Surge of Inequality: How Different Neighborhoods React to Flooding*

Ana Varela Varela †

2022

Abstract

Recovery scenarios after flooding vary by locality, from permanent declines in economic activity to capital gains. This paper shows that divergent post-flood trajectories at the neighborhood level increased preexisting spatial polarization along property value, racial, and income lines. Using evidence from property sales in four US states affected by Superstorm Sandy in 2012, combined with buyers' demographics, I find that flooded properties in neighborhoods with high preexisting income had more high-income white buyers and higher sale prices than comparable non-flooded coastal properties, seemingly capitalizing on the flood and offsetting average drops. Using machine learning algorithms, I conclude that of a rich set of preexisting place characteristics, neighborhood income best discriminates between most positively and most negatively affected properties. This evidence is consistent with a model of neighborhood segregation in which residential sorting—induced by credit-constrained households deriving higher disutility from flooding—rationally results in more high-income residents and higher property prices in initially higher-income neighborhoods. As coastal flooding is forecasted to increase, these results improve our understanding of the heterogenous impacts of floods, and on the existence of adaptive behavior, or lack thereof, after flooding.

JEL Codes: Q54, R31, H84, J15

Keywords: flood impacts, spatial inequality, segregation, real estate, climate change adaptation

* I am deeply grateful to Wolfram Schlenker, Geoffrey Heal, Dan O'Flaherty, and Eric Verhoogen for their guidance and support. I am thankful to Doug Almond, Laura Bakkensen, Sandra Baquié, Sandra Black, Anthony D'Agostino, Donald Davis, Eugenie Dugoua, Eyal Frank, Joséphine Gantois, Upmanu Lall, Ruiwen Lee, Jackie Klopp, Malgosia Madajewicz, Anouch Missirian, John Mutter, Francesc Ortega, Cristian Pop-Eleches, Jeffrey Shrader, Marco Tedesco, Anna Tompsett, Lex van Geen, and Reed Walker, as well as several presentation attendees for helpful comments and suggestions. I thank Christopher Mayer and the Paul Milstein Center for Real Estate at Columbia Business School for generous help in accessing real estate data from CoreLogic®. Any errors are mine.

†Assistant Professor, Amsterdam School of Economics, University of Amsterdam. Email: a.varelavarela@uva.nl

1 Introduction

Floods are the most frequent and costliest natural disaster in the US.[1] In this context, a large literature has emerged that investigates how affected areas evolve post-flood. This literature yields mixed evidence (Beltran et al., 2018): from permanent declines in real estate prices (Ortega and Taspinar, 2018) to prices that converge to pre-flood levels after an initial negative shock (Bin and Landry, 2013) to an increase in prices (Graff Zivin et al., 2020). These divergent recovery paths are present even within small geographies, such as the unequal recovery by neighborhoods observed in New Orleans after Hurricane Katrina (Bolin, 2007), in Houston after Hurricane Harvey (Fernandez, 2018), and in New York City after Superstorm Sandy (Bergren et al., 2013).[2] The potential causes of these differential outcomes remain underexplored in the economics literature. Gaps also remain in our understanding of how differential recovery paths post-flood contribute or are countercyclical to preexisting place dynamics, such as urban decay or gentrification.

This paper investigates the differential impacts of flooding on property price sales and buyers' demographics at the neighborhood scale. Specifically, I build on a model of segregation to derive and test a hypothesis that preexisting neighborhood income mediates responses to the flood , whereby places with high income levels could rationally experience an increase in sale prices post-flood. Using evidence from Superstorm Sandy, which caused extensive flooding across several US states in October 2012, I show that divergent post-flood changes at the neighborhood level increased preexisting spatial polarization along property value, racial, and income lines. Flooded properties in neighborhoods with high preexisting income had higher sale prices and more high-income white buyers than comparable non-flooded coastal properties, consistent with the model. To the best of my knowledge, this study is the first to document differential post-flood responses at the neighborhood scale that lead to increased spatial polarization using micro-level data on sale values and buyers' demographics. By doing so, it sheds light on previous disparate results in the flood-impacts literature.

Why would property prices go up in some places after a flood? To illustrate how residential sorting could lead to this result, I build on the model of segregation developed by Becker and Murphy (2000), which hinges on the existence of positive externalities to live among residents characterized by high income and other demographics related to income,

[1] It is estimated that 90% of all natural disasters involve some sort of flooding in the US (Insurance Information Institute, 2018), and floods were the number one disaster in terms of lives lost and property damage during the 20th century (Perry, 2000).

[2] Hurricane Sandy merged with a winter storm right before making landfall on the Atlantic coast of the US (Sobel, 2014). The resulting hybrid storm, no longer technically a hurricane, was labeled a "superstorm" by the media, a term that has since been used by the scientific community to describe this phenomenon (for instance, Barnes et al. (2013).

such as race.[3] Intuitively, low-income, credit-constrained households are assumed to derive a higher disutility from flooding, as they are less able to rebuild their homes, to retrofit them in preparation for the next flood, or to pay higher insurance costs. For a given household, these costs would be higher for more expensive properties. The residential sorting induced by this assumption—with low-income households retreating from places where the value derived from coastal amenities post-flood is not enough to compensate high residential prices—could lead to more high-income residents in some flooded places and in turn, to a residential price increase due to the positive externalities derived from living among them. This model helps to motivate why it would be rational to expect more high-income white residents and larger property prices post-flood in neighborhoods with higher preexisting average incomes.

To find empirical evidence of these predictions, I construct a rich dataset that captures the universe of property sales in the states most affected by Sandy's flooding (Toro, 2013): New Jersey, New York, Connecticut, and Rhode Island. Superstorm Sandy, which as of early 2022 remains the fourth most costly natural disaster in the history of the US,[4] is a compelling setting to explore heterogeneous flood effects, given that Sandy's flood affected a vast geographical area along the Atlantic coast of the US with very different preexisting place characteristics. I link a subset of the property sales observations with data on buyer's race, ethnicity, and income using publicly available data on anonymized individual mortgage lending transactions. I connect property sales with data on flooding extent, depth, and property-level damage caused by Sandy. Finally, the dataset includes a comprehensive set of 52 preexisting place variables compiled from diverse sources.[5]

I start the empirical analysis by exploring average flood impacts in a difference-in-differences setting, which serve as a benchmark for the heterogeneous effects evaluated later. The treatment group is composed of properties on Sandy's floodplain. Exploiting the richness of the data, together with the vast geographical area affected by Sandy, I can define a control group composed of non-flooded properties located no more than 500 meters (around 1,640 feet) from the coastline, which are credibly similar to properties in the floodplain in terms of access to coastal amenities and other unobservables.[6] Parallel trends in residential prices pre-Sandy increase confidence that this is indeed the case. Results show that on average, flooded properties are sold at lower prices, to buyers with less income and who are less

[3] Positive externalities could be related to job opportunities, networking, better schools, a prestige signal, etc.

[4] Sandy was estimated to cause 159 deaths and $80 billion USD in losses. These costs are only superseded by those incurred by Hurricane Katrina (2005), Hurricane Harvey (2017), and Hurricane Maria (2017) (NOAA National Centers for Environmental Information, 2022).

[5] These include the Opportunity Insights group at Harvard University, the US Census, the Stanford Education Data Archive, Open Street Map, etc.

[6] This choice of a control group, which mimics a boundary discontinuity design, is of particular relevance given that the literature has identified sorting along unobservable characteristics in coastal/non-coastal locations, e.g., Bakkensen and Barrage (2017).

likely to be white.

I then proceed to explore whether there is heterogeneity in the results on property prices according to preexisting neighborhood income in a triple-differences model. I find that the negative effect of the flood on prices and high-income white buyers decreases as preexisting place income increases; this result is robust to comparing properties with similar levels of damage. In fact, flooded properties in neighborhoods in the top decile of preexisting income increased property prices with respect to comparable non-flooded properties by 7.6%, and saw an increase of high-income white buyers equal to 23% of a standard deviation. These results are consistent with the predictions of the model: flooded properties in neighborhoods with higher preexisting average incomes increased in value, and these price increases coincided with changes in buyers' demographics.

Finally, I use data-driven machine learning algorithms as described in Chernozhukov et al. (2018) (henceforth, CDDF) to systematically evaluate heterogeneity in flood impacts according to preexisting place characteristics. Implementing this procedure serves two main goals. First, I am able to test and rule out that flood impacts were homogeneous, and hence boost confidence that the above results are not just spuriously capturing heterogeneity. Second, I corroborate that neighborhood income and secondary education, which are highly correlated in the data, best discriminate between most positively and most negatively affected properties among a rich set of preexisting place characteristics. Other variables that have been highlighted as positively affecting recovery after flooding, such as different indicators of social capital, do not vary as much between the most and least affected properties. Thus, the CDDF procedure serves as a check on the predictions of the model: flood impacts are not homogeneous, and preexisting place income is correlated with this heterogeneity.

The results have important policy implications. Coastal flooding is expected to increase due to sea level rise, an effect that will be compounded in the Northeast US by changes in hurricane activity induced by climate change (Hess et al., 2019).[7] Hence, after-flood dynamics such as those examined in this paper can lead to increased space polarization throughout the coast along racial and income lines. These dynamics could lead to the entrenchment of pockets of vulnerability with losses in the tax base that would make it harder to recover after the inevitable next storm. On the other hand, other coastal areas would become rich enclaves, which could regressively soak up public assistance for after-flood recovery. Moreover, other natural disasters that are also forecasted to worsen with climate change—such as fires— could lead to similar dynamics further from the coast (Fuller et al., 2019). A growing strand of the economics literature highlights the negative impacts of spatial segregation and

[7] In some areas in mid-Atlantic states in the US, flooding is 10 times more common now than in the 1950s (EPA, 2016). Annual property losses caused by hurricanes and other coastal storms could increase two- to fourfold by 2100 in the Northeast US compared to losses in 2014 (Gordon, 2014).

inequality.[8] Therefore, successful recovery policies and insurance schemes would need to consider these heterogeneous effects of natural disasters across space.

Contribution to the literature. This paper builds on and contributes to several strands of the literature. First, this paper relates to the literature on heterogeneous recovery after natural disasters. At the macro level, a large body of empirical work has documented diverging recovery patterns in different countries after a natural disaster (Bakkensen and Barrage, 2018; Jina and Hsiang, 2014; Mutter, 2015), from large positive to negative economic shocks. Some studies posit that a country high income per capita is correlated with fewer negative impacts (Kahn, 2005; Kellenberg and Mobarak, 2008; Lackner, 2019). At the micro level, the literature on the impacts of flooding events yields mixed evidence depending on the setting (for instance, Bin and Landry (2013), Graff Zivin et al. (2020), and Ortega and Taspinar (2018)). To the best of my knowledge, no paper in the economics literature has attempted to explain why these differential outcomes at the micro level are observed.

This paper bridges these two strands of the literature. Using property-level data, I am able to document that differential recovery paths at the micro, neighborhood level mimic those present at the country scale, with some flooded neighborhoods suffering long-lasting declines in real estate prices, while others seemingly capitalize on the flood. I focus on the fact that flooding does not occur in a vacuum, and that any heterogeneous effects are likely mediated by existing place characteristics and dynamics. Hence, this paper hypothesizes and finds evidence that there is heterogeneity in responses at the neighborhood level by preexisting income, as those present at the country level by income per capita. These results contribute by rationalizing previous disparate results found in the literature on flood impacts.

Second, this paper speaks to the literature on residential segregation, and specifically, how it is influenced by amenities and tastes for a particular type of neighbor (in line with work by Tiebout (1956) and Schelling (1969).) Several recent papers have investigated demographic shifts associated with changes in environmental amenities broadly (Andaloussi and Isaksen, 2017; Gamper-Rabindran and Timmins, 2013; Greenstone and Gallagher, 2011), and specifically to changes in flood risk and perception (Smith et al., 2006; Bakkensen and Ma, 2019; Graff Zivin et al., 2020; Keenan et al., 2018; Siders, 2018; van Holm and Wyczalkowski, 2018).[9] Part of this literature has focused on environmental justice concerns, and investigated how positive changes in public goods and environmental amenities may lead to

[8]For instance, work by Alesina and La Ferrara (2000), Bergman et al. (2019), Chetty et al. (2014), Chetty and Hendren (2017), Couture et al. (2021), Fogli and Guerrieri (2019), and Graham (2018).

[9]In the work most closely related to this study, Graff Zivin et al. (2020) focus on Florida and use a dataset of similar characteristics as the one used in this paper to show that real estate prices went up after flooding on average, and that buyers had higher incomes. They do not explore changes in the race or ethnicity of buyers, or investigate differential responses to the flood within the context of Florida.

sorting and displacement of low-income communities (Banzhaf and Walsh, 2013; Ma et al., 2019; Bento, 2013; Depro et al., 2015).

This paper posits and documents a novel mechanism that exacerbates spatial polarization and environmental justice concerns through a negative change in environmental amenities. Building on the model of Becker and Murphy (2000),[10] I show that an increase in a public *bad*, such as flooding, can lead to residential sorting and an increased spatial polarization, with low-income households retreating from places where the value derived from coastal amenities post-flood is not enough to compensate high residential prices. This paper then explores these predictions empirically and is the first to examine changes alongside buyers' race and income after flooding using micro-level data.

Finally, this paper builds on the growing literature on floods and flooding-risk impacts on real estate (Hallstrom and Smith, 2005; Bakkensen and Ma, 2019; Bin et al., 2008; Bernstein et al., 2019; Gallagher, 2014; Heal and Tedesco, 2018; Muller and Hopkins, 2019).[11] Some recent papers focus specifically on flood risk capitalization of Superstorm Sandy in New York City, using data on real estate transactions from the New York City Department of Finance (Ortega and Taspinar, 2018; Gibson and Mullins, 2020; Barr et al., 2017).[12] The fact that desirable coastal amenities and higher flooding risk are highly correlated makes disentangling flood risk capitalization a challenge that has been well documented (Beltran et al., 2018; Smith and Whitmore, 2019; Smith et al., 2006). Moreover, it is challenging to evaluate how much of the flood risk capitalization in flooded properties is mechanically driven by structural damage (Smith et al., 2006; Ortega and Taspinar, 2018).

In this paper, due to the richness of my dataset and the vast geographical area affected by Sandy, I am able to define a control group that mimics boundary discontinuity designs.[13] Identification hence relies on comparisons with non-flooded properties that are credibly similar in terms of access to coastal amenities and other unobservables.[14] Second, I construct

[10]Becker and Murphy (2000) present a model of neighborhood segregation in which residents value both endogenous and exogenous amenities, where the former are derived from neighbor types. When extended to several neighborhoods, this model implies the existence of "tipping" points—i.e., a minority ratio above which the neighborhood "tips" to become all-minority. Empirical work, notably by Card et al. (2008), has find evidence of the existence of these tipping points in US cities.

[11]This literature is itself is a subset of the vast literature initiated by Rosen (1974) that uses hedonic models to derive willingness to pay for goods for which there are not explicit markets (Davis, 2007; Greenstone and Gallagher, 2011; Chay and Greenstone, 2005; Greenstone, 2017)

[12]Ortega and Taspinar (2018) find that properties in Sandy's floodplain experienced an average price drop of 8%, using as controls other similar properties in the city. Two working papers explore different aspects of how much information about flood risk Sandy revealed with respect to FEMA's flood maps. Gibson and Mullins (2020) find that property prices went down on average for properties not affected by Sandy, but which were afterward included in updated flood risk maps. Barr et al. (2017) focus on the additional information shock Sandy imposed on non-flooded properties, given the relative distance between Sandy's flood surge and FEMA's flood risk area delimitation. Using locally weighted regressions, they conclude that non-flooded properties saw no effects in large areas of the city, although they did in areas closer to the city center.

[13]As Black (1999) and Duranton et al. (2011), for instance, do.

[14]Bernstein et al. (2019) and Bakkensen and Ma (2019) also control for distance to the coast in narrow bands.

a novel measure of vertical distance to the maximum flood surge for both flooded and non-flooded properties. This metric, which is correlated with damage, allows me to investigate flood impacts alongside the flood boundary in a flexible way.

Paper outline The remainder of the paper proceeds as follows. With the objective of motivating the empirical design, Section 2 lays out a model of segregation that rationalizes the existence of higher prices post-flood in neighborhoods with higher preexisting average incomes. Section 3 describes the constructed dataset, detailing data sources and processing steps, as well as presenting key summary statistics. Moving on to the empirical analysis, Section 4 describes key models and methodologies used to derive the main results, which are presented in Section 5. Section 6 summarizes further tests to fully contextualize the setting and to illustrate of potential mechanisms behind the main results. Robustness checks are presented in Section 7, and Section 8 concludes.

2 Conceptual model: flooding and neighborhood segregation

Why would residential prices go relatively up in some places after a flood? In this section, I build upon the segregation model developed by Becker and Murphy (2000) to illustrate that the premium to live in a particular place can go up after a disamenity shock, as long as the impacts are heterogenous across residents, and there is a preference to live among a certain type of resident. This section lays out the model's intuition and main predictions. A detailed description of the model is included in appendix C, and key mathematical derivations are in appendix D.

The set up of the model as defined by Becker and Murphy (2000)[15] considers that the willingness to pay to live in a certain place depends only on two variables: its exogenous amenities, and the type of neighbors. Specifically, it assumes that there are positive externalities to live among residents characterized by high-income and other demographics related to income, such as race.[16] [17] Hence, in equilibrium, residential prices would be higher

None of these papers investigate the impacts of direct flooding, but rather flood-risk (specifically, the impacts of sea level rise forecasts and presence on FEMA's Special Flood Hazard Area, respectively.)

[15] The model by Becker and Murphy (2000) has been used as theoretical underpinning for several empirical works, for instance, Card et al. (2008) and Banzhaf and Walsh (2013).

[16] Some work on the spatial segregation literature separate tastes for affluence and race of neighbors, for instance Sethi and Somanathan (2004) and Banzhaf and Walsh (2013). In this setting, for the sake of simplicity, I remain agnostic to the source of different tastes. Acknowledging that income and race are highly correlated in the data, I assume positive externalities of having high-income white neighbors.

[17] Positive externalities could be related to job opportunities, networking, better schools, a prestige signal, etc. This assumption is validated by recent empirical work. For instance, Bergman et al. (2019) use experimental evidence from Seattle to conclude that low-income families who move to higher income places report higher neighborhood satisfaction.

in places with more valued amenities—such as some coastal locations—and/or with more high-income residents.

From this set up defined by Becker and Murphy (2000), I focus on the model predictions when the system in equilibrium is shocked by a flood under two assumptions. First, I assume that flooding is a disamenity for all residents: everyone will be better off without increased flood risk. As a result, the bundle of amenities in flooded neighborhoods is less desirable than before for all residents. From this assumption alone, residential prices in flooded neighborhoods should decrease. Second, I interpret the flood as an income shock, and assume a decreasing marginal utility of income. This interpretation yields that (a) more credit-constrained, low income residents derive a higher disutility from flooding,[18] and (b) that the disutility shock would be positively correlated with pre-flood residential prices for a given type of resident. Intuitively, low-income residents would be less able to rebuild their homes, to retrofit them in preparation for the next flood, or to meet higher insurance costs; and, for a given type of resident, these costs would be higher for more expensive properties.

Conditional on these assumptions, the model yields two main predictions. First, low-income households could retreat from flooded places, notably those with higher pre-flood prices. The disutility derived from flooding makes the value derived from coastal amenities post-flood not enough to compensate high residential prices for them. Even if they value coastal amenities less too after the flood, high-income households are better able to cope and remain, leading to a higher proportion of high-income residents overall. Second, and as a consequence of this residential sorting, prices could increase for flooded properties. This result is a priori non-intuitive: as all residents derive disutility from the flood, economic theory would suggest a reduction in prices. But, intuitively, the increase in positive externalities derived from a larger rate of high income residents could override the negative utility shock from the flood.

Hence, this model provides key insights on adaptive behavior after a flood, or lack thereof. A seemingly irrational behavior of larger property prices after a flood could happen even if the flood disamenity is internalized, as long as flooding affects relatively more the lesser valued type of resident.

3 Data: sources, processing, and key statistics

This section summarizes data sources, processing and descriptive statistics for the key outcome variables (property sale prices and buyers' demographics), treatment variables (flood

[18] del Valle et al. (2019) provide empirical evidence for this assumption. In their study of household financial behavior after Hurricane Harvey, they find that there is a spike in new credit card originations after flooding, with the effect concentrated on borrowers with higher incomes, credit limits, and credit scores.

extent and depth), and place characteristics along which heterogenous impacts are evaluated. Further details about sources and processing are in appendix E.

Property sales Residential property sale data comes from CoreLogic®, a private supplier of US real estate data.[19] The dataset contains records of every property transaction in the states of New Jersey, New York, Connecticut, and Rhode Island from 2002 to 2016. It includes data on date of property sale transaction, sale amount, exact longitude and latitude of property, property characteristics, mortgage amount, mortgage lending institution, and name of buyer and seller. As described in section 4, my identification strategy relies on comparing flooded to non-flooded properties located on the coast. Hence, the analysis only considers properties which are located 500 meters (approximately 1640 feet) away from the coastline as defined by the National Oceanic and Atmospheric Administration (NOAA, 2016). Finally, I also obtain ground elevation by spatially matching each property longitude and latitude with Digital Elevation Models (DEMs) developed by the United States Geological Survey (USGS, 2019).[20]

Race and income of buyers I combine data on property transactions with buyers' demographics following the approach by Bayer et al. (2016), which is well used in the literature (for instance, by Bakkensen and Ma (2019) and Graff Zivin et al. (2020)). Hence, I use publicly available data on mortgage applications disclosed under the Home Mortgage Disclosure Act, which contain anonymized individual mortgage lending transactions, including self-reported race, ethnicity, and income of the borrower. I obtain buyers' demographics for 32% of the property transactions in my dataset[21] (which constitutes 57% of those properties with mortgage information.) This rate is comparable to other examples in the literature, which look at different time periods and locations.[22] Table 1, with descriptive statistics of the resulting dataset, shows that the subset of properties with socioeconomic demographics and the main sample are similar across several dimensions. Finally, I construct a summary variable of buyers' race and income using principal component analysis, as described in

[19] Access to the data was granted by the Paul Milstein Center for Real Estate at Columbia Business School.

[20] These models provide elevation, in meters, at a resolution of one-third of an arc-second (approximately 10 meters, or 33 feet) with respect to the North American Vertical Datum of 1988.

