relevance of sorting based on heterogeneity in people's perceived flood risk and/or aversion to flood risk.

Our paper speaks to two different strands of literature. First, previous papers (e.g Bin and Landry, 2013; MacDonald et al., 1990; Bin et al., 2008) have also looked at the relationship between house prices and flood risk. Our paper's main contribution here lies in the quality and extreme spatial detail of our dataset that entirely covers one of the most flood prone, yet also best protected, countries in the world. It allows us to provide quasi-experimental evidence on the effect of flood risk on house prices. In doing so, we relate to other papers that have used quasi-experimental research designs to identify the hedonic price of other housing attributes such as e.g. schooling (Black, 1999; Gibbons et al., 2013), hazardous waste sites (Greenstone and Gallagher, 2008), or air quality (Chay and Greenstone, 2005). Second, our paper relates to the recent literature on rare disasters. Barro (2009; 2014), Barro and Ursua (2012), Gabaix (2012) or Weitzman (2009) e.g. model the effect of rare disasters on asset prices (typically focussing on macroeconomic shocks, or catastrophic environmental change). In this paper we provide micro-evidence on the effect of one particular rare natural disaster, flooding, on the price of one particular asset, houses.

2. Identifying marginal willingness to pay to avoid flood risk using hedonic analysis.

In the Netherlands, no explicit market exists for avoiding the negative consequences of flood risk. You cannot buy flood insurance as an individual product. The only protection provided is that offered by the country's flood defenses that are upheld with public money. Moreover the government promises to compensate (an unspecified part of) any flood damage incurred were the defenses to fail. This makes it hard to directly identify people's willingness to pay to avoid flood risk[7].

In this paper, we use house prices to infer people's MWTP to avoid flood risk. However, when buying a house, being completely free from any flood risk is not the only attribute of the house determining its price. Hedonic analysis puts a price on each of the many different, yet inseparable, attributes, of a house based on the association of the overall price paid for the house and each of its attributes. In this section, we briefly set out the main idea behind the hedonic method. We focus specifically on how we incorporate flood risk into the standard hedonic model. We outline the assumptions under which we can identify the average MWTP to avoid flood risk, people's flood risk perception, and willingness to pay for private flood insurance on top of the public insurance provided by the Dutch government.

We follow the approach pioneered by Rosen (1974) and Freeman (1974)[8], and extend it in a way similar to Brookshire et al. (1985) to incorporate the uncertainty that comes with owning a house in a place facing the (rare) risk of a natural flood disaster.

Suppose, each house in the Netherlands is defined by a vector of n different attributes $\boldsymbol{H} = (h_1, h_2, \ldots, h_n)$. Some are directly related to the house itself, e.g. m^2 of floor space, number of rooms, size of the garden, etc. Others to neighborhood characteristics, e.g. quality of the nearby schools, type of neighbors, distance to parks, air quality, or other (dis)amenities. Besides these n different attributes, each house is also characterized by the flood risk it faces, s, where s denotes the centimetres of water that would enter the house in case the country's defenses fail. With probability $\rho(\tau,\mathrm{F}) \in [0,1)$ these defenses fail and a flood happens, and with probability $1 - \rho(\tau,\mathrm{F})$ no flood happens. $\rho(\tau,\mathrm{F})$ depends negatively on the share of total income, τ, that the government collects as taxes to uphold the country's flood defenses: $\partial\rho(\tau,\mathrm{F})/\partial\tau < 0$, and positively on the likelihood of a flood in the absence of any flood defenses, $\mathrm{F} = \rho(0,\mathrm{F})$: $\partial\rho(\tau,\mathrm{F})/\partial\mathrm{F} > 0$.

Flood risk effectively imposes uncertainty about people's future utility derived from housing[9]. In case a flood happens, a house located on unsafe ground ($s > 0$) is

[7] Otherwise we could look directly at flood insurance premia or at flood insurance take up rates. See e.g. Gallagher (2014), MacDonald et al. (1990) or Bin et al. (2008).

[8] See also Chay and Greenstone (2005), Greenstone and Gallagher (2008) or Heckman et al. (2010) for a more in depth discussion of a hedonic model.

damaged. The extent of this damage depends positively on how far (in cm) the house would be below the water line in case the country's defenses fail, s,[10] and negatively on the amount of damage compensation received from the government, which we again take to be dependent on the tax rate, τ. In particular, only a share $b_i(s,\tau) \in [0,1]$ of each housing attribute h_i, $i \in \{1,\ldots,n\}$ survives the flood, where: $\partial b_i(s,\tau)/\partial s < 0$ and $\partial b_i(s,\tau)/\partial\tau > 0\ \forall i$. In case the house is located on safe ground, it is not damaged even if a flood happens: $b_i(s = 0,\tau) = 1\ \ \forall i$. It means that moving to a house on safe ground is the only way to be fully free from any uncertainty regarding the utility derived from the house. Consistent with the situation in the Netherlands there is no possibility to take out private insurance that covers any potential future flood damages.

All of the house's attributes together determine its price:

$$P = P(\boldsymbol{H},s) \qquad (1)$$

The negative of the partial derivative of $P(\boldsymbol{H},s)$ with respect to the amount of water entering the house in case of a flood, $-\partial P(\boldsymbol{H},s)/\partial s$, then gives the marginal implicit price that people are willing to pay for an additional bit of flood safety.

In the hedonic model, (1), often called the hedonic price schedule (HPS), is determined by the (partial) equilibrium interactions between consumers and producers in a competitive housing market.

Each consumer's expected utility is the sum of their utility in case of a flood and their utility when no flood happens, each weighted by their respective likelihood:

$$V^k = \rho^k(\tau,\text{F})\, U^k(X,\boldsymbol{b}^k(s,\tau)\boldsymbol{H},\alpha^k) + (1-\rho^k(\tau,\text{F}))\, U^k(X,\boldsymbol{H},\alpha^k) \qquad (2)$$

, where $k \in \{1,\ldots,K\}$, and K is the total number of consumers. Note that we allow each consumer to have his/her own belief regarding both the likelihood of a flood $\rho^k(\tau,\text{F})$ as well as about the damage incurred to each housing attribute, $\mathbf{b}^k(s,\tau) = (b_1^k(s,\tau), b_2^k(s,\tau),\ldots,$

[9] Note that this does mean that we, as most hedonic house price model do, view a house purely as a consumption good. See Bayer et al. (NBER, 2011) for a recent model that considers houses as both a consumption as well as an investment good.

[10] The relationship between expected damage and the amount of water flowing into the house is not perfect into the house. It also depends on the time that the water is present in the house. This depends, among others, on the drainage properties of the soil, the spatial extent of the flood, as well as the time it takes to fix the breach in the flood defenses.

$b_n^k(s,\tau)$). In both the flood and no-flood state of the world, utility depends on the consumption of a vector of housing attributes, $\boldsymbol{H}$, as well as that of a numeraire good X with price normalized to 1. We assume that the utility derived from the other good, X, is the same in the flood or no-flood state of the world. This effectively means that this good does not incur any damages in case of a flood. Moreover it means that utility derived from good X is additively separable from that derived from housing: $U^k(X,\boldsymbol{H},\alpha^k) = U^k{}_X(X,\alpha^k) + U^k{}_H(\boldsymbol{H},\alpha^k)$. Do also note that we assume that each consumer's preference parameters, α^k, as well as the shape of his/her utility function does not differ in the flood and no-flood state of the world.[11] Finally, without any loss of generality, it holds that for each h_i: $\partial U^k{}_H(\boldsymbol{H},\alpha^k)/\partial h_i > 0$. For housing attributes negatively affecting utility it means that h_i measures the absence of these attributes, e.g. safety instead of crime levels, or air quality instead of pollution. Similarly $\partial U^k{}_X(X,\alpha^k)/\partial X > 0$.

Consumers maximize their expected utility subject to their individual budget constraint, $Y^k = P(\boldsymbol{H},s) + X + \tau Y^k$, where Y^k denotes consumer k's income, and each consumer pays a (nondistortionary) income tax, i.e. a share τ of his/her income, to the government for upholding the country's flood defenses and/or for a national flood damage compensation fund[12]. Note that, as is the case in the Netherlands, all consumers pay this tax, regardless of whether they actually live in a flood prone area or not.[13]

The first order conditions (FOCs) corresponding to each consumer's expected utility maximization determine what levels of h_i, X, and s, he/she chooses to consume. At this preferred point of consumption, the FOCs provide an expression for each consumer's marginal willingness to pay for an additional cm of flood safety, s:

[11] All these assumption our made for easy of exposition, and could be relaxed. Doing so does not add any substantial new insights, which is why we opted for the more parsimonious exposition. Also, it would be straightforward to add a term to (2) that captures any non-damage related disutility from the mere inconvenience experiencing a flood.

[12] We take this income as exogenous. It would be possible to make this income endogenous depending on e.g. hours worked and wages. It could even be done in a way to allow a potential flood to affect firms' production (and thus wages paid). This would however unnecessarily complicate things as people often do not work and live in the same location, each with a different degree of flood risk.

[13] There are differences in taxes paid between the 24 different so-called "Water boards" that are responsible for keeping up the flood defenses in their region, as well as safeguarding the quality and quantity of drinking water. However, each household within the same Water board pays the same tax, regardless of whether their house would actually flood in case the defenses fail or not.

$$p^k_s = -\partial P(H,s)/\partial s = -\rho^k(\tau,\text{Ғ}) \sum_{\mathrm{i}} \left[U^k_{Hi} h_i \frac{\partial b^k_i(s,\tau)/\partial s}{b^k_i(s,\tau)} \right] / U^k_{XX} \qquad (3)$$

where $U^k_{XX} = \partial U^k_X(X,\alpha^k)/\partial X$ and $U^k_{Hi} = \partial U^k_H(\boldsymbol{b}^k(s,\tau)\boldsymbol{H},\alpha^k)/\partial h_i$. His/her marginal willingness to pay for an additional bit of flood safety equals the expected marginal utility gain from moving to a house on safer ground in terms of the foregone marginal utility of X on which the consumer could have otherwise spent his/her (after tax) income. This expected marginal utility gain is the sum of the marginal utility gain derived from losing a lower share of each and every housing attribute h_i in case a flood happens, multiplied by the consumer's perceived likelihood of a flood happening $\rho^k(\tau,\text{Ғ})$.

Expression (3) clearly shows that $p^k_s > 0$, unless people perceive the likelihood of a flood to be zero ($\rho^k(\tau,\text{Ғ}) = 0$), or do not expect any difference in damage between their own house and an identical house with a lower degree of flood safety ($\partial b_i^k(s,\tau)/\partial s = 0$, $\forall i$). In the Dutch context this means that they are either fully confident that the country's defenses will keep them safe, and/or they expect the government to be able to fully compensate them for any flood related damage in case the defenses do fail.

By substituting the budget constraint in (3), and then inverting this equation, we can obtain each consumer's bid function for s:

$$B^k_s = B_s((1-\tau)Y^k - P(\boldsymbol{H}^*,s),\boldsymbol{H}^*,V^{k*}) \qquad (4)$$

, where V^{k*} is the highest level of expected utility attainable given the consumer's budget constraint and $\boldsymbol{H}^*$ the optimal quantities of the other housing attributes. Holding all other housing attributes constant except s, it reveals the maximum amount that he/she is willing to pay for different values of s, while holding utility constant. (4) clearly shows that this differs between consumers as a result of differences in preferences, income, and/or flood risk and flood damage perception.

On the supply side, suppliers of housing can also differ in their cost of supplying a unit of housing with attributes $\boldsymbol{H}$ and in an area facing s cm of water in case the Dutch defenses fail. Each supplier maximizes his/her profit, so that he/she provides houses to the market in locations with s cm of flood risk up to the point where his/her marginal willingness to sell the house equals his/her marginal cost of providing a house with s cm of flood risk: $\partial P(\boldsymbol{H},s)/\partial s = \partial C(\boldsymbol{H},s)^r/\partial s$, where $C(\boldsymbol{H},s)^r$ is supplier r's cost function. By inverting each supplier's profit function we can also derive his/her offer curve for s:

$$O^r_s = O_s(s, \boldsymbol{H}^*, \pi^{r*}) \qquad (5)$$

, where π^{r*} is the maximum profit available to producer r given its cost function and the HPS (1). Holding all other housing attributes constant at their optimal levels, $\boldsymbol{H}^*$, it reveals the minimum amount of s that he/she is willing to supply to the market while holding profits constant.

Figure 2: Equilibrium hedonic price schedule in the hedonic market for flood risk

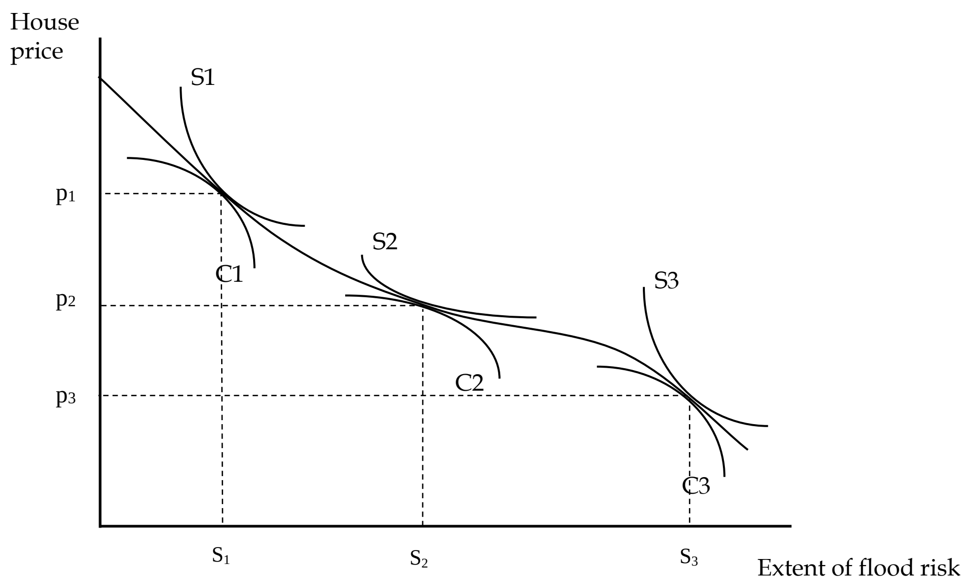

Consumers' bid functions and suppliers' offer functions together determine the equilibrium HPS for each particular housing attribute. At each point on this HPS an individual consumer's bid function is tangent to an individual supplier's offer function.

Figure 2 illustrates the equilibrium HPS in the hedonic market for the extent of flood risk, s. It also shows the offer curves of three different suppliers and three different consumers. As depicted, each bid function shows the standard increasing marginal rate of substitution between X and each housing attribute h_i (the loss of each h_i is increasing in s). Similarly, each offer function reflects the increasing marginal cost of providing houses with lower levels of flood risk. The gradient of the HPS with respect to the extent of flood risk reflects the equilibrium differential that allocates consumers across the locations. It compensates consumers for living in areas with a higher risk of flooding, and, similarly, it compensates suppliers for the cost of supplying houses facing less flood risk. House prices in places facing a higher risk of flooding must be lower in order to attract potential buyers.

2.1 MWTP to avoid flood risk, implied MWTP for additional flood insurance, and flood risk perception.

In our baseline regressions we identify the average MWTP to avoid flood risk as the log difference in house price between two otherwise equal homes, one facing flood risk, the other not. That is, α_{F} in the following hedonic house price regression:

$$ln\ P_i = -\alpha_{\mathrm{F}}\ I_i[s_i > 0] + \varepsilon_i \qquad (6)$$

, where $I_i[s_i>0]$ is a dummy variable taking the value 1 if the house floods if the Dutch flood defenses fail, and 0 otherwise. ε_i contains any other remaining unobserved components of house prices that our empirical design does not take account of[14].

