Estimating the flood discount: Evidence from a one-off national information shock[*]

Tom Gillespie[†] Ronan C. Lyons[‡] Thomas K.J. McDermott[§]

June 2023

Abstract

Flood risk is the most pervasive and costly natural hazard globally. With significant increases in flood risk expected over coming decades, future exposure to flood risk and associated costs will depend heavily on how private consumption decisions respond to new information about risk. We exploit a one-off national information treatment in the form of the release in 2011, for the first time, of detailed flood risk maps for Ireland, to test the effect of new information about flood risk on housing prices across an entire national housing market. We combine rich dwelling-level information on over 475,000 dwellings for the period 2006-2015 with detailed official data relating to flood risk, events and defences. Our core finding is that information matters. The price of housing responded dramatically to the release of flood risk maps at the end of 2011, with the emergence of a 4% price discount for dwellings at risk of flooding.

Keywords: Flood Hazards, Hedonic Prices, Urban Planning, Information Updating, Risk Assessments. **JEL codes**: H54, Q54, R21

[*]We thank Eoin Dignam and Patrick Howard for excellent research assistance, and the following for their assistance with data and related queries: daft.ie, the Central Statistics Office, the Office of Public Works, particularly Mark Adamson, An Taisce, Ordinance Survey Ireland and Ireland's Environmental Protection Agency, who also provided generous financial support. For helpful comments received, we also thank Matthew Collins, John Cullinan, Stephen Hynes, Matt Kahn, Danny Norton, Stephen O'Neill and seminar participants at the University of Galway and Trinity College Dublin. McDermott & Gillespie gratefully acknowledge financial support from the Irish Research Council (project ID: GOIPD/2017/1147 & GOIPG/2016/876) and from Ireland's Environmental Protection Agency (project ID: 2021-CE-1058). The authors declare no relevant conflict of interest, and have provided details here of all financial support received for this research.

[†]Centre for Economic Research on Inclusivity and Sustainability (CERIS) and Discipline of Economics, University of Galway.

[‡]Department of Economics and TRiSS, Trinity College Dublin, and Spatial & Economics Research Centre, LSE.

[§]Centre for Economic Research on Inclusivity and Sustainability (CERIS) and Discipline of Economics, University of Galway and corresponding author; email: thomas.mcdermott@universityofgalway.ie.

1 Introduction

Flood risk is the most pervasive and costly natural hazard, with an estimated one billion people in 155 countries exposed worldwide (European Commission Joint Research Centre, 2017). With the prospect of rising sea levels and more intense rainfall events due to climate change, flood risk is expected to increase in many locations over coming decades. Projections of the future costs of flooding depend not only on the risk of flood events but also on societies' exposure to those events (Hallegatte et al., 2013). This underlines the importance of the extent to which flood risk is taken into account in private decisions, especially where costs are borne, at least in part, by taxpayers (Lin et al., 2021).

Due to the immobile nature of real estate and its prevalence in the typical household's balance sheet, the housing market represents a unique and important window into how private actions reflect flood risk. Theory would suggest a price discount for dwellings at risk of flooding, given the associated costs. However, individual households may lack good information on flood risk (Bakkensen & Ma, 2020), and there is also the issue of moral hazard resulting from various forms of government interventions to protect households from disaster (Kydland & Prescott, 1977). In short, where market signals are weak, there may be a tendency towards over-exposure to flood risk.

There is a long literature estimating the effect of flood risk on housing prices, hereafter the *flood discount*. A recent review finds widely varying results, with estimates of the price effect ranging from −75% to +61% (Beltrán et al., 2018). Early contributions compared the value of dwellings within flood risk zones to those elsewhere, controlling for a range of dwelling attributes (e.g. Bin, Crawford, et al., 2008; Bin, Kruse, et al., 2008; MacDonald et al., 1990). However, interpreting the results of such studies as causal depends on the strong assumption that hazard risk is exogenous, conditional on other observable determinants of housing prices. As an alternative, many recent papers in this literature have exploited the more plausibly exogenous timing of particular flood events to estimate how flood discounts vary in response to the occurrence of flooding. A common finding in these studies is that there are significant discounts after flood events, which fade over time; see, for example, Atreya et al. (2013), Beltrán et al. (2019), Bin and Landry (2013), Bin and Polasky (2004), Gibson et al. (2017), Ortega and Taşpınar (2018), and Timar et al. (2018). As noted by Bosker et al. (2019), this strand of the literature estimates changes in households' risk perceptions following a recent flood event, rather than directly identifying their level of risk perceptions. There is also the possibility that these estimates reflect, at least partly, the damages from flooding, or indeed other impacts of flood events on local property markets.

A number of recent papers have turned to information treatments to recover flood discount estimates (e.g. Gibson & Mullins, 2020; Hino & Burke, 2020; Hsieh, 2021; Seo et al., 2021; Shr & Zipp, 2019; Votsis & Perrels, 2016). Several of these papers use updates to the U.S. Federal Emergency Management Agency (FEMA) flood risk maps as the basis of their estimates, but the FEMA maps are problematic on a number of levels. First, the rolling nature of updates to FEMA maps potentially creates a staggered treatment problem. It is now well-known in the applied econometrics literature that two-way fixed-effects (TWFE) estimations of difference-in-differences coefficients can lead to substantial biases when there are staggered treatment timing and heterogeneous/time-varying treatment effects (e.g. Baker et al., 2022; Borusyak et al., 2021; Callaway & Sant'Anna, 2021; De Chaisemartin & d'Haultfoeuille, 2020; Goodman-Bacon, 2021; Sun & Abraham, 2021).

Secondly, the process by which FEMA maps are updated creates endogeneity concerns, given that map updates are explicitly prioritised for areas with rapid housing development (Davlasheridze et al., 2017). Concerns have also been raised that the mapping process is amenable to local pressures and political influence (see, for example, Pralle, 2019). Papers exploiting map updates from other jurisdictions have typically been based on relatively small samples (e.g. Votsis & Perrels, 2016, using Finnish data) and/or found no information effect (e.g. Hsieh, 2021; Seo et al., 2021, for Taiwan and South Korea respectively).

In contrast to these existing studies, our paper exploits a one-off national information treatment in the form of the release in late 2011, for the first time, of detailed flood risk maps for Ireland. While our baseline specification includes a suite of property-level characteristics and highly localised spatial (and spatio-temporal) fixed effects, one may still be concerned that estimates of the flood discount could be confounded by unobservables that happen to correlate with flood risk in the cross-section.

To counter this concern we employ a difference-in-differences (DiD) style analysis exploiting the release in 2011 of detailed flood risk maps in Ireland. Our estimate of the flood discount, based on the DiD estimator, is -4%. We show that this estimate is robust to varying the set of included covariates, to varying the level of the spatial fixed effects, and that the flood discount only appears after the release of the risk information in 2011. We also show that prior to the information treatment, properties at risk of flooding and properties not at risk followed similar price trends in our data. Lastly, we show that our results are not driven by the use of listed prices: for a matched sample, switching between transactions and listing prices as the dependent variable has no significant effect on the estimated flood discount.