[21] The merge process is imperfect, as not all transactions are associated with a mortgage loan, and not all institutions are required to report under the HMDA. Moreover, lending institutions names are recorded as strings, and are subject to typographical errors or different abbreviations (e.g. "Bank of America" or "BK America") which complicate the merging process. Also, it is not possible to merge records which are not uniquely identified by the four merging variables (loan amount, name of lending institution, census tract, and transaction year.)

[22] Bayer et al. (2016), which study the San Francisco Bay Area from 1994 to 2004, match 55% of properties with mortgage information from Corelogic to mortgage applications; Liao and Panassie (2018), with a similar dataset for the state of Florida during 2000-2016, match 26% of the whole sample, 47% of the records with mortgage data.

appendix F. Buyers who rank high on this composite variable have higher income and are more likely to be white. This variable allows me to evaluate empirically predictions of the model described on 2, which predicts joint sorting along these dimensions.

Sandy's flooding To characterize Sandy's flooding, I use publicly available storm surge maps developed by the Federal Emergency Management Agency (FEMA) Modeling Task Force (FEMA, 2014).[23] I also obtain estimates of property-level damage developed by FEMA using aerial imagery and inundation assessments (FEMA, 2014). This dataset classifies each building damaged by Sandy into four damage levels (*affected*, *minor*, *major*, or *destroyed*) following a set of objective criteria.[24] [25]

Finally, I derive a metric of vertical distance to the maximum flood surge level for both flooded and non-flooded properties. As sketched in figure 2, this metric measures the difference in elevation between the properties' ground and the closest point in the flood extent boundary. The top panel of figure 3 shows that the distance metric is negatively correlated with damage, as expected. This metric allows me to advance the empirical analysis, as it provides a measure of vertical distance to flood level for control properties outside of the floodplain which allows me to investigate changes in outcomes alongside the flood boundary in a flexible way.[26]

Table 1 summarizes descriptive statistics of the resulting dataset, and Figure 1 delineates the area of study, and shows spatial extent of coastal property sales (top panel) and Sandy's flooding (bottom panel.)

[Figure 1 about here.]

[Table 1 about here.]

[Figure 2 about here.]

[23] These maps summarize flood depth at a 3 meter resolution using data from field-verified High Water Marks and Storm Surge Sensors from the USGS, as well as Civil Air Patrol and NOAA imagery assessments.

[24] A property whose exterior walls have collapsed is declared *destroyed*; if some exterior walls have collapsed or the flood depth is higher than 5 feet, the damage is *major*; if more than 20% of roof covering is missing or the flood depth is between 2 to 5 feet, it is classified as *minor*; if less than 20% of roof has been damaged or the flood depth is less than 2 feet, the property is *affected*.

[25] Properties in my transaction and damages datasets are both identified by points (that is, longitude and latitude coordinates.) To combine both datasets, I assign to each property point the closest damage level point up to 20 meters, beyond which I consider the property unaffected.

[26] Horizontal distances to the flood boundary on a plane would obviate changes in elevation which might be relevant in terms of flood perception. Also, the maximum elevation level reached by the flood changes locally due to complex ocean dynamics, wind patterns, soil geology, etc., which makes raw ground elevation a poor proxy for vertical distance to flood level for non-flooded properties.

[Figure 3 about here.]

Place characteristics With the objective of estimating heterogeneity with respect to preexisting place characteristics, I construct a comprehensive dataset to capture as many of the observable characteristics that might influence residential location choice as possible. To establish which variables might be relevant for inclusion, I refer to recent studies in the urban economics and public finance literatures which broadly consider the role of place characteristics, amenities, or neighborhood quality on some economic outcome of interest.[27] The resulting dataset is composed of 52 variables, summarized in table B.18.

Following the approach in Diamond (2016), I classify these 52 variables into overarching categories, and then summarize each of the categories using principal component analysis (PCA) – as described in appendix F with results summarized in figure A.6. The main objective of using PCA is to avoid spuriously capturing heterogeneity alongside single place variables: the composite variable obtained through with PCA captures the data signal within each of the overarching categories, and hence it is less noisy than a single indicator.

For instance, the place income category summarizes six different indicators indicative of the income and income distribution of a place: income per capita, poverty rate, fraction middle class, income share of the population with 1% highest income, Gini coefficient for income, and Gini coefficient for income without considering population with 1% highest income. The resulting place income variable, normalized to have a mean of zero and standard deviation equal to 1, is positively correlated to places with high income per capita and income inequality, and low fraction middle class.

4 Empirical Strategy

This section describes the difference-in-differences (DD) and triple differences (DDD) models used to infer average and heterogeneous marginal impacts of flooding, respectively. Then, it summarizes the approach followed to evaluate heterogeneity in flood impacts using machine learning algorithms as described in Chernozhukov et al. (2018).

4.1 Impacts of flooding: DD and DDD models

I estimate the average impact of flooding on each outcome of interest – either property sale value (in logs), a dummy variable indicating if the buyer identifies as a non-hispanic white, the buyer's income (in logs), or a composite variable indicating whether the buyer both has high income and is white – according to the following specification:

[27] Notably: Diamond (2016), Christensen and Timmins (2018), Chetty et al. (2018) , Chetty et al. (2014) , Chetty and Hendren (2017) , Collinson and Ganong (2018).

$$y_{pt} = \beta F_p \cdot A_t + \theta_1 F_p + \theta_2 A_t + \boldsymbol{X_p}\boldsymbol{\gamma} + \lambda_p + \lambda_t + \varepsilon_{pt} \tag{1}$$

y_{pt} is an outcome associated with property p sold on date t. A_t is a time indicator for sales recorded on or after November 1, 2012 (Sandy made landfall in the region between October 29 and November 1). F_p is a variable indicative of Sandy's flooding, either presence on the floodplain or vertical distance to flood level. Treatment status is then defined by the interaction term $F_p \cdot A_t$. The coefficient of interest, β, is the average effect of Sandy flooding on the outcome of interest.

The control group is composed of properties outside of Sandy's floodplain less than 500m away from the shoreline. This strategy ensures that properties on the control group are similar to those on the treatment group: among properties so close to the shoreline, Sandy flooding is more plausibly uncorrelated to unobservable characteristics that if the control group was further inland. This is of particular importance given that the literature has identified sorting along unobservable characteristics in coastal/non-coastal locations, e.g. Bakkensen and Barrage (2017) show that property owners closer to the coast tend to be less pessimist about the potential effects of climate change. Figure 4 zooms in the area of study to illustrate the identification strategy.

[Figure 4 about here.]

Time fixed effects λ_t control for any time shocks common to all properties,[28] and place fixed effects, λ_p, control for any time invariant unobserved place characteristics. My preferred specification controls for census block fixed effects.[29] This model is preferred to a repeat sales panel with property-level fixed effects as there might be concerns of any unobservables affecting properties that have been sold twice or more in the 15 year period.[30] To alleviate concerns that results might be driven for properties of different characteristics within blocks, I also include a set of property level characteristics $\boldsymbol{X_p}$ (year built, and square footage) on some models. For robustness, I also show that results are consistent to the inclusion of property-level fixed effects in a repeated sales panel.[31] Finally, robust standard errors, ε_{pt}, are two-way clustered, to allow errors to be correlated across census tracts within years, and across years within census tracts.

Identification of causal effects of flooding in equation 1 requires that non-flooded properties would have followed parallel trends to flooded properties in the absence of the flood.

[28] In the preferred specification, I control for month-of-year, although I show in section 7 that the results are robust to other time controls.

[29] Census blocks are the smallest geography level defined by the U.S. Census Bureau. In an urban setting, census blocks usually correspond to city blocks, surrounded on all sides by streets (Bureau, 2011)

[30] 82% of properties in the main sample had only been sold once in the period; 90% among those which have data on the race and income of the buyer

[31] In this case, property characteristics $\boldsymbol{X_p}$ will be embedded in the fixed effects.

Although this condition is not directly testable, I gain confidence that this would have been the case if the trends were parallel prior to the flood. Figure 5 confirms that prices for both flooded and non-flooded properties evolved concurrently prior to Sandy, diverging right after.

Finally, to estimate how flood impacts are marginally affected by different place characteristics, I use a triple differences (DDD) specification resulting from interacting the model in 1 with a variable Z_p summarizing a characteristic of the place property p is located in:

$$\begin{aligned} y_{pt} = {} & \beta F_p \cdot A_t \cdot Z_p + \\ & \theta_1 F_p + \theta_2 A_t + \theta_3 Z_p + \theta_4 F_p A_t + \theta_5 F_p Z_p + \theta_6 A_t Z_p + \\ & \boldsymbol{X_p} \gamma + \lambda_p + \lambda_y + \varepsilon_{pt} \end{aligned} \tag{2}$$

The coefficient of interest is β, which represents the marginal impact of the flood associated with one more unit of Z_p. For instance, when Z_p represents place income with a mean of 0 and a standard deviation of 1, β summarizes how much higher the flood impacts have been in places with 1 standard deviation above mean place income.

4.2 Estimating heterogeneous effects: machine learning procedure

I evaluate heterogeneity of flood impacts according to preexisting place characteristics using the machine learning procedures described in Chernozhukov et al. (2018) (henceforth, CDDF.) Machine learning methods such as those in CDDF provide a disciplined way to measure heterogeneity in the treatment ex-post, and and are becoming more common in the economics literature (Wager and Athey, 2018; Athey and Imbens, 2019; Rigol et al., 2017; Davis and Heller, 2017; Deryugina et al., 2019).[32]

Implementing the CDDF procedure in this context serves two main goals. First, I can use it to test and rule out that flood impacts were homogenous, and hence increase confidence that any potential heterogeneity captured on the DDD analysis is not spurious. Second, it yields the preexisting place characteristics that best discriminates between most positively and most negatively affected properties. Hence, I can use this procedure to derive a data-driven check for the predictions of the model in section 2, which concludes that places with higher preexisting place income should be those least affected by the flood.

Chernozhukov et al. (2018) provide a generic method to make inference on key features of the conditional average treatment effects.[33] Succinctly, CDDF relies on repeated data partitioning to estimate robust estimates of key features of the treatment effects. Then, con-

[32]They can be understood as an alternative to pre-registration of studies, while still guaranteeing that the selection of variables to measure heterogeneous effects has not been biased towards finding significant results.

[33]That is, the difference in expected outcomes between treatment and control groups, conditional on controls

fidence intervals and p-values are generated to consider uncertainty coming from both estimation and data splitting uncertainty.[34]

CDDF yields robust estimates for:

1. Best linear predictor for the average treatment effect (which can be compared with the DD estimate computed following equation 1 to check the fitness of the procedure.)
2. Average treatment effects for properties on the top and bottom quintiles, according to their estimated treatment effect.
3. Average characteristics of the properties on the top and bottom quintiles, according to their estimated treatment effect, as well as differences in these characteristics between the two quintiles.

Even if the methods in Chernozhukov et al. (2018) are defined for randomized control trials, they are generalizable to other non-experimental settings as long as it is possible to construct an unbiased signal of the treatment effect conditional on controls. Hence, the required assumption for inference is the same as for the DD model described above: after controlling for place and time fixed effects, the difference in expected potential outcomes between properties inside and outside the floodplain yields the unbiased effect of the flood.

5 Main Results

This section describes the main flood impacts estimated following the empirical models described above in section 4. It starts by documenting average effects of the flood on property values and buyers' demographics measured with the DD model. Then, results from DDD reveal that these impacts were heterogeneous according to preexisting place income, as the model in section 2 would predict. Finally, it presents results from a machine learning procedure which confirm that preexisting place income best discriminates between most positively and most negatively affected properties.

5.1 Average impacts: lower prices, fewer white and high-income buyers

I find that, on average, flooded properties were sold for lower values to poorer buyers who were less likely to be white. These drops were larger for properties which experienced a higher flood depth. However, I rule out that structural damage is solely driving the results, as properties which were just affected by the flood or with low flood depths saw significant decreases in prices.

[34]For completeness, I describe in the Appendix G the steps followed to implement the CDDF procedure. The original code from Chernozhukov et al. (2018) is available for downloading here

Binary treatment: presence in the floodplain Table 2 reports results on the average impacts of flooding. In the top panel, the treatment variable is a binary variable equal to 1 for properties in Sandy's floodplain, interacted with an indicator of sales post-Sandy. All models include census blocks and month fixed effects. Results from this panel show that property prices of flooded properties have decreased almost 9% on average after Sandy. Homebuyers were 5 percentage points less likely be white, and their income was 2.6% lower. They were also less likely to be both white and high-income (drop equal to 10% of a standard deviation).

The relative average drop in prices after Sandy for flooded properties is also evident in the top panel of figure 5. This figure also shows that prices before Sandy followed parallel trends for properties in the treatment and the control groups. This pattern increases confidence that trends would have continued to evolve in parallel in the absence of the flood, and make the results above causally interpretable.

[Table 2 about here.]

[Figure 5 about here.]

Continuous treatment: vertical distance to flood level The bottom panel of table 2 presents results using a flood treatment variable equal to the absolute value of vertical distance to flood level for flooded properties.[35] These results show that the average drops in prices reported in the top panel were concentrated in properties with higher flood depths. Hence, one more meter of distance to the maximum flood level decreased prices by 5.4% on average. Similarly, one more meter of flood depth decreased the likelihood of a buyer being white by 2.8 percentage points, and to be both white and high-income by 5% of a standard deviation. Buyer's income does not change significantly (at least at the 10% confidence level) with respect to vertical distance to flood.

The bottom panel of figure 5 provides graphical intuition for the changes in prices for properties inside and outside the floodplain with respect to vertical distance to flood level, before and after Sandy. There is no apparent discontinuity around the flood boundary prior to Sandy, which is consistent with a lack of prior sorting along both sides of the floodplain. Prices for properties on the floodplain change with respect to flood depth remarkably differently after Sandy. The figure shows that property prices post-Sandy are below those prior to the flood even for properties at lower distances to the flood (e.g. 0.5 meters or lower.) The drop in prices decreases with vertical distance to flood level, consistently with the results in table 2 above. The figure also suggests an inflection point in the relation between

[35] The steps taken to compute the metric of the vertical distance to flood level are described in section 3.

flood depth and price changes, with properties with flood depths around 1 meter experiencing the largest drops, although the results for higher flood depths are noisier. Non-flooded properties' prices remain unchanged before and after the flood.

Property damage The evidence above seems to suggest that negative shocks to prices are not just driven mechanically by structural damage, given that price drops do not increase monotonically with flood depth, and that properties at low flood depths experience price drops. To gain further confidence, table 3 interacts the treatment variable (a dummy indicating if a property is in the floodplain and the sale was after the flood) with a set of dummy variables indicating the level of damage according to FEMA's evaluation (as described in section 3.) The omitted category includes properties that were deemed not damaged by FEMA. Besides regression results, the table also presents p-values of pairwise testing the equality of the dummies coefficients. These results show that the point estimate of flood impacts increases with the category of damage. However, the drop in prices for merely *affected* properties is significant at the 99% and of large magnitude, equal to 8.2%. I also fail to reject that the coefficient for properties with *minor* and *major* damage are equal (p-value equal to 0.428.)

[Table 3 about here.]

5.2 Heterogeneous impacts according to place income

In this section, I show that the negative impact of the flood on prices and high-income white buyers decreases as preexisting place income[36] increases. This result is robust to comparing properties with similar levels of damage. The elasticity of property prices with respect to flood depth also decreases with preexisting place income.

Binary treatment: presence in the floodplain The top panel of table 4 shows that the magnitude of flood impacts changes according to preexisting place income. Properties located on places with a mean value of place income (equal to 0, by construction) experience a decrease in property prices equal to 4.5%. With respect to these, flooded properties in places one standard deviation above mean income see an *increase* in prices of 7.3%. This coefficient is significantly different both from zero and from the coefficient for properties at the mean place income (p-value of both differences is smaller than 1%.)

[36] As described in section 3, place income is a composite variable that describes six different indicators indicative of the income and income distribution of a place. It is normalized to have a mean of zero and standard deviation equal to 1, and it is positively correlated to places with high income per capita and income inequality, and low fraction middle class.

Regarding buyers' demographics, table 4 reports that one standard deviation of place income increases flood impacts on the likelihood of a white buyer by 4.1 percentage points, and on the composite variable of white high-income buyers by 5.8% of a standard deviation, effectively offsetting average drops. As with the case of property prices, these coefficients are significantly different from zero (95% confidence level), as well as different from those for properties at the mean place income (p-value of the differences equal to 0.023 and 0.015, respectively.) Flood impacts on buyers' income does not change significantly with respect to preexisting place income.[37]

Figure 6 provides visual confirmation of the divergent evolution in prices for properties with different values of preexisting place income. This figure plots the evolution of property price log residuals, net of block fixed effects, separately for flooded and non-flooded properties located in the bottom decile of preexisting place income (bottom panel) and in the top decile (top panel.) It shows that flooded and control properties followed parallel trends prior to Sandy in both low and high income places (more clearly so after the housing boom ended in 2008.) In the bottom decile, there is a stark drop in prices after Sandy, which persists after five years. On the other hand, there is no evidence of a negative shock for flooded properties in the top decile after Sandy. Not only that, but prices for flooded properties are consistently above those in the control group after Sandy.

[Table 4 about here.]

[Figure 6 about here.]

Continuous treatment: vertical distance to flood level The bottom panel of table 4 reports results of using the absolute value of vertical distance to flood level[38] as flood treatment. These results show that one standard deviation of place income positively affects the impacts of one more meter of flood depth on property prices (by 4.8%), the likelihood of having a white buyer (by 2.3 percentage points), and on buyers being high-income white (by 7.7% of a standard deviation.) Significance tests of the sum of the DD and DDD coefficients[39] yield that, in places one standard deviation above mean preexisting income, one

[37] An increase in property prices with respect to income, as experienced on average by properties on high income places, has been related in recent literature in urban economics with gentrification processes (Bunten, 2018; Dragan et al., 2019). It should be noted that most authors, e.g. (Bunten, 2018), label a process as gentrification if the property price increases happen in originally low-income areas. According to this definition, the effect observed in this case could not be categorized as gentrification per se, as the increase in property prices takes place in high-income places. Some other definitions of gentrification focus on the displacement of minorities by high-income, usually white, population (Brooks et al., 2012; Reeves et al., 2019). The higher probability of a white buyer in high income places will be consistent with these definitions of gentrification as well.

[38] The steps taken to compute the metric of the distance to the flood are described in section 3

[39] Specifically, I test whether the linear combination of $After_t$ x $Flooded_p$ + $After_t$ x $Flooded_p$ x $Place\ Income_p$ is equal to zero.

more meter of flood depth does not significantly change prices, the likelihood of having a white buyer, or the buyers' income. Hence, the sensitivity of these outcomes to flood depth, present on average and in low income places, is not replicated in high-income places.

Figure 7 provides further confirmation of the divergent evolution of prices in low and high income places. It plots how prices change with respect to vertical distance to flood level before and after Sandy for properties located in the bottom decile of place income (bottom panel) and in the top decile (top panel.) In low income places, property prices post-Sandy decrease for all flood depth bins. Property prices are negatively correlated with flood depth. On the other hand, flooded property prices in high income places are *higher* than those on the control group for flood depths below 1 meter. There is no evidence of changes in prices before and after the flood for properties which experienced flood depths higher than 1.5 meters.

[Figure 7 about here.]

Differential property damage Differential responses by preexisting neighborhood income could be in part mechanically driven by differential damage. To explore this, figure 3 plots how the likelihood of a property being damaged (with a category of damage equal to *affected* or above) changes with respect to vertical distance to flood level for different quantiles of the place income variable. This figure shows that the likelihood of being damaged by the flood does not increase monotonically with place income. However, properties in the top quintile of place income do experience a smaller likelihood of being affected than those on the lower quintiles. To rule out that this difference is solely driving heterogeneity, I replicate the results on the first column of table 4 restricting the sample to those properties with the same category of damage. Results are summarized in table 5, and show that the effect of place income on flood impacts remains positive. For instance, *affected* properties in places with mean preexisting income experienced a decrease in prices equal to 3.1% post-flood on average. This negative impact is offset by 7.3% on average in equally *affected* properties located in places one standard deviation above mean preexisting place income.[40] Finally, differential damage according to place income could explain why some properties experience different levels of negative shocks due to the flood. However, it will not on its own explain why some flooded properties actually increased their value with respect to those on the control.

[Table 5 about here.]

[40] Moreover, table B.15 shows that results for property prices and white buyers in table 4 are robust to dropping from the sample properties in the top quintile of place income (although the DDD coefficient for buyers' income becomes negative.)

5.3 Estimating heterogeneity with machine learning algorithms

This section summarizes the results of applying the procedure in Chernozhukov et al. (2018) (CDDF) with a random forest algorithm to systematically evaluate heterogeneity of flood impacts in property prices according to preexisting place characteristics. Specifically, I test and rule out that flood impacts on property prices were homogeneous. Second, I corroborate that neighborhood income and secondary education, which are highly correlated in the data, best discriminate between most positively and most negatively affected properties among a rich set of preexisting place characteristics.

[Figure 8 about here.]

Is there heterogeneity in the results? The top panel of figure 8 summarizes the results of estimating heterogeneous effects of flood impacts using the CDDF procedure.[41] The model implemented uses property prices (in logs) as an outcome variable, a binary variable indicating presence in the floodplain as the treatment variable, and the preexisting place characteristics described in section 3 as covariates.

First, the figure shows that the estimated best linear predictor of average treatment effect is equal to -0.084. This value is very similar to the DD estimate presented in table2, which is equal to -0.089. This similarity serves as a check of the fitness of the procedure.

Then, the figure plots the estimated average effects for groups of properties with different levels of predicted average treatment effects.[42] Hence, group 1 is composed of 20% of the properties with the lowest values of predicted treatment effect, while group 5 is composed of 20% of the properties with the highest values of predicted treatment effect. If the results were homogeneous, there should be no significant differences in the average treatment effects between the groups.

However, in this setting, the group averages are remarkably different. The average flood impacts for group 1 were -31.2 log points, significantly below average. By design, the average treatment effect increases with group. Flood impacts in group 2 were equal to -20 log points. The average treatment effect for group 3 is not statistically different from the treatment effect for the whole sample, while the average treatment effect for group 4 is not statistically different from zero. Finally, group 5 experiences an average treatment effect equal to +17.8 log points. The difference between the point estimates for groups 1 and 5 is equal to 49 log points, and it is significant at the 1% level. Hence, these results show that the average flood effect of a 8.4% drop in prices is masking significant heterogeneities, where

[41] The procedure in Chernozhukov et al. (2018) is described in section 4.2 and appendix G

[42] In brief, the CDDF procedure predicts a treatment effect per unit of observation in the main sample, which enables the grouping of observations by their predicted treatment effect.

some properties have lost in value significantly more than this, and others have instead increased their sale prices. The hypothesis that results were homogeneous can be rejected.