Assuming that flood risk perceptions can be summarized at the 6PPC level, it follows from (3) and (6) that, in the absence of taste-based sorting our estimate identifies the following:

$$\alpha_{\mathrm{F}} \approx -\frac{1}{K}\sum_k \left[\rho^k(\tau, \mathrm{F}) \frac{\sum_{\mathrm{i}} \left[U_{Hi}^k h_i (b_i^k(s,\tau) - 1) \right] / U_{XX}^k}{P(H^k, s^k = 0)} \right] \qquad (7)$$

where K is the size of the total Dutch population and the approximation comes from the fact that we include a "flood or no flood risk" dummy variable in our regression instead of an actual measure of the cm of water entering the house in case of flood, s (in extensions we do further specify our dummy into five different flood risk categories, each with increasing s). With our estimate of α_{F} in hand we answer the following questions that lie at the heart of our paper:

[14] See the next section for much more detail on the exact empirical specification that we estimate, and the empirical design that we use to obtain a consistent estimate of α_{F}.

1. *Do the Dutch feel so safe behind their publicly provided dikes that house prices no longer reflect flood risk?*

The answer to this question can be directly inferred from the significance of our estimate of the average MWTP to avoid flood risk. An insignificant estimate would imply that the Dutch are fully confident in the publicly provided protection against flooding.[15] (7) shows that such a finding implies that they either fully trust in the country's defenses so that their perceived likelihood of a flood approaches zero: $\rho^k(\tau,\text{F}) = 0$ for all k[16], and/or they are confident that the government will be able to fully compensate any damages they incur if a flood does happen: $b_i^k(s,\tau) = 1$ for all k,i. By contrast, if we find a positive and significant MWTP to avoid flood risk, this immediately implies that people are both not fully confident in the country's defenses, as well as in the government's ability to fully compensate them for any incurred flood damage.[17]

2. *What is the average flood risk perception in the Netherlands?*

(7) shows that our estimate of average MWTP to avoid flood risk could in principle also be directly used to infer people's flood risk perception conditional on the existing defenses in place. Our identified average MWTP does however also depend on the uncompensated damage people expect to their house in case of a flood, as well as on people's exact preferences for each and every housing attribute and the outside good. Generally, it is impossible to separate people's flood risk perception from these other factors influencing average MWTP.[18]

However, under some (strong) assumptions on consumer preferences, and assuming that each consumer expects that a flood would destroy the same share of each and every housing attribute ($b_i^k(s,\tau) = b^k(s,\tau)$), our estimated α_F is a direct measure of the the average future (uncompensated) flood damage people expect to their house:

[15] Another explanation for such a finding could be that people are unaware of the flood risk that they and their property face, i.e. they are unaware that for their house s > 0 and wrongly believe that s = 0.
[16] Note that this does not depend on people actually having the correct perception about the likelihood of a flood without any defenses. It holds as long as they believe that the public defenses in place fully reduce this likelihood to zero.
[17] Finding a negative and significant MWTP to avoid flood risk is of course also possible, but this would be hard to explain.
[18] Note that this becomes even more complicated when specifying a fully dynamic version of our simple static hedonic model, see e.g. equation (10) below. In that case our estimated MWTP e.g. also depends on people's discount rate, and their expectations of future income and prices of the outside good.

$$\alpha_F \approx \; - \Sigma_k[\rho^k(\tau,F)\,(b^k(s^k,\tau)\text{-}1)] \,/\, K \qquad (8)$$

In particular this requires us to make the, likely invalid, assumption that consumer preferences are homogenous and linear in each and every housing attribute.[19] Using (8) we can even put a number on the average perceived likelihood of a flood for which all people expect the exact same uncompensated damage to their house (i.e. by further assuming that $b^k(s^k,\tau) = b$ for all k). For example, people's average perceived likelihood of an all-destructive flood ($b^k(s^k,\tau) = 0$ for all k) would be directly identified by our estimated α_F. Such a number can be compared to the officially stated flood risk levels that the country's defenses are, by law, supposed to keep up (e.g. the defenses of the economic heartland of the Netherlands should fail only once every 10,000 years). Also, we can compare it to people's perceived flood risk reported in earlier survey evidence.

3. *How does the average MWTP for flood protection relate to the public money already spent on flood protection?*

Our estimate directly reveals average MWTP to avoid flood risk as the average percentage people are willing to pay more for an otherwise equal house without any flood risk. The Dutch government currently spends an average €400 per year on each of the 2.8 million houses in the Netherlands facing flood risk. In order to compare this yearly number to our estimated average MWTP, we adopt our static hedonic model in the simplest way possible. Assume that now each, infinitely lived, consumer maximizes expected utility over his/her lifetime:

$$W^k = \Sigma^{\infty}_{m=0}\,[(\delta^k)^m V^k_{t+m}] \qquad (9)$$

, where $0 < \delta^k < 1$ reflects the rate at with each consumer discounts utility in future periods, and V^k_t specified as in (2) with a subscript t added to X only[20]. At the start of

[19] Note that a similar assumption is not uncommon in other hedonic studies that aim to gauge the the welfare effects of non-marginal changes in their housing attribute of interest, although they typically only need to assume that preferences are homogenous and linear in their housing attribute of interest only (see e.g. Chay and Greenstone, 2005).

[20] This means that we assume that people's flood perception, the damage to each and every housing attribute that they expect in case of a flood, as well as their preferences do not change over time. Also, no subscript t is added to **H** as people buy their house at the start of their lifetime and live there throughout their entire life.

his/her lifetime, each consumer buys one house that he/she will occupy for the rest of his/her lifetime, paying $rP(\boldsymbol{H},s)$ as mortgage payment in each period[21]. The rest of his/her after tax income, $(1\text{-}\tau)Y_t$, is spent on the outside good X so that the budget constraint in each period looks like: $(1\text{-}\tau)Y_t = rP(\boldsymbol{H},s) + P_{xt}X_t$. For simplicity we further normalize the price of the outside good to 1 in each period and also assume that income is the same in each period.[22] The FOCs of the consumer's lifetime expected utility maximization problem subject to his/her budget constraint in each period now give the following *yearly* MWTP for an additional bit of flood safety:

$$p^k{}_s = -\partial P(H,s)/\partial s = -\rho^k(\tau,\text{F})\sum_{\text{i}}\left[U^k_{Hi}h_i\,\frac{\partial b^k_i(s,\tau)/\partial s}{b^k_i(s,\tau)}\right]/\,rU^k_{XX} \qquad (10)$$

Combining (10) and (7) we can thus calculate the implied yearly amount a household living in the median house at risk of flooding is willing to pay (over and above the taxes paid for the public protection already provided) to be fully insured (privately or publicly) against flood risk:

$$P_{ins} = \alpha_{\text{F}}\, m[P(\boldsymbol{H},s>0)]\,/\,\text{r} \qquad (11)$$

, where $m[P(\boldsymbol{H},s>0)]$ is the median price of a house at risk of flooding in the Netherlands and r is the interest paid on an interest only mortgage.

The resulting P_{ins} can be directly compared to the €400 publicly spent on protecting each house at risk, or alternatively to the yearly amount people currently pay for their private home insurance or contents insurance (both of which explicitly exclude flood damage from their coverage). Do note that (11) shows that these calculations require us to make an assumption on r.

[21] Note that most people in the Netherland actually took out such an interest only mortgage over our sample period. Up until 2013, the tax system in the Netherlands encouraged people to take out such a mortgage since all mortgage interest payments were fully tax deductible.

[22] We could relax these two assumptions. They do not affect our conclusion regarding MWTP for additional flood safety. It would only mean that this MWTP depends additionally on people's time preference parameter δ^k, and on their expectations of both their income and of the price of the outside good in each future period. Making these simplifying assumptions allows us to write MWTP in a way that is almost (up to $1/r$) equivalent to that in our baseline static hedonic model, see (3).

4. *How much housing wealth will be lost due to future climate change?*

Our estimated average MWTP also allows us to provide a tentative estimate of the impact of the expected rise in sea levels due to future climate change on housing wealth in the Netherlands. Based on predictions of, among others, the Royal Dutch Meteorological Institute's (KNMI) for the centimeters rise in sea levels by 2100 (van den Hurk et al., 2006), we first obtain the number and price of the currently safe houses that a future rise in sea levels would put at risk of flooding in a best-, medium- and worst-case climate change scenario. Under the very strong assumption of an unchanging MWTP to avoid flood risk, and assuming that the state of the Dutch flood defenses is unaffected by this climate change, we can then easily get an estimate of the loss of housing wealth due to rising sea levels by multiplying the value of these houses by our estimated MWTP to avoid flood risk, α_F.[23] Interestingly this loss of housing wealth would be realized without a single drop of water coming over the dikes. It purely arises by increasing the number of houses at risk of flooding as well as increasing the expected damage to the houses currently already at risk of flooding.

3. Empirical identification strategy

3.1 Identification issues when estimating MWTP for flood protection using the hedonic approach

Answering our main research questions requires a consistent estimate of the average MWTP to avoid flood risk in the Netherlands[24]. Consistent estimation of the HPS for

---

[23] Note that an ideal measure of welfare change would also take into account how consumers and suppliers of housing respond to the change in the extent of flood risk induced by these rising sea levels by moving or changing the amount or quality of housing they supply. Most likely these changes will mitigate the welfare losses that we report here. In order to put a number on these changes one would need to have consistent estimates of each consumer's offer and each supplier's bid function, something that (see also footnote 25) lies beyond the scope of this paper.

[24] In some sense this aim is less ambitious than trying to estimate individual consumers' bid functions, primitive preference parameters, or recovering the entire MWTP function (see e.g. Rosen, 1974 for one of the first attempts to do this). Being able to do this would be very helpful as it would allow one to assess the welfare effects of nonmarginal changes in flood protection. However, doing this has met with empirical challenges that we cannot credibly claim to be able to address using our data (see e.g. Brown and Rosen, 1982; Epple (1987); Bartik (1987); Deacon et al., (1998); or for more recent contributions Ekeland, et al., 2004, or Heckman et al., 2010). However, as pointed out in Chay and Greenstone (2005), consistent estimation of the HPS for a particular housing attribute is as important from a practical perspective. Not being able to obtain a consistent estimate of the HPS invalidates any subsequent welfare analysis because it will lead to an inconsistent estimate of the MWTP function; irrespective of the subsequent method used to

flood risk (or any housing attribute) is however extremely difficult (see also Gibbons et al., 2013; Black, 1999; Chay and Greenstone, 2005; or Greenstone and Gallagher, 2008 for insightful discussions on this). Two reasons for this stand out.

The first is misspecification of the HPS in (1). The most prominent reason for this is unobserved determinants of house prices that are correlated with flood risk[25]. For example, areas that have a high risk of flooding may also offer better opportunities for water recreation, or nicer views. Also, the west of the Netherlands is both most at risk of flooding, as well as the country's heavily urbanized economic heartland. As such, estimates of the HPS for flood risk may be biased because of omitted variables.

Taste- and/or income based sorting is the second reason making it difficult to infer average MWTP for a particular housing attribute. As shown in (4), consumers' bid functions may differ when they have different preferences, different incomes, and or different flood risk or flood damage perceptions. As a result different individuals may have a different MWTP to avoid flood risk, so that e.g. individuals having a higher flood risk perception and/or higher incomes will sort into areas without any flood risk. Income-based sorting makes it more difficult to separate MWTP to avoid flood risk from MWTP for wealthier neighbors (or characteristics of these neighbors correlated with income). If people sort based on their tastes and/or perception of flood risk, it would mean that we are no longer sure to identify the average MWTP of the entire Dutch population. Instead, our estimate would represent the average MWTP of some non-random subpopulation of the overall Dutch population only.

3.2 Our data and empirical design

In light of the above-raised identification issues, we adopt an empirical identification strategy that provides us with the best possible consistent estimate of the average Dutch MWTP to avoid flood risk. It relies heavily on the extreme spatial detail of our dataset.

---

recover the "deeper" preference or technology parameters determining individual consumers' bid functions or suppliers' cost functions.

[25] The other is an incorrect choice of functional form relating observed housing attributes to house prices. See Cropper et al. (1988) for a discussion.

3.2.1 Data

We use the most detailed and comprehensive dataset ever compiled on Dutch house prices and flood risk. It contains information on official recorded house prices, various measures of flood risk, and a host of other possible house price determinants for each of the 459.279 six-digit postal code areas (6PPC-area) in the Netherlands in each year covering the period 1999-2011. The median size of a 6PPC-area is only 60 x 60 meters, containing on average 20 housing units.

The use of aggregate 6PPC level data instead of considering individual houses may induce some biases in the presence of heterogeneity within 6PPC areas in house prices, flood risk and/or other house price determinants (see Chay and Greenstone (2005) or Greenstone and Gallagher (2008) that also use aggregated data when estimating a hedonic house price regression). However, we expect any bias, if present, to be small given that the variation in all our variables is much smaller within 6PPC areas than between them. Moreover, given that most median 6PPC house prices reported in our data set (about 70%) are actually based on the sale of a single house, this bias is mostly due to measurement error in flood risk (and/or other house price determinants). Unless this measurement error is negatively correlated with actual flood risk[26], our estimated average MWTP will be biased towards finding that the Dutch are satisfied with the current levels of public flood protection, i.e. towards finding $\alpha_F = 0$ in (7).

House prices

For each 6PPC-area we know the median house price of all houses sold in each year during the period 1999-2011. If no houses were sold in a particular year, we exclude the 6PPC in that year from the analysis[27]. These median house prices are based on all

[26] In the worst case, our aggregate measure of flood risk at the 6PPC level would hide that all the least expensive houses in a 6PPC area classified as facing flood risk are actually located in the safe part of the 6PPC, and, similarly, all the most expensive houses in a 6PPC area classified as not at risk actually being located in the flood prone parts of the 6PPC. Such systematic measurement error in our data is very unlikely.

[27] We discuss selection bias arising from this in detail when interpreting our findings. It is only an issue when this selection is based on an unobservable variable that is not controlled for in our empirical design that is, moreover, correlated with flood risk. We also show that there is no evidence that either the number of houses sold nor the likelihood of observing a sale is correlated with flood risk. Finally, we note that the average time that a house is on sale before being sold is less than a year. Given that we observe 13 years of

property transactions in the Netherlands as registered by Het Kadaster. Het Kadaster records the exact price of all property transactions in the Netherlands as stated in the official purchase agreement registered at the notary. Each price therefore represents the actual price paid and not the ask-price nor the sometimes different selling price claimed by the broker selling the house. For each transaction Het Kadaster also records the type of property involved in the transaction.[28] We exclude 6PPC-areas from the analysis where non-residential property was sold. The universal coverage of all property transactions in the Kadaster data does result in one minor issue. When people sell part of their residential property (e.g. a piece of their garden, a garage or a boathouse), this is recorded as a property transaction of the same type under which a transaction of the entire property would be listed. It is impossible to distinguish these transactions from the sale of the entire residential property. They result in some unrealistically low "house" prices in the data set. We deal with these observations by excluding those 6PPC areas from our analysis that report a median house price below the 5pct-quantile of all median 6PPC house prices in the same town that the 6PPC is part of.[29] In the end this leaves us with 1,274,629 observations of median 6PPC house prices over our sample period.

Flood risk

As our main flood risk measure we use the expected maximum amount of water that would flood each and every 6PPC area in case the Dutch primary river or sea defenses fail. This information is publicly available on the website *www.risicokaart.nl*, the main interactive website of the Dutch government where it informs citizens about many different types of risks that they and their property face (flood risk, environmental

---

data, it is therefore very likely that all houses that were put up for sale during this period [except for those in the latest year(s)] were (eventually) sold.

[28] It distinguishes six different categories of property: non-residential, apartments, detached houses, semi-detached houses (two houses under one roof), attached houses at the end of a housing block (attached to one neighboring house), and attached houses in the middle of a housing block (attached to two neighboring houses).