Hedonic price estimates reveal the marginal willingness to pay (WTP) of the marginal buyer (Greenstone, 2017; Kuminoff & Pope, 2014; Pinchbeck et al., 2021). Given the potential for selection into risk based on idiosyncratic household preferences and budget constraints (sorting), we interpret our estimates cautiously as "local" treatment effects. We return to this point in our discussion in Section 6.

Our analysis is based on a large and unusually rich dataset, with nationwide coverage of the Irish property market over a long time period (2010-2015 for transactions, 2006-2015 for listings). The richness of our data enables us to control for important covariates that are typically excluded from similar studies, including a sophisticated measure of sea views, independently assessed building quality, and time-varying and spatially explicit measures of flood events and areas protected by flood defences.

Most closely related to our paper in setting is Pilla et al. (2019), who compare directly the effects of assessed risk and a large flood event for the case of Dublin, Ireland, after severe flooding in 2011. They find evidence that flood events had a bigger impact on housing prices than assessed flood risk. This reinforces the idea that actors in housing markets are not always well informed about flood risk. Pilla et al. (2019), however, is limited to cross-sectional analysis and uses a more limited set of control variables, which presents challenges to a causal interpretation of their results.

The rest of the paper is structured as follows. The next two sections describe the local context and the data used in our analysis, in particular the data relating to flood risk, flood defences and flood events. Section 4 outlines our empirical framework and our rationale for interpreting the analysis as causal. The results of our empirical analysis are presented in Section 5, before the final section concludes.

2 Context

Flood risk is an important public policy issue in Ireland, and a cause for concern amongst the general public.[1] Most of Ireland's urban areas are coastal, and the country regularly experiences flooding events along major rivers and in coastal areas, as a result of heavy rainfall and winter storms. These floods have been costly, with roughly €1bn, or close to €800 per household, in insured losses over the period 2000-2014.[2] Moreover, the Irish government has committed to spending large sums on flood relief schemes: 68 flood relief schemes included in our analysis cost €226.6 million in total, with an additional €1 billion of planned public expenditure, or roughly 0.5% of national income, on flood relief schemes over the next 10 years (Office of Public Works, 2018).[3]

Prior to 2011, there were no detailed maps available on flood risk in Ireland. This changed with the release in August 2011 of Preliminary Flood Risk Assessment (PFRA) maps for the entire country. This followed a period of investment in detailed flood risk assessment by the Irish government, to comply with EU Directive 2007/60/EC, which requires all member states to assess and manage flood risk. The Preliminary Flood Risk Assessment (PFRA) maps were published in August 2011 and made widely available by early 2012 with the launch of myplan.ie, a central repository for spatial information related to planning.

Importantly for our purposes, these new risk maps represent a relatively clean information treatment. In particular, the maps are not used to inform flood insurance decisions. In Ireland, flood insurance is provided by the private insurance industry without any government involvement and is usually bundled with general household insurance. As part of the government's investment in flood risk mapping, an agreement was made with the insurance industry that the new risk maps would not be used as the basis for assessing risk for insurance purposes. Private insurers use their own independent risk assessments to price flood risk, and these did not change with the information treatment.

The maps that we use are the official primary source of flood risk information for the entire country and have been made readily accessible to the general public via online interactive mapping tools. Given their role in helping prioritize investment in flood defences, we control in our analysis for flood defences using detailed time-varying and spatially precise data on properties protected. The official risk maps were updated with more detailed risk maps for 300 Areas of Further Assessment starting in 2016; for this reason, we restrict our analysis to the period up to the end of 2015.

One additional piece of context is in relation to the Irish housing market during the period of our study. We start our analysis in 2006, when national online listings become available – but it is worth noting the dramatic change in housing market conditions between 2007 and 2012, when prices fell by over 50% on average, followed by growth

[1]As part of this research, we also conducted an online survey of public attitudes to flood risk, which found that the general public in Ireland is concerned about flooding, that those concerns have increased for many over the last 10 years, and that a large majority of people expect the problem to get worse in the coming decades. Further details about the survey available on request.

[2]Information provided directly from Insurance Ireland (the representative body for private insurance companies operating in Ireland).

[3]The stated intention is that these schemes will provide protection to 80% of the 34,500 dwellings in Ireland assessed as having a 1% chance of experiencing a significant flood event in any year. In scaling by national income, the measure used in Ireland is modified Gross National Income (GNI*), which was valued at €197.5 billion in 2018, according to data from the Central Statistics Office (CSO), available here https://www.cso.ie/en/releasesandpublications/ep/p-nie/nie2018/ (last accessed in June 2023).

again in prices from 2013. Our method focuses on differences, not levels, and in all our analysis we control for detailed region-specific trends, as detailed further in Section 4 below.

3 Data

3.1 Flood data

The principal source of information related to flooding in Ireland is the Office of Public Works (OPW), the agency with responsibility for flood risk management in Ireland. The OPW provided us with data on scientifically assessed risk and on flood defences, as detailed below.

Flood risk The flood risk maps that we use in our analysis come in the form of high-resolution polygons depicting both fluvial and coastal flood risk zones for the entire country of Ireland. The maps show areas with at least a 1% risk of flooding per year for fluvial flood risk and at least a 0.5% risk per year from coastal flooding. In our analysis we combine the coastal and fluvial maps together and generate an indicator variable for whether or not a property is inside either of these risk categories, which we label $Risk$.

Flood defences The OPW provided us with polygon data related to 68 existing flood defence schemes completed between 1996 and 2017. Attributes include the date of completion, spatial extent of protection, whether the defence was permanent or demountable, and the cost of each scheme. To ensure that our estimate of the flood discount is based on dwellings at risk of flooding and not protected by flood defences at the time of their listing or sale, we distinguish dwellings at risk of flooding and protected by flood defences from dwellings at risk of flooding but not protected by flood defences.

Flood events We also control for past flood events, using an extensive archive of historical flooding in Ireland, compiled by the OPW. Information in the archive is drawn from various sources, such as reports by local authorities, engineers' reports, newspaper articles, and photos. We extract location and timing information from this archive. The dataset contains a total of 1,947 dated flood points and 84 dated flood polygons dating from 1763 to 2015. We construct indicators for dwellings affected based on the distance of the dwelling from the most recent flood event: dwellings within 100m of a flood event (or within a flood event polygon), and dwellings between 100m and 250m from a flood event. The time since the most recent flood within a 100 (or 250) meter radius of a dwelling was modelled as a categorical variable of which there were four categories: >30 years since a flood, 10-30 years, 2-10 years, and <2 years, with a base of no flood event recorded.

3.2 Dwelling & Location data

Our baseline housing dataset is drawn from the Daft.ie National Listings Database, the archive of sale listings from the leading real estate website in Ireland. The daft.ie listings dataset is both long, covering the period from 2006Q1, at the height of a real estate boom, through to 2015Q4, and broad, covering the national market in its entirety.

It is also deep, with an estimated coverage of over 95% of all listings in the Irish market, and rich, in terms of the information available for each listing. This includes both structured information on dwelling attributes (including: property type; bedrooms; bathrooms; and building energy rating, BER) and the text of the ad, which is mined for a range of variables, including site size and orientation.