It should be noted that this procedure it is agnostic regarding the source of heterogeneity: it only reveals that impacts were hot homogeneous across properties, but it does not impose heterogeneity along any variable on any functional form.

What are the characteristics of the most positively and most negatively affected properties? The bottom panel of figure 8 summarizes the average place characteristics for those properties in group 1 (the most negatively affected properties, in blue), and for those in group 5 (the most positively affected, in orange.) This figure shows the five preexisting place characteristics, of the 12 considered, that yield the highest differences between groups.[43] All preexisting place characteristics have been normalized to have a mean of zero and a standard deviation of one, so the coefficients are comparable to each other.

This figure shows that the variable along which both groups differ the most is secondary education, followed by place income. The difference between the two groups averages along these variables is around one standard deviation. Secondary education and place income are very highly correlated in the data: the bottom panel of figure A.1 shows that secondary education and income have the highest correlation among all possible pairs of variables. Other place characteristics in which the most and least affected places differ more are retail, migration, and segregation.[44]

Some of the place characteristics that the literature highlights might positively affect recovery after a flood, such as social capital (Aldrich, 2012), do not seem to be as relevant as income to distinguish between the most positively and most negatively affected properties. Hence, one of the composite variables that summarizes social capital – which according to figure A.6 represents places with low crime, high social capital index, high fraction religious, low fraction of children with single mothers, and high number of overall occupied housing – has the second smallest difference in average between the two groups, as noted in B.19.

Overall, these results show that flood impacts were not homogeneous in the sample, and that place income, together with a variable highly correlated with it – secondary education, is correlated with differences in treatment effects. This latter result then corroborates predictions from the model in section 2, which concluded that flooded properties in places with high preexisting income could increase in value post-flood.

[43] Results for the 12 preexisting characteristics are shown in table B.19. Table B.17 in turn reports the original place variables, not PCA composites, that yield the largest differences.

[44] As described in figure A.6, the secondary education variable describes places with large numbers of college graduates and colleges per capita; the income variable, places with high income and high income inequality, and low levels of middle class; the retail variable, places with large number of retail opportunities; the migration variable, places with high levels of in and out migration and foreign born residents; and the segregation variable, high racial and income segregation, and large population density.

6 Further results

This section summarizes further results to comprehensively contextualize the setting in which the results presented in the previous section take place, and to illustrate potential mechanisms. I find evidence consistent with properties in low preexisting income places relying less on flood insurance post-flood, even if flood insurance premiums have decreased. I rule out that places with high preexisting income obtain larger public assistance for reconstruction, experience changes of magnitude in the type of houses being sold after the flood, or see a change in the proportion of renters or in the likelihood of a property being a primary residence. Finally, I also present evidence against post-flood impacts being driven by properties that had not been recently flooded.

6.1 Flood insurance

A higher reliance in flood insurance, for those who could afford it, might incentivize living in risky areas (Silvis, 2017). To evaluate changes in flood insurance composition before and after Sandy, I use two datasets on the universe of anonymized individual flood insurance policies and claims underwritten by FEMA between 2010 and 2018.[45] [46]

Figure 9 and table B.4 describe the evolution of the number of flood insurance policies (in logs) for places in the bottom and top deciles of preexisting place income before and after Sandy. In both types of places, the number of policies increases after Sandy (although not significantly in the top decile), decreasing one year post-Sandy. The decline is sharper in low income places. As a result, 3 years after Sandy places in the bottom decile have fewer insurance policies than prior to Sandy. There are not significant changes (at least at the 90% confidence level) for the number of policies on the top decile.

[Figure 9 about here.]

Table B.5 explores the characteristics of the flood insurance policies.[47] Column (1) shows that insurance premiums have decreased after Sandy in census tracts with a mean level of income (-5.7%), and more so if they were flooded (further drop of -7.6%.) However, premiums have increased in relative terms for high income census tracts. One standard deviation above mean place income is correlated with 13.9% increase in premiums. There are not significant changes by preexisting place income regarding building and content insurance

[45] FEMA's National Flood Insurance Program (NFIP) is the largest provider of residential flood insurance in the United States. In 2016, 90% of all flood insurance premiums were written for the NFIP, while only 10% were for private flood insurance (Insurance Information Institute, 2018).

[46] These datasets are described in further detail in appendix E

[47] The anonymized individual policies are only identified anonymized by tract, so I construct a treatment variable equal to the percent of area covered by the floodplain by tract.

coverage of the policies. Fewer policies report elevated buildings[48] in high income places after the flood, which could partly explain the relatively higher premiums.

Finally, table B.6 reports that there have been no differential changes in terms of insurance claims related to Sandy for building and content damage with respect to preexisting place income, or that places with higher income received more claims per capita.

Hence, these results show that residents of places with high preexisting income maintain similar levels of flood insurance take-up post-flood despite increasing premiums. However, drops in property prices in low income places reported in the previous sections are happening against a backdrop in which fewer properties are insured against the flood, even with lower average premiums. This evidence is consistent with residents in low-income places opting out of the most expensive policies.

6.2 Housing stock

Property price increases and residential sorting post-flood could also be brought about by changes in the housing stock, along the lines of the "housing stock age" hypothesis as developed by Brueckner and Rosenthal (2009). They suggest that high-income households have a preference for newer housing stock, and that this determines patterns of segregation between low- and high-income households within cities. According to this hypothesis, a natural disaster such a flood could sweep outdated housing, and make room for more desirable, newer housing.[49] Under this hypothesis, observed changes in neighborhoods with higher of preexisting place income would be mostly driven by an upgrade of housing stock characteristics, rather than neighbor type, as in the model in section 2.

I do not find evidence supporting the housing renewal hypothesis in this setting. Fist, table B.1 shows that one standard deviation higher preexisting place income is not correlated with a higher likelihood of new construction among flooded properties either pre- or post-Sandy. Second, figure A.3 plots the fraction of properties in the sample that are sold in a given month[50] [51] for the whole sample, as well as separately for properties in the low and top quintiles of preexisting place income. On average, properties are not more likely to be sold post-Sandy, and this result does not change significantly with preexisting place income. The likelihood of sale for flooded properties is on a slight upward trend (every month after Sandy, the probability of a flooded property being sold increases by 0.002 percentage points

[48] A building is identified as elevated if it has no-basement, and was constructed so as to meet certain flood prevention criteria (e.g. lowest flood is above ground level, building is elevated by columns, etc.)

[49] Hornbeck and Keniston (2017), for instance, provide evidence that the Boston Fire of 1812 generated an opportunity for widespread building upgrades in burned plots. This reconstruction generated positive externalities which contributed to increased land and building prices.

[50] After removing month fixed effects, to control for seasonal variation common across years.

[51] To generate this figure, I construct a balanced panel of properties, in which the unit of observation is property-month. A dummy sale variable indicates whether a particular property was sold in a given month.

with respect to those on the control, result significant at the 90% level), but this trend does not change significantly with place income either.[52] Finally, table B.3 shows that there are statistically significant changes in the characteristics of properties being sold post-flood in high income places, but these changes are of small magnitude (flooded properties in places one standard deviation above mean place income post-Sandy have 43 more square feet and are 1 year newer on average than those on places with mean preexisting income.)

6.3 Neighborhood composition

Table B.10 shows that the likelihood of a property being occupied by its owner does not change significantly post-flood on average, nor does it change significantly with respect to place income.[53] Using data from anonymized individual insurance policies[54], I find that overall there are fewer properties that are used as primary residence in the sample, and this drop does not change significantly for properties in more extensively flooded census tracts or in places with different preexisting income (results summarized in column (5) of table B.5.) However, the proportion of primary residences prior to Sandy was higher for high income places. Hence, figure A.2 shows that properties in the top decile of place income have the largest proportion of properties used as a primary residence in the sample (point estimate is close to 80%.) These results constitute evidence against property price increases in high income places being driven by absentee landlords.

6.4 Expectations about flood risk

Different prior expectations about the likelihood of flooding could also affect the magnitude of responses to the flood. For instance, properties that are regularly flooded could be hypothesized to experience smaller price shocks as the new flood does not add much new information on the likelihood of future flooding. This could in turn yield heterogeneity in results, if places with low- and high- preexisting income have a different flood history. To gain clarity on this channel, I use data on the flood extent caused by Hurricane Irene, which made landfall in the Northeast US in August 2011. Table B.8 show that average drops in property prices, buyers' likelihood of being white or have a high income are not statistically significantly different between properties that were flooded both by Sandy and Irene, or just by Sandy.

I also investigate whether properties that are in FEMA's Special Flood Hazard Area see differential flood impacts in table B.9. It could be that the priors about future flood risk

[52]Statistical significance of these results is reported in table B.2.

[53]Data on owner-occupied status of properties comes from mortgage applications, and hence it is only available for those properties that were successfully matched with buyers' demographics

[54]As described in section 6.1.

for owners of properties in the SFHA were less affected by Sandy's flood than for those outside.[55] Table B.9 shows that in fact that drops in property prices are significantly higher for properties both in Sandy's floodplain and on the SFHA, and that properties in the SFHA post-Sandy had lower prices even if outside Sandy's floodplain. Hence, in this setting, the price effect of being in the SFHA and on Sandy's floodplain are compounded. These results need to be interpreted with caution, however. As described in more detail in appendix E, the available SFHA maps do not cover the whole area of study, and are current as of October 2018. SFHA designation could have endogenously changed after Sandy. This endogeneity on the treatment variable could bias these results.

6.5 Public Assistance

In order to evaluate whether high-income places are receiving relatively larger funds for recovery, I use data from FEMA's Public Assistance program, "FEMA's largest grant program providing funds to assist communities responding to and recovering from major declared disasters or emergencies" FEMA (2019a). Table B.7 shows that higher preexisting income is not correlated with larger quantities of recovery funds at the county level, after controlling for population and extent of area flooded in a county. However, these results should be interpreted with caution: public assistance funds are aggregated at the county level (the smallest geographical unit at which projects are identified), so inference in table B.7 is based in only 33 observations.

7 Robustness checks

This section summarizes further tests to check robustness of the results to alternative specifications, choice of control group, sample restrictions, and use of raw preexisting place characteristics rather than PCA composites.

Are results robust to the specification of fixed effects and clustering of standard errors? The preferred DDD model presented in table 4 includes census block and month-year fixed effects, and two-way clusters standard errors along census tracts and years. The impact of place income on flooded property prices is robust to using different location fixed effects (notably, to adding property controls or parcel fixed effects) as shown in table B.11; to using alternative time fixed effects (county-month, quarter, and week) as shown in table B.12; and to using different clusters of standard errors (tract month-year, tract, block month-year, zip code year, zip code month-year, and zip code) as shown in table B.13.

[55] along the lines of the mechanism proposed by Gibson and Mullins (2020).

Is there evidence of spillovers on the control group? Table B.14 runs an event study using only on properties on the control group to evaluate whether there have been any changes post-flood that could be signalling spillovers. This table shows that properties in the control group have not seen any significant changes before and after the flood with respect to property prices, buyers' likelihood of being white, their income, or their propensity to be high-income white. To gain further evidence against spillovers, figure A.4 shows that the chosen control group of properties located no more than 500 meters from the coastline does not evolve differently with respect to other properties located further inland (either between 1.5 and 2km, or between 4km or 5km.) For reference, it should be noted that no point in the island of Manhattan is located more than 4km away from the coastline.

Are results robust to the choice of sample? Results from the DD and DDD models presented in tables 2 and 4 are robust to using only observations that have been identified as single homes by CoreLogic®, the provider of property sales data. Results are summarized in table B.16. The point estimate of the effect of preexisting place income on flood impacts is larger than when using the whole sample. Results on the propensity of buyers being high-income white are noisier than with the whole sample. It should be highlighted that this single homes indicator is predicted by CoreLogic, and does not come directly from the property sales as recorded by county clerks.

Table B.15 shows that DDD results for property prices and white buyers in table 4 are robust to dropping from the sample properties in the top quintile of place income. However, the DDD coefficient for buyers' income changes sign.

Results using original place variables I replicate results of key models in the paper using raw original variables, rather than PCA composites. Hence, figure A.5 summarizes the results of including the variables in the place income composite in the DDD model described in equation 2. Table B.17 summarizes the average difference in raw place characteristics of most negatively and most positively affected properties as computed by the machine learning procedure described in section 5.3. Half of the 12 variables that yield the largest differences between the most negatively and most positively affected properties are related to preexisting income or secondary education levels of a place (e.g. Gini coefficients for income, income per capita, college graduation rate, etc.)

8 Conclusions

This paper shows that divergent post-flood changes at the neighborhood level increased preexisting spatial polarization along property value, racial, and income lines. Building on the model of segregation by Becker and Murphy (2000), I derive and test a hypothesis that preexisting income at the neighborhood scale mediates responses to the flood, whereby places with high income levels could rationally experience an increase in sale prices post-flood. To find empirical evidence of this prediction, I combine data from residential property sales, buyers' demographics, a rich set of neighborhood characteristics, and data on Superstorm Sandy's flooding extent and damage across the coastlines of the states of New Jersey, New York, Connecticut, and Rhode Island in 2012. Consistent with the model, the paper's main empirical result shows that flooded properties in neighborhoods with high preexisting income had higher sale prices and more high-income white buyers than comparable non-flooded coastal properties, offsetting average drops. Using machine learning algorithms, I corroborate that neighborhood income, of a rich set of preexisting place characteristics, best discriminates between most positively and most negatively affected properties. To the best of my knowledge, this study is the first to document differential responses after flooding at the neighborhood scale that lead to increased polarization in space using micro-level data on sale values and buyers' race and income. This paper sheds light on and helps rationalize disparate results previously found in the literature.

This paper contributes to our understanding of heterogeneous impacts of flooding across space, which is key to inform successful flood insurance schemes, and after-flooding policies more broadly. As of 2022, the National Flood Insurance Policy is undergoing reforms with the objective of determining insurance rates that are "equitable" (Service, 2022). With this equity objective in mind, it is imperative that we take into account how different communities react post-flood. In some places, lower property prices and changes in neighborhood demographics after flooding can speed the process of urban decay. The resulting community post-flood may have less capacity to economically respond or bounce back after the inevitable next storm, or even to evacuate in case of immediate danger. This dynamic could lead to the entrenchment of pockets of social vulnerability along an increasingly risky shoreline. In other places, flooded properties increase in value post-flood. For these, insurance rate discounts based on property construction year and location[56] could actually serve as regressive subsidies for high-income households.

This paper opens up several avenues for future research. The paper focuses exclusively on flood impacts on owned property. However, homeowners and renters differ along many dimensions; they have different incentives and discount rates and are characterized by different demographics (Goodman and Mayer, 2018) and social capital levels (Glaeser et al., 2002). Renters can play an important role in neighborhood change (fostering gentrification or decay.) Hence, future research could use micro-level data to explore the post-flood behavior of renters and to what extent it mimics, or fails to, that of homeowners. Also, future

[56] For instance, discounts that FEMA currently implements for properties built before flood risk maps were developed in their respective communities and for properties included in a flood risk area after a map update.

research could use micro-level data to track residents that move in and out of the floodplain to fully characterize neighborhood dynamics and provide further evidence of an increased flood risk. Who gets displaced after the flood, and where do they move? What is the original residence of those who come to live in the floodplain?

List of Figures

Figure 1: Study area: spatial extent

(a) Property sales location

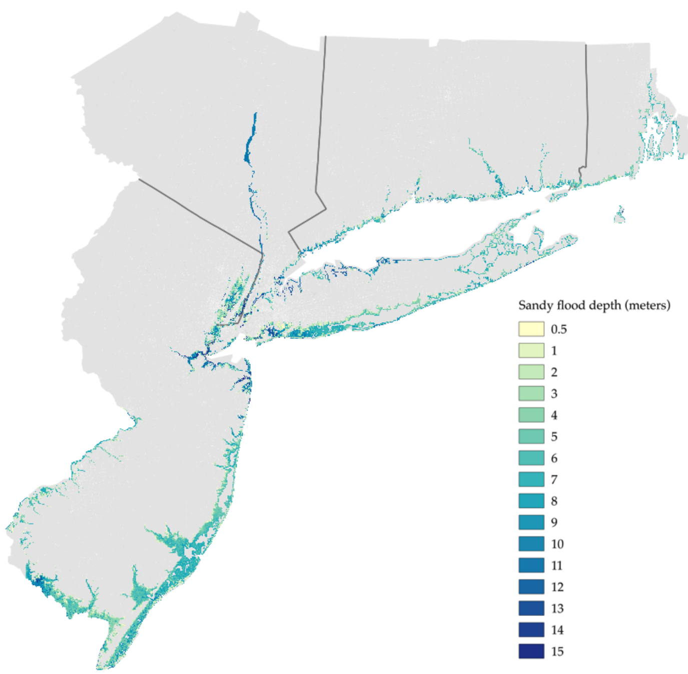

(b) Sandy's flooding

Notes: Figure delineates the area of study covering the states of New Jersey, New York, Connecticut, and Rhode Island. The top panel shows the location of the sample of property sales data, fully covering the coastline in the area of study. Each orange dot represents a property sale record in the data. The bottom panel delineates Sandy's flood extent.

Figure 2: Measuring vertical distance to flood level

(a) Flood extent boundary: layout

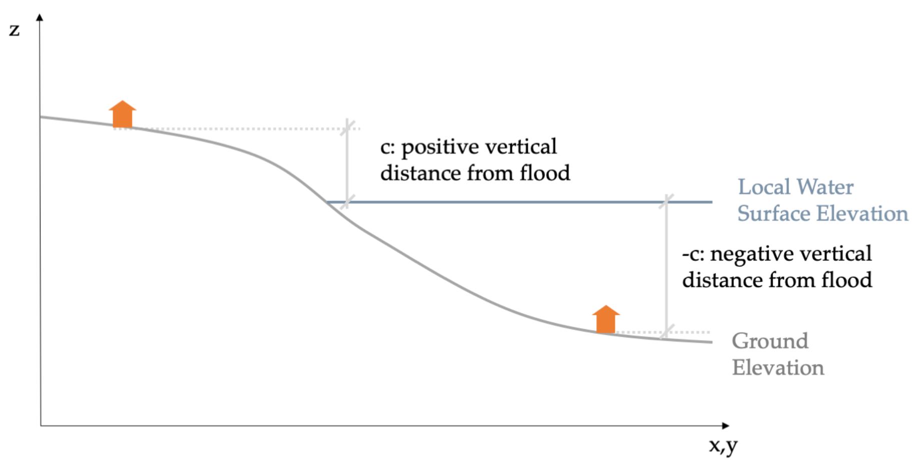

(b) Distance to the flood: metric

Notes: The top panel shows the position of the maximum flood extent boundary (blue line) with respect to property sale records (orange dots.) The bottom panel illustrates how the metric that measures vertical distance to flood level is computed. Specifically, for each property (1) I determine which point in the flood extent boundary is the closest, and (2) I compute the vertical distance to flood level metric as the difference in elevation between the properties' ground and the closest point in the flood extent boundary. It will be positive for properties whose ground stood above the flood maximum water elevation, and negative for those below.

Figure 3: Damage and vertical distance to flood level

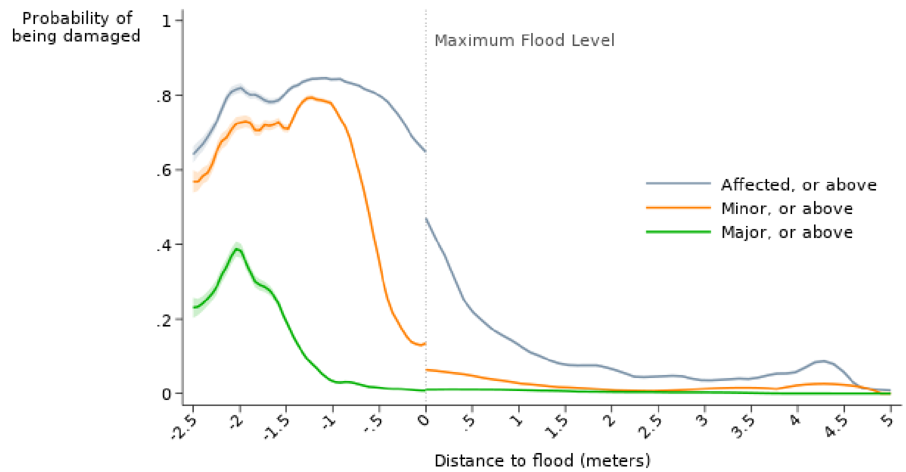

(a) Likelihood of damage with respect to vertical distance to flood level

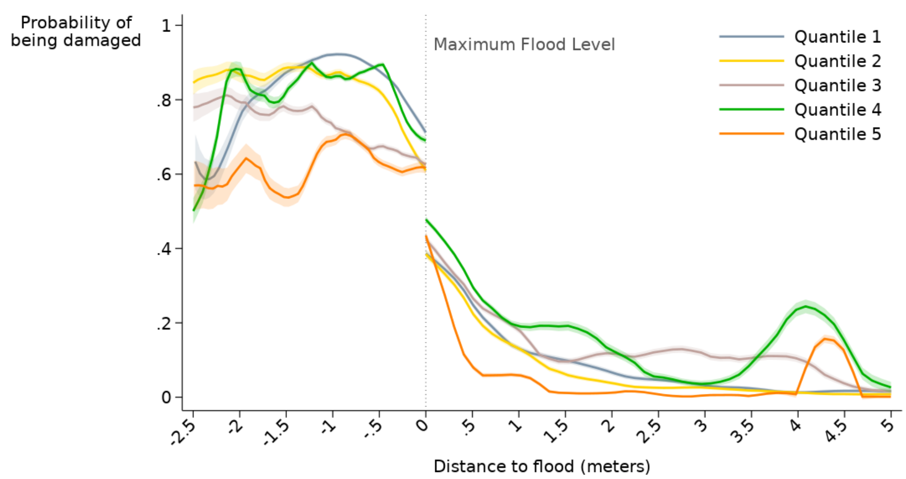

(b) Likelihood of being *affected* with respect to vertical distance to flood level

Notes: The top panel shows that the vertical distance to flood level, as measured by the metric described in section 3, is negatively correlated with damage. The probability of a property suffering damage of category *affected* or above changes discretely at the flood boundary. The bottom panel shows the likelihood of a property having a category of damage equal to *affected* or higher by quantiles of the place income variable. All lines represent smooth values from a kernel-weighted local polynomial regression, and shaded areas indicate 95% confidence intervals.