[29] Our results do not depend on excluding these observations, but the inclusion of these small scale property transactions introduces additional noise that makes our estimates less precise. We show results without using any correction for these transactions or using a different correction in Appendix A (using e.g. the 2.5pct quantile as cutoff, or the 5pct-quantile of all 6PPC median house prices in each of the 2345 districts ("wijken") into which the towns of the Netherlands can be divided). Our town and district definitions are those of 2008, there were only minor changes in these definitions over time. In 2008, there were 443 towns in the Netherlands.

hazards, storage and/or transportation of toxic materials, etc.). The extent of flood risk is classified in seven different categories that increase in the expected amount of water flowing into the 6PPC in case the country's primary defenses fail: no flood risk (0 cm), 0-20cm, 20-50cm, 50-80cm, 80cm-2m, 2-5m, >5m. These numbers are based on software developed by Deltares (e.g. Sobek or Delft3D), the leading Dutch institute for applied research in water management. It uses a host of surface characteristics (altitude, natural barriers, etc), in combination with sophisticated water flow modeling to provide detailed predictions on the maximum water levels in each 6PPC in case of a flood. Figure 1a already showed the distribution of flood risk over the country.

In our main regressions, we focus on a dummy variable constructed using this data that takes the value one if the 6PPC area floods in case the primary defenses fail (i.e. s > 0cm). The main reason for doing this is that there is more uncertainty about the exact amount of water entering the 6PPC than about whether or not a 6PPC area will flood in case the primary defenses fail. It depends e.g. on the size and type of the breach (see Klijn et al., 2010; or Asselman et al., 2009 for detailed discussion on this). Given the accuracy of Deltares' flood predictions, the possibility that our flood dummy misclassifies whether a 6PPC area floods or not is basically only present in the 3% of 6PPC areas that are expected to receive 0-20cm of water. We will use the more refined categories detailing the amount of water flowing into the 6PPC in extensions to our results. Our main flood dummy is directly related to (3). If it takes the value one, we know two things, $\digamma > 0$ and $s > 0$, i.e. the likelihood of a flood in the 6PPC area without any defenses in place is nonzero and water will damage the house in case a flood happens.

Besides this main variable, we have also collected information on the level at which the government claims to keep up the primary defenses. The Dutch flood defenses are organized in so-called dike-ring areas. Each dike-ring area has an official acceptable flood risk level up to which its defenses should, by law (Water Act, 2009), be kept up. Figure 3 below shows the variation in accepted flood risk in each of the dike-rings. Acceptable flood risk levels range from once per 10,000 years for the economic heartland in the west of the country, to once per 1,250 years along the main rivers in the more eastern parts of the country.

Figure 3. The Dutch defenses: accepted flood risk levels by dike-ring area

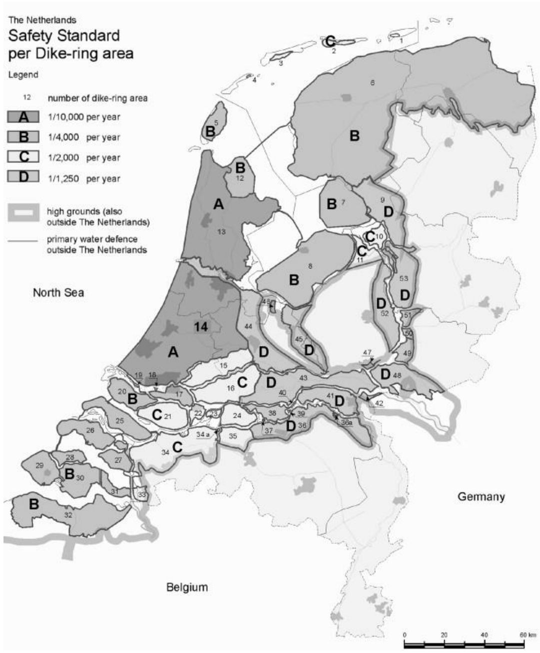

Notes: one dike-ring, nr.40 Heerewaarden has a lower safety standard of 1 flood per 500 years. Figure taken from Botzen et al, 2009.

This information can be directly related to people's perceived flood risk, $\rho(\tau,F)$. If people know, and trust, the government's claims about the upheld levels of flood protection paid for by their taxes (τ), we can expect our estimate of MWTP to be systematically related to these safety standards.

House and neighbor(hood) attributes

Besides flood risk, we have also collected a host of other variables for each 6PPC area possibly affecting house prices. They can be distinguished between house and neighborhood characteristics. Here we discuss the most important of these control variables (see Appendix B, for a complete list of all 139 variables).

The information on housing attributes comes from two different sources. From the Kadaster data we know the type of homes (see footnote 29) sold in the 6PPC. Moreover, from the centralized town administration of the Netherlands (Basis Administratie Gemeenten (BAG)), and the Rijksdienst voor Cultureel Erfgoed (RCE), we know the median floor space of the houses in each 6PPC, the year in which they were built, and the percentage of homes that qualify as a national monument. These latter two are important, as the year a house was built is correlated with its likelihood of facing flood risk. Historic houses built before the main defenses were in place are more often found on safe ground. They are also often desired for their aesthetic appeal. If not controlled for, this may therefore confound our estimate of MWTP to avoid flood risk.

The neighborhood characteristics come from a wide range of sources (see the Appendix for the full listing). From e.g. the BAG, the Dutch Central Bureau of Statistic (CBS) and the Integrale Veiligheidsmonitor Rijk (IVR), we have information on the availability of many different (dis)amenities in the neighborhood of each 6PPC. For example: distance to parks, restaurants, shops, scrapyards, airports, or main roads, but also littering, burglary etc. Importantly, we have information on the distance to fourteen different types of water bodies (sea, (main) rivers, lakes, estatuaries, open wetlands, etc). This group of variables is of particular importance to us as flood prone areas may also offer better access to water recreation, nicer views, etc, that, if not controlled for, may confound our estimates of people's MWTP to avoid flood risk. The same can be said of elevation. Places on higher ground tend to face a lower risk of flooding[30], but they may also provide people with nicer views. We take each 6PPC's average altitude from the Actueel Hoogtebestand Nederland (*www.ahn.nl*), where the Dutch government informs citizens about the exact altitude of their 6PPC. It is measured in meters Normaal Amsterdams Peil (NAP).

Finally, the Dutch Central Bureau of Statistics (CBS) provides us with information on the characteristics of households living in each 6PPC (e.g. age, income, % immigrants, % households with children). The coverage of this data is less complete then that of our other variables.[31] For our purpose, the information on household income is of particular importance as it allows us to say something about the relevance of income based sorting in our context. Also, we employ the information on the presence of

[30] Especially compared to the 25% of Dutch land below sea level. This correlation is however far from perfect. A substantial part of the flood risk in the Netherlands does not come from places below sea level, but from areas facing flood risk from the main rivers and lakes (compare Figure 1a to Figure 1b).

[31] This information is only available at the 6PPC level in two years of our sample (2004 and 2010; for income this is 2003 and 2008). Also, for privacy reasons, information is only available if the aggregate numbers reported can not be traced back to individual households (mostly this means that information is missing in 6PPC areas with less than 5 or 10 households). We fill in the missing numbers for years before 2004 using the 2004 numbers, for the years in between 2004 and 2010 we take the average of the reported 2004 and 2010 numbers, and from 2010 onwards we take the 2010 numbers.

households with children and on the number of foreign born inhabitants as proxies for individuals' tastes for flood safety to shed light on possible taste-based sorting.

3.2.2 Empirical design

We identify average MWTP to avoid flood risk using a simple hedonic regression of the form already specified in equation (6):

$$ln\ P_{ijt} = -\alpha_F\ I_i[s_i>0] + \boldsymbol{X}_{it}\boldsymbol{\beta} + \varepsilon_{ijt} \qquad (12)$$

, where, given that we do not observe individual house prices but the median house price per 6PPC, *i* from now stands for a 6PPC area. We add subscripts *t* denoting the year of observation, and *j* denoting the 5PPC area that the 6PPC area belongs to. $\boldsymbol{X}_{it}$'s are other determinants of house prices at the 6PPC level that possibly vary over the years, and $\boldsymbol{\beta}$ is a vector capturing the effect of these controls on house prices. $I_i[s>0]$ is our main indicator of flood risk. It is a time invariant dummy variable taking the value 1 if water flows into the 6PPC in case the Dutch primary defenses break. α_F is our main parameter of interest. If identified correctly, it represents the average Dutch MWTP to avoid flood risk.

Estimating (12), we obtain a consistent estimate of α_F only if the following condition holds:

$$\mathrm{E}[I_i[s_i>0]\ \varepsilon_{ijt} \mid X_{it}] = 0 \qquad (12)$$

That is, only if, conditional on the observed control variables included in our regression, a 6PPC's flood risk is uncorrelated to any unobserved determinants of house prices. Basically this means that we need to ensure that assignment of houses to flood risk areas is random conditional on the observable house price determinants we include in our model. Despite the fact that we can control for more than one hundred other house price determinants at the 6PPC level, we can never be sure that we have correctly controlled for all possible unobserved determinants of house prices also related to flood risk[32]. Possible income based sorting for example would make it difficult to distinguish MWTP

[32] Any misspecification in the additively linear way in which we control for the observed house price determinants in our regressions could also result in a violation of condition (12).

to avoid flood risk from MWTP for wealthy neighbours (or other characteristics of these neighbors correlated to their income).

To make our identification as credible as possible, we adopt an empirical design that is heavily based on the spatial detail at which our data is available. It consists of three different steps that aim to get us as close as possible to identifying our effect of interest off of quasi-experimental (random) variation in flood risk only. They are each aimed at making the houses in our sample as comparable as possible, also when it comes to their unobservable house and neighborhood characteristics.

First, we restrict our sample to 6PPC areas where only terraced houses, attached single family houses, were sold (*rijtjeshuizen* in Dutch). Terraced houses are the most often sold type of property in the Netherlands. Two randomly selected terraced houses have much more comparable housing characteristics than e.g. two detached houses, let alone a detached house and an apartment. For example, the standard deviation in m^2 median floor space of houses sold in 6PPC areas reduces from 44.2 m^2 to 28.9m^2 when considering only terraced houses[33]. Basically, this selection of terraced houses means that we use a sample of flood and non-flood 6PPC areas that are matched on having the same type of housing stock (see also Gibbons et al. (2013) or Black (1999) for a similar (implicit) matching on property type). The condition under which we consistently estimate average MWTP to avoid flood risk effectively changes from (13) to:

$$\mathrm{E}[I_i[s_i>0]\ \varepsilon_{ijt} \mid X_{it}, I_i[T_i = 1]] = 0 \qquad (14)$$

where T_i indicates whether or not the house sold was a terraced house. This focus on 6PPC areas with a housing stock consisting exclusively of terraced houses halves our sample to 619,605 observations[34]. This first step is primarily aimed at ensuring that we are comparing areas with a similar housing stock. It does not control for any other neighborhood characteristics unless they are perfectly correlated with the likelihood of a 6PPC's housing stock consisting solely of terraced houses.

[33] This reduction is even larger when considering median plot size. There the standard deviation reduces by more than a factor four when only considering 6PPC areas consisting of terraced houses.
[34] For comparison, the percentage of 6PPC areas with a housing stock consisting solely of detached houses, two-under-one-roof houses, or apartments is 11.9%, 10.1%, 16.5% respectively. The other 12.9% are 6PPC areas with a mixed housing stock.

This is what our second and third step aim for. They are the backbone of our empirical strategy. As a second step we include 5PPC-year fixed effects in (12):

$$ln\ P_{ijt} = -\alpha_{\mathrm{F}}\ I_i[s_i>0] + \boldsymbol{X}_{it}\boldsymbol{\beta} + \mu_{jt} + v_{it} \qquad (15)$$

So that we now get a consistent estimate of α_{F} under the assumption that

$$\mathrm{E}[I_i[s_i>0]\ v_{it}\ |\ X_{it}, I_i[T_i=1], \mu_{jt}] = 0 \qquad (16)$$

Basically, we identify α_{F} only off of variation in house prices and flood risk between 6PPC areas located in the same 5PPC area. The median size of a 5PPC-area is only 294 by 294 meters. It means that, when estimating (14), only very localized, 6PPC specific, unobservables (v_{it}) can confound our estimate of MWTP for flood risk, i.e. those varying within 5PPC areas. We argue that virtually all neighborhood characteristics, such as availability of restaurants, shops, or work, the quality of public service provision, or community taxes, and also, to a large extent, the types of neighbors, are implicitly controlled for by the inclusion of 5PPC-year fixed effects. Moreover, the terraced houses built in the same 5PPC area are very often identical, so that our focus on within-5PPC variation also controls for almost all unobserved characteristics of each house (nr. of rooms, size of windows, roof type, etc.). One could question whether we are left with enough observations providing us the necessary variation to estimate (15). In our baseline terraced house sample, 8% of all observations (47,308 in total) are in 5PPC areas exhibiting, within the year of observation, variation in both house prices and flood risk. Figure A1 in the Appendix illustrates this localized variation in flood risk in case of the city of Dordrecht and its immediate surroundings.

Finally, as a third step in our identification strategy, we adopt a BDD design[35] inspired by earlier contributions of among others Black (1999), Bayer et al. (2007), Ries and Somerville (2010), and Gibbons et al., (2013). In particular, we restrict our sample to

[35] Note that our BDD design is different from a true regression discontinuity (RD) design. The fact that different types of people (or houses) may sort into flood or non-flood prone areas invalidates the crucial assumption underlying any RD design that people have imperfect control over their assignment into treatment or not. See Lee and Lemieux (2010, p.347) for a detailed discussion of the issues arising from the use of geographic borders in an RD design. Our BDD design is meant to ensure that we are comparing as similar as possible homes, some of which flood and some do not (e.g. having the same neighborhood characteristics, being similar in type, size, and quality, etc). We discuss possible issues arising from sorting when interpreting our main findings.

only 6PPC areas that are located within 100 meters of what we call the *flood line*.[36] More specifically, we define 6PPC areas along the flood line as either flood prone 6PPC areas located within 100 meters from a 6PPC area without any flood risk, or vice versa. This focus on the flood line aims to further control for unobserved housing and neighborhood characteristics varying *within* each 5PPC area that are correlated with both flood risk and house price (variation that is not controlled for by the inclusion of 5PPC-year fixed effects). The condition under which we can claim to consistently estimate average MWTP to avoid flood risk, changes to (16) below because of our BDD design:

$$\mathrm{E}[I_i[s_i>0]\ v_{it} \mid X_{it}, I_i[T_i=1], \mu_{jt}, I_i[D_{FL,i} < 100m]] = 0 \qquad (17)$$

, where $D_{FL,i}$ denotes distance to the nearest 6PPC area with a different flood risk than i. The focus on the flood line does significantly reduce our sample, leaving us with 42,998 observations of 6PPC areas along the flood line. However, not surprisingly, half of these observations now provide us with the necessary within 5PPC-variation in house prices and flood risk.

The fact that we reduce our data set according to two different selection criteria, see (14) and (17), plus the fact that we identify our effect of interest relying only on within-5PPC variation in house prices and flood risk, does strictly speaking mean that we are no longer sure to be identifying average MWTP to avoid flood risk for the whole Dutch population of home-owners. Our baseline estimate provides us with the average MWTP to avoid flood risk for owners of a terraced house located along the above defined flood line. If people e.g. sort into either terraced houses or into areas along the flood line based on their flood risk and/or flood damage perceptions, this means that we are no longer able to infer MWTP for people owning different types of houses than a terraced house and/or living farther away from the flood line.

One issue that immediately comes to mind regarding our focus on areas along the flood line is the visibility of the flood risk. Are we not simply focusing on those areas in the Netherlands immediately adjacent to the water (rivers, lakes, sea) threatening to flood them? Effectively, it could mean that we select exactly that sample with the highest chance of finding an effect of flood risk on house prices. It turns out that the average

[36] Using a 100 meter cutoff is arbitrary, we show results using different cutoffs in our robustness section.

distance to the nearest large water body (definitely protected by a dike) is indeed smaller for our flood line sample, however the difference in terms of visibility of the water posing the flood risk is not that different[37]. The median [5th percentile] distance to this water is about 5.8km [250m] compared to 6.8km [570m] for houses not located along the flood line.