In addition to dwelling attributes, the presence of information on dwelling location allows for individual dwellings to be assigned to "at risk" of flooding, as described above, but also for the calculation of a variety of other location-specific amenities. These include nearest city centre, transport facilities, schools, and natural amenities. Based on Gillespie et al. (2023), we also include a location-specific measure of sea-view breadth and depth, unique in the flood discount literature, the omission of which may be a source of downward bias in existing flood discount estimates.

We exclude dwellings with prices above €2,000,000 or below €30,000, as these are either atypical properties or an error in the data. We also exclude dwellings listed as having fewer than one or more than five bedrooms, and fewer than one or more than eight bathrooms, for similar reasons. After applying size and price filters, we have a sample of 475,436 residential sale listings between 2006Q1 and 2015Q4. The data are illustrated in Appendix Figures A1 and A2. The breakdown of observations by risk and timing (pre- and post-information treatment) is provided in Table 1. Further detail on the data are included in the Data Appendix and in Tables A1 and A2.

	No. of Listings	% of pre/post sample
Pre-treatment:		
Outside PFRA (not at risk)	294,760	92.8
Inside PFRA (at risk)	22,862	7.2
Defended	84	0.03
Post-treatment:		
Outside PFRA (not at risk)	145,721	92.4
Inside PFRA (at risk)	11,330	7.2
Defended	679	0.4
Total	475,436	100.00

Table 1: Listings sample (2006-2015) by treatment category

Additional analysis is performed on a dataset of transactions matched to listings, using Ireland's Property Price Register (PPR), the official listing of residential real estate transactions in Ireland. This dataset has several major drawbacks for our purposes: it begins in January 2010, shortly before the information treatment, restricting the set of pre-treatment observations; the PPR data contain no dwelling attributes other than price, address and date of transaction; and finally, there is the issue of non-unique addresses in Ireland (prior to the introduction of postcodes, known in Ireland as Eircodes, in 2015). For these reasons, we use the PPR data as an additional robustness check, rather than as our main source of housing data. To do so, we match observations in the PPR data with listings where possible. The total number of transactions matched to listings was 45,161.[4]

[4]Further details on our housing data and the matching process between transactions and listings are

4 Empirical Strategy

The standard methodology in this literature builds on Rosen's (1974) theoretical framework of hedonic prices; a recent review of best practice in using hedonic property value models is given in Bishop et al. (2019). Conceptually, the value of a dwelling takes the following form:

$$Price = f(S, L, F) + \epsilon \tag{1}$$

where S are structural characteristics (bedrooms, bathrooms, garden etc.), L are location and environmental characteristics (amenities, neighborhood etc.), and F are flood-related variables. The dwelling price is thus a function of all the attributes relating to the dwelling and the resulting coefficients are the implicit marginal prices of the attributes.

In our context, the key identification challenge with estimating such a model is that even with a rich set of controls to capture S and L, flood risk is not randomly assigned and therefore may be correlated with unobservables. To counter this concern we employ a difference-in-difference (DiD) style analysis exploiting the release in 2011 of detailed flood risk maps in Ireland, as follows: each listing is assigned to one of four categories, based on $Risk$ and the timing of the listing (pre- and post-information treatment).[5] We then estimate the following regression:

$$\begin{aligned} log(price)_i = \beta_0 + \beta_1 Risk * pre_i + \beta_2 NoRisk * post_i + \beta_3 Risk * post_i \\ + \beta_4 X_i' + \beta_5 Z_i' + \mu_t + \lambda_{ed,t} + \epsilon_i \end{aligned} \tag{2}$$

where the omitted risk category (base) is $NoRisk * pre$ (dwellings not located in flood risk zones and listed pre-information treatment), X_i' is a vector of dwelling-specific attributes (structural characteristics), and Z_i' are location-specific amenities. The $\lambda_{ed,t}$ are location fixed effects at the level of Ireland's official electoral districts (EDs), which represents a very fine spatial disaggregation of the data. In our main specifications these are included flexibly as ED-by-year fixed effects. The μ_t account for within-year quarterly trends at the national level, to capture variation in the Irish housing market over time. Robust standard errors ϵ_i in all specifications are clustered at the ED level. The estimate of the flood discount, conditional on information, is then the difference-in-differences: $(\beta_3 - \beta_2) - \beta_1$.

A number of identifying assumptions are required in interpreting the estimates of Equation (2). First, the treatment should not be correlated with unobservables. While flood risk is potentially correlated with many location-based factors that affect property values, e.g. coastal amenities, these do not change with the treatment in our context, given that our treatment refers to both the timing and spatial extent of risk information. Our estimate is still subject to bias, however, if treated and control dwellings experience differential observable or unobservable trends. To test for parallel trends, we estimate versions of Equation (2) separately for dwellings inside and outside of the official flood risk zones, omitting the $Risk$ categories from the regression. In Figure 1 we report the coefficients on the μ_{qt} year-quarter dummies for each sub-sample. The

given in the Data Appendix.

[5] In practice we have two additional categories, for properties in official risk zones but protected by flood defences, listed either before or after the information treatment. These categories are included in the estimating equation, but not used in the calculation of the flood discount.

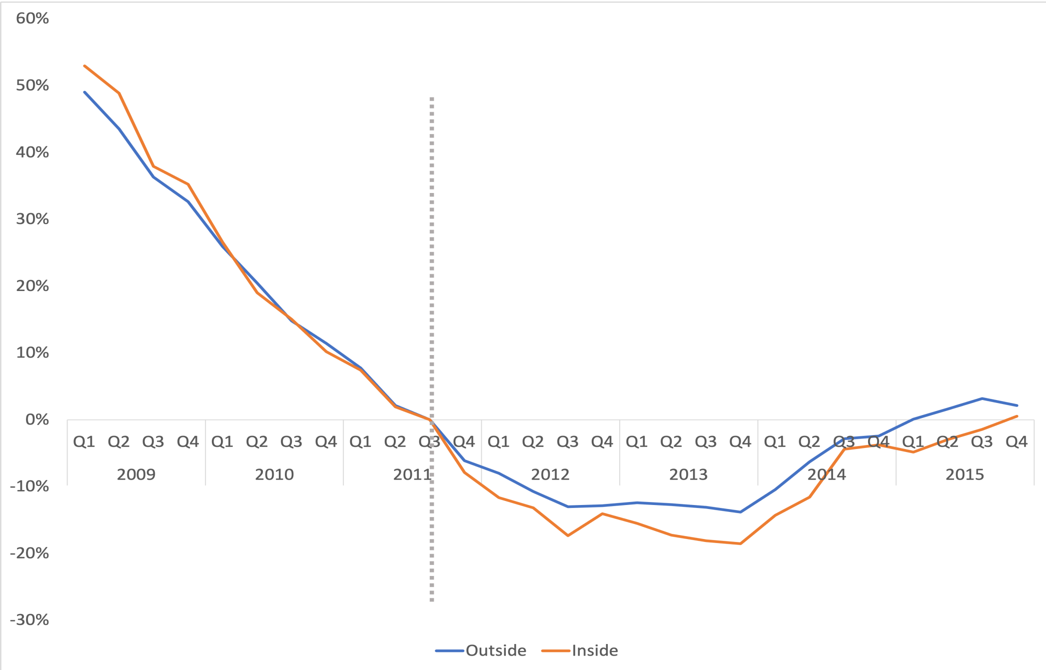

Figure 1: Parallel trends

Note: The figure shows a mix-adjusted price index over time, constructed from the coefficients on the year-quarter dummies from a regression similar to equation (2), omitting the flood risk indicators, and estimated separately for properties inside and outside the PFRA flood risk zones. The dashed vertical line indicates the timing of the release of the flood risk maps in August 2011.

results suggest that at-risk and not-at-risk dwellings followed similar price trends in the years immediately prior to the release of the risk information in 2011, with a gap emerging between the two series after the release of the flood risk maps and persisting thereafter for the remainder of our sample period.