Figure 4: Illustration of identification strategy

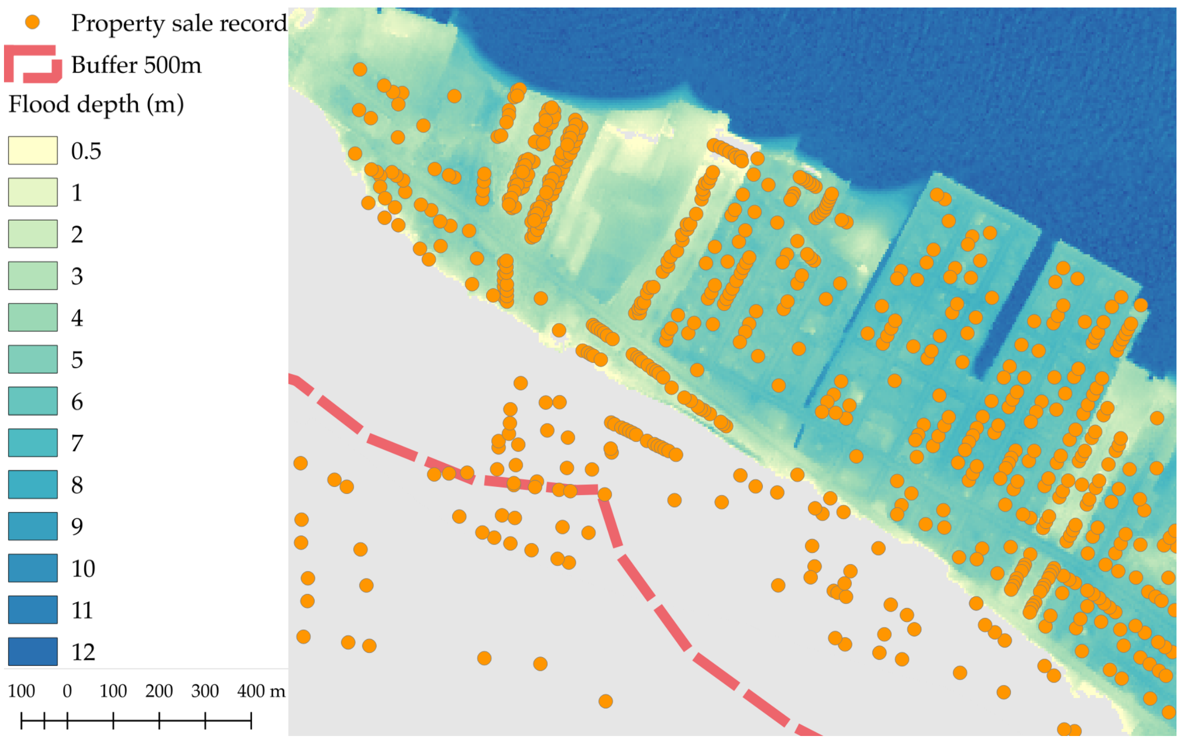

Notes: This figure zooms in the area of study to illustrate the identification strategy. Property sale records are marked in orange. The treatment group is composed of property sales which are on Sandy's floodplain (marked in yellow to blue hues, depending on the maximum depth the flood reached.) The control group is formed by property sales outside the floodplain, but within 500 meters from the coastline (500m buffer marked with a red dashed line.)

Figure 5: Average impacts of flooding on property prices

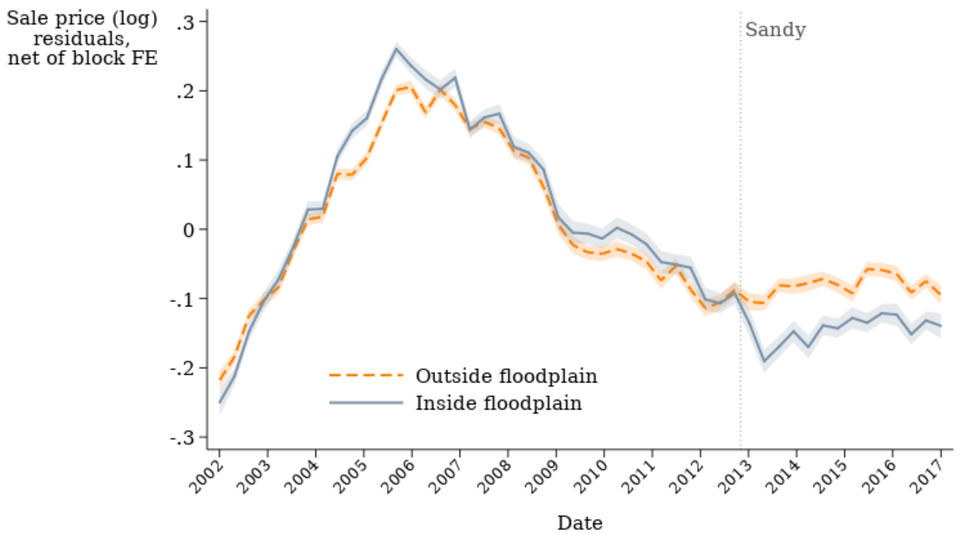

(a) Property prices over time

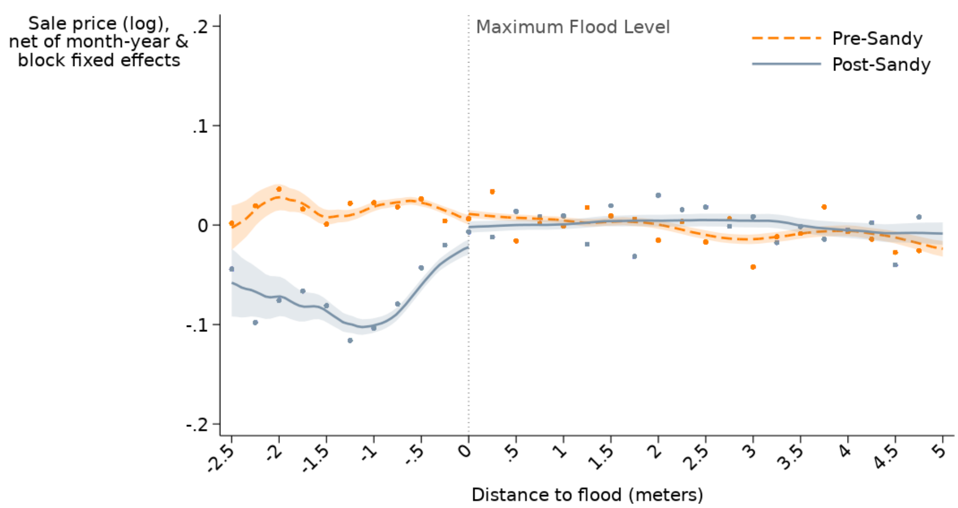

(b) Property prices over vertical distance to flood level

Notes: Figure summarizes the average impacts of flooding on property prices. Top panel plots the evolution of property sale prices – in logs, net of block fixed effects – over time. The dotted line at the end of October 2012 indicates the time of Sandy's landfall. The blue line represents properties in the Sandy floodplain (the treatment group), while the dashed orange line is for properties located within 500m of the shoreline, but not on the floodplain (the control group.) The bottom panel shows how prices evolve with vertical distance to flood level. The dashed orange line plots data before the flood, and the blue line shows data post-flood. The y-axis represents property price sale residuals – in logs, net of block and month fixed effects. The x-axis shows the difference between a property's ground level and the maximum level the flood reached closest to the property, which is negative for flooded properties, and positive for non-flooded properties. The dotted line at zero marks the maximum flood surge level. In both panels, lines represent smooth values from a kernel-weighted local polynomial regression, and shaded areas indicate 95% confidence intervals. In the bottom panel, dots represent price averages at 0.25m-wide distance bins.

Figure 6: Property prices over time, by place income

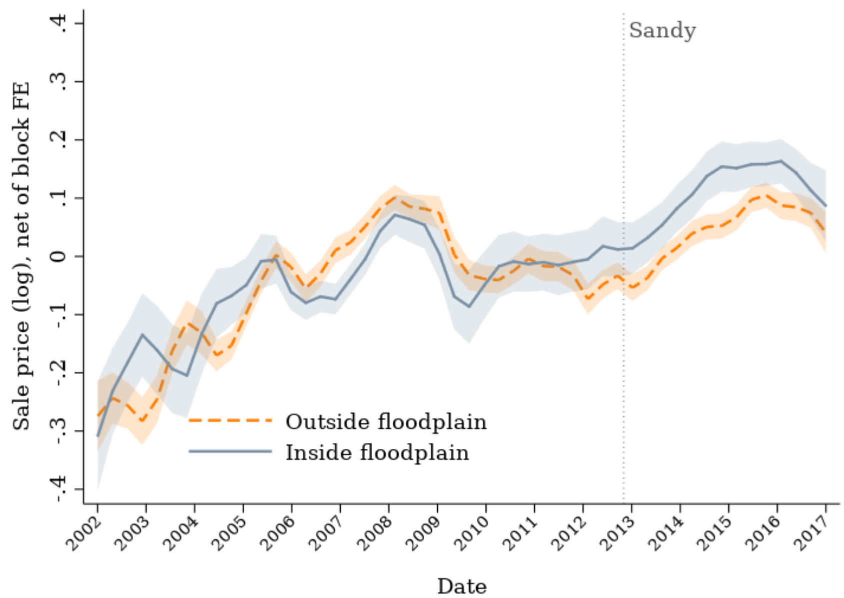

(a) Highest-income places (top 10%)

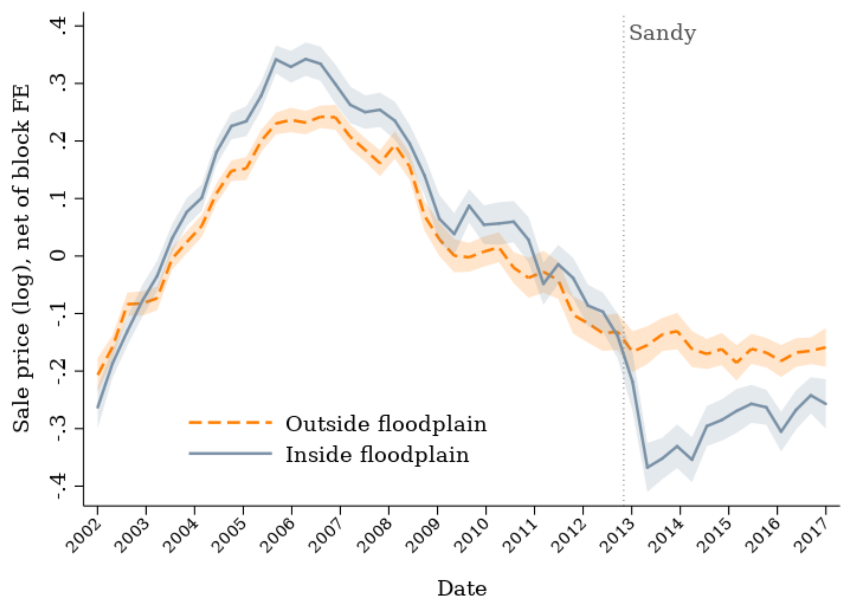

(b) Lowest-income places (bottom 10%)

Notes: Figure plots the evolution of property sale prices – in logs, after removing block fixed effects – for properties on places in the bottom decile of preexisting place income (bottom panel), and for properties on the top decile of preexisting place income (top panel). The dotted line at the end of October 2012 in both panels indicates the time of Sandy's landfall. The blue line represents properties in the Sandy floodplain (the treatment group), while the orange line is for properties located within 500m of the shoreline, but not on the floodplain (the control group.) Lines represent smooth values from a kernel-weighted local polynomial regression, shaded areas indicate 95% confidence intervals.

Figure 7: Property prices over vertical distance to flood level, by place income

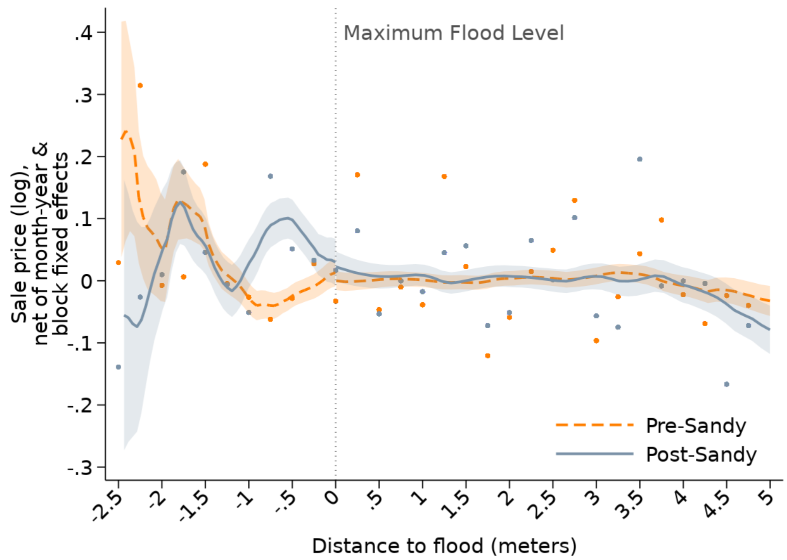

(a) Top quintile of variable INC

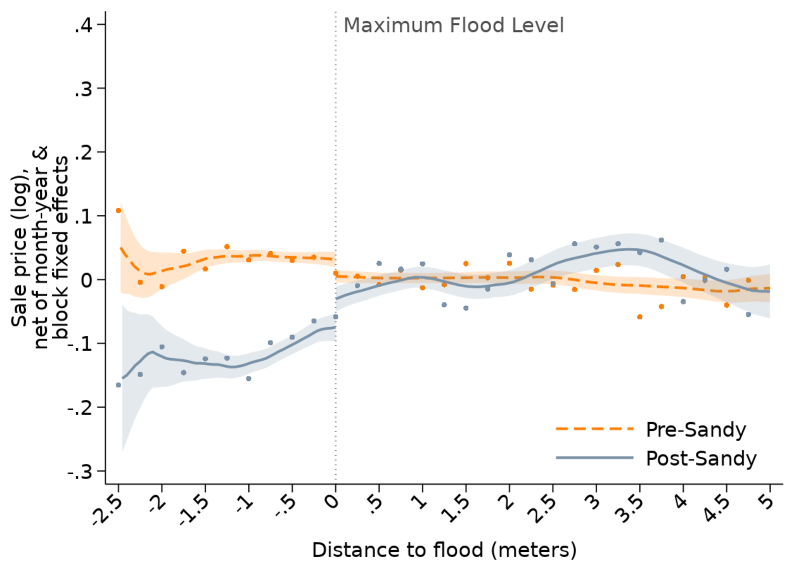

(b) Bottom quintile of variable INC

Notes: Figure shows how prices evolve with vertical distance to flood level for properties on places in the bottom decile of preexisting place income (bottom panel), and for properties on the top decile of preexisting place income (top panel). The dashed orange line plots data pre-flood, and the blue line shows data post-flood for properties on places in the bottom decile of place income (bottom panel), and for properties on the top decile of place income (top panel). Specifically, the y-axis represents property price sale residuals, net of census block and month fixed effects. The x-axis shows the difference between a property's ground level and the maximum level the flood reached closest to the property, which is negative for flooded properties, and positive for non-flooded properties. The dotted line at zero marks the maximum flood surge elevation. Lines represent smooth values from a kernel-weighted local polynomial regression, shaded areas indicate 95% confidence intervals, and dots plot price averages at 0.25m-wide elevation bins.

Figure 8: Evaluating heterogeneity using machine learning

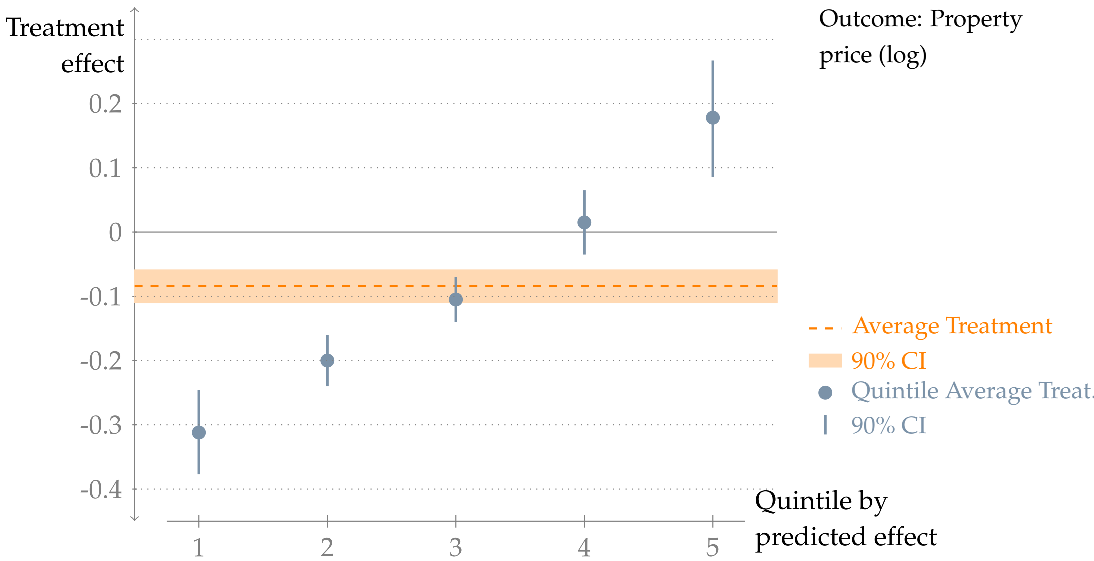

(a) Group average treatment effect

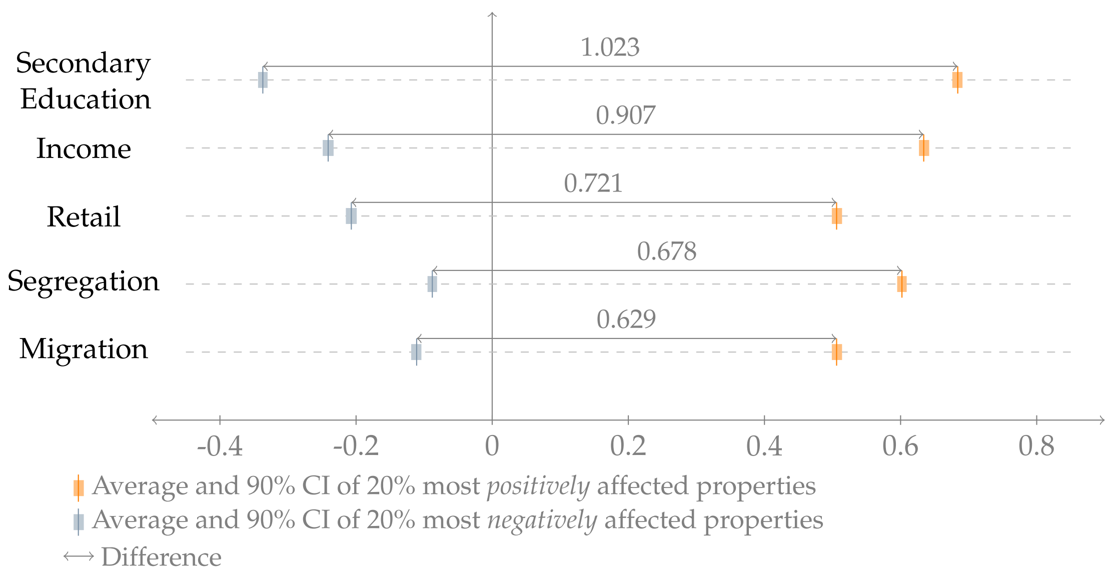

(b) Average characteristics of most positively and negatively affected properties

Notes: Top panel summarizes results from estimating heterogeneity of flood impacts according to preexisting place characteristics using the machine learning procedures described in Chernozhukov et al. (2018). Bottom panel summarizes the average difference in place characteristics for the most positively and most negatively affected properties.

Figure 9: Evolution of number of NFIP insurance policies over time

Notes: Count of flood insurance policies under the National Flood Insurance Program over time, for census tracts with a value of preexisting place income in the bottom decile (dashed, orange) and in the top decile (blue.) The vertical dotted line at the end of October 2012 represents the date Sandy made landfall. The other two lines mark the first and third anniversary of Sandy, respectively. Lines represent smooth values from a kernel-weighted local polynomial regression, shaded areas indicate 95% confidence intervals.

List of Tables

Table 1: Summary statistics

		Full sample			With HMDA data		
	Units	Mean	SD	N	Mean	SD	N
Sale price	2010$	548,314	519,988	505,113	523,108	475,728	159,145
Square feet	-	1,750	3,968	399,018	1,740	3,307	138,541
Year built	-	1964	37	354,499	1962	37	122,831
In Sandy floodplain	-	0.31	0.46	505,113	0.30	0.46	159,145
In SFHA	-	0.32	0.48	482,021	0.31	0.47	153,031
Sandy flood depth	meters	2.53	1.78	157,590	2.47	1.71	47,426
Elevation	meters	8.25	11.36	494,142	7.71	10.05	155,139
Non-Hispanic White	-				0.68	0.46	122,547
Income	2010$ (thousands)				198	299	153,043

Notes: This table shows descriptive statistics for the whole sample (leftmost three columns), and those properties that were matched with a mortgage application disclosed under the Home Mortgage Disclosure Act (HMDA), and hence had data for race and income of buyers (rightmost three columns.)