Figure 4. The flood line

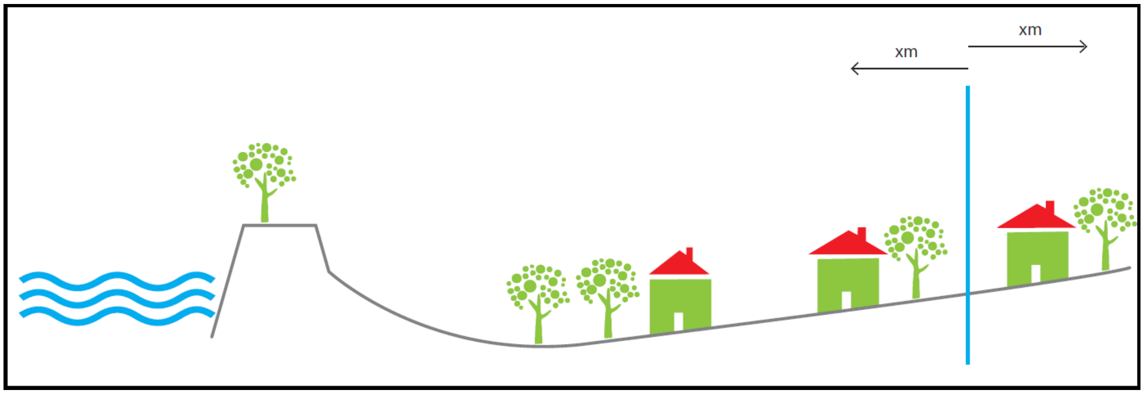

Notes: the extent of the flood line depends on the distance x (in meters) set to delimit it. In our baseline results x is set to 100.

Figure 4 illustrates where the flood line we focus on in our BDD is typically located relative to the water posing the flood risk. Typically, the dike protects an area that extends far into its hinterland. The area immediately surrounding that point where the water would reach in case of a flood, is what we call our flood line. Often this point is located far away from the water posing the risk.

In extensions to our baseline results we also shed light on the possibility of differences in flood awareness among people, as well as on the likelihood that our baseline findings may be affected by income and/or taste-based sorting. We do this in two different ways. First, using different variables that can be argued to affect flood awareness (visibility of the risk, official state of the defenses, flood depth) we extend (15) as follows:

$$ln\ P_{ijt} = -\alpha_{F}\ I_i[s_i>0] + (I_i[s_i>0]\boldsymbol{Z}_i')\ \boldsymbol{\alpha}_{\mathbf{AF}} + \boldsymbol{Z}_i\boldsymbol{\gamma} + \boldsymbol{X}_{it}\boldsymbol{\beta} + \mu_{jt} + v_{it} \qquad (18)$$

[37] In extensions to our baseline findings, we also explicitly verify whether MWTP varies among people depending on the visibility of the water that threatens to flood them.

, where $\boldsymbol{Z_i}$ denotes the variable(s) possibly affecting people's flood awareness. A significant estimate of $\boldsymbol{\alpha_{AF}}$ would be an indication that MWTP differs between people depending on the particular variables included in $\boldsymbol{Z_i}$.

Second, to probe into the likelihood of sorting, we adapt (15) by replacing our main dependent variable, ln house prices, by household characteristics on the basis of which one can expect sorting to occur. We use median household income when looking for evidence of income based sorting. Taste-based sorting is more difficult to identify, especially so since individual tastes for flood safety are generally unobserved.[38] In the absence of such data, we follow Greenstone and Gallagher (2008), and look for evidence of taste-based sorting using different proxies for households' risk preferences (families with (young) children, and being foreign born). If people do sort into flood safe areas based on one of these characteristics, they should be systematically related to flood risk.

4 Results

4.1 Descriptives

We start by showing some descriptive statistics of the key variables used in our study. Table 1 shows the mean and standard deviation of these variables for four different samples. From left to right, we focus on the *full sample* of all 459,279 6PPC areas (in the 13 years of our sample), 6PPC areas with at least one house was sold, 6PPC areas where only terraced houses were sold, and finally 6PPC areas in this terraced house sample located along our previously defined flood line.

Comparing column (1) and (2) shows no substantial differences in flood risk characteristics between 6PPC areas where no houses were sold and 6PPC areas where houses (or only terraced houses) were sold. They are very similar when it comes to the distribution of the extent of flood risk (cm of water), their elevation, distance to water, location below sea level, and the official state of their dike-ring defenses. The biggest

[38] Bayer et al. (2007) propose a solution to this issue when individual tastes for the housing attribute of interest vary with observable characteristics of these individuals (e.g their income, education levels, sex, etc). If tastes vary for other unobserved reasons, their method would however still not fully address it. Others have proposed a correlated random-coefficient model to deal with this issue (see e.g. Chay and Greenstone, 2005). Their approach relies on the availability of a credible instrument that, in our case, would be correlated with the extent of flood risk facing a 6PPC, but not to any other (unobserved) determinant of median house price in the 6PPC. We lack a credible candidate for such an instrument.

difference between 6PPC areas with and without any home sales is, not surprisingly, the size of the housing stock.

Table 1. Descriptive statistics

	(1) all 6PPC		(2) 6PPC any house(s) sold		(3) 6PPC terraced house(s) sold		(4) 6PPC < 100m flood line terraced house(s) sold	
	mean	s.d.	mean	s.d.	mean	s.d.	mean	s.d.
price	-	-	231283	147827	210723	100748	223828	116473
nr. sold	-	-	1.5	0.9	1.4	0.8	1.4	0.8
floor space (m2)	-	-	120	44	116	29	113	31
nr. houses	16	15	21	14	20	11	20	11
flood risk?	0.36	0.48	0.37	0.48	0.42	0.49	**0.58**	0.49
0-20cm	0.03	0.17	0.03	0.17	0.03	0.16	**0.16**	0.36
20-50cm	0.04	0.20	0.05	0.21	0.04	0.21	**0.14**	0.35
50-80cm	0.04	0.20	0.04	0.20	0.04	0.20	0.08	0.27
80cm-2m	0.14	0.34	0.14	0.35	0.17	0.38	0.15	0.36
2-5m	0.11	0.31	0.11	0.31	0.14	0.34	**0.06**	0.23
> 5m	0.003	0.05	0.002	0.05	0.003	0.05	0.002	0.04
NAP (m) - elevation	9.1	20.4	8.9	20.5	8.3	20.5	**0.8**	5.3
< 0m NAP?	0.26	0.44	0.26	0.44	0.32	0.47	**0.51**	0.50
dist. to water (m)	582	669	562	642	505	577	**284**	294
dist. to *large* water body (m)	9982	10129	9908	10080	9594	9573	**6698**	5762
off. state of defenses								
not in dike-ring	0.32	0.47	0.31	0.46	0.28	0.45	**0.02**	0.15
1/10000 year	0.27	0.45	0.28	0.45	0.28	0.45	**0.49**	0.50
1/4000 year	0.16	0.37	0.16	0.36	0.16	0.36	0.20	0.40
1/2000 year	0.07	0.25	0.08	0.27	0.09	0.29	0.08	0.27
1/1250 year	0.17	0.38	0.17	0.38	0.19	0.39	0.20	0.40
1/500 year	0.0001	0.01	0.0001	0.01	0.00003	0.01	-	-
distance to flood line (m)	5385	9448	5194	9213	4892	8975	67	19.7
observations	5970627		1274629		619605		42998	

The difference between 6PPC areas with only terraced house sales and those with any house sales are also small when it comes to their flood risk characteristics. 6PPC areas with only terraced house sales are slightly more often at risk of flooding, as well as located below sea level. The average price paid for a house in the terraced house sample is however about 10% lower than when considering all home sales, which can be explained by the fact that detached and “two-under-one-roof” homes are typically sold at much higher prices. Do note that the standard deviation in house prices, as well as in m^2 floor space, is also much lower in the terraced house sample. Both reflect the fact that terraced houses are a much more homogenous group of houses in terms of their housing characteristics than e.g. the detached homes or apartments included in the sample of all house sales.

When further restricting our sample to 6PPC areas along the flood line, house prices, m^2 floor space, and the number of houses sold are very similar to 6PPC areas located further away from this line. The flood risk characteristics of this sample are, not surprisingly, quite different from the other samples. Now about 58% (51%) of 6PPC areas are at risk of flooding (located below sea level), compared to around 40% (30%) in the other samples. However, the extent of the risk faced tends to be smaller around the flood line: flood prone 6PPC areas there are facing lower expected amounts of water in case of a flood than in the other samples. Also, virtually all 6PPC areas along the flood line fall under the protection of one of the dike-ring areas, compared to only 70% in the three country wide samples. Finally, both the average distance to any type of water as well as that to a large water body (definitely protected by a dike) is smaller for the flood line sample. Based on these descriptives, our focus on the flood line, while having the virtue of being much more sure to be comparing otherwise similar homes (see Table 2 for descriptive evidence on this), does raise the concern that we are identifying MWTP to avoid flood risk based on a part of the country where the flood risk that people face is more visible than in other parts of the country (see also our earlier discussion of this issue in section 3.2). In section 4.2 we show evidence that this selection issue appears to be of limited concern (see e.g. Figure 4).

Table 2 below shows important additional descriptives that are useful in assessing the validity of our empirical design. They show how well our empirical design manages to balance the house price determinants that we do observe across flood and non flood prone 6PPC areas. If we do not manage to balance these observables this would shed doubt on the ability of our design to do so for any possible remaining unobservable housing attributes. Moreover, if the observable housing attributes are balanced, consistent estimation of MWTP to avoid flood risk does not depend on the functional form assumptions that we make regarding the relation between these observable attributes and house prices.

Table 2. Evaluating our empirical design – differences in house and neighborhood characteristics (along the flood line)

	1	2	3	4	5	6	7	8
					within 5PPC variation only			
terraced houses	All		< 100m of floodline		all		< 100m of flood line	
Observations	357640	261965	17895	25103	357640	261965	17895	25103
	no flood risk	flood risk	no flood risk	flood risk	no flood risk	flood risk	no flood risk	flood risk
ln median house price	12.17	12.20***	12.25	12.23***	12.18	12.17***	12.24	12.23***
% houses built in:								
< 1549	0.04	0.01***	0.10	0.04**	0.04	0.01*	0.1	0.1
1550 - 1749	0.2	0.1***	0.7	0.4***	0.2	0.02**	0.6	0.5*
1750 - 1849	0.3	0.2***	1.1	0.6***	0.4	0.02***	0.9	0.7***
1850 - 1879	0.3	0.3	1.1	0.7***	0.4	0.3	0.9	0.9
1880 - 1899	1.1	0.9***	1.4	1.7	1.1	0.8***	1.6	1.6
1900 - 1909	2.1	1.6***	3.3	2.9	2.0	1.7	3.3	3.0
1910 - 1919	2.2	1.5***	2.4	2.6	1.7	2.1**	2.1	2.8***
1920 - 1929	4.6	2.8***	4.9	5.4	4.2	3.5***	5.3	5.1
1930 - 1939	7.4	5.0***	5.6	5.8	6.4	6.3	5.5	5.9
1940 - 1949	1.8	0.9***	1.8	1.5	1.5	1.4	1.6	1.7
1950 - 1959	7.5	4.4***	5.0	5.6	6.0	6.5*	5.2	5.4
1960 - 1969	13.8	12.0***	10.4	12.2***	12.5	13.7***	11.4	11.4
1970 - 1979	23.1	24.5***	19.8	19.8	23.1	24.5***	19.7	19.9
1980 - 1989	18.4	22.9***	21.2	20.1	20.6	19.9	20.9	20.3
1990 - 1999	12.6	17.1***	15.7	15.3	14.6	14.3	15.3	15.6
2000 - 2011	4.4	5.8***	5.1	5.2	5.1	4.9	5.3	5.1
% monumental houses	0.4	0.2***	1.6	0.9***	0.5	0.1***	1.4	1.1**
ln median floor space (m2)	4.74	4.71***	4.70	4.69***	4.73	4.72**	4.70	4.69
% monumental building in 6PPC?	1.1	0.6***	2.3	1.7**	1.1	0.6***	2.1	1.8
ln area (m2)	8.25	8.15***	8.07	8.01***	8.21	8.20	8.06	8.02**
ln built-up area (m2)	7.14	7.09***	7.07	7.03***	7.13	7.11**	7.05	7.04
elevation (m NAP)	13.8	0.7***	1.2	0.5***	8.4	8.1***	0.9	0.8***
# houses	19.6	20.3***	19	19	20	20	19	19
# inhabitants	51	56***	51	52	53	54*	51	52
% non-western foreign born	1.8	1.9***	1.8	1.9***	1.8	1.8	1.9	1.9
% 0-14yr	20.3	21.6***	20.7	21.0	20.9	20.8	20.8	20.8
% 15-24yr	11.5	11.7***	11.5	11.5	11.6	11.5	11.4	11.6
% 25-44yr	31.7	32.3***	31.7	32.3**	31.9	32.0	32.0	32.1
% 45-64yr	26.2	25.8***	26.9	26.1***	26.1	25.9	26.5	26.4
% 65+	10.4	8.5***	9.2	9.1	9.5	9.8	9.2	9.1
% 1pp households (hh)	23.2	20.3***	22.6	22.1	22.1	21.8	22.4	22.2
% >1pp hh (children)	45.2	49.2***	47.0	47.4	46.9	46.9	47.0	47.3
% >1pp hh (no children)	31.6	30.6***	30.5	30.5	31.0	31.3	30.6	30.5
mean hh size	2.5	2.7***	2.6	2.6	2.6	2.6	2.6	2.6
mean hh income rel. to Dutch average	1.04	1.09***	1.12	1.11**	1.06	1.05***	1.12	1.11***

	1	2	3	4	5	6	7	8
Table 2 continued					**Within 5PPC variation only**			
	All		< 100m flood line		All		< 100m flood line	
	no flood risk	flood risk	no flood risk	flood risk	no flood risk	flood risk	no flood risk	flood risk
Distance to (m)								
rail	2298	2566***	2488	2358*	2406	2419***	2410	2414*
road	143	151***	136	139	143	150***	135	140***
shops/restaurants	814	934***	875	831*	864	865	849	850
Park	469	392***	405	395	437	437	398	400
public facilities	1312	1312	1233	1221	1308	1317**	1226	1225
recreational area	2492	2189***	2054	2020	2366	2362	2034	2035
agricultural land	503	488***	507	533**	500	492***	523	521
forest	944	1138***	1241	1283**	1025	1027	1264	1267
open plain (dry)	4397	7466***	7259	8091***	5689	5703***	7747	7743
Water bodies:								
open plain (wet)	3366	2659***	2868	2757**	3065	3069	2804	2802
Ijssel-/Markermeer	67466	44460***	44363	44073	57745	57731***	44195	44192
estatuary (closed)	68238	54246***	50665	48305***	62322	62321	49287	49287
Rhine / Meuse river	31211	18413***	24813	23084***	25801	25799	23802	23805
lake	64113	45094***	47797	48192	56077	56065***	48029	48026
storage / watershed	44217	26621***	28968	26121***	36780	36775	27307	27306
recreational inland waterways	1798	1373***	1357	1284**	1622	1613***	1315	1313
water for nat. resource extraction	14530	18055***	19046	19422	16019	16022	19264	19267
mud flat	21544	17792***	17059	16656*	19964	19949***	16825	16823
other small water	814	372***	370	347***	629	625**	358	356
Waddenzee, Eems, Dollard.	109826	95100***	89732	91049	103605	103593***	90503	90499
Oosterschelde	110661	84227***	84716	77952***	99488	99481*	80767	80768
Westerschelde	123219	102347***	105019	98800***	114396	114392	101388	101389
Noordzee	70159	47056***	37313	34603***	60396	60384***	35734	35729

Notes: numbers are averages of each reported variable in 6PPC areas with and without flood risk respectively. *,**,*** denotes significantly different from 6PPC areas without flood risk at the 10%, 5%, 1% respectively. Numbers in the first four columns are based on regressions of flood risk on the reported variable. Numbers in the last four columns are based on regressions of flood risk on the reported variable controlling for 5PPC-year fixed effects. The no flood risk numbers in these columns report the average of all 5PPC-year effects. All reported significance levels are based on standard errors clustered at the 6PPC level.