Further reassurance on the exogeneity of our treatment is provided by additional analysis that shows robustness of our estimate of the flood discount to varying the set of included covariates and to variations in the level of the spatial fixed effects (reported in Section 5 below). Our baseline specification also allows for unobservable trends at a local level: In all specifications, we include very local trends ($\lambda_{ed,t}$). While many papers assume a national or city-level trend – we allow for each of over 3,000 EDs to have their own path in housing prices (or experience their own individual annual shocks). We also separately control for within-year variation in prices at the national level (by including quarterly fixed effects).

A further identifying assumption required in our context is that the treatment is relevant – both in the sense that buyers and sellers are aware of the risk information and use it in their decisions (once the maps become available), and that the risk as captured by the maps represents the actual risk experienced by households (or market participants). In terms of the former assumption – awareness – we are confident that market participants had access and motivation to use the risk information, as appears to be borne out in our empirical results.[6] As for the latter assumption – that the maps represent actual risk – given their basis in scientific risk assessment, it seems reasonable

[6] Access to the risk maps was made relatively easy with their inclusion in the online platform myplan.ie.

to assume the maps are a good guide to the true risk. However, it is of course possible that households are insulated from the consequences of this risk, for example through various forms of government intervention, including subsidised insurance, compensation schemes and flood protection measures, as is common in the context of flood risk around the world. In our setting, flood insurance is provided entirely by the private sector in Ireland with no government involvement. Subsidised insurance is therefore not relevant. We also control for existing flood defence investments in our analysis. However, we cannot entirely rule out the possibility that households anticipate future flood defence investments or compensation from government in the event of experiencing flooding. To the extent that either of these are prevalent in our context, they would represent an attenuation bias of our estimates of the flood discount, meaning our results understate the true effect.

Finally, intrepreting our estimates as average WTP would require further assumptions – in particular that there is no sorting based on risk (Greenstone, 2017; Kuminoff & Pope, 2014; Pinchbeck et al., 2021). We return to this point in the discussion of our findings in Section 6.

5 Results

The results of our baseline estimate of the flood discount are reported in Table 2. The specification estimated here includes the full suite of controls – dwelling-specific attributes and location-based amenities, as well as ED-by-year fixed effects and national-by-quarter fixed effects, and is estimated on the full listings sample 2006-2015. The results in Column (1) show that prior to the release of the flood risk maps in 2011, there was no significant price discount for dwellings located in flood risk zones. In Column (2) we see that after the release of the risk maps, dwellings in flood risk zones attract a 3.5% price discount, relative to dwellings outside the flood risk zones. Finally, taking the difference in the differences (in Column 3) we arrive at our headline estimate of the flood discount of -4%.

The $Risk$ variable is defined here as $= 1$ for dwellings located within one of the flood risk zones defined in the official PFRA risk maps released in 2011, and 0 otherwise. However, it is possible that the information treatment could also have effects on neighboring properties. For example, if demand is displaced by the new risk information, it's possible that this could push up prices for neighboring properties. On the other hand, the risk information might spillover to neighboring properties if buyers are concerned about being located near to places with flood risk.

In additional specifications, reported in Table A3, we test for these neighbor effects by allowing for buffers of either 100m or 250m around the flood risk zones as defined in the PFRA maps.[7] The results show that the estimated flood discount remains negative and statistically significant, albeit declining in magnitude, as we add 100m and 250m buffers, respectively to the definition of $Risk$. These results suggest that our baseline findings are not being driven by locally displaced demand. Instead, they are consistent with market participants treating flood risk as continuous across space, rather than a spatially discrete phenomenon defined by the boundaries of the mapped flood risk zones.

[7] For urban areas, a 100m buffer corresponds roughly to including dwellings on the next street adjacent to a flood risk zone, while a 250m buffer corresponds to including entire city blocks that are adjacent to flood risk zones. See the data illustrations presented in Figures A1 and A2.

	Pre-maps (1)	Post-maps (2)	(3)
$NoRisk$	$(Omitted)$	-0.142	
	-	(-19.4)	
$Risk$	0.005	-0.177	
	(0.705)	(-17.6)	
Difference	0.005	-0.035	-0.040
	(0.705)	(-4.8)	(-5.24)
Controls		Y	
Spatio-temporal FE		Y	
Year-Quarter FE		Y	
Obs.		475,436	
R-Sq		0.792	
RMSE		0.281	

Table 2: Baseline estimate

Note: The table shows the results of a regression estimate of equation (2) using the full sample of listings, with dwelling and location controls as described earlier; t-statistics in parentheses.

Tables 3 and 4 present the results of a set of robustness tests. In Table 3 we vary the set of included covariates – removing, in turn, sets of control variables that might be expected to correlate with flood risk in the cross section; flood defences, flood events, blue space variables, and amenities. In each case the estimated flood discount, based on the difference-in-difference style estimator, remains remarkably insensitive to the inclusion or exclusion of these groups of control variables. This exercise provides further reassurance that the empirical setup is capturing the effect of flood risk on property values, conditional on availability of flood risk maps, and not the effects of some other unobservables.

Similarly, in Table 4 we test the robustness of our findings to varying the level of the spatial fixed effects. In our baseline specification we include ED-by-year fixed effects. There are over 3,000 EDs in our data. With 10 years of data, this results in the inclusion of over 30,000 fixed effect units in our baseline specification. In Table 4, we experiment with specifications using both more and less granular definitions of the location fixed effects: markets (of which there are 54 in the data), micro-markets (nearly 400) and Census Small Areas (over 18,000). Again the results show that our estimate of the flood discount is remarkably robust to the level of these spatio-temporal controls.