Table 2: Difference in Difference: Main results

	(1) Price (log)	(2) White	(3) Income (log)	(4) White-Income
Panel A: Binary Treatment				
Flooded_p x After_t	-0.089***	-0.050*	-0.026**	-0.101**
	(0.018)	(0.026)	(0.010)	(0.040)
Mean	548,319	0.68	193,613	0
Observations	505,113	122,547	153,043	118,979
Panel B: Continuous Treatment				
Vertical distance to flood_p x After_t	-0.054***	-0.028*	-0.002	-0.046**
	(0.015)	(0.015)	(0.009)	(0.021)
Mean	548,319	0.68	193,613	0
Observations	501,973	115,313	146,600	111,677

Notes: Difference in Difference model results. Dependent variable is the log of property price sale in column (1), a dummy variable equal to 1 if the buyer is non-Hispanic white in column (2), log income of the buyer in column (3), and a composite variable indicating whether the buyer both has high income and is white (with mean 0 and standard deviation of 1) in column (4). *After* is a dummy equal to 1 for sales at time *t* after Sandy. In the top panel, *Flooded* is a dummy equal to 1 for a property *p* in Sandy's floodplain, 0 otherwise. In the bottom panel, *Vertical distance to flood* measures the absolute value of the difference between a a property's ground level and the geographically closest maximum level the flood reached for flooded properties *p*, and is equal to 0 for non-flooded properties. All models include census block and month-year fixed effects, as well as the non-interacted variables *Flooded* or *Distance to flood* and *After* (coefficients not shown for clarity.) Standard errors two-way clustered at the census tract and year level. * $p< 0.1$, ** $p< 0.05$, *** $p< 0.01$.

Table 3: Difference in Difference: flood impacts according to damage levels

	(1) Price (log)	p-value = Affected_p	p-value = Minor_p	p-value = Major_p
Flooded_p x After_t x Affected_p	-0.082***	-	-	-
	(0.018)			
Flooded_p x After_t x Minor_p	-0.138***	[0.025]	-	-
	(0.023)			
Flooded_p x After_t x Major_p	-0.171***	[0.083]	[0.428]	-
	(0.047)			
Flooded_p x After_t x Destroyed_p	-0.409***	[0.005]	[0.019]	[0.054]
	(0.098)			
Mean	548,314			
Observations	501,973			

Notes: Difference in Difference model results, according to damage level. Column (1) regresses the log of property sale prices on a treatment variable (*Flooded* a dummy equal to 1 for properties *p* in the floodplain, and *After* is a dummy equal to 1 for sales at time *t* after Sandy) interacted with a set of dummy variables denoting the level of damage according to FEMA's evaluation, as follows: *Affected* if less than 20% of roof has been damaged or the flood depth is less than 2 feet; *Minor* if more than 20% of roof covering is missing or the flood depth is between 2 to 5 feet; *Major* if some exterior walls have collapsed or the flood depth is higher than 5 feet; *Destroyed* if exterior walls have collapsed. The model in column (1) also includes census block and month-year fixed effects, as well as variables *Flooded* and *After* and their interactions with the damage dummies (coefficients not shown for clarity.) Standard errors two-way clustered at the census tract and year level. The second column shows the p-value of the difference between the coefficient in each row with that of the interaction between the treatment variable and the *Affected* category; the third column, with that of the *Minor* category; and the fourth column, with that of *Major* category. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Difference in Difference in Differences by place income: Main results

	(1) Price (log)	(2) White	(3) Income (log)	(4) White-Income
Panel A: Binary Treatment				
$After_t$ x $Flooded_p$	-0.045***	-0.043*	-0.021*	-0.083**
	(0.015)	(0.024)	(0.010)	(0.036)
$After_t$ x Place $Income_p$	0.102***	0.000	0.061***	0.056**
	(0.016)	(0.012)	(0.009)	(0.022)
$After_t$ x $Flooded_p$ x Place $Income_p$	0.073***	0.041**	-0.010	0.058**
	(0.014)	(0.016)	(0.016)	(0.026)
Mean	548,306	0.68	193,626	0
Observations	504,971	122,533	153,022	118,965
Panel B: Continuous Treatment				
$After_t$ x V. distance to $flood_p$	-0.037**	-0.025	0.004	-0.035*
	(0.013)	(0.014)	(0.010)	(0.019)
$After_t$ x Place $Income_p$	0.108***	0.005	0.056***	0.058**
	(0.016)	(0.013)	(0.008)	(0.023)
$After_t$ x V. distance to $flood_p$ x Place $Income_p$	0.048***	0.023*	0.033	0.077**
	(0.014)	(0.011)	(0.021)	(0.027)
Mean	548,306	0.68	193,626	0
Observations	504,971	122,533	153,022	118,965

Notes: Difference in Difference in Difference model results, according to preexisting place income. Dependent variable is the log of property price sale in column (1), a dummy variable equal to 1 if the buyer is non-Hispanic white in column (2), log income of the buyer in column (3), and a composite variable indicating whether the buyer both has high income and is white (with mean 0 and standard deviation of 1) in column (4). *After* is a dummy equal to 1 for sales at time *t* after Sandy. *Place income* is a composite variable that describes preexisting income of the place property *p* is located. It is normalized to have a mean of zero and standard deviation equal to 1, and it is positively correlated to places with high income per capita and income inequality, and low fraction middle class. In the top panel, *Flooded* is a dummy equal to 1 for a property *p* in Sandy's floodplain, 0 otherwise. In the bottom panel, *V. distance to flood* measures the absolute value of the difference between a a property's ground level and the geographically closest maximum level the flood reached for flooded properties *p*, and is equal to 0 for non-flooded properties. All models include census block and month-year fixed effects, as well as the non-interacted variables *Flooded* or *V. distance to flood*, *After*, *Place Income*, and all of their interactions (some coefficients not shown for clarity.) Standard errors two-way clustered at the census tract and year level. * $p< 0.1$, ** $p< 0.05$, *** $p< 0.01$.

Table 5: Difference in Difference in Differences by place income and damage category

	(1) Affected	(2) Minor	(3) Major	(4) Destroyed
$After_t$ x Damage $Level_p$	-0.031*	-0.082***	-0.074	-0.578
	(0.016)	(0.018)	(0.053)	(0.347)
$After_t$ x Place $Income_p$	0.098***	0.099***	0.098***	0.098***
	(0.016)	(0.016)	(0.016)	(0.016)
$After_t$ x Damage $Level_p$ x Place $Income_p$	0.073***	0.104***	0.182*	-0.186
	(0.020)	(0.015)	(0.086)	(0.603)
Observations in treatment	77,625	66,532	9,532	237
Observations	425,075	413,901	356,975	347,860

Notes: Difference in Difference in Difference model results, by preexisting place income and damage level. The dependent variable in all models is the the log of property sale prices. Each column restricts the sample of treated units to those with the same level of damage according to FEMA's evaluation: *affected* in column (1), *minor* in column (2), *major* in column (3), and *destroyed* in column (4). The control group in all models is composed by properties declared unaffected by FEMA's evaluation. *Damage level* is a dummy equal to 1 for properties *p* assigned the level of damage on the column title. *After* is a dummy equal to 1 for sales at time *t* after Sandy. *Place income* is a composite variable that describes preexisting income of the place property *p* is located. It is normalized to have a mean of zero and standard deviation equal to 1, and it is positively correlated to places with high income per capita and income inequality, and low fraction middle class. All the models include census block and month-year fixed effects, as well as the non-interacted variables *Damage Level*, *After*, *Place Income*, and all of their interactions (some coefficients not shown for clarity.) Standard errors two-way clustered at the census tract and year level. * $p< 0.1$, ** $p< 0.05$, *** $p< 0.01$.

References

Aldrich, D. P. (2012). *Building Resilience*. Social Capital in Post-Disaster Recovery. University of Chicago Press.

Alesina, A. and La Ferrara, E. (2000). Participation in Heterogeneous Communities. *The Quarterly Journal of Economics*, pages 1–59.

Andaloussi, M. B. and Isaksen, E. T. (2017). The environmental and distributional consequences of emissions markets: evidence from the Clean Air Interstate Rule. *Working Paper*, pages 1–71.

Athey, S. and Imbens, G. W. (2019). Machine Learning Methods Economists Should Know About. *Working Paper*, pages 1–62.

Bakkensen, L. A. and Barrage, L. (2017). Flood Risk Belief Heterogeneity and Coastal Home Price Dynamics: Going Under Water? . *NBER Working Paper Series*, pages 1–55.

Bakkensen, L. A. and Barrage, L. (2018). Climate Shocks, Cyclones, and Economic Growth: Bridging the Micro-Macro Gap . *NBER Working Paper Series*, pages 1–28.

Bakkensen, L. A. and Ma, L. (2019). Distributional Impacts of National Flood Insurance Program Reform. *Working Paper*, pages 1–61.

Banzhaf, H. S. and Walsh, R. P. (2013). Segregation and Tiebout sorting: The link between place-based investments and neighborhood tipping. *Journal of Urban Economics*, 74(C):83–98.

Barnes, E. A., Polvani, L. M., and Sobel, A. H. (2013). Model projections of atmospheric steering of Sandy-like superstorms. *Proceedings of the National Academy of Sciences*, 110(38):15211–15215.

Barr, J., Cohen, J., and Kim, E. (2017). Storm Surges, Informational Shocks, and the Price of Urban Real Estate: An Application to the Case of Hurricane Sandy. *Working Paper*, pages 1–47.

Bayer, P., Ferreira, F., and McMillan, R. (2007). A Unified Framework for Measuring Preferences for Schools and Neighborhoods. *Journal of Political Economy*, 115(4):588–638.

Bayer, P., McMillan, R., Murphy, A., and Timmins, C. (2016). A Dynamic Model of Demand for Houses and Neighborhoods. *Econometrica*, 84(3):893–942.

Becker, G. and Murphy, K. (2000). *Social Economics*. Market Behavior in a Social Environment.

Beltran, A., Maddison, D., and Elliott, R. J. R. (2018). Is Flood Risk Capitalised Into Property Values? *Ecological Economics*, 146:668–685.

Bento, A. M. (2013). Equity Impacts of Environmental Policy. *Annual Review of Resource Economics*, 5(1):181–196.

Bergman, P., Chetty, R., DeLuca, S., Hendren, N., Katz, L. F., and Palmer, C. (2019). Creating Moves to Opportunity: Experimental Evidence on Barriers to Neighborhood Choice . *NBER Working Paper Series*, pages 1–89.

Bergren, E., Coffey, J., Aldana Cohen, D., Crowley, N., Koslov, L., Liboiron, M., Merdjanoff, A., Murphree, A., and Wachsmuth, D. (2013). A Tale of Two Sandys. Technical report, Superstorm Research Lab.

Bernstein, A., Gustafson, M. T., and Lewis, R. (2019). Disaster on the horizon: The price effect of sea level rise. *Journal of Financial Economics*, 134(2):253–272.

Bin, O., Brown Kruse, J., and Landry, C. (2008). Flood Hazards, Insurance Rates, and Amenities: Evidence From the Coastal Housing Market. pages 1–20.

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

Black, S. (1999). Do Better Schools Matter? Parental Valuation of Elementary Education. *The Quarterly Journal of Economics*, 114(2):577–599.

BLS (2019). Consumer Price Index, Northeast Region.

Bolin, B. (2007). Chapter 7: Race, Class, Ethnicity, and Disaster Vulnerability. In Rodríguez, H., Quarantelli, E. L., and Dynes, R. R., editors, *Handbook of Disaster Research*.

Brooks, N., Donaghy, K., and Knaap, G.-J. (2012). *The Oxford Handbook of Urban Economics and Planning*.

Brueckner, J. K. and Rosenthal, S. S. (2009). Gentrification and Neighborhood Housing Cycles: Will America's Future Downtowns Be Rich? *Review of Economics and Statistics*, 91(4):725–743.

Bunten, D. (2018). When the Neighborhood Goes: The Effects of Gentrification on Individuals and Communities.

Bureau, U. C. (2011). What are census blocks?

Calonico, S., Cattaneo, M. D., and Titiunik, R. (2014). Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs. *Econometrica*, 82(6):2295–2326.

Card, D., Mas, A., and Rothstein, J. (2008). Tipping and the Dynamics of Segregation. *The Quarterly Journal of Economics*, 123:177–218.

Chay, K. Y. and Greenstone, M. (2005). Does Air Quality Matter? Evidence from the Housing Market. *Journal of Political Economy*, 113(2):376–424.

Chernozhukov, V., Demirer, M., Duflo, E., and Fernandez-Val, I. (2018). Generic Machine Learning Inference on Heterogenous Treatment Effects in Randomized Experiments . *NBER Working Paper Series*, pages 1–40.

Chetty, R., Friedman, J. N., Hendren, N., Jones, M. R., and Porter, S. R. (2018). The Opportunity Atlas: Mapping the Childhood Roots of Social Mobility. *NBER Working Paper Series*, pages 1–95.

Chetty, R. and Hendren, N. (2017). The Impacts of Neighborhoods on Intergenerational Mobility II: County-Level Estimates. *The Quarterly Journal of Economics*, pages 1–95.

Chetty, R., Hendren, N., Kline, P., and Saez, E. (2014). Where is the land of Opportunity? The Geography of Intergenerational Mobility in the United States *. *The Quarterly Journal of Economics*, 129(4):1553–1623.

Christensen, P. and Timmins, C. (2018). Sorting or Steering: Experimental Evidence on the Economic Effects of Housing Discrimination . *NBER Working Paper Series*, pages 1–66.

Collinson, R. and Ganong, P. (2018). How Do Changes in Housing Voucher Design Affect Rent and Neighborhood Quality? *American Economic Journal: Economic Policy*, 10(2):62–89.

Couture, V., Gaubert, C., Handbury, J., and Hurst, E. (2021). Income growth and the distributional effects of urban spatial sorting. *Review of Economic Studies*.

Davis, J. M. V. and Heller, S. B. (2017). Using Causal Forests to Predict Treatment Heterogeneity: An Application to Summer Jobs. *American Economic Review: Papers & Proceedings*, 107(5):546–50.

Davis, L. W. (2007). The Effect of Health Risk on Housing Values: Evidence from a Cancer Cluster. *American Economic Review*, 94(5):1693–1704.

del Valle, A., Scharlemann, T., and Shore, S. (2019). Household Financial Behavior After Hurricane Harvey. *Working Paper*, pages 1–36.

Depro, B., Timmins, C., and O'Neil, M. (2015). White Flight and Coming to the Nuisance: Can Residential Mobility Explain Environmental Injustice? *Journal of the Association of Environmental and Resource Economists*, 2(3):439–468.

Deryugina, T., Heutel, G., Miller, N. H., Molitor, D., and Reif, J. (2019). The Mortality and Medical Costs of Air Pollution: Evidence from Changes in Wind Direction. *American Economic Review*, pages 1–103.

Diamond, R. (2016). The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980–2000. *American Economic Review*, 106(3):479–524.

Dragan, K., Gould Ellen, I., and Glied, S. A. (2019). Does Gentrification Displace Poor Children? New Evidence from New York City Medicaid Data. *NBER Working Paper Series*, pages 1–40.

Duranton, G., Gobillon, L., and Overman, H. G. (2011). Assessing the Effects of Local Taxation using Microgeographic Data. *The Economic Journal*, 121(55):1017–1104.

EPA (2016). Climate Change Indicators: Coastal Flooding. Technical report.

FEMA (2014). Hurricane Sandy Impact Analysis .

FEMA (2019a). FEMA - Data Feeds.

FEMA (2019b). FEMA Flood Map Service Center.

FEMA (2019c). FIMA NFIP Redacted Claims Data Set.

FEMA (2019d). FIMA NFIP Redacted Policies Data Set.

Fernandez, M. (2018). A Year After Hurricane Harvey, Houston's Poorest Neighborhoods Are Slowest to Recover. *The New York Times*.

Fogli, A. and Guerrieri, V. (2019). The End of the American Dream? Inequality and Segregation in US Cities. *NBER Working Paper Series*, pages 1–59.

FRED St Louis (2019). Consumer Price Index for All Urban Consumers: Housing in U.S. City Average.

Fuller, T., Turkewitz, J., and Del Real, J. A. (2019). Despair for Many and Silver Linings for Some in California Wildfires. *The New York Times*.

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

Gamper-Rabindran, S. and Timmins, C. (2013). Does cleanup of hazardous waste sites raise housing values? Evidence of spatially localized benefits. *Journal of Environmental Economics and Management*, 65(3):345–360.

Geddes, J. A., Martin, R. V., Boys, B. L., and Van Donkelaar, A. (2017). Global 3-Year Running Mean Ground-Level Nitrogen Dioxide (NO2) Grids from GOME, SCIAMACHY and GOME-2.

Gelman, A. and Imbens, G. (2017). Why High-Order Polynomials Should Not Be Used in Regression Discontinuity Designs. *Journal of Business & Economic Statistics*, 37(3):447–456.

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

Glaeser, E. L., Laibson, D., and Sacerdote, B. (2002). An Economic Approach to Social Capital. *The Economic Journal*, 112(483):F437–F458.

Goodman, L. S. and Mayer, C. (2018). Homeownership and the American Dream. *Journal of Economic Perspectives*, 32(1):31–58.

Gordon, K. (2014). Risky Business: The Economic Risks of Climate Chante in the United States. Technical report.

Graff Zivin, J., Liao, Y. L., and Panassie, Y. (2020). How hurricanes sweep up housing markets: Evidence from florida. *NBER Working Paper Series*.

Graham, B. S. (2018). Identifying and Estimating Neighborhood Effects. *Journal of Economic Literature*, 56(2):450–500.

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.

Greenstone, M. and Gallagher, J. (2011). Does Hazardous Waste Matter? Evidence from the Housing Market and the Superfund Program. *The Quarterly Journal of Economics*, 123(3):951–1003.

Hallstrom, D. G. and Smith, V. K. (2005). Market responses to hurricanes. *Journal of Environmental Economics and Management*, 50(3):541–561.

Heal, G. and Tedesco, M. (2018). Rising Oceans: Economics and Science. *Working Paper*, 115(9):2022–2025.

Hess, H., Delgado, M., Hamidi, A., Houser, T., Kopp, R., Bolliger, I., Hsiang, S., and Greenstone, M. (2019). New Jersey's Rising Coastal Risk. Technical report.

Horn, D. P. and Webel, B. (2019). Introduction to the National Flood Insurance Program (NFIP). *Congressional Research Service*, pages 1–31.

Hornbeck, R. and Keniston, D. (2017). Creative Destruction: Barriers to Urban Growth and the Great Boston Fire of 1872. *American Economic Review*, 107(6):1365–1398.

Insurance Information Institute (2018). Spotlight on: Flood insurance.

James, G., Witten, D., Hastie, T., and Tibshirani, R. (2015). *An Introduction to Statistical Learning*. with Applications in R.

Jina, A. and Hsiang, S. (2014). The Causal Effect of Environmental Catastrophe on Long-Run Economic Growth: Evidence From 6,700 Cyclones. *NBER Working Paper Series*, pages 1–69.

Kahn, M. E. (2005). The Death Toll from Natural Disasters: The Role of Income, Geography, and Institutions . *The Review of Economics and Statistics*, 87(2):271–284.

Keenan, J. M., Hill, T., and Gumber, A. (2018). Climate gentrification: from theory to empiricism in Miami-Dade County, Florida. *Environmental Research Letters*, 13(5):054001–12.

Kellenberg, D. K. and Mobarak, A. M. (2008). Does rising income increase or decrease damage risk from natural disasters? *Journal of Urban Economics*, 63(3):788–802.

Lackner, S. (2019). Earthquakes and Economic Growth. *Job Market Paper*, pages 1–61.

Lee, D. S. and Lemieux, T. (2010). Regression Discontinuity Designs in Economics. *Journal of Economic Literature*, 48(June):281–355.

Liao, Y. and Panassie, Y. (2018). How Hurricanes Sweep Up Housing Markets: Evidence from Florida. *Job Market Paper*, pages 1–63.

Ma, L., Timmins, C., and Banzhaf, H. S. (2019). Environmental Justice: The Economics of Race, Place, and Pollution. *Journal of Economic Perspectives*, 33(1):185–208.

MTA (2016). Sandy-Related Repairs Coming to Clark St 2/3 Tunnel.

MTA (2019). L Project.

Muller, N. Z. and Hopkins, C. (2019). Hurricane Katrina Floods New Jersey: The Role of Information in the Market Response to Flood Risk. *NBER Working Paper Series*, pages 1–96.

Mutter, J. C. (2015). *The Disaster Profiteers*. How Natural Disasters Make the Rich Richer and the Poor Even Poorer. St. Martin's Press.

NCES (2019). National Center for Education Statistics - Data.

NOAA (2016). NOAA Medium Resolution Shoreline.

NOAA National Centers for Environmental Information (2022). U.S. Billion-Dollar Weather and Climate Disasters.

Open Street Map (2019). OpenStreetMap.

Ortega, F. and Taspinar, S. (2018). Rising sea levels and sinking property values: Hurricane Sandy and New York's housing market. *Journal of Urban Economics*, 106:81–100.

Perry, C. A. (2000). Significant Floods in the United States During the 20th Century USGS Measures a Century of Floods.

PRISM (2019). PRISM - Climate Data.

Ramey, C. (2015). New York Disputes FEMA on Flood Risk. *The Wall Street Journal*.

Reeves, J., Boustan, L. P., Margo, R. A., Miller, M. M., and Steil, J. P. (2019). Does Condominium Development Lead to Gentrification? pages 1–31.

Rigol, N., Hussam, R., and Roth, B. (2017). Targeting High Ability Entrepreneurs Using Community Information: Mechanism Design In The Field. *Job Market Paper*, pages 1–82.

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

Schelling, T. C. (1969). Models of segregation. *American Economic Review*, 59(2):488–493.

Service, C. R. (2022). National flood insurance program: The current rating structure and risk rating 2.0. Technical report, US Congress.

Sethi, R. and Somanathan, R. (2004). Inequality and Segregation. *Journal of Political Economy*, 112(6):1296–1321.

Shafroth, F. (2017). The Before and After Effects of Flooding on Property Taxes.

Siders, A. R. (2018). Social justice implications of US managed retreat buyout programs. *Climatic Change*, 152(2):239–257.

Silvis, V. G. (2017). Flooding by Design: A Look at the National Flood Insurance Program. *Risk, Hazards & Crisis in Public Policy*, 9(1):82–99.

Smith, V. K., Carbone, J. C., Pope, J. C., Hallstrom, D. G., and Darden, M. E. (2006). Adjusting to natural disasters. *Journal of Risk and Uncertainty*, 33(1-2):37–54.

Smith, V. K. and Whitmore, B. (2019). Amenities, Risk, and Flood Insurance Reform . *NBER Working Paper Series*, pages 1–62.

Sobel, A. (2014). *Storm Surge*. Hurricane Sandy, Our Changing Climate, and Extreme Weather of the Past and Future. Harper Collins.

Stanford CEPA (2019). Stanford Education Data Archive.

Tiebout, C. M. (1956). A pure theory of local expenditures. *Journal of Political Economy*.