Table 2 shows the average of each observable housing attribute in 6PPC areas with and without flood risk, also indicating whether they differ significantly from each other. It does so for four different samples.[39] Columns 1 and 2 compare these averages when

[39] We only show these averages for a subset of all our observable housing attributes. For all other observable housing attributes observed at the 6PPC level that we include as controls in our analysis (see the

considering the full sample of all 6PPC areas in our terraced house sample. Next, columns 3 – 8 put perspective on the two ways in which our empirical design controls for unobservables. Columns 3 and 4 report these averages restricting the sample to 6PPC areas within 100m from the flood line, and columns 5 and 6 when controlling for any unobserved house price determinants at the 5PPC level by including 5PPC-year fixed effects. Finally, columns 7 and 8 report these averages employing both these steps in our empirical design.

Column 1 and 2 show that most of the observables differ significantly between flood and non flood risk 6PPC areas in the raw data. Each of the two strategies used to balance the other (un)observables across flood and non flood risk 6PPC areas manages to substantially reduce the number of significant differences (see column 3 – 6). However, only when employing both at the same time, we find hardly any remaining significant differences (column 7 and 8). Moreover, without any exception, the differences in means are greatly reduced (often to a less than 1% or 1ppt difference)[40].

Notable significant differences of concern are those related to the year the house was built, elevation and household income[41]. As briefly mentioned before, houses in the Netherlands built before the major defences were in place are mostly found in places without any flood risk. Many people also find them aesthetically appealing and might therefore pay a higher price for them (although keeping them up is also generally more expensive). If the age dummies we include in our regressions do not adequately control for this, this may confound our estimate of MWTP for flooding. Do note however, that only less than 4% of home sales in our sample were of houses built before 1900, which is why we believe that this is an insignificant concern.

Appendix for a complete list), these descriptives are available upon request. They are also nicely balanced in our “flood line”-sample.

[40] A strong indication that the higher p-values do not merely reflect the fact that our empirical design greatly reduces the variation in the sample by focusing on within-5PPC variation only and/or restricting the sample to 6PPC areas within 100m of the flood line.

[41] Flood prone 6PPC areas are also located significantly further away from a main road or railways. The difference is however about 5m which is very little compared to an average distance of 140m or 2414m to a main road or railway respectively. To address this issue further we include, besides each distance variable proper, also a dummy variable based on each respective distance variable that takes the value 1 if the 6PPC area lies within 25m of each respective (dis)amenity that we have distance information on. This is aimed to tackle the issues that it is mostly location right next to a particular amenity that matters for people’s location decision.

Elevation is another concern. It may, besides affecting the extent of flood risk, also affect the view from the house. Although the difference in average elevation between flood and non flood prone 6PPC areas is only 10 centimeters in our flood line sample, this may also confound our estimates. To address this issue we show results including a fourth order polynomial in elevation as well as dummy variables indicating whether the 6PPC area's elevation is more than 0.5, 1, 1.5, or 2m above the average elevation of all 6PPC areas in the 5PPC[42].

Household income is the only of our neighbour characteristics that remains significantly different between 6PPC areas at risk and not at risk of flooding. It points to possible issues with income based sorting, that we already touched upon in the previous sections (and will explicitly discuss in the next section).

Overall, however, our empirical design greatly reduces the likelihood of confounding in our analysis.

4.2 Average Dutch MWTP to avoid flood risk

Table 3 builds up to our main estimate of the average MWTP to avoid flood risk in the Netherlands. It provides crucial perspective on the three ways our empirical design aims to get us as close as possible to identifying our effect of interest off of quasi-experimental (random) variation in flood risk only.

We start in column 1 by relying on assumption (14), and find a significant positive effect of flood risk on house prices. This puzzling positive effect however disappears as soon as we control for the full set of all 113 other possible house price determinants that we observe at the 6PPC-level (column 2). They consist of characteristics of the housing stock (year built, floor space, plot size, etc), (dis)amenities present in the neighborhood (restaurants, parks, railways, etc), and, importantly a set of 14 controls related to the presence of different types of natural, water-related, amenities. We now find a significant negative effect of flood risk on house prices of 1.8%. However, this result depends on the two strong assumptions that any remaining unobservable house price determinants are uncorrelated with flood risk, and that we specified the correct functional form for the way

[42] These dummy variables are based on the idea that you will only have a nice view when your own house is elevated more than a certain amount of cm above your neighbors' houses.

the observed controls influence house prices[43]. Column 8 (to which we come back later) illustrates that these two assumptions are unlikely to be met: when including also neighbor characteristics as controls, the point estimate of the flood risk discount, although still significant, drops to 0.4%. It clearly illustrates the need for an empirical design that controls for any other possible (unobserved) determinants of house prices as convincingly as possible.

Table 3. Baseline results

	1	2	3	4	5	6	7	8
Flood Risk	0.029	-0.018	-0.010	-0.023	-0.013	-0.010	-0.007	-0.004
	[0.000]	[0.000]	[0.003]	[0.002]	[0.000]	[0.001]	[0.011]	[0.004]
house characteristics	no	yes	no	no	no	yes	yes	yes
6ppc - amenities	no	yes	no	no	no	yes	yes	yes
6ppc - water	no	yes	no	no	no	yes	yes	yes
6ppc - neighbors	no	no	no	no	no	no	yes	yes
FE	-	-	5ppc/yr	-	5ppc/yr	5ppc/yr	5ppc/yr	-
< X m floodline	-	-	-	< 100m	< 100m	<100m	< 100m	-
Observations	619605	616646	619605	42998	42998	42760	38858	563985

Notes: p-values based on standard errors clustered at the 6PPC level in brackets. See Appendix B for an overview of all control variables observed at the 6PPC level that we include in our regressions.

Columns 3 – 5 provide crucial perspective on the success of the two additional steps in our empirical design to achieve just that. In these columns we do not include any observable control variables and solely rely on our empirical design to get an estimate of the effect of flood risk on house prices. In column 3 we include 5PPC-year fixed effects and identify the effect of flood risk only off of variation between the price of houses sold in 6PPC areas within the same 5PPC and in the same year. In column 4 we restrict attention to houses located within 100 meters of the flood line only. And, in column 5, we employ both these steps. In all three columns, we find a significant negative flood risk effect. Employing both steps of our empirical design, we find a flood risk discount on house prices of 1.3%, substantially lower than the 1.8% discount we found when only including a full set of controls in column 2.

[43] The fact that the more flood-prone western part of the Netherlands is also its economic heartland is one very important example of a confounder of the effect of flood risk on house prices we find in column 1. Only including province-year fixed effects based on the twelve Dutch provinces already turns the positive effect of flood risk that we find in column 1 into a negative one.

Next, when also including the full set of observable controls in our baseline specification, we find a significant negative effect of flood risk on house price of 1% (or about € 2400 for the median house at risk of flooding). Importantly, this point estimate is similar to that when solely relying on our BDD design in column 5. It gives us a lot of confidence that our empirical design is indeed getting us close to the (quasi-) experimental ideal.

In our baseline specification we do not control for neighbor characteristics. Although they could be important confounders of the effect of flood risk on house prices that we identify, they are, in our view, clear examples of "bad controls" (see Angrist and Pischke, 2009 p.64-66 for a general discussion; see Gibbons et al. (2013) or Greenstone and Gallagher (2008) for other examples when estimating hedonic house price regressions). Households choose whether to buy a house at risk or not at risk of flooding. A 6PPC's flood risk status may therefore also determine the characteristics of the households in the 6PPC area, making household characteristics a bad control in a regression of flood risk on house prices. In fact, in many hedonic applications such household characteristics are often excluded from the hedonic house price regression as they are seen as demand shifters, used only when estimating the entire MWTP function (see also Chay and Greenstone (2005) for a discussion). Moreover, the coverage of the data that we have available on these neighbor characteristics, especially those on income and household composition, is much worse than that of our other data (see our discussion in section 3.1 and footnote 32). Nevertheless, column 7 shows that our result holds up to also adding these 26 "bad controls" to the regression. The effect of flood risk is only slightly lower.[44]

Column 2 and 8 also provide some perspective on the possibility that we, given the focus of our empirical design on 6PPC areas around the flood line only (see section 3.2), are identifying the average MWTP to avoid flood risk for a specific part of the Dutch population only, i.e. those living close to the flood line. The column 2 and 8 results are based on variation between flood and non-flood risk 6PPCs in all parts of the Netherlands. The significantly negative effect of flood risk on house prices in these two

[44] Note that the difference between a regression with or without these "bad controls" is also much smaller when relying on all steps in our empirical design (compare column 2 vs. 8 to column 6 vs. column 7). We take it as another indication of the success of our empirical design.

columns therefore shows that the price discount for houses at risk of flooding appears not to be restricted to places right along the flood line[45].

Figure 5 provides additional evidence on this. It plots the average median 6PPC house price in 100m bins around the flood line up to a distance of 2500m[46]. The picture clearly shows a discrete negative jump in house prices as soon as we "cross the flood line". Furthermore, if anything, when moving deeper into flood prone territory (i.e. to the left in the figure), house prices keep falling. For houses that do not run any risk of flooding (to the right of the flood line in Figure 4), we do not observe any clear relationship between prices and distance to the flood line.

Figure 5. Average median 6PPC house prices up to 2500 meter of the flood line

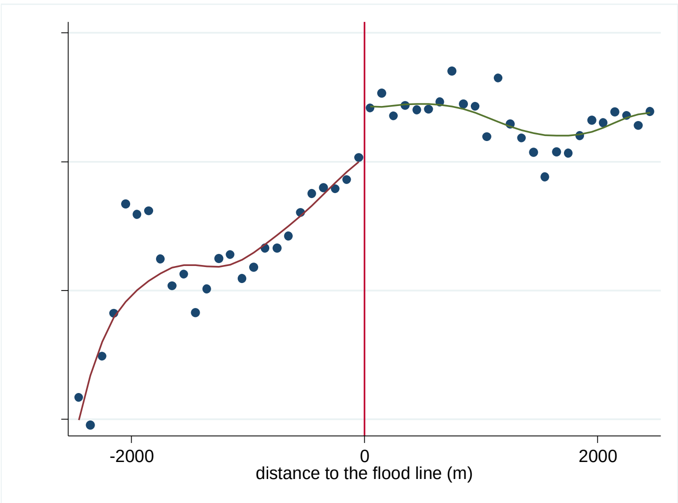

Notes: a negative distance refers to the distance of 6PPC areas **with** flood risk to the flood line, a positive distance to the distance of 6PPC areas without any flood risk to the flood line.

These patterns suggest that the focus on the flood line that is the crucial feature of our empirical design, is unlikely to result in an overestimate of people's MWTP to avoid flood risk. It may even be an underestimate given the much lower house prices in flood prone 6PPCs that are further away from the flood line. A prominent explanation for the

[45] Of course this is not conclusive evidence. This would only be so if all other unobservable house price determinants were uncorrelated with flood risk. Moreover, it requires having chosen the correct functional form for the way the observables influence house prices. Both these assumptions are unlikely to be met (see the earlier discussion in the main text).

[46] 98% of the observations in our terraced house sample that are at risk of flooding are also located within 2.5km from the flood line (256891 observations in total). For the observations not at risk this number is 48% (170630 observations in total).

decreasing price of houses at risk of flooding that are located further away from the flood line is that distance to the flood line is positively associated with flood depth, i.e. the amount of water entering the house in case of a flood (see section 4.3.2 for more on this).

Before discussing the implications of the significantly negative 1% discount that we identify on house prices in flood prone areas, we first show that this finding holds up to a host of robustness checks.

4.2.1 Robustness

The results of the different robustness checks are shown in Table 4a and 4b.[47] Columns 1 – 3 in Table 4a show a first very important robustness check to our baseline findings. The fact that our main variable of interest, the flood risk dummy, is 6PPC specific and does not vary over time can result in overstating the significance of our estimated effect flood risk effect (see Angrist and Pischke, 2009 p.311). For this reason we always cluster our standard errors at the 6PPC level[48]. However, one could argue that clustering our standard errors still does not entirely solve this issue (Angrist and Pischke, 2009 p.313). An alternative way to deal with this issue is to first calculate 6PPC averages of all our variables, and then use these 6PPC averages to estimate (15) using WLS with the number of years a particular 6PPC is present in our sample as weights. Column 1 – 3 show that our results hold up to this important robustness check. The size, sign and significance of our flood risk effect is very similar to that found in columns 6 and 7 in Table 3a respectively.

[47] In addition, Table A1 in Appendix A show the sensitivity of our results to the way we deal with the fact that the official house price data of the Kadaster reports the sale of a piece of residential property (a garage, piece of garden or a boathouse for example) as if it were the sale of an entire property.

[48] All our results are robust to using standard errors clustered at the 5PPC, 4PPC or town level instead. These less strict specifications also take account of possible spatial dependence between the error terms of 6PPC areas in the same 5PPC, resp. 4PPC, area. We decided to show the standard errors clustered at the 6PPC level in the main text. The inclusion of 5PPC-year fixed effects in all our main specifications implies that we already control for most of the spatial dependence in unobservable house price determinants. Moreover, clustering at the 5PPC, 4PPC or town level reduces the number of clusters used substantially (by a factor 3, 10 or 50 respectively), with possibly unwanted consequences for the reliability of inference (Angrist and Pischke, 2009 section 8.2.3).

Table 4a. Robustness I

	6PPC averages			1-1 matching to nearest across FL		all houses no apartments	all houses	no houses built 2004-2010	median 6PPC = 1 sale	entire 5PPC terraced houses	detailed elevation
	1	2	3	4	5	6	7	8	9	10	11
Flood Risk	-0.011	-0.012	-0.007	-0.009	-0.013	-0.005	-0.005	-0.010	-0.015	-0.012	-0.010
	[0.012]	[0.005]	[0.042]	[0.001]	[0.000]	[0.055]	[0.052]	[0.001]	[0.000]	[0.001]	[0.001]
house characteristics	yes	yes	yes	yes	no	yes	yes	yes	yes	yes	yes
6ppc - amenities	yes	yes	yes	yes	no	yes	yes	yes	yes	yes	yes
6ppc - water	yes	yes	yes	yes	no	yes	yes	yes	yes	yes	yes
6ppc - neighbors	-	-	yes	-	-	-	-	-	-	-	-
FE	-	5ppc	5ppc	5ppc/yr	5ppc/yr	5ppc/yr	5ppc/yr	5ppc/yr	5ppc/yr	5ppc/yr	5ppc/yr
< X m floodline	< 100m	< 100m	< 100m	<100m	<100m	<100m	<100m	<100m	<100m	<100m	<100m
nr.observations	9645	9645	8980	18041	18112	53130	76441	42029	29543	16679	42760

Notes: p-values based on standard errors clustered at the 6PPC level in brackets. Columns 1 – 3 report results of regressing each variable's average over the 13 years in our sample on average ln house prices over the same period using WLS with the number of years a 6PPC area is observed over our sample 13-year sample period as weights. P-values in these last three columns are based on robust standard errors. In all columns except 3 and 4 we focus on the sale of terraced houses only. In column 8 our detailed elevation specification involves including a fourth order polynomial in elevation (m NAP), as well as four different dummy variables indicating whether the 6PPC's elevation is more than 0.5m, 1m, 1.5m, 2m higher than the average elevation of all 6PPC areas in the same 5PPC. In column 3 and 4 we also include controls for the percentage of each house type sold in each 6PPC (terraced houses, detached, "two-under-one-roof", and apartments).