Unsurprisingly, the estimated discount is attenuated (and less precise) in Column (4) of Table 4, where we use the most granular version of spatial fixed effects at the Census Small Area (SA) level. With over 18,000 SAs in our data, some of these are geographically very small (in some cases individual apartment blocks), meaning that the unit of the fixed effect and treatment may be correlated: for fixed effects to aid identification, it is required that there is variation within geographical units. Additionally, as the specification in Column (4) includes over 120,000 fixed effects units, leaving relatively few degrees of freedom to estimate the coefficients on the $Risk$ categories. Again, the robustness of the estimated flood discount to these variations provides additional reassurance that location-based unobservables – potentially correlated with flood risk

	Baseline (1)	Excl defences (2)	Excl events (3)	Excl blue space (4)	Excl amenities (5)
DiD	-0.040 (-5.24)	-0.043 (-5.57)	-0.040 (-5.22)	-0.040 (-5.11)	-0.041 (-5.30)
Controls	Y	Y	Y	Y	Y
Spatio-temporal FE	Y	Y	Y	Y	Y
Year-Quarter FE	Y	Y	Y	Y	Y
Obs.	475,436	475,436	475,436	475,436	475,454
R-Sq	0.792	0.792	0.792	0.792	0.792
RMSE	0.281	0.281	0.281	0.282	0.282

Table 3: Robustness (1) Varying the set of controls

Note: The table shows the results of regression estimates of equation (2) using the full sample of listings, with dwelling and location controls as described earlier, with the exception of the controls omitted as per the column headings; t-statistics in parentheses.

	Mkt*year (1)	MicroMkt*year (2)	ED*year (3)	SA*year (4)
DiD	-0.046 (-4.63)	-0.051 (-6.41)	-0.040 (-5.74)	-0.025 (-2.77)
Controls	Y	Y	Y	Y
Spatio-temporal FE	Y	Y	Y	Y
Year-Quarter FE	Y	Y	Y	Y
Obs.	475,494	475,405	475,436	475,436
R-Sq	0.712	0.753	0.792	0.86
RMSE	0.321	0.298	0.281	0.259
No. of FE units	540	3,878	30,036	121,700

Table 4: Robustness (2) Varying the level of the FE

Note: The table shows the results of regression estimates of equation (2) using the full sample of listings, with dwelling and location controls as described earlier; t-statistics in parentheses.

– are not affecting our estimates.

One might wonder if our estimate of the discount reflects risk perceptions or simply changes in the salience of flood risk – for example, due to media coverage of the new flood risk maps. Indeed, a common finding in the existing literature is a temporary spike in flood discounts following flood events, that fades over time. Firstly, our empirical specification includes past flood events and, due to flexible functional form, allows more recent events (and those closer to a property) to have a greater impact on prices than more distant ones. To the extent that any concerns about salience of flood risk beyond events is important, we estimate a version of Equation (2) where we interact $Risk$ with a set of annual dummies, and omit the pre- and post-information distinction. The coefficients (and 95% confidence intervals) from the $Risk * year$ interactions are displayed in Figure 2. As the figure shows, the estimated discount is indistinguishable from zero in the years prior to the release of the risk maps, turns negative and significant with the release of the maps (in 2011), and remains consistently

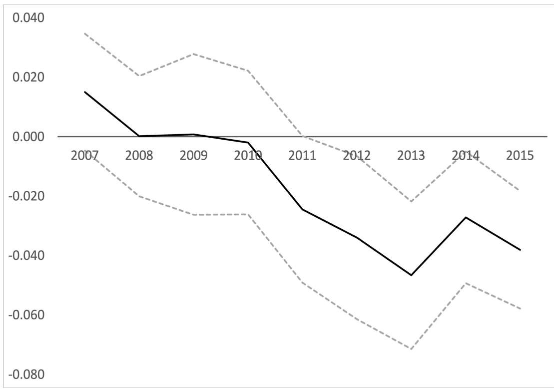

Figure 2: Annual estimate of the flood discount

The figure shows coefficient estimates and 95% confidence intervals from a regression similar to equation (2) but with $Risk$ interacted with year dummies. This regression uses the full sample of listings (2006-2015).

negative and of a similar magnitude for the remainder of our sample period (to the end of 2015). While we cannot rule out salience as a potential mechanism, the results of this exercise suggest that what we observe in the data is more than a transient spike in the salience of flood risk.

Finally, readers might be concerned about the use of list prices to estimate the flood discount. It could be, for example, that sellers of risky homes are optimistic about their property's value and list it for a price that does not reflect flood risk. Alternatively, it could be that agents are motivated to advise owners of risky properties to list for a relatively low price in order to achieve a quicker sale. In Table A4, we report the results of estimates based on a matched sample of listings and transactions. The two sets of results are based on an identical specification, estimated on the same set of dwellings, but swapping out (log) list prices in Column (1) for (log) transaction prices in Column (2). Given the relatively short pre-treatment period in the transaction price data (the PPR dataset), the coefficients reported here are the simple cross-sectional coefficients on $Risk$ in the post-information treatment period. The estimated flood discounts reported here are somewhat smaller than in our preferred specifications, likely reflecting an urban bias in the (smaller) matched transactions sample. More importantly, the results suggest that moving from list to transaction prices has little or no bearing on our estimate of the flood discount.

6 Discussion and Conclusions

In this paper, we examine the relationship between flood risk and housing market outcomes, using the case of Ireland in the decade around the release of flood risk information in 2011. In particular, we exploit rich housing data – including a dataset of over 475,000 sale listings 2006-2015, supplemented by an additional set of 45,000 transactions 2010-2015 – and detailed official data relating to flood risk, previous flood events, and completed flood defences. Our estimation strategy exploits a one-off national information treatment to estimate the flood discount. Our baseline estimate, using a difference-in-differences style estimator is -4%. We show that this estimate is robust to varying the set of included controls and to varying the level of our spatio-temporal fixed effects. We also find that the discount for at risk dwellings is not present prior to the release of the risk maps in 2011, emerges quickly after the information treatment and persists for the remainder of our sample period.

Overall, we believe that our findings represent the causal effect of flood risk on housing prices, conditional on the availability of flood risk information. Nonetheless, our results have limitations. First, given the potential for sorting into risk, as has been observed in other contexts (e.g. Bakkensen & Barrage, 2022; Bakkensen & Ma, 2020), we interpret our finding as a local treatment effect. The estimated flood discount represents the market price of flood risk, given the prevailing market conditions and population preferences at the time of the information treatment (Pinchbeck et al., 2021).

In terms of the magnitude of our estimate, a 4% discount would be consistent with damages per flood event in the region of €50,000 to €60,000, for reasonable parameter values.[8] That value would be in line with, albeit somewhat higher than, the (limited) available evidence on average flood costs for residential property in Ireland. This estimate is perhaps best thought of as an estimate of the welfare costs of living with flood risk – including potential damages to property, as well as more hidden costs of flooding, such as disruptions to daily life and mental health costs – for the marginal household that moves in to a flood risk zone.

It is also possible that the maps we use do not reflect the actual risk experienced by households (or market participants) for example to the extent that they are shielded from the consequences of flooding by government compensation schemes or investments in flood protection (including in the future). Our estimates also exclude other non-residential costs of flooding such as damage to public infrastructure or commercial real estate. On this basis, the true costs of flooding to society are likely larger than what we estimate here.

While our discussion and our estimation strategy emphasises the information treatment, we cannot fully rule out salience as an alternative mechanism, for example due to media coverage of the new flood risk maps. The evidence we present shows the information treatment coincided with the emergence for the first time of a flood discount, which persists for the remainder of our sample period. This suggests that the effect we observe is more than a transient spike in flood risk salience.