Toro, R. (2013). Hurricane Sandy's Impact (Infographic).

US Census (2019). County Business Patterns (CBP).

USGS (2019). 3D Elevation Program (3DEP).

Van Donkelaar, A., Martin, R. V., Brauer, M., Hsu, N. C., Kahn, R. A., Levy, R. C., Lyapustin, A., Sayer, A. M., and Winker, D. M. (2018). Global annual PM2.5 grids from MODIS, MISR and SeaWiFS Aerosol Optical Depth (AOD) with GWR, 1998-2016.

van Holm, E. J. and Wyczalkowski, C. K. (2018). Gentrification in the wake of a hurricane: New Orleans after Katrina. *Urban Studies*, 9(1):004209801880044–16.

Villarreal, C. (2015). Where the Other Half Lives: Evidence on the Origin and Persistence of Poor Neighborhoods from New York City 1830-2012. *Working Paper*, pages 1–64.

Wager, S. and Athey, S. (2018). Estimation and Inference of Heterogeneous Treatment Effects using Random Forests. pages 1–16.

Woods, C. and Edwards, M. C. (2011). Factor Analysis and Related Methods. *Essential Statistical Methods for Medical Statistics*, pages 174–201.

Appendices

Appendix A Additional figures

Figure A.1: Preexisting place characteristics

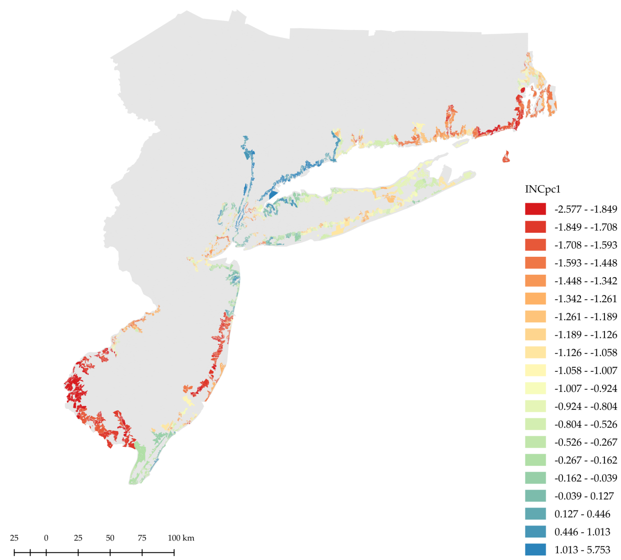

(a) Spatial variation of preexisting place income

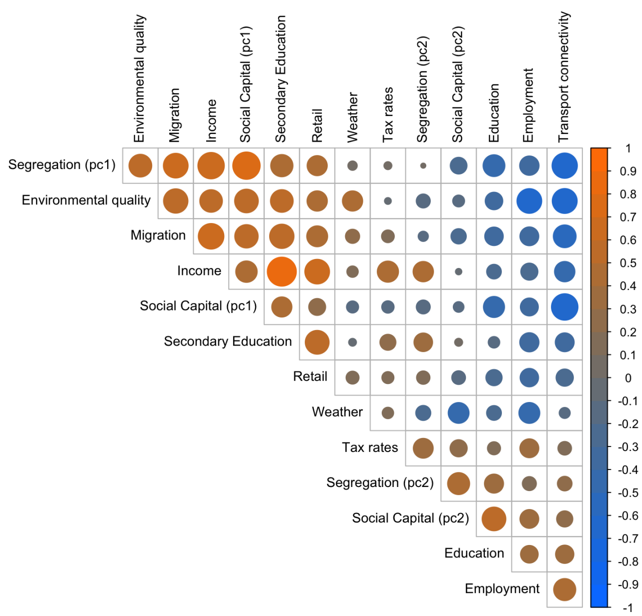

(b) Correlation among preexisting place characteristics

Notes: The top panel shows spatial variation of preexisting place income, from high values (in blue, green hues) to low (in red, orange hues.) Places which rank high in this variable have high income per capita and high income inequality, and a low fraction of middle class households. The bottom panel summarizes correlation among preexisting place characteristics. Positive correlation is plotted in orange hues, and negative in blue hues. The area of each circle is proportional to the Pearson correlation coefficient between the relevant row and column variables.

Figure A.2: Primary residencies by decile of place income

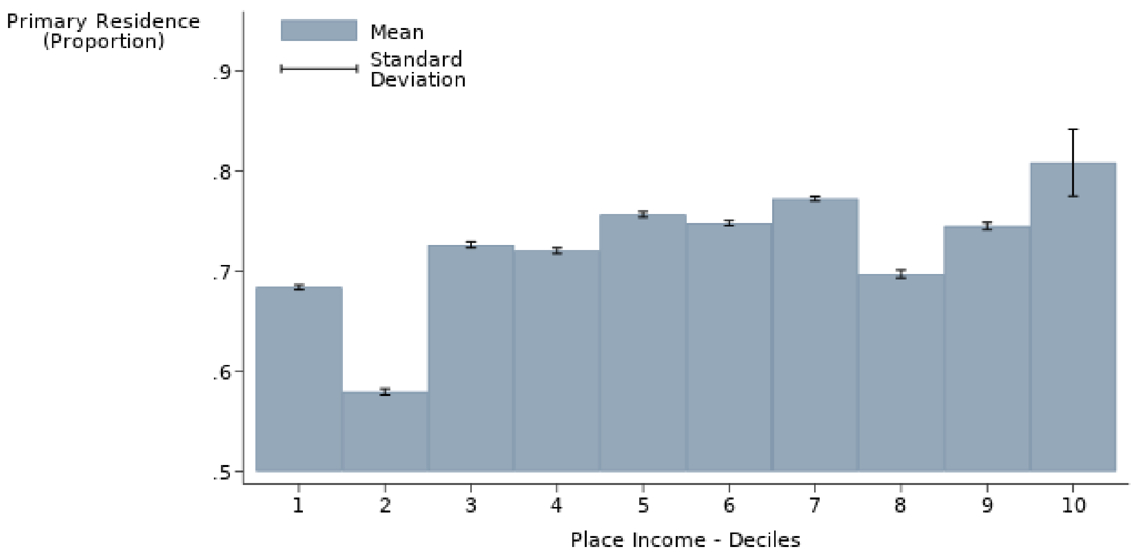

Notes: Figure shows average and standard deviation of residences that are primary residences by decile of preexisting place income. Data source are flood insurance policies underwritten by FEMA between 2010 and 2012 (prior to Sandy.)

Figure A.3: Evolution over time of fraction of properties sold

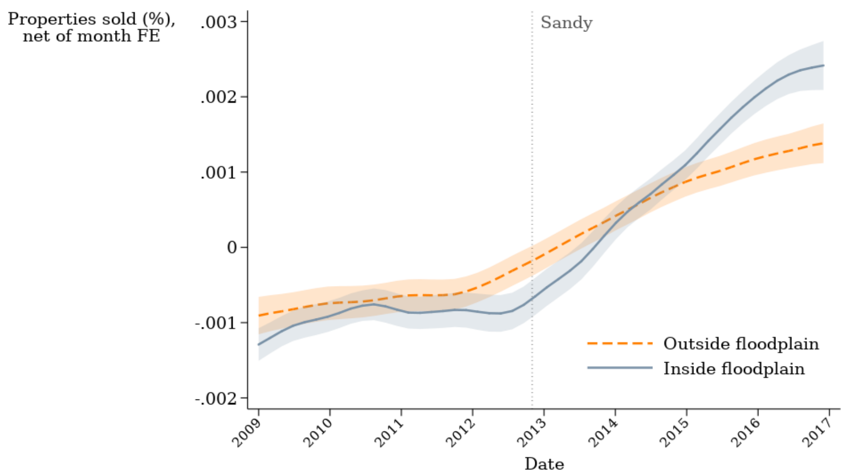

(a) Whole sample

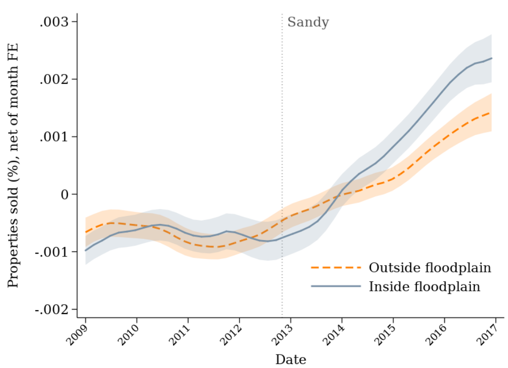

(b) Bottom quintile of place income

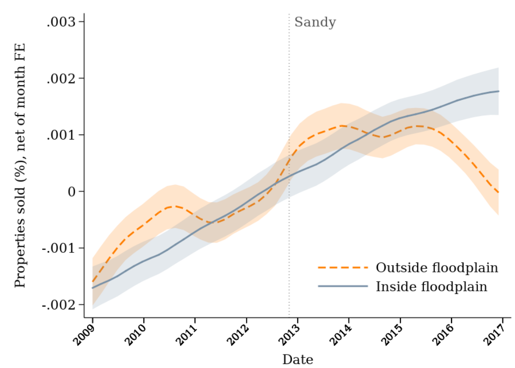

(c) Top quintile of place income

Notes: Figure shows the fraction of properties in the sample sold in a given month, net of month fixed effects to remove seasonal variation common across years. Lines show smooth values from a kernel-weighted local polynomial regression for properties in the floodplain (in blue) and those outside (in orange.) The top panel plots the whole sample, the bottom left panel uses data from properties on places in the bottom quintile of preexisting place income, and the bottom right panel uses data from for properties on the top quintile of preexisting place income. The dotted line at the end of October 2012 indicates the time of Sandy's landfall.

Figure A.4: Inland control group

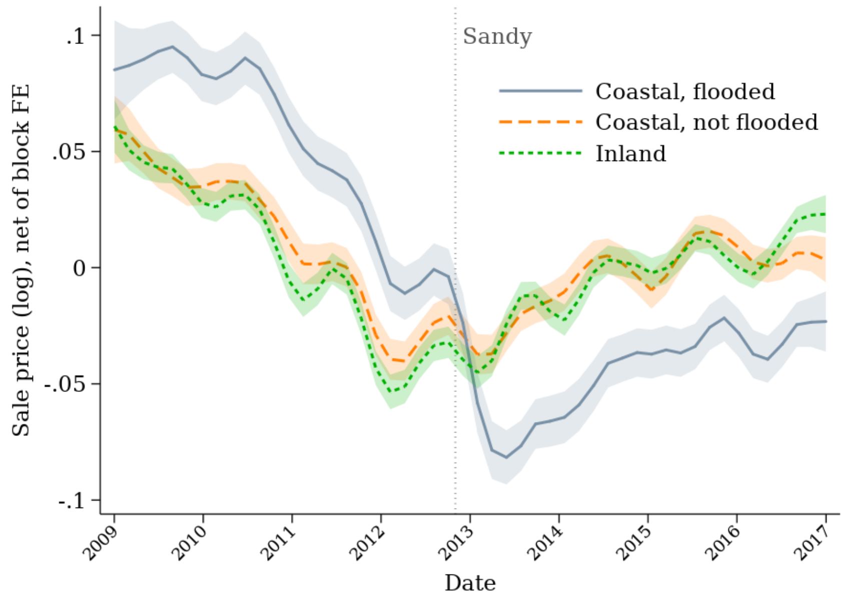

(a) Inland control: between 1.5 and 2 km from coast

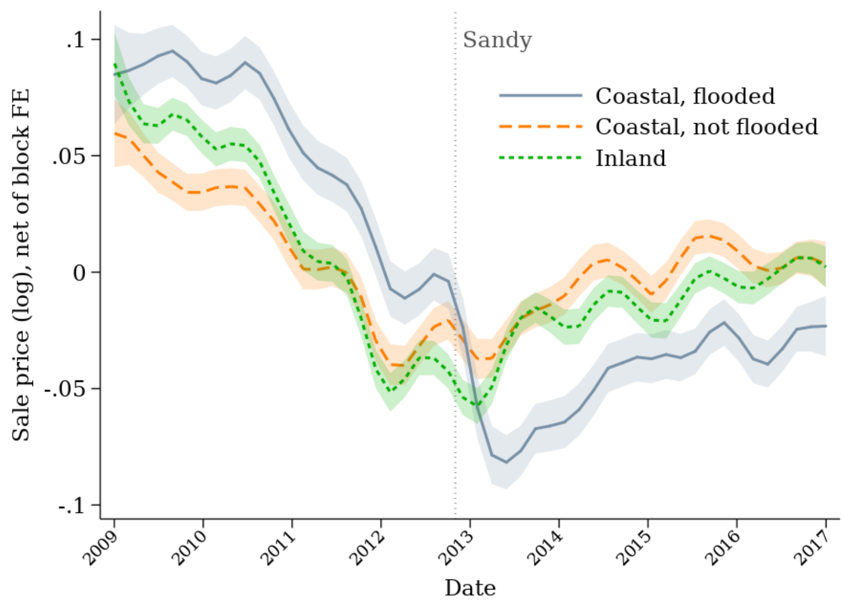

(b) Inland control: between 4 and 5 km from coast

Notes: Figure plots the evolution of property sale prices (logs, net of block fixed effects) for properties in Sandy's floodplain (in blue), properties outside the floodplain located no more than 500m from the coast (in orange), and a further inland control (in green.) Lines represent smooth values from a kernel-weighted local polynomial regression, shaded areas indicate 95% confidence intervals.

Figure A.5: Original single indicators in preexisting place income

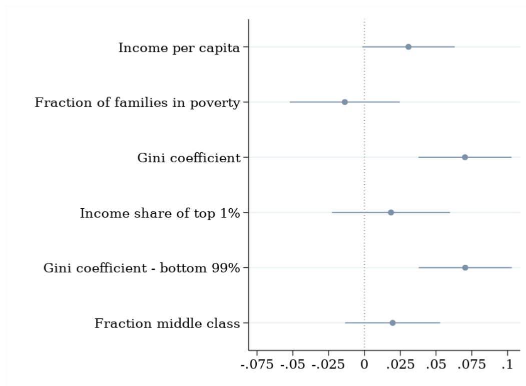

(a) Property price sale (log)

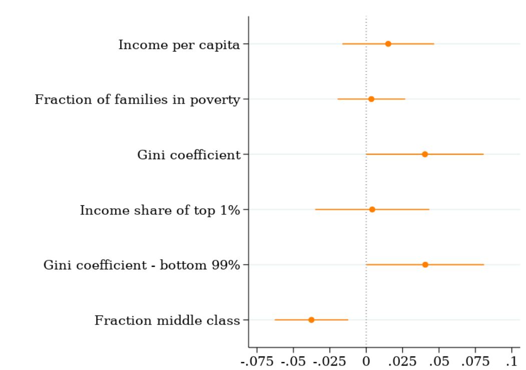

(b) Non-Hispanic White

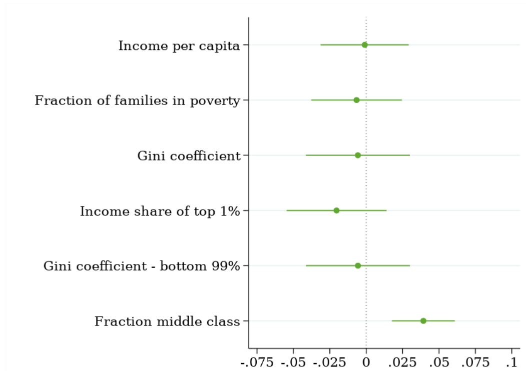

(c) Income (log)

Notes: Figure plots the point estimate and 95% confidence intervals of the coefficients on the triple interaction in the DDD model (equation 2), among (1) *Flooded*, a dummy equal to 1 for properties p inside Sandy floodplain; (2) *After*, a dummy equal to 1 for sales at time t after Sandy; and (3) each place characteristic on the y-axis. Each line on the y-axis represents a different model. Dependent variable is the log of property price sale in the top panel, a dummy variable equal to 1 if the buyer is non-Hispanic white in the middle panel, and the log income of the buyer in the bottom panel.

Figure A.6: Principal Component Analysis - Summary

Social Capital (pc1): (*Total variance explained:* 32%): High number of owner-occupied housing, low number of overall occupied or seasonal use housing units, high fraction religious, low census mail return rate.

Social Capital (pc2): (*Total variance explained:* 28%): Low crime, high social capital index, high fraction religious, low fraction of children with single mothers, high number of overall occupied housing.

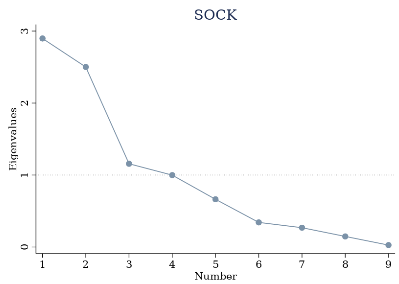

	1st	2nd
Owner-occupied housing units	0.49	0.20
Housing units for occasional use	-0.52	-0.14
Occupied housing units	-0.35	0.36
Fraction of children with single mothers	0.07	-0.34
Social capital index	0.17	0.32
Fraction religious	0.35	0.34
Violent crime	0.27	-0.41
Total crime	0.03	-0.52
2010 Census mail return rate	-0.38	0.20

Income (pc1) (*Total variance explained:* 63%): High income per capita and income inequality, low fraction middle class.

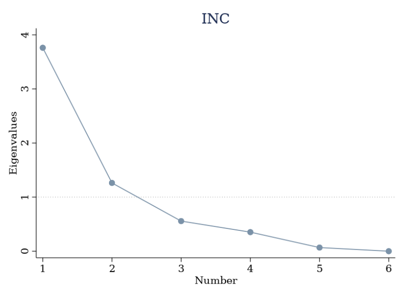

	1st
Income per capita	0.35
Poverty rate	-0.07
Gini coefficient for income	0.49
Top 1% income share	0.49
Gini Bottom 99%	0.49
Fraction Middle Class	-0.37

Migration (pc1): (*Total variance explained:* 66%): Large migrants inflow and outflow, and share of foreign born residents

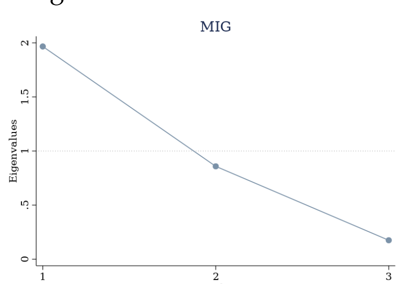

	1st
Migration inflow rate	0.58
Migration outflow rate	0.68
Fraction of foreign born residents	0.45

Notes: Principal Component Analysis - Summary. This figure shows the following elements for each of the twelve overarching place categories: (1) narrative description of the resulting principal components (PC) variable(s) and the percentage of variance explained by each PC; (2) scree plot, which plots the eigenvalue associated with each of the PCs, and hence, provides a metric of the proportion of variance explained by each; and (3) the loadings with which each single indicator enters into the 1st PC (and also 2nd, when applicable.)

Segregation (pc1): (*Total variance explained:* 41%): High racial and income segregation, large population density, low share of white population, long commutes
Segregation (pc2): (*Total variance explained:* 22%): Large share of white population, large income segregation

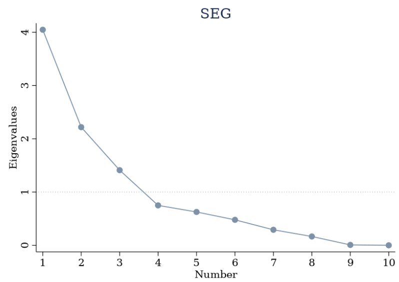

	1^{st}	2^{nd}
Share with commute less than 15 min	-0.27	0.16
Racial shares (white)	-0.29	0.49
Population density	0.39	-0.10
Racial segregation	0.37	-0.02
Income segregation	0.39	0.38
Segregation of poverty	0.36	0.38
Segregation of affluence	0.39	0.35
Racial shares (black)	0.19	-0.41
Racial shares (hispanic)	0.22	-0.35
Racial shares (asian)	0.19	-0.18

Education (pc1) (*Total variance explained:* 68%): High language and math grades, low expenditure per student.

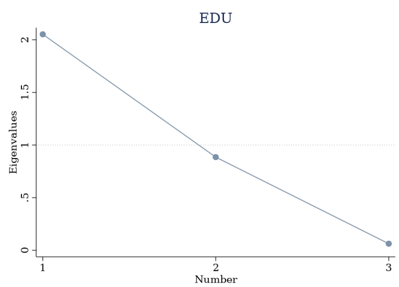

	1^{st}
School expenditure per student	-0.31
Average language score 3-8 grade	0.67
Average math score 3-8 grade	0.67

Secondary Education (pc1) (*Total variance explained:* 74%): High fraction of college graduates, and high number of colleges per capita.

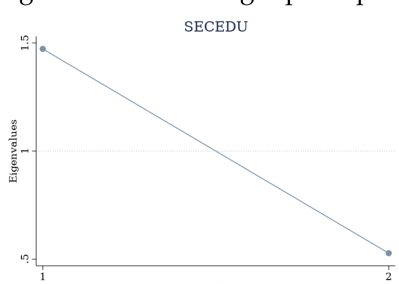

	1^{st}
Number of colleges per capita	0.71
College graduation rate	0.71

Notes: Principal Component Analysis - Summary. This figure shows the following elements for each of the twelve overarching place categories: (1) narrative description of the resulting principal components (PC) variable(s) and the percentage of variance explained by each PC; (2) scree plot, which plots the eigenvalue associated with each of the PCs, and hence, provides a metric of the proportion of variance explained by each; and (3) the loadings with which each single indicator enters into the 1^{st} PC (and also 2^{nd}, when applicable.)

Tax rates (pc1) (*Total variance explained:* 56%): High local tax rate, and high local government expenditures per capita.

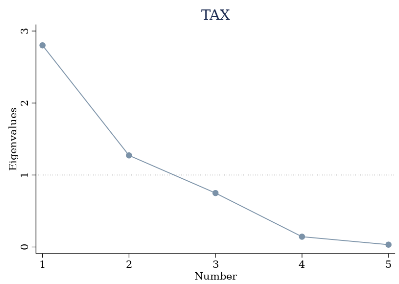

	1st
Local tax rate	0.57
Local tax rate per capita	0.58
Local government expenditures per capita	0.56
State EITC exposure	-0.12
State income tax progressively	0.10

Employment (pc1) (*Total variance explained:* 27%): High labor force participation, high fraction of manufacturing workers, and high teenage labor force participation rate.