Column 4 and 5 show a second robustness check. Instead of restricting the sample to all 6PPC areas within 100 meter of the flood line, we now match each 6PPC area within 100 meter from the flood line to the nearest 6PPC area in the same 5PPC area with a different flood risk status. The resulting sample is much smaller than our baseline sample due to this 1-1 matching.[49] Results when using this "matched sample" are very similar to those in our baseline findings, both when relying solely on our empirical design (in column 5) as well as when also including all observable controls (in column 4).

As a third robustness check, column 6 and 7 put some perspective on the first step in our empirical design, i.e. our exclusive focus on terraced houses. In column 6 we also include 6PPC areas with a housing stock consisting solely of detached or "two-under-one-roof" houses; and in column 7 we also consider apartments.[50] We still find a negative price discount for houses in flood prone areas when widening the sample this way. However, the point estimate of the discount decreases to about 0.5%. Moreover it is less precisely estimated. The focus on a sample of 6PPC areas with a much less homogenous housing stock clearly comes at the cost of losing some precision. Unobserved differences in the characteristics of the housing stock between 6PPC areas within the same 5PPC area are the least well accounted for by steps 2 and 3 of our empirical design. Although our results do come through when considering sales of any house type, we take our focus on terraced houses only as a very important first step in our empirical design aimed at getting the best possible estimate of the effect of flood risk on house prices.

Next, column 8, 9 and 10 show results when considering an even more selective sample than our baseline 6PPC terraced house sample. In column 8 we exclude all 6PPC areas where new houses were built during our sample period. This aims to address possible concerns with sample selection arising from people making an active choice about where to build new homes that may confound our estimate of MWTP to avoid

[49] The typical 6PPC area in this 1-1 matched sample is also, not surprisingly, located 6 meters closer to the flood line than in our baseline sample. The uneven number of observations in this matched sample is explained by the fact that each 6PPC area may be the nearest 6PPC area with a different flood risk status for more than one other 6PPC areas.

[50] We show results with and without apartments as we do not have information on the exact floor of the apartment building that the sold apartment is located on. The effect of flood risk can be expected to be different for apartments located on the higher floors of an apartment building. They will not experience any damage if a flood would happen, only the inconvenience of a flooded first, and at most second floor, of the apartment building.

flood risk. In column 9 we only consider 6PPC areas where the median house price reported in a particular year is based on the sale of single house. As already mentioned in section 3.2.1, our use of aggregate 6PPC-specific measures of house prices and house price determinants could result in issues with measurement error. When the median house price reported for a particular 6PPC is based on a single sale however, this measurement error is confined to the independent variables only. Unless this measurement error is negatively correlated with actual flood risk (see also footnote 27), this will only bias our estimated effect of flood risk towards zero. Finally in column 10, we not only restrict attention to 6PPC areas with a housing stock consisting entirely of terraced houses, but further restrict the sample to those "terraced house 6PPC areas" located in a 5PPC area whose entire housing stock consists of terraced houses only (i.e. 5PPC areas consisting solely of "terraced house 6PPC areas"). The results in column 8 – 10 show that our baseline findings hold up to using any of these more selective samples.

Finally, column 11 takes explicit note of the fact that elevation was one of the most prominent examples of a variable that was not nicely balanced between places with and without flood risk, even when employing all three steps in our empirical design (see Table 2). Places at risk of flooding are, not surprisingly, typically located at lower altitudes. Although on average this difference is only 10 centimeters, one could be worried that our estimated flood risk effect is confounded by the effect of having a nicer view at higher altitudes. In our baseline regressions we simply include elevation as one of the controls in an additive linear way. As the effect of elevation need not be linear, column 11 shows that our results do not change when instead adding a fourth order polynomial in elevation in combination with four different "view dummies" that indicate whether a 6PPC's elevation is more than 0.5, 1, 1.5, or 2 meters higher than the average elevation of all 6PPC areas in the same 5PPC.[51]

Table 4b contains the two final robustness checks. Column 1 – 4 show that flood risk is not significantly related to the likelihood of observing a house sale nor to the number of houses sold in a 6PPC area. Column 1 and 2 show this when considering the sale of any type of house, and column 3 and 4 when only considering the sale of terraced

[51] We always find the same result regardless of the order of the polynomial in elevation used and/or the definition and number of "view dummies" included.

houses. Although this is, of course, not conclusive evidence that we do not face any selection problems, it makes it quite unlikely. Only an unobservable variable that is not already controlled for by our empirical design and that is related to both the likelihood of a sale and to house prices can bias our estimated effect of flood risk on house prices. In this regard it is also good to note that that the average time that a house is on sale before being sold is less than a year. Given that we observe 13 years of data, it is therefore very likely that all houses that were put up for sale during this period were sold, except maybe, those in the latest year(s).

Table 4b. Robustness II

	Quantity instead of price?				Definition of the flood line				False "Flood line" at +1m NAP - "Flood risk" = <1m NAP
	sale?	nr sold	sale terraced houses?	nr sold terraced houses	< 300m	< 200m	< 70m	< 60m	
	1	2	3	4	5	6	7	8	9
Flood Risk	0.002	0.001	0.003	0.012	-0.004	-0.006	-0.010	-0.005	0.006
	[0.388]	[0.862]	[0.428]	[0.131]	[0.097]	[0.028]	[0.008]	[0.241]	[0.174]
house characteristics	yes	yes	yes	yes	yes	yes	yes	yes	yes
6ppc - amenities	yes	yes	yes	yes	yes	yes	yes	yes	yes
6ppc - water	yes	yes	yes	yes	yes	yes	yes	yes	yes
FE	5ppc/yr	5ppc/yr	5ppc/yr	5ppc/yr	5ppc/yr	5ppc/yr	5ppc/yr	5ppc/yr	5ppc/yr
< X m floodline	< 100m	< 100m	< 100m	< 100m	< 300m	< 200m	< 70m	< 60m	< 100m
nr.observations	367029	367029	105053	105053	147074	100183	22289	15227	13348

Notes: p-values based on standard errors clustered at the 6PPC level in brackets. In column 1 and 2, we consider sales of any house type. In column 3 and 4 we restrict our sample to 6PPC areas where the housing stock consists of terraced houses only. See Appendix B for a detailed description of all control variables observed at the 6PPC level that we include in our regressions.

Next, columns 5 – 8 verify the robustness of our findings to the exact definition used for the "flood line". Instead of considering only 6PPC areas located within 100 meters from the nearest 6PPC with a different flood risk status, we widen this range to 300 or 200 meters in column 5 and 6. We find that widening the distance band reduces the estimated flood discount, moreover its estimate becomes increasingly imprecise[52]. It underlines the importance of the extreme spatial detail of our data, allowing us to consider a sample of only very nearby houses. In column 7 and 8, we instead use an even narrower range of 70

[52] When further increasing the distance band to 400, 600, 800 or 1000 meters, the point estimates for the flood risk discount always lies within the 0.4 - 0.5% range, but is never significant at even the 10% level.

or 60 meters. We find the same significant 1% flood risk discount as in our baseline sample when using a cutoff distance to define the flood line up to 70m[53]. When we reduce this distance even further (see column 8), we still find a negative flood risk discount, but it is no longer significant. Using an ever stricter definition of the floodline results in a substantial loss of observations, at some point simply leaving us with too little variation between 6PPC areas within the same 5PPC and year to precisely identify the flood risk discount.[54] Note that using an ever smaller distance cutoff also implies looking at an increasingly selected sample of the smaller 6PPC areas.

Finally, column 9 shows results of a "false experiment": we first limit the sample to 6PPC areas facing *absolutely no* flood risk and mark those 6PPC areas located below +1m above sea level as places "artificially at risk of flooding". Next, we further limit the sample to the "artificial flood line" by considering only those 6PPCs "artificially at risk of flooding" that are located within 100m of a 6PPC "not artificially at risk", and vice versa; and run our hedonic house price regression (15). A significant "artificial flood risk" effect would shed some doubt on the success of our design. We do not find this (results are similar when using e.g. +2m above sea level instead). Moreover, besides being insignificant, the point estimate is now positive.

4.3 Interpretation

We now turn to the implications of the 1% flood risk discount that we have identified, and answer the five different questions that we set out in detail at the end of section 2.

1. *Do the Dutch feel so safe behind their publicly provided dikes that house prices no longer reflect flood risk?*

Our estimated 1% flood risk discount immediately implies that the Dutch do not feel completely safe behind their dikes. The rare risk of flooding is still reflected in house prices. It means that, see (7), the average Dutch person does not fully trust the publicly built flood defenses to keep him/her safe from flooding ($\rho(\tau,Ғ) > 0$). Also, he/she does not

[53] Results when using a 90 or 80 meter cutoff are available upon request.

[54] When not including 5PPC-year fixed effects nor any observable controls, we do find a significantly negative flood risk discount when using a 60 meter cutoff to define the flood line. Also, do note that the 100 meter cutoff used to define the flood line in our baseline sample is already much smaller than the 0.15 mile and 0.1 mile cutoff used in the BDD in Black (1999) and Bayer et al. (2007) respectively.

believe that the government will be able to fully compensate the flood damage they expect to their house when a flood would happen ($b_i(s,\tau)$ < 0, for at least one housing attribute).

2. What is the average flood risk perception in the Netherlands?

If we are willing to make the additional assumptions discussed in section 2, (8) shows that our estimated 1% flood risk discount is a direct measure of the average expected future (uncompensated) flood damage to houses located in flood prone areas. That is, the average Dutch person believes that a house at risk of flooding is expected to incur a future (uncompensated) flood damage amounting to 1% of the house.

If one is also willing to assume that all people have the exact same perception of the (uncompensated) share of one's house lost in case a flood actually happens, α_F can be directly interpreted as the average perceived likelihood of a major flood, i.e. one that completely destroys one's house and that would inflict so much damage that the Dutch government will be unable to provide any effective compensation (i.e. $b^k(s^k,\tau) = 0$ for all k). Our estimated α_F of 1% thus implies that people expect such a major flood to happen once every 100 years. The average perceived likelihood of more minor floods with an expected uncompensated loss of e.g. half or a third of the house is even higher: once every 50 or 33 years respectively. These probabilities are much higher than the official protection levels at which the Dutch government claims to uphold the country's defenses (see Figure 3). In the best protected places they should reduce the likelihood of a flood to happen to once every 10,000 years. But even in the dike-ring with the highest acceptable flood risk, its defenses should be built to fail only once every 500 years. This finding is in line with recent survey evidence. The self-reported likelihood of a flood in these surveys is typically even higher (e.g. 7%-19% chance in the next 10 years (TNS Nipo, 2006), or >20% chance in the next 50 years (Bockarjova et al., 2010)).

An explanation for these much higher perceived risks is that few people trust and/or are aware of the official state of the flood defenses defending their home (see the next section for more evidence on this). In fact, the government itself may not know whether the claimed levels of defense are actually attained by the current defenses in place. Recent evidence showed that only about two thirds of all dikes and dunes, and half of the other primary waterworks were up to standard (Inspectie Verkeer en Waterstaat, 2011). Instead of basing their perceived flood risk levels on the official government standard for their particular location, people appear to be much more influenced by the actual occurrence of flood events. In the 61 years after the last major flood disaster in 1953, there were at least another five instances of (near-) failure of one or more of the country's flood defenses that received widespread (media) attention.[55]

3. *How does the average MWTP for flood protection relate to the public money already spent on flood protection?*

To keep up the flood defenses at their required level, the Dutch government currently spends about € 400 per year on each house at risk of flooding. Using (10), we can directly relate our estimate to this amount of public money spent each year to keep the country safe. Assuming a 3% mortgage interest rate (r = 0.03)[56], and using the fact that the median price of a house at risk of flooding was €243,000 in 2011, our estimated 1% MWTP to avoid flood risk corresponds to a MWTP a yearly flood insurance premium of €69. About 18% more than the annual amount of public money spent per house at risk to protect the country from flooding. It can also be directly compared to the average annual insurance premium paid for home insurance. In 2011 this was €131 (COELO, Woonlastenmonitor 2011). Our findings thus imply that people living in flood risk areas would be willing to pay about 50% more for their home insurance policy if its coverage were extended to also insure against flood risk[57].

[55] In 1960 a dike near Amsterdam broke over a length of 80 meters, flooding the entire neighborhood of Tuindorp-Oostzaan. About 15000 people had to leave their flooded homes. In 1984, several villages along the Meuse flooded. In 1993 about 8% of the province of Limburg flooded due to exceptional high water levels in the river Meuse. Fortunately, mostly farmland; 8000 people were evacuated. In 1995 250,000 people were evacuated out of fear that some of the main dikes along the Meuse and Rhine would break. In the end the dikes held up and flood damage was limited to several small villages in Limburg. In 2003, a dike near Wilnis gave way over a length of 60 meters, flooding the town. 1500 people had to be rapidly evacuated. Nobody lost his/her life in any of these events. See Gallager (2014) for the effect of media attention on people's perceived flood risk.

[56] This is based on the typical ***after-tax*** interest rate on an interest only mortgage in our sample period.

[57] Note that this interpretation of our estimate is of particular interest for the recent policy debate in the Netherlands about whether or not part of the costs of the country's public flood protection should no longer be paid from the government's general tax income, but, instead, from taxes levied on people living in flood prone areas only. Our estimate indicates that people in flood prone areas would be willing to pay such additional taxes up to an amount of €69 per year. But only if it credibly buys them full insurance against (the negative consequences of) future floods.

4. *How much housing wealth will be lost due to future climate change?*

Finally, under the strong assumption of an unchanging MWTP to avoid flood risk, our identified MWTP allows us to put some tentative numbers on the cost of a future rise in sea levels in terms of lost housing wealth. To do this we assume that the state of the Dutch defenses remains at their current levels. We base our calculations on scenarios for sea level rise of e.g. the Royal Dutch Meteorological Institute (KNMI) or the Deltacommissie (2008). In particular, we consider sea levels to rise by 24cm, 100cm and 150cm in a best-, medium-, and extreme-case climate change scenario respectively[58]. These sea level rises would mean that an additional 275,450, 1,233,620, 1,562,545 houses respectively would face the risk of being flooded in case the defenses fail. The median price of these, currently flood safe, houses was about €250,000 in 2011. As a result, our estimated MWTP to avoid flood risk of 1%, implies a tentative total loss of housing wealth due to a future sea level rise of €0.7, €3.1, €4 billion in the best-, median-, worst-case climate change scenario respectively, even if the Dutch defenses were upgraded to offer the same protection level at these increased sea levels as they do today.

4.4 Heterogeneity in MWTP to avoid flood risk – sorting and awareness

4.4.1 Sorting

Our baseline regressions identify the *average* MWTP to avoid flood risk of Dutch households. However, as previously discussed, heterogeneity in household income or in household preferences for flood safety can result in a biased estimate of this average MWTP when estimating (14).

In this section, we probe into the likelihood of such income- and taste-based sorting. We do this, similar to Greenstone and Gallagher (2008), by replacing the dependent variable in (14) by ln household income, and three different proxies of households' preferences for flood risk respectively. Two of these proxies, the percentage of households in the 6PPC with children and the percentage of 0 – 14 year olds in total

[58] Future climate change will most likely also increase the risk of river floods due to more erratic rainfall patterns in upstream areas. However, contrary to rising sea levels, these river floods will, although increasing in frequency, not lead to many currently flood safe areas becoming flood prone. Moreover, it is much harder to identify in the data which currently flood safe houses, if any, would switch to being flood prone due to these more frequent river floods. In case of a sea level rise this can be readily inferred from elevation data.

6PPC population, are based on the idea that families with young children have a stronger taste for flood-safety. The other proxy, the percentage of non-western foreign born in total 6PPC population, is based on the idea that immigrants might be more sensitive to the risk of flooding as they are less accustomed to living with this flood risk. Also, flood risk may be one of the first things people think about when moving to the Netherlands[59].