Our findings have important policy implications for flood risk management, insurance and flood defences, as well as for projections of future flood losses in a world of increasing flood risk. Perhaps most importantly for policymakers, our results present compelling evidence on the effectiveness of public investments in flood risk information provision. Better information results in more awareness and a clear price signal, which should translate into less exposure to flooding in future.

Finally, we show that, on a like-for-like basis, list prices act as a good proxy for transaction prices, where those are unavailable. Given the prevalence of flood risk in lower-income settings, where formal housing statistics are typically weaker, this is a useful finding for both researchers and policymakers.

[8]The flood discount should represent the capitalised value of future flows of expected flood damages, absent risk aversion. For a dwelling valued at €300,000, a risk of flooding per year of 1%, a time horizon of 30 years and a time discount rate of between 2% and 5%, a flood discount of 4% is consistent with damages per flood of €50,000 to €60,000.

References

Atreya, A., Ferreira, S., & Kriesel, W. (2013). Forgetting the flood? an analysis of the flood risk discount over time. *Land Economics*, *89*(4), 577–596.

Baker, A. C., Larcker, D. F., & Wang, C. C. (2022). How much should we trust staggered difference-in-differences estimates? *Journal of Financial Economics*, *144*(2), 370–395.

Bakkensen, L. A., & Barrage, L. (2022). Going underwater? flood risk belief heterogeneity and coastal home price dynamics. *The Review of Financial Studies*, *35*(8), 3666–3709.

Bakkensen, L. A., & Ma, L. (2020). Sorting over flood risk and implications for policy reform. *Journal of Environmental Economics and Management*, *104*, 102362.

Beltrán, A., Maddison, D., & Elliott, R. (2018). Is flood risk capitalised into property values? *Ecological Economics*, *146*(100), 668–685.

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

Bin, O., Crawford, T. W., Kruse, J. B., & Landry, C. E. (2008). Viewscapes and flood hazard: Coastal housing market response to amenities and risk. *Land economics*, *84*(3), 434–448.

Bin, O., Kruse, J. B., & Landry, C. E. (2008). Flood hazards, insurance rates, and amenities: Evidence from the coastal housing market. *Journal of Risk and Insurance*, *75*(1), 63–82.

Bin, O., & Landry, C. E. (2013). Changes in implicit flood risk premiums: Empirical evidence from the housing market. *Journal of Environmental Economics and management*, *65*(3), 361–376.

Bin, O., & Polasky, S. (2004). Effects of flood hazards on property values: Evidence before and after hurricane floyd. *Land Economics*, *80*(4), 490–500.

Bishop, K. C., Kuminoff, N. V., Banzhaf, H. S., Boyle, K., von Gravenitz, K., Pope, J. C., Smith, V. K., & Timmins, C. (2019). Best practices in using hedonic property value models for welfare measurement.

Borusyak, K., Jaravel, X., & Spiess, J. (2021). Revisiting event study designs: Robust and efficient estimation. *arXiv preprint arXiv:2108.12419*.

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

Callaway, B., & Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, *225*(2), 200–230.

Davlasheridze, M., Fisher-Vanden, K., & Klaiber, H. A. (2017). The effects of adaptation measures on hurricane induced property losses: Which fema investments have the highest returns? *Journal of Environmental Economics and Management*, *81*, 93–114.

De Chaisemartin, C., & d'Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, *110*(9), 2964–2996.

European Commission Joint Research Centre. (2017). *Atlas of the human planet 2017: Global exposure to natural hazards*. Publications Office. https://doi.org/10.2760/932505

Gibson, M., & Mullins, J. T. (2020). Climate risk and beliefs in new york floodplains. *Journal of the Association of Environmental and Resource Economists*, *7*(6), 1069–1111.

Gibson, M., Mullins, J. T., & Hill, A. (2017). Climate change, flood risk, and property values: Evidence from new york city. *Work. Pap., Dep. Econ., Williams Coll., Williamstown, MA*.

Gillespie, T., Hynes, S., & Lyons, R. C. (2023). The value of a sea view: Hedonic estimates using 3d simulation and natural language processing. *Land Economics*.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, *225*(2), 254–277.

Greenstone, M. (2017). The continuing impact of sherwin rosen's "hedonic prices and implicit markets: Product differentiation in pure competition". *Journal of Political Economy*, *125*(6), 1891–1902.

Hallegatte, S., Green, C., Nicholls, R. J., & Corfee-Morlot, J. (2013). Future flood losses in major coastal cities. *Nature climate change*, *3*(9), 802–806.

Hino, M., & Burke, M. (2020). *Does information about climate risk affect property values?* (Tech. rep.). National Bureau of Economic Research.

Hsieh, L.-H. C. (2021). Is it the flood, or the disclosure? an inquiry to the impact of flood risk on residential housing prices. *Land Use Policy*, *106*, 105443.

Kuminoff, N. V., & Pope, J. C. (2014). Do "capitalization effects" for public goods reveal the public's willingness to pay? *International Economic Review*, *55*(4), 1227–1250.

Kydland, F. E., & Prescott, E. C. (1977). Rules rather than discretion: The inconsistency of optimal plans. *Journal of political economy*, *85*(3), 473–491.

Lin, Y., McDermott, T. K., & Michaels, G. (2021). Cities and the sea level.

MacDonald, D. N., White, H. L., Taube, P. M., & Huth, W. L. (1990). Flood hazard pricing and insurance premium differentials: Evidence from the housing market. *Journal of Risk and Insurance*, 654–663.

Office of Public Works. (2018). *Kevin 'boxer' moran publishes flood risk management plans and announces $257 million for 50 new flood relief schemes, as part of a ten-year $1 billion programme of investment in flood relief measures?* Retrieved June 2, 2023, from https://www.gov.ie/en/press-release/5f18ca-kevin-boxer-moran-publishes-flood-risk-management-plans-and-announce/

Ortega, F., & Taṣpınar, S. (2018). Rising sea levels and sinking property values: Hurricane sandy and new york's housing market. *Journal of Urban Economics*, *106*, 81–100.

Pilla, F., Gharbia, S. S., & Lyons, R. (2019). How do households perceive flood-risk? the impact of flooding on the cost of accommodation in dublin, ireland. *Science of The Total Environment*, *650*, 144–154.

Pinchbeck, E. W., Roth, S., Szumilo, N., & Vanino, E. (2021). The price of indoor air pollution: Evidence from risk maps and the housing market.

Pralle, S. (2019). Drawing lines: Fema and the politics of mapping flood zones. *Climatic Change*, *152*(2), 227–237.

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

Seo, J., Oh, J., & Kim, J. (2021). Flood risk awareness and property values: Evidences from seoul, south korea. *International Journal of Urban Sciences*, *25*(2), 233–251.

Shr, Y.-H. J., & Zipp, K. Y. (2019). The aftermath of flood zone remapping: The asymmetric impact of flood maps on housing prices. *Land Economics*, *95*(2), 174–192.

Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, *225*(2), 175–199.