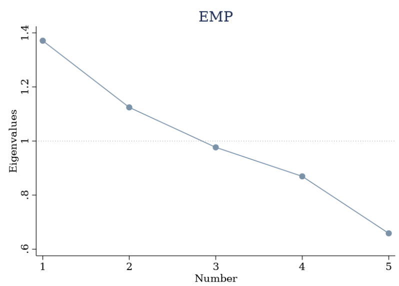

	1st
Unemployment	0.25
Fraction working in manufacturing	0.52
Labor force participation rate	0.55
Teenage (14-16) labor force participation rate	0.51
Job growth rate (2004-2013)	-0.32

Retail (pc1) (*Total variance explained:* 85%): high number of retail, movie theaters, and restaurants and bars.

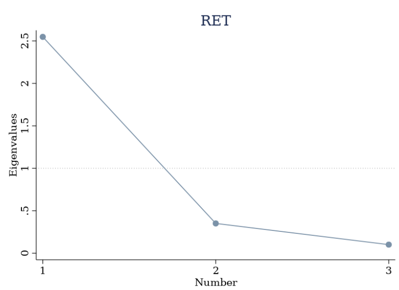

	1st
Number of movie theaters	0.55
Number of restaurants and bars	0.60
Number of retail trade	0.58

Notes: Principal Component Analysis - Summary. This figure shows the following elements for each of the twelve overarching place categories: (1) narrative description of the resulting principal components (PC) variable(s) and the percentage of variance explained by each PC; (2) scree plot, which plots the eigenvalue associated with each of the PCs, and hence, provides a metric of the proportion of variance explained by each; and (3) the loadings with which each single indicator enters into the 1st PC (and also 2nd, when applicable.)

Transport connectivity (pc1) (*Total variance explained:* 51%): Large distances to bus and train stations, and motorway crossings.

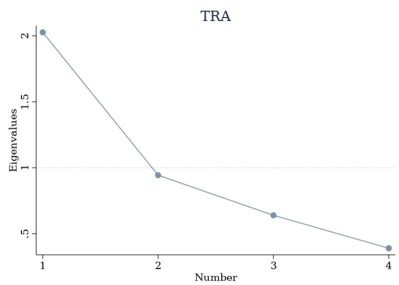

	1^{st}
Average commuting time	-0.25
Distance to bus station	0.59
Distance to motorway crossing	0.49
Distance to train station	0.59

Environment - Pollution (pc1) (*Total variance explained:* 80%): Highly polluted

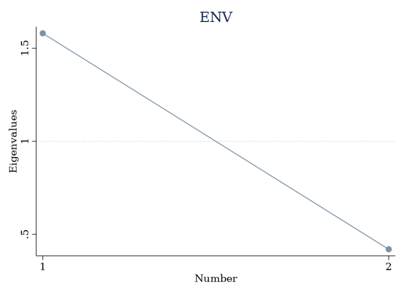

	1^{st}
Pollution - $PM_{2.5}$	0.71
Pollution - NO_2	0.71

Weather (pc1) (*Total variance explained:* 68%): High January and July temperatures, and low precipitation

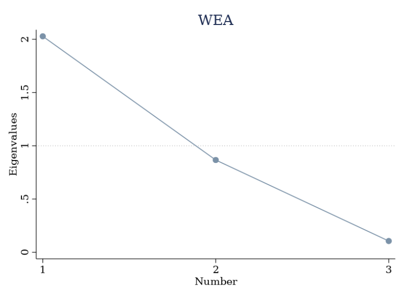

	1^{st}
Average daily precipitation	-0.52
January average temp	0.68
July average temp	0.51

Notes: Principal Component Analysis - Summary. This figure shows the following elements for each of the twelve composite variables describing preexisting place characteristics: (1) narrative description of the resulting principal components (PC) variable(s) and the percentage of variance explained by each PC; (2) scree plot, which plots the eigenvalue associated with each of the PCs, and hence, provides a metric of the proportion of variance explained by each; and (3) the loadings with which each single indicator enters into the 1^{st} PC (and also 2^{nd}, when applicable.)

Appendix B Additional tables

Table B.1: New construction

	(1) 2013-2016	(2) 2009-2012
Flooded_p	0.013	0.011
	(0.006)	(0.011)
Place income_p	-0.001	-0.011**
	(0.004)	(0.003)
Flooded_p x Place Income_p	0.002	0.003
	(0.005)	(0.005)
Sample	13-16	09-12
Observations	150,469	98,279

Notes: Table shows how the probability that the sale of property p corresponds to a newly constructed property changes with respect to being in Sandy's floodplain (*Flooded*) and *Place income*, a composite variable that describes preexisting income of the place property p is located. It is normalized to have a mean of zero and standard deviation equal to 1, and it is positively correlated to places with high income per capita and income inequality, and low fraction middle class. Column (1) shows results for the post-Sandy (2013-2016) period, and column (2) for the period immediately pre-Sandy (2009-2012). Dependent variable is a dummy equal to one if the property sold was built within the 4-year period considered. Standard errors two-way clustered at the census tract and year level. * $p< 0.1$, ** $p< 0.05$, *** $p< 0.01$.

Table B.2: Fraction of properties in sample sold by period

	(1)	(2)	(3)	(4)
Flooded_p x After_t	0.00039	0.00039		
	(0.00027)	(0.00025)		
Flooded_p x After_t x Place income_p		0.00017		
		(0.00021)		
Trend after_t x Place income_p				-0.00001**
				(0.00001)
Flooded_p x Trend after_t			0.00002*	0.00002
			(0.00001)	(0.00001)
Flooded_p x Trend after_t x Place income_p				0.00001
				(0.00001)
Location fixed effects	property	property	property	property
Observations	43,491,525	43,478,085	43,491,525	43,478,085

Notes: Table shows results for changes in the probability of being sold of a property. Dataset is a balanced panel for properties, in which the unit of observation is at the property-month level. Dependent variable is a dummy showing whether property p was sold in a given month. *Flooded* is a dummy equal to 1 for properties p in Sandy's floodplain. *After* is a dummy equal to 1 for sales at time t after Sandy. *Trend after* is a continuous time variable which increases one unit every month after Sandy. *INCpc1* is the first principal component for the Income category for the place property p is located in. Places which rank high in this variable have high income per capita and high income inequality, and a low fraction of middle class households. Standard errors two-way clustered at the census tract and year level. * $p< 0.1$, ** $p< 0.05$, *** $p< 0.01$.

Table B.3: Flood impacts on property characteristics

	(1) Square Feet	(2) Year Built
Panel A: DD model		
Flooded_p x After_t	-22.550	-1.051**
	(19.479)	(0.453)
Mean	1750	1964
Observations	396,340	351,578
Panel B: DDD model		
Flooded_p x After_t	-22.153	-0.874*
	(19.969)	(0.439)
After_t x Place Income_p	-41.302***	-0.223
	(10.647)	(0.153)
Flooded_p x After_t x Place Income_p	42.942**	1.087***
	(16.750)	(0.354)
Mean	1750	1964
Observations	396,330	351,570

Notes: Table reports changes in house characteristics on average (top panel), and by place income (bottom panel.) *Flooded* is a dummy equal to 1 for properties *p* in Sandy's floodplain. *After* is a dummy equal to 1 for sales at time *t* after Sandy. *Place income* is a composite variable that describes prexisting income of the place property *p* is located. It is normalized to have a mean of zero and standard deviation equal to 1, and it is positively correlated to places with high income per capita and income inequality, and low fraction middle class. Standard errors two-way clustered at the census tract and year level. * $p< 0.1$, ** $p< 0.05$, *** $p< 0.01$.

Table B.4: Count of FEMA insurance policies, changes after Sandy

	(1) All	(2) 1^{st} decile	(3) 10^{th} decile
1 year after$_m$	0.126***	0.047***	-0.010
	(0.037)	(0.012)	(0.049)
Between 2 and 3 years after$_m$	0.156**	0.007	0.040
	(0.062)	(0.019)	(0.082)
Beyond 3 years after$_m$	0.026	-0.154***	-0.047
	(0.063)	(0.028)	(0.099)
Observations	145,908	15,444	5,724

Notes: Dependent variable is the average number of active flood insurance policies underwritten by FEMA by month in the census tracts included in the study between 2010 and 2018. Column (1) includes the whole sample, column (2) restricts the sample to properties located in places in the bottom decile of place income, and column (3) restricts the sample to properties located in places in the top decile of place income. *1 year after* is a dummy equal to one for months *m* 1 year after Sandy. *Between 2 and 3 years after* is a dummy equal to one for months *m* 2 and 3 years after Sandy. *Beyond 3 years after* is a dummy equal to one for months *m* more than 3 years after Sandy. Standard errors clustered at the county level. * p< 0.1, ** p< 0.05, *** p< 0.01.

Table B.5: FEMA insurance policies characteristics, changes by place income

	(1) Policy Premium (log)	(2) Building Coverage (log)	(3) Contents Coverage (log)	(4) Building Elevation	(5) Primary Residence
After$_m$	-0.057***	-0.048***	0.038***	0.044***	-0.074***
	(0.007)	(0.009)	(0.014)	(0.005)	(0.006)
After$_m$ x Extent$_t$	-0.076*	0.030	0.084***	0.056***	-0.014
	(0.043)	(0.019)	(0.025)	(0.016)	(0.019)
After$_t$ x Place income$_t$	-0.031*	-0.009	-0.026	0.005	0.017*
	(0.017)	(0.013)	(0.023)	(0.011)	(0.008)
After$_m$xExtent$_t$xPlace income$_t$	0.139**	0.035	0.003	-0.074**	-0.025
	(0.068)	(0.025)	(0.039)	(0.033)	(0.024)
Mean	991	205,217	62,553	0.29	0.67
Observations	2,107,935	2,079,552	1,488,291	2,107,940	2,107,944

Notes: The dependent variable is the log of insurance premiums in column (1), the log of coverage for building in column (2), the log of coverage for contents in column (3), a dummy variable whether the property is elevated in column (4), and whether the property is the primary residence of the policyholder in columns (5). *After* is a dummy equal to 1 for month *m* after Sandy. *Extent* is tract *t* area covered by the flood, in percentage terms. *Place income* is a composite variable that describes average preexisting income at the tract *t* level. It is positively correlated to places with high income per capita and income inequality, and low fraction middle class. Standard errors clustered at the county level. * p< 0.1, ** p< 0.05, *** p< 0.01.

Table B.6: Flood insurance claims related to Sandy, by place income

	(1) Building (log)	(2) Content (log)	(3) Cost Compliance (log)	(4) Claims per capita
Place income$_{tract}$	-0.050	-0.099	0.030***	-0.030
	(0.143)	(0.117)	(0.008)	(0.018)
Depth$_{tract}$	0.067	0.044	-0.010	0.031
	(0.048)	(0.036)	(0.011)	(0.020)
Extent$_{tract}$ (%)	-1.888***	-1.240**	-0.009	0.296***
	(0.474)	(0.533)	(0.038)	(0.089)
Depth$_{tract}$ x Extent$_{tract}$ (%)	0.587***	0.434***	0.005	-0.070
	(0.113)	(0.140)	(0.020)	(0.042)
Mean	36,731	11,702	27,763	0.062
Observations	109,161	69,186	102,700	665

Notes: Table shows how flood insurance claims with a reported flooding date between October 26th and November 2nd 2012, from census tracts with at least one property sale record in main dataset, change with respect to place income. *Place income* is a composite variable that describes average preexisting income at the county level. It is positively correlated to places with high income per capita and income inequality, and low fraction middle class. The dependent variable is the log of the amount claimed for building damage in column (1), the log of the amount claimed for content damages in column (2), the log of the claim made for an increased cost of compliance in column (3), and the number of claims made by census tract divided by the total number of population on the tract, according to 2010 census, in column (4). Then, the unit of observation in columns (1)-(3) is individual policy, and in column (4) is census tract. All models control for the average flood depth (*Depth*), the tract area covered by the flood (*Extent*), and an interaction of both. Standard errors clustered at the county level. * $p< 0.1$, ** $p< 0.05$, *** $p< 0.01$.

Table B.7: Public Assistance by average place income at county level

	(1) Public Assistance (log)
Place Income$_{\text{county}}$	0.553
	(0.524)
Population (million)$_{\text{county}}$	11.025***
	(2.070)
Area flooded (%)$_{\text{county}}$	5.242***
	(1.562)
Mean	1,733,123
Median	29,980
Observations	7,932

Notes: Public assistance funds by average place income. Dependent variable is the federal share obligated for disaster assistance projects after Sandy. The unit of observation is the county. *Place income* is a composite variable that describes average preexisting income at the county level. It is positively correlated to places with high income per capita and income inequality, and low fraction middle class. * p< 0.1, ** p< 0.05, *** p< 0.01.

Table B.8: Flood impacts, by prior flood experience

	(1) Price (log)	(2) White	(3) Income (log)	(4) White-Income
After$_t$ x Sandy$_p$	-0.070***	-0.055*	-0.020*	-0.102**
	(0.018)	(0.028)	(0.010)	(0.042)
After$_t$ x Irene$_p$	-0.043	-0.050	-0.039*	-0.099
	(0.051)	(0.031)	(0.022)	(0.058)
After$_t$ x Sandy$_p$ x Irene$_p$	-0.025	0.056	0.012	0.087
	(0.049)	(0.034)	(0.030)	(0.063)
Observations	501,973	115,313	146,600	111,677

Notes: *Irene* is a dummy equal to 1 for properties p in Hurricane Irene's floodplain. Hurricane Irene made landfall in the Northeast US in August 2011, fourteen months before Sandy. *Sandy* is a dummy equal to 1 for properties p in Sandy's floodplain. *After* is a dummy equal to 1 for sales at time t after Sandy. Standard errors two-way clustered at the census tract and year level. * p< 0.1, ** p< 0.05, *** p< 0.01.

Table B.9: Flood impacts, by presence in a flood risk area

	(1) Price (log)	(2) White	(3) Income (log)	(4) White-Income
$After_t$ x $Sandy_p$	0.021	0.023	-0.038**	-0.000
	(0.025)	(0.021)	(0.017)	(0.033)
$After_t$ x Flood risk $area_p$	-0.080***	-0.081**	-0.002	-0.129*
	(0.019)	(0.038)	(0.018)	(0.063)
$After_t$ x $Sandy_p$ x Flood risk $area_p$	-0.067**	-0.032	0.020	-0.027
	(0.030)	(0.030)	(0.027)	(0.042)
Observations	475,949	110,178	140,343	106,729

Notes: *Flood risk area* is a dummy equal to 1 for properties *p* located in FEMA's Special Flood Hazard Area. *Sandy* is a dummy equal to 1 for properties *p* in Sandy's floodplain. *After* is a dummy equal to 1 for sales at time *t* after Sandy. Standard errors two-way clustered at the census tract and year level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B.10: Flood impacts on likelihood of owner occupying property

	(1) Owner - occupied
$Flooded_p$ x $After_t$	0.003
	(0.008)
$After_t$ x Place $Income_p$	-0.008
	(0.005)
$Flooded_p$ x $After_t$ x Place $Income_p$	-0.013
	(0.011)
Observations	152,344

Notes: Dependent variable is a dummy equal to one if a property is occupied by its owner. *Flooded* is a dummy equal to 1 for properties *p* in Sandy's floodplain. *After* is a dummy equal to 1 for sales at time *t* after Sandy. *Place income* is a composite variable that describes prexisting income of the place property *p* is located. It is normalized to have a mean of zero and standard deviation equal to 1, and it is positively correlated to places with high income per capita and income inequality, and low fraction middle class. Standard errors two-way clustered at the census tract and year level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B.11: Robustness to different location fixed effects

	(1)	(2)	(3)	(4)
Flooded_p x After_t	-0.045***	-0.050***	-0.048***	-0.043**
	(0.015)	(0.016)	(0.015)	(0.017)
After_t x Place Income_p	0.102***	0.108***	0.109***	0.131***
	(0.016)	(0.013)	(0.013)	(0.011)
Flooded_p x After_t x Place Income_p	0.073***	0.048***	0.045***	0.047***
	(0.014)	(0.012)	(0.012)	(0.013)
Location FE	block	block	block	parcel
Property controls	no	no	yes	–
Observations	501,845	324,486	324,486	166,882

Notes: Robustness to different location fixed effects of results in table 4, panel (A), column (1). Column (1) replicates preferred model with block fixed effects and no property controls. Column (3) also includes some property characteristics as controls (square footage and year built.) Column (2) does not include controls, but is run on the same data that model (3), that is, on properties that do have data on property controls (around 35% of the observations in the main sample are missing data on one or more of property controls.) Finally, the model in column (4) includes parcel level fixed effects. All models include month-year fixed effects, and cluster standard errors at the tract and year levels. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B.12: Robustness to different time fixed effects

	(1)	(2)	(3)	(4)
Flooded_p x After_t	-0.045***	-0.042***	-0.044***	-0.045***
	(0.015)	(0.012)	(0.015)	(0.015)
After_t x Place Income_p	0.102***	0.289***	0.102***	0.102***
	(0.016)	(0.051)	(0.016)	(0.016)
Flooded_p x After_t x Place Income_p	0.073***	0.049***	0.073***	0.074***
	(0.014)	(0.011)	(0.014)	(0.014)
Time FE	month	county x month	quarter	week
Observations	501,845	501,796	501,845	501,845

Notes: Robustness to different time fixed effects of results in table 4, panel (A), column (1). All models include census block fixed effects, and cluster standard errors at the tract and year levels. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B.13: Difference in Difference: robustness to different clustering of standard errors

	(1)	(2)	(3)	(4)
$Flooded_p$ x $After_t$	-0.045***	-0.045***	-0.045***	-0.045***
	(0.015)	(0.013)	(0.013)	(0.009)
$After_t$ x Place $Income_p$	0.102***	0.102***	0.102***	0.102***
	(0.016)	(0.007)	(0.006)	(0.006)
$Flooded_p$ x $After_t$ x Place $Income_p$	0.073***	0.073***	0.073***	0.073***
	(0.014)	(0.014)	(0.014)	(0.010)
Cluster SE	tract year	tract month-year	tract	block month-year
Observations	501,845	501,845	501,845	501,845

	(5)	(6)	(7)
$Flooded_p$ x $After_t$	-0.044**	-0.044***	-0.044***
	(0.018)	(0.017)	(0.017)
$After_t$ x Place $Income_p$	0.102***	0.102***	0.102***
	(0.017)	(0.009)	(0.008)
$Flooded_p$ x $After_t$ x Place $Income_p$	0.074***	0.074***	0.074***
	(0.016)	(0.017)	(0.017)
Cluster SE	zip year	zip month-year	zip
Observations	501,507	501,507	501,507

Notes: Robustness to different clusters of standard errors of results in table 4, panel (A), column (1). All models include census block and month-year fixed effects. * p< 0.1, ** p< 0.05, *** p< 0.01.

Table B.14: Event study with properties in the control group

	(1) Price (log)	(2) White	(3) Income (log)	(4) White-Income
$After_t$	-0.003	-0.025	0.031	-0.003
	(0.008)	(0.016)	(0.019)	(0.042)
Observations	154161	44078	48290	43168

Notes: Sample in all models include only properties in the control group, that is, those not on the floodplain, but located less than 500m from the shoreline. *After* is a dummy equal to 1 for sales at time *t* after Sandy. Standard errors two-way clustered at the census tract and year level. * p< 0.1, ** p< 0.05, *** p< 0.01. * p< 0.1, ** p< 0.05, *** p< 0.01.

Table B.15: Flood impacts by place income, removing higher income places

	(1) Price (log)	(2) White	(3) Income (log)	(4) White-Income
$\text{Flooded}_p \times \text{After}_t$	-0.037	-0.016	-0.048***	-0.066
	(0.022)	(0.028)	(0.013)	(0.043)
$\text{After}_t \times \text{INC}_p$	0.178***	-0.057**	0.074***	-0.024
	(0.031)	(0.026)	(0.019)	(0.046)
$\text{Flooded}_p \times \text{After}_t \times \text{INC}_p$	0.099**	0.091**	-0.083**	0.068
	(0.040)	(0.037)	(0.029)	(0.060)
Mean	459,154	0.69	171,780	-0.05
Observations	405,801	103,705	128,728	100,811

Notes: Sample includes only observations from properties with place income below the 80$^{\text{th}}$ percentile. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B.16: Flood impacts, for properties identified as single homes

	(1) Price (log)	(2) White	(3) Income (log)	(4) White-Income
Panel A: DD model				
$\text{Flooded}_p \times \text{After}_t$	-0.099***	-0.039**	-0.025**	-0.084***
	(0.018)	(0.017)	(0.010)	(0.025)
Mean	418,646	0.62	140,516	0.02
Observations	268,417	68,666	88,503	66,473
Panel B: DDD model				
$\text{Flooded}_p \times \text{After}_t$	-0.061***	-0.014	-0.042**	-0.058***
	(0.017)	(0.012)	(0.016)	(0.018)
$\text{After}_t \times \text{Place Income}_p$	0.076***	-0.056**	0.044**	-0.049
	(0.014)	(0.021)	(0.015)	(0.028)
$\text{Flooded}_p \times \text{After}_t \times \text{Place Income}_p$	0.106***	0.065**	-0.044	0.067
	(0.027)	(0.030)	(0.033)	(0.058)
Mean	418,646	0.62	140,516	0.02
Observations	268,410	68,666	88,503	66,473

Notes: Sample includes only observations from properties identified as single homes by Corelogic®. Results of the DD model are in the top panel, and DDD model are in the bottom panel. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B.17: Difference between average place characteristics of most positively and negatively affected properties

Indicators	Category	Difference
Gini for income (Bottom 99%)	INC	0.927
Gini for income	INC	0.926
Colleges per capita	SECEDU	0.899
College graduation rate	SECEDU	0.839
Migration outflow rate	MIG	0.814
Daily precipitation	WEA	0.805
Violent crime	SOCK	0.779
Food and drink business	RET	0.759
Population density	SEG	0.751
Top 1% income share	INC	0.775
Income per capita	INC	0.742
Retail business	RET	0.682

Notes: Difference between average place characteristics of the 20% most positively affected and the 20% most negatively affected properties. Table shows the 12 variables (out of 52) that yield the largest differences. All variables have been normalized to have a mean equal to 0 and a standard deviation equal to 1. All differences have a p-value smaller than 0.001.