Table 5. Sorting: income and risk preferences

	ln hh Income	% non-western foreign-born	% hh with children	% 0-14yr	ln hh Income	% non-western foreign-born	% hh with children	% 0-14yr
	1	2	3	4	5	6	7	8
Flood Risk	-0.007	0.003	0.506	0.079	-0.003	-0.012	-0.018	0.014
	[0.035]	[0.889]	[0.135]	[0.637]	[0.227]	[0.561]	[0.100]	[0.336]
house characteristics	yes	yes	yes	yes	yes	yes	yes	yes
6ppc – amenities	yes	yes	yes	yes	yes	yes	yes	yes
6ppc – water	yes	yes	yes	yes	yes	yes	yes	yes
6ppc – neighbors (incl. ln house price)	-	-	-	-	yes	yes	yes	yes
FE	5ppc/yr	5ppc/yr	5ppc/yr	5ppc/yr	5ppc/yr	5ppc/yr	5ppc/yr	5ppc/yr
< x m floodline	< 100m	< 100m	< 100m	< 100m	< 100m	< 100m	< 100m	< 100m
nr.obs	41756	42455	42455	42398	38858	38858	38858	38858

Notes: p-values based on standard errors clustered at the 6PPC level in brackets. See Appendix B for a detailed description of all control variables observed at the 6PPC level that we include in our regressions.

Table 5 shows the results of these "sorting-regressions", always employing all three steps in our empirical design. In column 1 – 4 we do not include other neighbor characteristics as controls. In column 5 – 8 we do, and we also include a 6PPC's ln median house price (reflecting the quality of the housing stock) despite the fact that these variables can be argued to be "bad controls" (see our earlier discussion in section 4.2).

We find some evidence for income-based sorting in column 1. Household income is 0.7% higher in places without any flood risk. Interestingly, this appears to be entirely capitalized into the housing market. Controlling for house prices in column 5, we no longer find any significant effect of flood risk on household income. The wealthier households are the ones that can afford to pay the more expensive houses in these areas[60]. It means that our baseline estimate of MWTP to avoid flood risk might be (partly)

[59] Note that one could also well expect the contrary: immigrants being less aware of the flood risk they face when moving to the Netherlands.

[60] Not that surprising maybe, given the fact that Dutch households spend on average about 24% of their net income on their house (CBS, 2012). Only controlling for house prices already turns the effect of flood risk on household income insignificant. Result available upon request.

picking up MWTP to live next to wealthy neighbors. However, this seems to be of limited concern. Column 7 and 8 in Table 3 and column 3 in Table 4a already showed that our main finding is robust to controlling for household income when estimating (15).

We find no evidence of taste-based sorting based on the three proxies that we use for households' flood safety preferences. Of course this is no exclusive evidence that taste-based sorting is not an issue. Indeed, one could even take the evidence for income-based sorting that we find as being (partially) driven by differences in flood safety preferences between high and low income households. Without information on the exact distribution of flood safety preferences over the Dutch population it is very difficult to provide conclusive evidence on the relevance of taste based sorting, and to infer whether we, if at all, over- or underestimate average MWTP to avoid flood risk[61]. In the next subsection we come back to this issue one last time when looking at the differences in MWTP between consumers facing different extents of flood risk.

4.3.2 Awareness

Our baseline result provides the best possible estimate of the average MWTP to avoid flood risk in the Netherlands. Earlier papers have however shown that people's flood awareness may differ based on previous flood experience (Gallagher, 2014; Hallstrom and Smith 2005; or Bin and Landry, 2013), on the visibility of the threat in their immediate surroundings, or on the extent of the risk that they and their property face (Troy and Romm, 2004).

In this section, we further explore possible heterogeneity in people's MWTP to avoid flood risk[62]. In particular, we use detailed information on the visibility of water in the neighborhood, the official flood protection levels (claimed to be) upheld by the government, as well as the actual expected water levels in case of a flood, and verify whether MWTP to avoid flood risk depends on these variables. Each of these variables affects MWTP by possibly affecting people's flood risk and/or flood damage perception (7). The first two primarily affect people's perceived likelihood of a flood happening,

[61] This depends on the exact distribution of preferences over the population, as well as on the availability of houses at different levels of flood risk.

[62] Do note that we already briefly touched upon this issue in section 4.2, where we showed that MWTP to avoid flood risk is not restricted to places close to the flood line only (see e.g. Figure 5).

$\rho^k(\tau,\mathrm{F})$, whereas expected water levels in case of a flood directly relate to the damage people expect in case of a flood, $b_i^k(s,\tau)$.[63] Table 6 shows our results.[64]

Table 6. Differences in flood risk awareness?

Official state of defenses		Visibility of water			Flood severity	
	1		2	3		4
Flood Risk	-0.008**	Flood risk	-0.009***	-0.008**	Flood risk	-
(1/10000 yrs)	[0.034]		[0.004]	[0.032]		-
official state of defenses:		Distance to:	large water body	water	water level if flood:	
1/4000 yrs	-0.011	< 250m	-0.001	-	0 - 20cm	-0.005
	[0.179]	250 - 500m	-0.008	-		[0.230]
1/2000 yrs	-0.013	500 - 750m	0.013	-	20 - 50cm	-0.004
	[0.269]	750 - 1000m	-0.036**	-		[0.395]
1/1250 yrs	0.006	1000 - 1250m	0.007	-	50 - 80cm	**-0.014**
	[0.434]	< 30m	-	-0.013		**[0.025]**
no dike-ring	0.002	30 - 60m	-	0.004	80cm - 2m	**-0.020**
	[0.938]	60 - 90m	-	-0.003		**[0.000]**
		90 - 120m	-	-0.010	> 2m	**-0.020**
		120 - 150m	-	0.003		**[0.043]**
p-value equal?	[0.416]	p-value equal?	[0. 287]	[0. 842]	p-value all 0?	[0.002]
house characteristics	yes		yes	yes		yes
6ppc - amenities	yes		yes	yes		yes
6ppc - water	yes		yes	yes		yes
FE	5ppc/yr		5ppc/yr	5ppc/yr		5ppc/yr
< X m floodline	< 100m		< 100m	< 100m		< 100m
median (q25) distance to water (m)			5859 (1511)	197 (102)		
nr.obs	42760		42760	42760		42760

Notes: p-values based on standard errors clustered at the 6PPC level in brackets. ***, **, * denotes significant at the 10%, 5%, 1% level respectively. Estimated coefficients are for each respective variable *interacted* with the flood risk dummy. See Appendix B for a detailed description of all control variables observed at the 6PPC level that we include in our regressions.

First, in column 1, we allow MWTP to avoid flood risk to depend on the officially published flood protection levels at which the country's defenses are claimed to be

[63] This distinction is not perfect, we will come back to this when discussing the results.

[64] We also find no evidence that MWTP to avoid flood risk is different in 6PPC areas that were evacuated during the large-scale near river floods in 1995. However, data at the 6PPC level on house prices and many other house price determinants is unfortunately only available from 1999 onwards. This complicates the interpretation of this finding substantially, which is why we do not report it in the main text. It could e.g. mean that people have not updated their flood expectations following the 1995 evacuations. Since no actual flooding happened, the evacuations only reinforced people's trust in their flood defenses. Instead people's flood expectations may have changed immediately after the evacuations, but, four years later, they have reverted back to what they were before 1995 (see Gallagher, 2014 or Bin and Landry, 2013 for evidence of such temporary effects of actual flood events on people's flood expectations). Alternatively, it could also be that people living in these areas at the time of the evacuations have updated their flood expectations, but have sold their homes to unaware people that lived in non-evacuated areas in 1995.

upheld. As shown in Figure 3 there is substantial heterogeneity in these protection levels in different parts of the country. They could lead to differences in people's flood risk perception $\rho^k(\tau,\text{Ғ})$. We take locations located in a dike-ring offering the highest protection level (a chance of failing only once every 10,000 years) as benchmark, and verify whether MWTP to avoid flood risk is higher in places located in dike-rings protected by weaker flood defenses. Results show no significant differences in MWTP to avoid flood risk depending on these official risk levels. The easiest explanation is that people do not know the officially published protection levels of the Dutch defenses. Alternatively, people are aware of them, but doubt that they accurately reflect the actual state of the Dutch flood defenses.

The latter is in our view most likely. There are in fact three good reasons to doubt the officially published protection levels. First, recent research shows that 25-50% of the Dutch primary flood defenses are not offering the protection they are supposed to (Inspectie Verkeer en Waterstaat, 2011). Second, current protection levels were set around 1960 and were based on the likelihood of a flood due to water levels *higher* than what the defenses are built for. Since then it has become clear that overflow constitutes only about 40% of the flood risk, the rest coming from e.g. piping, a dike subsiding or breaking, or the failure of a lock or weir. Taking these into account results in a sometimes drastic reduction of the protection levels offered by many of the Dutch defenses (see Veiligheid Nederland in Kaart, 2014). Finally, there exists substantial heterogeneity in actual protection levels within the same dike-ring area (see also Veiligheid Nederland in Kaart, 2014).

The second factor that we look at that may result in heterogeneity in MWTP to avoid flood risk is the visibility of water in the near surroundings of a house. People's flood risk perception, $\rho^k(\tau,\text{Ғ})$, or their expected uncompensated damage in case of a flood, $b^k_i(s,\tau)$, may be higher in places where the water that threatens to flood them is very visible. We look at this possibility in two different ways. In column 2 of Table 6 we verify whether MWTP to avoid flood risk depends on the distance to a large water body that is visibly protected by the country's primary flood defenses (the sea, the main rivers, and lakes). In column 3 we instead consider the distance to any type of water, also including streams, creeks, or ditches that are visible, but not always visibly protected[65]. This distance is typically much smaller than that to larger water bodies (which is also the reason why we use different distance bands in each of the two cases). In both cases, we do not find any evidence that people living closer to the water that threatens them are more aware of the threat.

Finally, we verify whether MWTP to avoid flood risk depends on the expected extent of the flood. In particular, instead of including a single flood risk dummy in our regressions, we now estimate MWTP to avoid flood risk separately for each category of expected water levels in case of a flood reported on www.risicokaart.nl (0 – 20cm, 20 – 50cm, 50 – 80cm, 80 – 200cm, > 200cm)[66].

We find that house prices depend systematically on the actual extent of the flood risk facing the house, *s*[67]: the higher expected water levels in case of a flood, the more people are willing to pay to avoid this flood risk. Moreover, although we find a price discount on all houses facing flood risk, it is not significant in 6PPCs that would receive less than 50cm of water in case of a flood.

These results can be interpreted in several different ways however, see (3). The most likely and straightforward interpretation is that people simply expect more (uncompensated) damage to their house the higher the expected water level in case of a flood. In fact, (8) shows that, if we, in addition to the assumptions outlined in section 2, assume that all people have the same perceived likelihood of a flood, $\rho^k(\tau,\mathrm{F}) = \rho(\tau,\mathrm{F})$, our estimates can be directly used to infer how much more damage people expect on average in places where the water would reach higher levels in case of a flood. People for example expect three times as much damage when water levels would reach up to 80cm compared to water levels of only 20cm. Or, five times as much damage when the house would be flooded by more than 2 meters of water compared to water levels of only 50cm.

A second explanation for our findings is that people's perceived flood risk is larger in 6PPC areas that are expected to be flooded by more water in case of a flood. The larger this amount of water, the more visible it may be that your house is running the risk

---

[65] For example because their water level is kept at acceptable levels by a weir, lock or pumping station.

[66] We group the 2 – 5m, and > 5m categories together given that only about 0.25% of observations fall into this latter category.

[67] Results are very similar when not restricting the sample to the flood line only. They are available upon request.

of being flooded. Survey evidence suggests that this may indeed be part of the explanation of our findings in column 4. The share of respondents correctly stating that their house will be flooded in case the defenses break is higher, the higher the actual expected water levels in case of a flood (see Watermonitor, 2009).

Finally, these findings could be taken as an indication of taste-based sorting. Disregarding the insignificant effects up to 50cm of flood water, they imply a MWTP to avoid flood risk that is *higher at lower* flood water levels.[68] This is difficult to reconcile with homogenous tastes for flood safety. Do note that this interpretation is, although a real possibility, very tentative. It is based on a very rough estimate of the way the house price – flood risk gradient depends on the expected severity of the flood. In fact, these differences in MWTP to avoid flood risk are statististically insignificant.

Taking the estimates in column 4 seriously, we can also fine-tune our predicted loss of housing wealth due to rising sea levels in the best-, median- and worst-case climate change scenario of the KNMI. To do this we collected two pieces of additional information. First, the number of houses currently facing the risk of being flooded by no, or less than 50cm of water that would start facing the risk of being flooded by 50 – 80cm, or by more than 80cm of water in each of these scenarios respectively. And second, the number of houses currently facing the risk of being flood by 50-80cm that would start facing more than 80cm of water in each of the three scenarios. Using these numbers, and the significant point estimates in column 4 of Table 6, we come to a total loss of housing wealth of €1.1 billion, € 4.9 billion, and € 8.1 billion in case of a 24cm, 1m, 1.5m sea level respectively. Especially in the latter two scenarios these numbers are substantially higher than that based on our baseline average MWTP in section 4.3.

5. Conclusions

The Netherlands is one of the most flood prone countries in the world. It is also the best protected country against floods in the world. The Dutch government spends over € 1.1 billion per year (about € 400 per house at risk) on protecting the country against floods.

[68] I.e. 1.4% to move from a house facing a flood risk of 50-80cm of water to an, otherwise equal, house facing a lower, or no flood risk; 0.6% to move from a house facing a flood risk of 80cm-2m of water to an, otherwise equal, house facing a flood risk of 50-80cm of water; and 0% to move from a house facing a flood risk of >2m of water to an, otherwise equal, house facing a flood risk of 80-2m of water).

Without these defenses 36% of the Netherlands floods; containing approximately 2.8 million houses/households.

We show in this paper that, despite these world-class flood defenses in place, the average Dutch citizen still prices in the (rare) risk of flooding into the amount they are willing to pay for their house. Using a border-discontinuity type design, that relies heavily on our unique, extremely detailed dataset on house prices and flood risk in all 459,279 six-digit Dutch postal-code areas, we find that houses that would flood in case the country's defenses fail cost on average 1% less than otherwise equal homes that do not run any flood risk. Furthermore, we show that this flood risk discount depends positively on the expected amount of water that would enter the house in case the Dutch defenses fail.

Our results imply that the average Dutch person does not feel fully protected by the publicly provided flood defenses in place, nor does he/she believe that the government is able to live up to its claim to compensate people for the flood damage they would suffer in case the defenses fail. In fact, this estimate can be taken to imply that the average Dutch citizen expects a major flood to happen once every 100 years, much more often than what the government claims the Dutch defenses are built for.

Our estimate also puts an interesting perspective on the recent (heated) debate on whether or not the costs of the Dutch flood defenses should continue to be borne by all Dutch citizens, regardless of the actual risk of flooding they, and their property, face. Our results indicate that those living in a house at risk of flooding would, on average, be willing to spend an additional €69 per year to be fully insured against any future flood risk. In the absence of any private insurance market for flood risk in the Netherlands, it opens the scope for a property "flood tax" paid only by the 2.8 million households currently living in areas facing actual flood risk. If the government would offer credible flood protection in return such a tax could increase public spending on the country's defenses by up to €200 million. The future costs of keeping the Netherlands flood safe are only likely to further increase as a result of rising sea levels. Funding these costs by such a property "flood tax" would, in our view, be a good idea that also, based on our results, may meet less resistance than the government thinks.

REFERENCES

Angrist, J and S. Pischke, (2009), *Mostly Harmless Econometrics: An Empiricist's Companion,* Princeton University Press, Princeton.