Timar, L., Grimes, A., & Fabling, R. (2018). That sinking feeling: The changing price of urban disaster risk following an earthquake. *International Journal of Disaster Risk Reduction*, *31*, 1326–1336.

Votsis, A., & Perrels, A. (2016). Housing prices and the public disclosure of flood risk: A difference-in-differences analysis in finland. *J. Real Estate Fin. Econ.*, *53*(4), 450–471.

For Online Publication - Appendices for:

Estimating the flood discount: Evidence from a one-off national information shock

Tom Gillespie (University of Galway)
Ronan C. Lyons (Trinity College Dublin)
Thomas K.J. McDermott (University of Galway)

Online Appendix A: Data

Dwelling attributes from the Daft.ie listings data

As noted in the main text, our main source of housing data is the Daft.it National Listings Database. Key property-level attributes for inclusion in a hedonic housing price model include the property?s type and size. In our data, we distinguish between apartments and houses dwelling types, with apartments further segmented between duplexes and regular apartments. For houses, there are additional distinctions in the data between terraced, semi-detached, detached and bungalow houses. These distinct property types are captured in our regressions using categorical variables. To capture a property's size in the listings data, indicator variables are included for number of bedrooms (one to five) and for number of bathrooms (one to seven) relative to number of bedrooms.

In addition to detailed systematic information on dwelling attributes such as size and type, a large number of other attributes were also available from the text of the ad placed on daft.ie for each dwelling. This includes a vector of categorical variables for specific features (such as "built-in wardrobes", "patio", "red brick" or "balcony"), as well as information relating to the age or condition of the property.

The full set of controls included in the main regression specifications reported in the paper includes: A categorical variable that flexibly controls for proximity and time since flood events; dwelling quality (Building Energy Rating); building size and plot size; dwelling type; distance to a range of amenities including: roads, nearest CBD, schools, golf course, powerlines, forests (of different types), nature reserve, canals, rivers, lakes, coast, and transitional water bodies; a measure of sea view breadth and depth; and a series of phrase controls based on the text of the sales listing on daft.ie, as follows: "south, west, southwest, period, edwardian, balcony, baywindow, cornerh, utility, conservatory, grannyflat, culdesac, jacuzzi, wardrobefitted, wardrobewalkin, wetroom, underfloor, ensuite, fireplace, stove, aga, burner, solarpanels, victorian, georgian, terrace, endofterrace, detached, semid, mews, garden, garage, frenchdoors, highceiling, corniced, refurb, doubleglazed, pvc, pvcu, brands, beach, luxury, penthouse, views".

Transactions data and matching

To complement our listings dataset, we also draw on information from Ireland's official Residential Property Price Register (PPR).[9] This register is a comprehensive database of all property transactions in Ireland since 2010, based on transaction tax

[9] The data are available at https://www.propertypriceregister.ie/Website/NPSRA/pprweb.nsf/page/ppr-home-en (last accessed March 2023).

returns made by solicitors. The register only contains the property's address as entered by the solicitor, the date of its transfer, its contractually agreed price, whether the dwelling is newly built, and whether the price is a full market price or not (i.e. whether the transaction is arm's length). We exclude transactions with prices that are not full-market.

In order to accurately map these transactions, and to add the dwelling characteristics needed for hedonic housing price regressions, transactions data were mapped to Ireland's official Eircode dwelling-level identifiers. This was undertaken by daft.ie, using an iterative process of automatic scripts, reviewed manually to ensure accuracy. Successful matches between the PPR and Eircodes enabled the identification of the exact dwelling, based on its location, and subsequent matching to the daft.ie data. This matching process resulted in 45,161 dwellings for which we have both transaction and listings information.

Spatial fixed effects units

As noted in Section 4 in the main text, spatial fixed effects are included in all our specifications, to capture the impact on housing prices of factors that are not otherwise captured by our suite of controls, including location-specific and population-specific attributes. Four options are considered: local markets, micro-markets, Electoral Divisions, and Small Areas. The first two are based on daft.ie's breakdown of real estate markets nationwide. Local markets refer to cities, postal districts (within Dublin city), and counties elsewhere in the country; there are a total of 54 markets in Ireland and 25 within Dublin. Micro-markets refer to collections of named areas on the daft.ie system. They are aggregated up from approximately 2,500 areas around the country into micro-markets, based on the volume of market activity, and geographical and socio-economic coherence. There are 375 micro-markets included in the dataset, of which 118 are in Dublin. The latter two options for spatial fixed effects are based on Census divisions of the country. There are just over 3,500 Electoral Divisions (EDs) in the country. Lastly, Census Small Areas (SAs) are a new spatial categorization of Ireland, introduced in the 2011 Census and with an average of 180 dwellings per SA. There are a total of 18,641 SAs in the country.

	N	mean	sd	min	max
Listed Sale Price (EUR)	475,559	275,650	179,666	30,000	2,000,000
Transacted Sale Price (EUR)	45,161	241,435	176,148	16,000	2,725,000
Listed Sale Price Matched Sample (EUR)	45,161	247,627	177,527	30,000	2,000,000
Distance to nearest CBD (m)	475,559	39,613	31,455	16	138,049
Distance to Major Road (m)	475,559	2,511	3,507	0	37,244
Distance to Primary School (m)	475,559	907.9	903.6	0	20,069
Distance to Secondary School (m)	475,559	2,889	3,438	0	31,146
Distance to Coastline (m)	475,559	23,857	24,245	0	94,307
Distance to Transitional Water Body (m)	475,559	17,163	20,972	0	86,823
Seaview Share	475,559	0.00275	0.0255	0	0.791

Table A1: Summary statistics for main continuous variables

Note: The summary statistics presented in this table refer to the full (baseline) listings sample, with the exception of rows 2 and 3 of the table, which refer to the (smaller) matched listings and transactions sample.

Variable	No. of listings	% of sample
Flood variables		
Inside PFRA (at risk)	34,192	7.2%
Defended	763	0.16%
Ever flooded	17,303	3.6%
Dwelling type		
Apartment, duplex, townhouse	62,194	13.1%
Detached, bungalow	213,651	45.0%
Semi-D, terraced, end-of-terrace	199,002	41.9%
No. of bedrooms		
1-bed	15,073	3.17%
2-bed	74,493	15.67%
3-bed	192,077	40.40%
4-bed	153,500	32.29%
5-bed	40,293	8.48%
Urban/rural		
Urban	304,586	64.06%
Rural	170,850	35.94%

Table A2: Frequencies for key dwelling characteristics

Notes: The frequencies reported here are for the full (baseline) listings sample. Defended here refers to dwellings inside the flood risk zone according to PFRA maps (at risk), but protected by flood defences, prior to their listing for sale. Ever flooded here refers to dwellings that are either inside historical flood event polygons or within 250m of a historical flood point, for flood events that occurred prior to their listing for sale. Urban here includes cities, satellite urban towns and independent urban towns.