Asselman, N., J. ter Maat, A. de Wit, G. Verhoeven, S . Soares Frazão, M. Velickovic, L. Goutiere, Y. Zech, T. Fewtrell, and P. Bates, P., (2009), "Flood inundation modelling: model choice and application". in P. Samuels, S. Huntington, W. Allsop, and J. Harrop (eds.) *Flood Risk Management: Research and Practice*, Taylor and Francis Group, London, 211-219

Klijn, F., N. Asselman, and H. Van der Most, (2010), "Compartmentalisation: flood consequence reduction by splitting up large polder areas", *Journal of Flood Risk Management,* 3, 3-17.

Bayer, P., F. Ferreira and R. McMillan. (2007): "A Unified Framework for Measuring preferences for Schools and Neighborhoods", *Journal of Political Economy*, 115, 588-638.

Barro, R, (2009), "Rare Disasters, Asset Prices, and Welfare Costs," *American Economic Review*, 99, 243–264.

Barro, R. and J.F. Ursua, (2012), "Rare Macroeconomic Disasters", *Annual Review of Economics*, 4, 83-109.

Barro, R. (2014), "Environmental Protection, Rare Disasters and Discount Rates", *Economica,* forthcoming.

Bartik, T.J, (1987), "Estimating Hedonic Demand Parameters with Single Market Data: The Problems Caused by Unobserved Tastes", *Review of Economics and Statistics,* 69(1), 178-180

Bayer, P, R. McMillan, F. Ferreira, (2007), "A unified framework for measuring preferences for schools and neighborhoods", *Journal of Political Economy,* 115(4), 588-638.

Bayer, P, R. McMillan, A.Murphy and C. Timmins, (2011), “A Dynamic Model of Demand for Houses and Neighborhoods”, *NBER Working Paper*, no 17250, NBER, Cambridge Mass.

Black, S., (1999), “Do Better Schools Matter? Parental Valuation of Elementary Education”, *Quarterly Journal of Economics,* 114, 577-599.

Bin, O., J.B. Kruse, and C.E. Landry, (2008), “Flood Hazards, Insurance Rates, and Amenities: Evidence from the Coastal Housing Market”, *Journal of Risk and Insurance,* 75(1), 1-252

Bin, O. and C.E. Landry, (2013), “Changes in Implicit Flood Risk Premiums: Empirical Evidence from the Housing Market”, *Journal of Environmental Economics and Management,* 65(3), 361-376.

Bockarjova, M., P. Geurts, M. Oosterhaven, A. van der Veen, 2010: Mag het wat kosten?, in: H. van der Most, S. de Wit, B. Broekhans, W. Roos, Kijk op waterveiligheid (Uitgeverij Eburon, Delft), pp. 56-73.

Botzen, W.J.W., J.C.J.H. Aerts, and J.C.J.M Van den Bergh, (2009). “Willingness of homeowners to mitigate climate risk through insurance” *Ecological Economics, 68*(8-9), 2265-2277.

Brookshire, D., M. Thayer, J. Tschirhart, and W. Schulze. 1985. “A Test of Expected Utility Model: Evidence from Earthquake Risks” , *Journal of Political Economy,* 93, 369-89

Brown, J.N. and H.S. Rosen, (1982), “On the Estimation of Structural Hedonic Price Models”, *Econometrica,* 50(3), 765-768.

CBS, (2012), Welvaart in Nederland. Inkomen, vermogen en bestedingen van huishoudens en personen. *Central Bureau of Statistics*, The Netherlands.

Chay, K. and M. Greenstone, (2005), “Does Air Quality Matter? Evidence from the Housing Market,” *Journal of Political Economy*, 113(2). 1121–1167.

COELO Woonlastenmonitor, (2011), COELO, University of Groningen, The Netherlands.

Cropper, M. L., L.B. Deck, and K. E. McConnell, (1988), "On the Choice of Functional Form for Hedonic Price Functions," *Review of Economics and Statistics*, 70(4), 668-675.

Deacon, R., C. Kolstad, A. Kneese, D. Brookshire, D. Scrogin, A. Fisher, M. Ward, K. Smith, and J. Wilen, (1998), "Research Trends and Opportunities in Environmental and Natural Resource Economics", *Environmental and Resource Economics*, 11(3), 383-397.

Deltacommissie, (2008). Working Together with Water: A Living Land Builds for its Future, Findings of the Deltacommissie. *Hollandia printing*, Den Haag.

Ekeland, I., J. J. Heckman, and L. Nesheim, (2004), "Identification and Estimation of Hedonic Models," *Journal of Political Economy*, 112(1), S60–S109.

Epple, D., (1987), "Hedonic Prices and Implicit Markets: Estimating Demand and Supply Functions for Differentiated Products," *Journal of Political Economy*, 95(1), 59–80.

Freeman, A. Myrick III, (1974), "On estimating air pollution control benefits from land value studies", *Journal of Environmental Economics and Management*, 1, 74-83.

Gabaix, X. (2012), "Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance", *Quarterly Journal of Economics*, 127(2), 645-700.

Gallagher, J. (2014), "Learning About an Infrequent Event: Evidence from Flood Insurance Take-Up in the US." *American Economic Journal: Applied Economics*, 6(3), 206-233

Gibbons, S, S. Machin, and O. Silva, (2013), "Valuing School Quality Using Boundary Discontinuities", *Journal of Urban Economics*, 75, 12-28.

Greenstone, M and J. Gallagher, (2008), "Does Hazardous Waste Matter? Evidence from the Housing Market and the Superfund Program," *Quarterly Journal of Economics*, 123(3), 951-1003

Hallstrom D. G. and V. K. Smith (2005), “Market Responses to Hurricanes*”, Journal of Environmental Economics and Management,* 50, 541–562.

Heckman, J. R.M. Matzkin and L. Nesheim, (2010), “Nonparametric Identification of Nonadditive Hedonic Models,” *Econometrica*, 78(5), 1569–1591. .

Inspectie Verkeer en Waterstaat, (2011), “Derde toets primaire waterkeringen. Landelijke toets 2006-2011”, *Ministry of Transport, Public Works and Water Management*, The Netherlands.

Lee, D.S., and Th. Lemieux*, (*2010), “Regression Discontinuity Designs in Economics." *Journal of Economic Literature*, 48(2), 281-355.

MacDonald, D. N., H. L. White, P.M. Taube, and W. L. Huth, (1990), “Flood Hazard Pricing and Insurance Premium Differentials: Evidence From the Housing Market”, *Journal of Risk and Insurance* 57(4): 654-63

Ries, J. and T. Somerville, (2010), “School Quality and Residential Property Values: Evidence from Vancouver Rezoning”, *The Review of Economics and Statistics*, 92(4), 928-944.

Rosen, S., (1974), “Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition,” *Journal of Political Economy*, 82(1), 34–55

Troy, A. and J. Romm, (2004), “The Role of Disclosure in the flood zone: Assessing the price effects of the California Natural Hazard Disclosure Law (AB 1195)”, *Journal of Environmental Planning and Management.* 47(1), 137-162.

TNS Nipo, (2006), Risicoperceptie bij overstromingen in relatie tot evacuatiebereidheid, Amsterdam.

De Vries, E.J., (1998), “Vergoeding van Rampschade”, *Nederlands Juristenblad*, 42, 1908-1915.

Water Act, (2009), *Ministry of Transport, Public Works and Water Management*, The Netherlands.

Watermonitor, (2009), “Inzicht in waterbewustzijn van burgers en draagvlak voor beleid”, Intomart GfK bv, projectnummer 22265.

Weitzman M., (2009), “On Modeling and Interpreting the Economics of Catastrophic Climate Change”. *The Review of Economics and Statistics*, 91(1), 1-19.

Van den Hurk, B., Klein Tank, A., Lenderink, G., van Ulden, A., van Oldenborgh, G.J.,Katsman, C., van den Brink, H., Bessembinder, J., Hazeleger, W., Drijfhout, S., (2006). KNMI Climate Change Scenarios 2006 for the Netherlands. *KNMI Scientific Report WR 2006-01*. KNMI, De Bilt, The Netherlands.

Van Winsum-Westra, M., A. Buijs, M. de Groot, (2010), “Pilot: Tevreden met hoogwaterbescherming? Een studie naar de tevredenheid met hoogwaterbescherming onder de bevolking”, *Alterra-rapport 2051*, Wageningen, The Netherlands.

Van der Woude, A.M., (1972) “Het Noorderkwartier. Een regionaal historisch onderzoek in de demografische en economische geschiedenis van westelijk Nederland van de late middeleeuwen tot het begin van de negentiende eeuw”, *A.A.G. Bijdragen*, 16.

Veiligheid Nederland in Kaart, (2014). *Ministry of Transport, Public Works and Water Management*, The Netherlands.

World Risk Report, (2014), *United Nations University - Institute for Environment and Human Security* and *Bündnis Entwicklung Hilft*.

Appendix A

Figure A1. Variation in the extent of flood risk around the city of Dordrecht

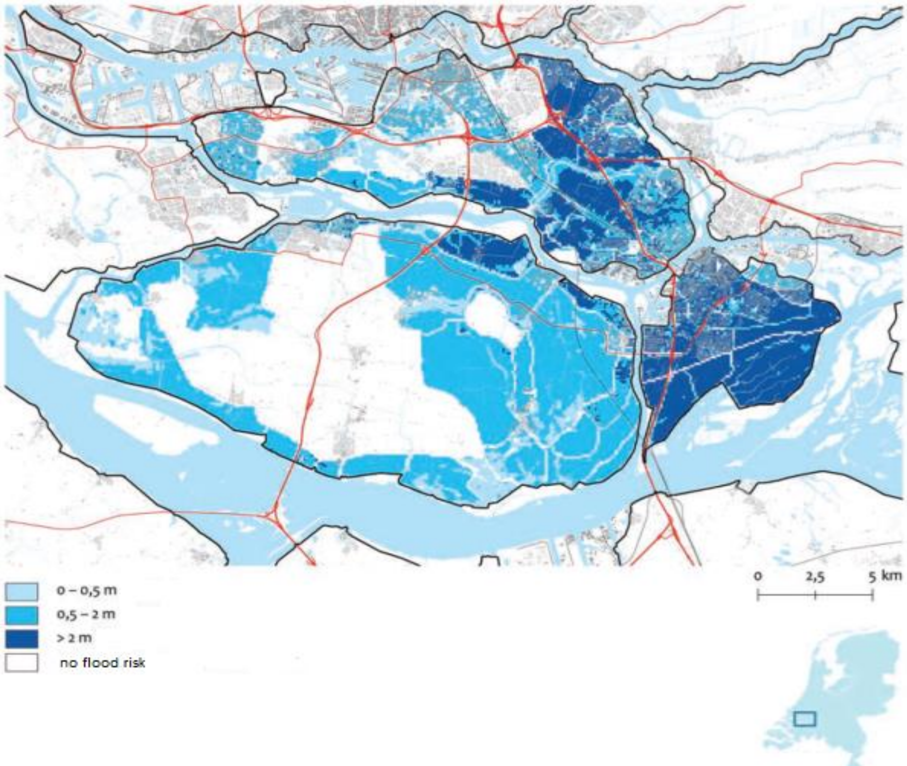

Notes: Figure taken from PBL 2009 "Overstromingsrisicozonering in Nederland"

Table A1. Robustness to the cutoff used to restrict the sample to sales of an entire property only

	Cutoff to identify sale of entire property			
	2.5% town	5% neighborhood	no cutoff	no cutoff
	1	2	3	4
Flood Risk	-0.008	-0.009	-0.005	-0.008
	[0.005]	[0.002]	[0.104]	[0.042]
house characteristics	yes	yes	yes	no
6ppc - amenities	yes	yes	yes	no
6ppc - water	yes	yes	yes	no
FE	5ppc/yr	5ppc/yr	5ppc/yr	5ppc/yr
< X m floodline	<100m	<100m	<100m	<100m
nr.obs	43392	42770	43921	44168

Notes: p-values based on standard errors clustered at the 6PPC level in brackets. See Appendix **B** for a detailed description of all control variables observed at the 6PPC level that we include in our regressions.

As explained in section 3.2.1., our data set contains some unrealistically low house prices due to the fact that the Kadaster records the sale of a piece of residential property as if it were the sale of an entire house. In our baseline results, we deal with this by excluding those 6PPC areas from our analysis that report a median house price below the 5pct-quantile of all median 6PPC house prices in the same town that the 6PPC is part of recorded in the same year. In column 1 of Table A1 we use the 2.5pct quantile as cutoff instead. And, in column 2, we use the 5pct-quantile of all 6PPC house prices in each of the 2345 districts (“wijken”) into which the towns of the Netherlands are divided. Finally, in column 3 and 4, we use no cutoff at all, and simply also include these “unrealistically low” house prices in the analysis. Using a different cutoff leaves our main result unaffected. When using no cutoff we find a less precisely estimated flood risk discount. This is especially so when also controlling for the other observed house price determinants in our data set (see column 3), which is not surprising given that these controls are included as house price determinants, and not as garage, piece of garden, boathouse, or any other “piece of residential property”-price determinants. Taking account of the presence of sales of “pieces of property” by excluding the lowest prices in our sample avoids this unnecessary noise.

Appendix B. Data sources and overview of all control variables

Controls variables		"Bad controls" (see discussion in section 4.2)
% monumental houses [ix]	**Distance to (m) & dummy within 25 meters of: [iv]**	**% inhabitants reporting neighborhood problems with: [v]**
monumental building? [ix]	rail	alcohol
% buildings non houses [vi]	road	drugs
iconic building? [ix]	airport	bicicyle theft
museum? [viii]	residential area	graffiti etc
theatre? [vii & x]	shops/restaurants	harassment
ln area (m2) [iv]	public facilities	loitering
ln built-up area (m2) [iv]	socio-cultural facilities	car content theft
ln median floor space (m2) [iii]	business area	littering
elevation (m) [i]	dumping-ground	burglary
distance to town centre: [ii]	scrapyard (cars, boats)	damaged public facilities
0 - 1km	cemetery	
1 - 2km	nat.resource extraction	**(less complete) data from the**
2 - 3km	building site	**Dutch Central Bureau of Statistics [iv]**
3 - 4km	other semi-paved area	# houses
4 - 5km	park	address density
5 - 10km	sports-ground	# inhabitants
home type PC6: [vi]	kitchen-garden	% non-western foreign born
% apartments	recreational area	% 0-14yr olds
% detached homes	residential recreational area	% 15-24yr olds
% terraced house	greenhouse horticulture	% 25-44yr olds
% half house	agricultural land	% 45-64yr olds
% terraced house (corner)	forest	% 65+ olds
% houses built in period: [iii]	open plain (dry)	# men
< 1549	national border	# women
1550 - 1749	*water-related:*	% 1pp households (hh)
1750 - 1849	open plain (wet)	% >1pp hh (no children)
1850 - 1879	IJsselmeer/Markermeer	% >1pp hh (children)
1880 - 1899	estatuary (closed)	mean hh size
1900 - 1909	Rhine / Meuse rivers	mean hh income rel. to Dutch average
1910 - 1919	lake	
1920 - 1929	storage basin / watershed	
1930 - 1939	recreational inland waterways	
1940 - 1949	water for nat.resource extraction	
1950 - 1959	mud flat	
1960 - 1969	other water	
1970 - 1979	Waddenzee, Eems, Dollard.	
1980 - 1989	Oosterschelde	
1990 - 1999	Westerschelde	
2000 - 2011	Noordzee	

Notes: Roman numerals behind a variable name or bold face category of variables refers to the data source for this variable. See the list on the next page for the different sources.

Data sources

i. Actueel Hoogtebestand Nederland (AHN, www.ahn.nl)

ii. Atlas voor gemeenten (www.atlasvoorgemeenten.nl)

iii. Basis Administratie Gemeenten (BAG)

iv. Centraal Bureau voor de Statistiek (CBS)

v. Integrale Veiligheidsmonitor Rijk (IVR)

vi. Kadaster

vii. Nederlands Uitburo (NUB)

viii. Nederlandse Museumvereniging (NMV)

ix. Rijksdienst voor het Cultureel Erfgoed (RCE)

x. Vereniging voor Schouwburg- en Concertgebouwdirecteuren (VSCD)