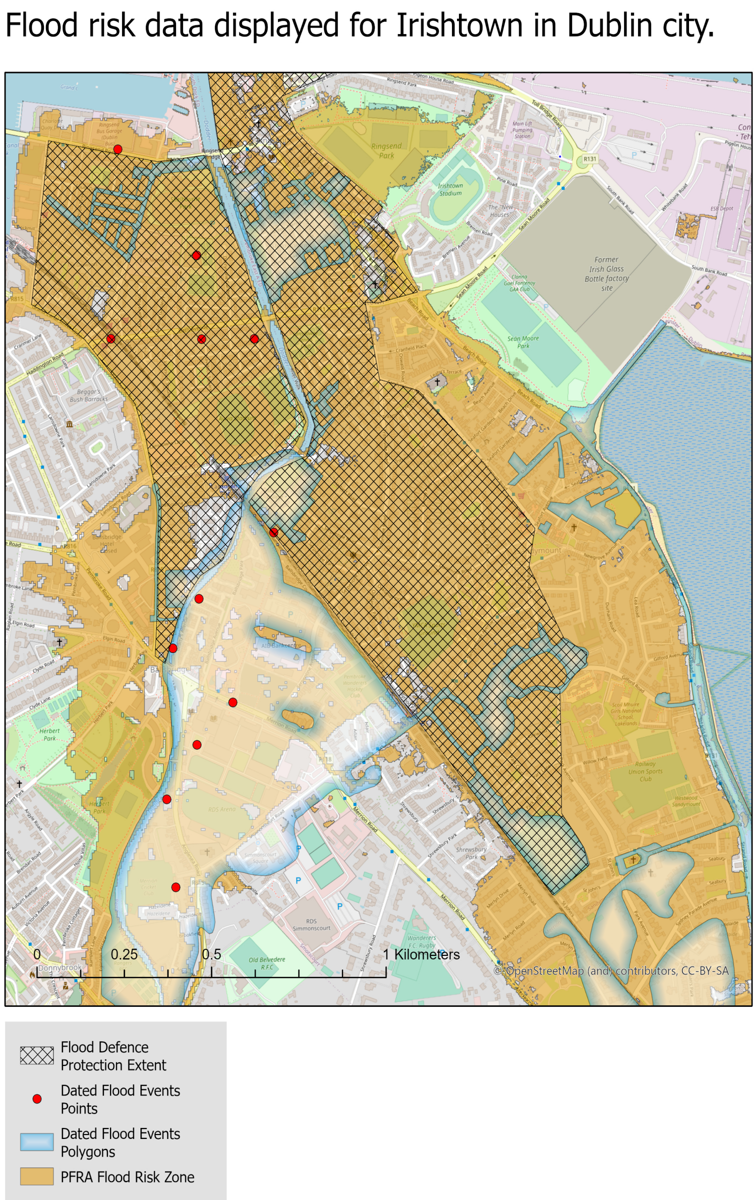

Figure A1: Illustration of flood-related data

Flood risk data, sales sample, and electoral district boundaries for Irishtown area in Dublin city.

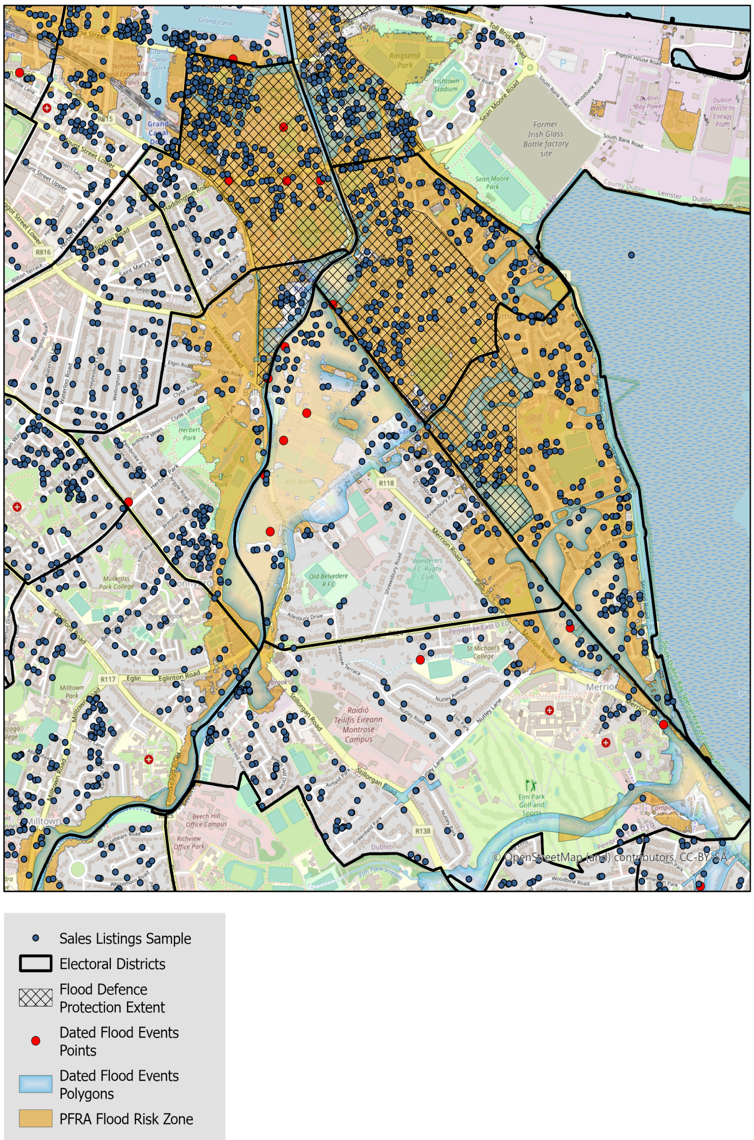

Figure A2: Illustration of data - adding listings and ED boundaries

	Baseline (1)	100m buffer (2)	250m buffer (3)
DiD	-0.040	-0.029	-0.016
	(-5.24)	(-6.21)	(-3.93)
Controls	Y	Y	Y
Spatio-temporal FE	Y	Y	Y
Year-Quarter FE	Y	Y	Y
Obs.	475,436	475,436	475,436
R-Sq	0.792	0.792	0.792
RMSE	0.281	0.281	0.281

Table A3: Flood boundary buffers

Note: The table shows the results of regression estimates of equation (2) using the full sample of listings, with dwelling and location controls as described earlier; t-statistics in parentheses. Column (1) is our baseline specification, where $Risk$ is defined according to the PFRA flood risk maps as described in the text. In Columns (2) and (3) $Risk$ is defined to include dwellings located inside flood risk zones, according to the PFRA flood risk maps, and dwellings located within 100m and 250m, respectively, of the boundary of a flood risk zone.

	Listings (1)	Transactions (2)
$Risk$	-0.029	-0.024
	(-4.63)	(-6.41)
Controls	Y	Y
Spatio-temporal FE	Y	Y
Year-Quarter FE	Y	Y
Obs.	45,161	45,161
R-Sq	0.903	0.904
RMSE	0.216	0.211

Table A4: Comparing estimates based on list vs transaction prices

Note: The table shows the results of regression estimates of cross-sectional versions of equation (2) using the matched sample of listings and transactions, with dwelling and location controls as described earlier; t-statistics in parentheses.