Table B.18: Place characteristics: data sources, geography, and year

Group & Indicators	Geography	Year	Source
SEG: Segregation			
Racial shares	Tract	2010	Opp Lab
Share with commute less than 15 min	Tract	2010	ACS
Racial segregation	County	2000	Opp Lab
Income segregation	County	2000	Opp Lab
Segregation of poverty	County	2000	Opp Lab
Segregation of affluence	County	2000	Opp Lab
Population density	Tract	2010	Census SF1
INC: Income			
Income per capita	Tract	2010	ACS
Gini coefficient for income	County	2000	Opp Lab
Top 1% income share	County	2000	Opp Lab
Gini Bottom 99%	County	2000	Opp Lab
Fraction Middle Class	County	2000	Opp Lab
Poverty rate	Tract	2010	ACS
EDU: Education			
School expenditure per student	ScDistrict	2010	NCES
Average math score 3-8 grade	ScDistrict	2009-2011	SEDA
Average language score 3-8 grade	ScDistrict	2009-2011	SEDA
SOCK: Social Capital			
Social capital index	County	2000	Opp Lab
Fraction religious	County	2000	Opp Lab
Violent crime	County	2000	Opp Lab
Total crime	County	2000	Opp Lab
Occupied housing units	Tract	2010	Census SF1
Housing units for occasional use	Tract	2010	Census SF1
Owner-occupied housing units	Tract	2010	Census SF1
Fraction of children with single mothers	Tract	2010	Census SF1
2010 Census mail return rate	Tract	2010	Opp Lab
TAX: Tax rates			
Local tax rate	County	2000	Opp Lab
Local tax rate per capita	County	2000	Opp Lab
Local government exp pc	County	2000	Opp Lab
State EITC exposure	County	2000	Opp Lab
State inc tax progress	County	2000	Opp Lab

Group & Indicators	Geography	Year	Source
EMP: Employment			
Labor force particip rate	Tract	2010	Opp Lab
Fraction manufactur work	Tract	2010	ACS
14 16yo in labor force	County	2000	Opp Lab
Unemployment	Tract	2010	ACS
Job growth rate	Tract	04 - 13	Opp Lab
SECEDU: Secondary Education			
Colleges per capita	County	2000	Opp Lab
College graduation rate	Tract	2010	Opp Lab
MIG: Migration			
Migration inflow rate	County	2000	Opp Lab
Migration outflow rate	County	2000	Opp Lab
Fraction foreign born residents	Tract	2010	Opp Lab
RET: Retail			
Number of retail trade	Zip code	2010	NAICS
Number of restaurants and bars	Zip code	2010	NAICS
Number of movie theaters	Zip code	2010	NAICS
TRA: Transport connectivity			
Distance to bus station	Point	2018	OpenMap
Distance to motorway crossing	Point	2018	OpenMap
Distance to train station	Point	2018	OpenMap
Average commuting time	Tract	2010	ACS
ENV: Environment - Pollution			
Pollution - $PM_{2.5}$	Grid	2008-2010	NASA
Pollution - NO_2	Grid	2008-2010	NASA
WEA: Weather			
January average temperature	Grid	2000-2010	Prism
July average temperature	Grid	2000-2010	Prism
Average daily precipitation	Grid	2000-2010	Prism

Notes: This table classifies the original place descriptive variables into twelve overarching categories. For the original variables, it lists the the geography the variables were measured in, the time in which they were measured, and the data sources.

Table B.19: Random Forest - Heterogeneity - Characteristics

	(1) 20% Least Affected	(2) 20% Most Affected	(3) Difference (1) - (2)
Secondary Education (pc1)			
Estimate	0.684	-0.337	1.023
90% CI	(0.676,0.691)	(-0.344,-0.330)	(1.012,1.034)
p-value	-	-	[0.000]
Income (pc1)			
Estimate	0.634	-0.241	0.907
90% CI	(0.627,0.642)	(-0.249,-0.233)	(0.897,0.918)
p-value	-	-	[0.000]
Retail (pc1)			
Estimate	0.506	-0.207	0.721
90% CI	(0.499,0.514)	(-0.215,-0.199)	(0.711,0.731)
p-value	-	-	[0.000]
Segregation (pc1)			
Estimate	0.602	-0.088	0.678
90% CI	(0.595,0.609)	(-0.095,-0.081)	(0.669,0.688)
p-value	-	-	[0.000]
Migration (pc1)			
Estimate	0.506	-0.111	0.629
90% CI	(0.499,0.514)	(-0.119,-0.104)	(0.619,0.639)
p-value	-	-	[0.000]
Social Capital (pc1)			
Estimate	0.463	-0.040	0.508
90% CI	(0.457,0.469)	(-0.046,-0.035)	(0.500,0.516)
p-value	-	-	[0.000]

	(1) 20% Least Affected	(2) 20% Most Affected	(3) Difference (1) - (2)
Environment (pc1)			
Estimate	0.289	-0.071	0.382
90% CI	(0.282,0.296)	(-0.078,-0.064)	(0.372,0.393)
p-value	-	-	[0.000]
Segregation (pc2)			
Estimate	0.203	-0.159	0.364
90% CI	(0.196,0.210)	(-0.167,-0.152)	(0.354,0.375)
p-value	-	-	[0.000]
Transport (pc1)			
Estimate	-0.378	-0.017	-0.352
90% CI	(-0.384,-0.372)	(-0.023,-0.011)	(-0.360,-0.344)
p-value	-	-	[0.000]
Education (pc1)			
Estimate	-0.270	0.051	-0.326
90% CI	(-0.277,-0.263)	(0.045,0.058)	(-0.336,-0.316)
p-value	-	-	[0.000]
Employment (pc1)			
Estimate	-0.181	0.105	-0.282
90% CI	(-0.189,-0.174)	(0.098,0.113)	(-0.293,-0.272)
p-value	-	-	[0.000]
Weather (pc1)			
Estimate	-0.156	0.135	-0.281
90% CI	(-0.163,-0.149)	(0.129,0.142)	(-0.290,-0.271)
p-value	-	-	[0.000]

	(1) 20% Least Affected	(2) 20% Most Affected	(3) Difference (1) - (2)
Social Capital (pc2)			
Estimate	-0.057	0.068	-0.122
90% CI	(-0.063,-0.050)	(0.061,0.075)	(-0.132,-0.112)
p-value	-	-	[0.000]
Taxes (pc1)			
Estimate	0.041	0.018	0.035
90% CI	(0.034,0.048)	(0.010,0.025)	(0.025,0.045)
p-value	-	-	[0.000]

Notes: Table presents the place characteristics along which the places in which most and least affected properties are located differ the most, according to the analysis following the procedure in Chernozhukov et al. (2018) with a random forest method. The place characteristics (in rows) have been normalized to have a mean of 0 and a standard deviation of 1, so their coefficients are comparable. Column (1) shows the average value and 90% confidence interval for each place variable for properties in the *least* affected quintile. Column (2) shows those values for properties in the *most* affected quintile. Column (3) shows the difference between the coefficients in (2) and (3), as well as a 90% confidence interval and the p-value of the difference.

Appendix C Conceptual model: detailed description

This appendix describes the full set up of the conceptual model outlined in section 2.

Set up The basic model set up considers two neighborhoods, A and B, and two types of households, H and L. Half of the households are of type H, and the other half are type L. s_i summarizes the proportion of residents who are of type H in each neighborhood i:

$$s_i = \frac{H_i}{L_i + H_i} \qquad i = A, B$$

H-type households can be characterized as having high-income, and/or other kind of socioeconomic traits – e.g. race, ethnicity – as long as those are correlated with a higher income. The externalities derived from having H types as neighbors are higher than L types for all households,[57] which make H types the preferred neighbors of all residents.

Both neighborhoods A and B have the same number of houses, and each household lives in one house. Each neighborhood i is endowed with a bundle of exogenous amenities Z_i. I assume that one of the amenities in the bundle is the absence of flood risk, $-F_i$:

$$-F_i \in Z_i \qquad i = A, B$$

In this setting, I assume neighborhood A represents a coastal location, whose amenity bundle overall is more valued by residents of both types with respect to those in B. B could be inland, or also a coastal location with fewer amenities.[58]

Willingness to pay to live in neighborhood i A household prefers to live in a neighborhood that has more residents of the preferred type, more valued amenities, or both. How these preferences change is allowed to vary by household type. Following Becker and Murphy (2000), I assume that the willingness to pay for amenities is separable from the willingness to pay for neighborhood type. Then, the overall willingness to pay for living in neighborhood i for a resident of type j is defined by:

$$V^j = u^j(Z_i) + f^j(s_i) \qquad \text{with} \quad u', f' > 0; \;\; u'' = f'' = 0$$ [59]

Equilibrium There is a competitive market for housing in both neighborhoods, so each house is sold to the highest bidder, regardless of her type. Hence, neighbors cannot choose which type moves into their neighborhood (in other words, the model as formulated does not allow active discrimination.)

House prices in equilibrium in neighborhoods A and B, P_a and P_b respectively, are pinned down by the different resident types' willingness to pay functions. In particular, an equi-

[57] Positive externalities from neighbors could be related to job opportunities, networking, increased school achievement, a prestige signal, etc.

[58] A could have higher amenities than B for historical reasons and path dependency, e.g. B could be on a coastal site where there were historical marshes (Villarreal, 2015)

[59] Becker and Murphy (2000) allow u to be concave or linear, while they do not impose any constraints on the curvature of f. For simplicity, I assume a linear willingness to pay for amenities and preferred neighbors.

Figure C.1: Conceptual model set up: premium for living in A in equilibrium

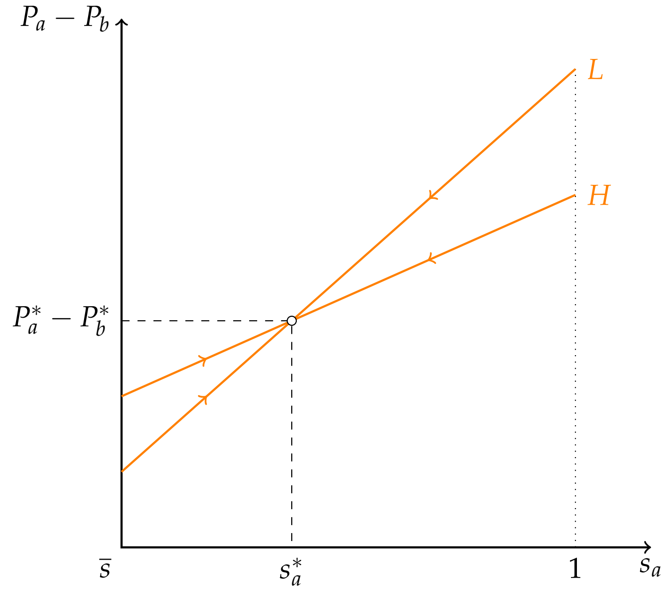

Notes: Figure graphically depicts the set up for the segregation model described in section 2. Its plots the willingness to pay for living in neighborhood A with respect to B (y-axis) versus the ratio of H types living in neighborhood A (s_a, in the x-axis.) for both types of residents (L and H curves.) It shows that, if a partial segregation equilibrium exists, that equilibrium is stable.

librium yields sufficiently higher prices in neighborhoods with larger shares of H and more valued amenities (in this case, neighborhood A), so no one would be better off by moving.

If both types H and L are indifferent between neighborhoods in equilibrium, it must be that:

$$\begin{aligned} P_a^* - P_b^* &= f_h(s_a^*) + u_h(Z_a) - f_h(s_b^*) - u_h(Z_b) \\ &= f_l(s_a^*) + u_l(Z_a) - f_l(s_b^*) - u_l(Z_b) \end{aligned} \tag{3}$$

where s_a^* and s_b^* are the equilibrium ratios of H types in both neighborhoods; $P_a^* > P_b^*$, and $s_a^* > s_b^*$.

Figure C.1 graphically shows the premium for living in A with respect to B is an stable equilibrium with partial segregation.[60] [61] [62]

In this figure, I assume that H-types put a higher relative value on the coastal amenities of A than L-types. This assumption would be consistent with H and L valuing these ameni-

[60]Note that for the existence of an interior partial segregation solution, a required condition is that the slope of the willingness to pay for the preferred neighbor is higher for L-types than for H-types.

[61]This equilibrium is stable: To the left of s_a^*, the willingness to pay of H-types is above that of L-types. Hence, if an L type were to move from A to B, moving s to the left of equilibrium, a type H – with a higher willingness to pay at this point – would outbid a type L for the empty house, raising s back to s_a^*. The opposite would happen to the right of s_a^*: the higher willingness to pay of type L at this point would bring s down.

[62]Note there are two other possible equilibria besides the one drawn in figure C.1. One fully integrated (with $s_a^* = 0.5$) if L's WTP for living in A is always above H's; and another one fully segregated (with $s_a^* = 1$) if H's WTP is always above L's.

ties differently because of their intrinsic characteristics (e.g. wealth, education levels, etc), but implications from the model below would hold regardless of this assumption.[63]

Flooding shock as a disamenity in coastal location: assumptions Starting from an equilibrium scenario, I assume that neighborhood A is hit by a flood. All of the residents dislike flooding, as a result, the bundle of amenities in A is less desirable than before for all residents. I further make the following assumption:

Assumption 1: *Utility loss after the flood is larger for L types than H types, that is:*

$$u_h(Z_{i,1}) - u_h(Z_{i,2}) < u_l(Z_{i,1}) - u_l(Z_{i,2}) \qquad \text{with} \quad i = A, B$$

where 1 represents the pre-flood period, and 2 is the post-flood period.

A change in the risk of flooding can be understood as an income shock (e.g. as in Smith and Whitmore (2019).) Responding to a flood could mean structure reconstruction, raised flood insurance premiums, expenses incurred in preparation for the next flood, etc. Under the assumptions of diminishing returns to income and that L-types are more credit constrained, the disutility from the flood would be higher for low-income L-types than for H-types.

Further, if B is also a coastal location impacted by a flood, I also assume that the loss of utility within each resident type would be lower in B than in A. Specifically, I carry on the interpretation of the flood as an income shock, and assume that the utility loss is proportional to pre-flood residential prices.[64]

Assumption 2: *Utility loss after the flood for a given resident type is proportional to the pre-flood residential prices, and hence, higher in neighborhood A than in B in this setting.*

$$u_j(Z_{a,1}) - u_j(Z_{a,2}) = \alpha \cdot \left[u_j(Z_{b,1}) - u_j(Z_{b,2})\right] \qquad \text{with} \quad \alpha \propto \frac{P_a}{P_b} > 0 \text{ and } \quad j = H, L$$

where 1 represents the pre-flood period, and 2 is the post-flood period.

Comparative statics Given the assumptions above, the model yields two main implications on the ratio of the preferred type of resident in A, and the evolution of premiums to live in A with respect to B after a flood:

Prediction 1: *In the new equilibrium after flooding, neighborhoods are more segregated, i.e. $s^*_{a,2} > s^*_{a,1}$, were $s^*_{a,2}$ is the equilibrium ratio of H types in neighborhood A after flooding in period 2, and $s^*_{a,1}$ is the pre-flood s in equilibrium.*

Appendix D shows that a negative shock to amenities would yield the same level of segregation, s_a, as long as the drop in utility derived from this negative shock is equal for

[63] If H- and L- types valued amenities in A equally, I would need to assume some concavity on the willingness to pay for preferred neighbors for the partial segregation equilibrium to exist

[64] e.g. taxes could be levied to protect properties in anticipation of future storms (Shafroth, 2017), flood insurance premiums could be more expensive for more expensive properties, etc.

both H and L types. If, as it assumed in this context, the drop in utility is rather relatively higher for L-types than H-types, the outcome is necessarily an increase in s_a^* to reach a new equilibrium.

Intuition from this result follows up directly from Becker and Murphy (2000), who argue that differences in the evaluation of amenities should contribute to segregation. After the flood, amenities in A with respect to B are relatively more valuable for H than L. This leads to an influx of H-types to A, in detriment of L type residents.

Prediction 2: *In the new equilibrium after flooding, the premium for living in A with respect to B could increase, stay constant, or decrease with respect to pre-flood levels, depending on the relative drops in utility derived from the flood and the slopes of $f_j(\cdot)$ for both types of residents.*

Appendix D shows that, mathematically, the premium for living in A with respect to B would stay constant after a flood as long as:

$$\frac{f_h'(s)}{f_l'(s)} = \frac{\delta_h}{\delta_l} \tag{4}$$

Where $f_j'(s_a)$ is the slope of the willingness to pay for preferred neighbors and δ_j is defined as the difference in utility derived from amenities in A and B before and after the flood for resident type $j = H, L$, that is:

$$\delta_j = u_j(Z_{a,1}) - u_j(Z_{b,1}) - \left[u_j(Z_{a,2}) - u_j(Z_{b,2})\right]$$

Then, the premium for living in A will increase if $\frac{f_h'(s_a)}{f_l'(s_a)} > \frac{\delta_h}{\delta_l}$, and will decrease if the opposite is true. The more H values to be around other H-types than L does (i.e. the larger the ratio $\frac{f_h'(s_a)}{f_l'(s_a)}$), and the larger the drop in utility after the flood for L-types with respect o H-types (i.e. the smaller the ratio $\frac{\delta_h}{\delta_l}$), the more likely it is that the premium to live in A increases after a flood.

Intuitively, the new equilibrium after a flood would necessarily have a higher level of s_a^* given the model assumptions, as the model prediction 1 notes. Both L- and H-types derive a higher utility from having more H-type neighbors, even if they find the flood a disamenity. For the premium for living in A to stay constant after a flood (i.e. for condition 4 to be satisfied), it would be necessary that the relative increase in the willingness to pay for higher H is enough to compensate for the flood disamenity.

Graphically, figure C.2 plots two potential scenarios on the change of the premium for living in neighborhood A after a flood. In the top panel, $P_a^* - P_b^*$ increases after the flood, while in the bottom panel decreases.

In the bottom panel of figure C.2, and even if s_a^* is higher at time 2, premium to live in A decreases as the disamenity from the flood for both types is large enough given the slopes of the willingness to pay for a higher s_a^*. In the top panel however, the utility loss derived from the flood is not large enough in relative terms to overcome the utility gain of having more H-type neighbors. The premium to live in A with respect to B increases.

The second result is a priori less intuitive. All of the residents derived disutility from the flood, and still, prices in A relatively *increase* after the flood. This result holds if only A is hit by a flood and not B, or if A is more severely hit than B. In fact, the result holds even

Figure C.2: Conceptual model - Premium changes after flood

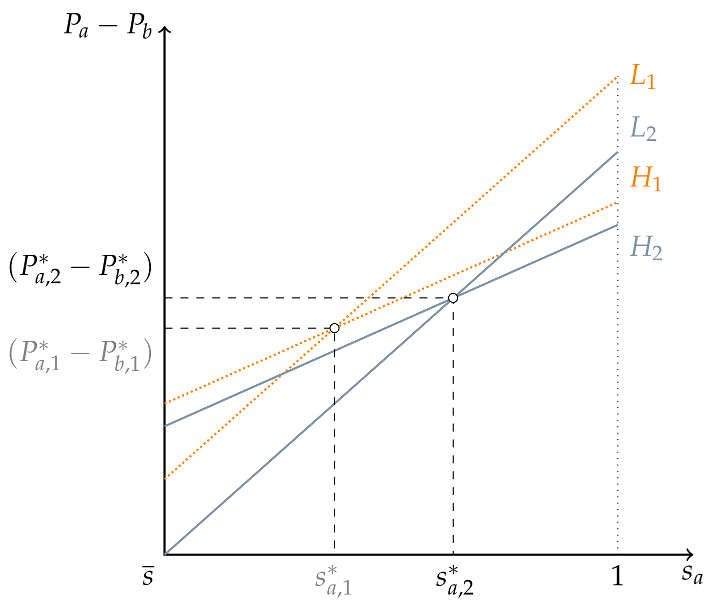

(a) Premium for neighborhood A increases

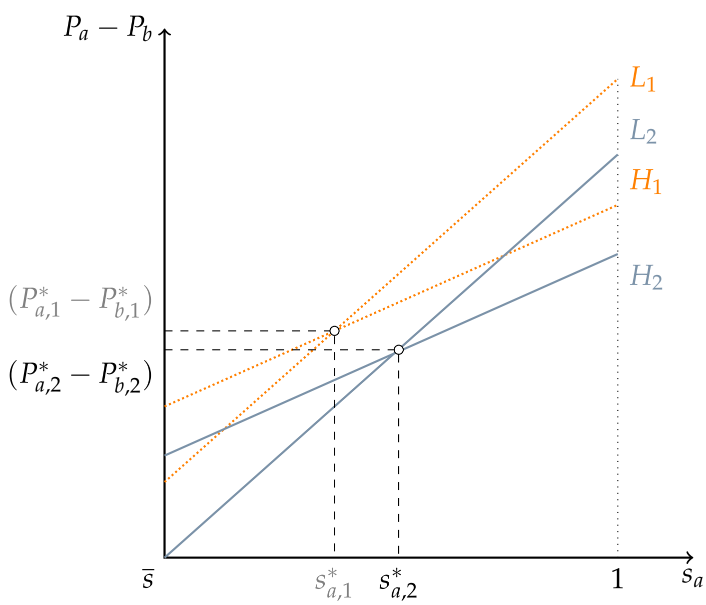

(b) Premium for neighborhood A decreases

Notes: These figures show two potential after-flood scenarios for the segregation model described in section 2. In both panels, the ratio of H-types living in A (s_a) increases after the flood. In the top panel, the premium to live in A with respect to B also increases, as the utility derived from more H-type neighbors more than compensates the flood disamenity. In the bottom panel, which only differs from the top panel in that the disamenity drop for H-types is larger than in the top panel, the premium to live in A decreases after the flood.

if the disutility from the flood was so large as to make the bundle of amenities in A less desirable than B.[65] In this case, B would have more valued amenities and cheaper houses than A. However, without the ability to form coalitions to coordinate a move to B en-masse, H-types will stay in A where s_* is higher.

Results from this model provide key insights on adaptive behavior after a flood, or lack thereof. A seemingly irrational behavior of relatively increased property prices after a flood could happen even if the flood disamenity is internalized, as long as flooding affects relatively more the lesser valued type of resident.

[65]Graphically, it would mean that the premium to live in A, $P_a^* - P_b^*$, evaluated at $s = 0.5$, fully integrated equilibrium, is negative.

Appendix D Conceptual model: key mathematical derivations

Model prediction 1: *In the new equilibrium after flooding, neighborhoods are more segregated, i.e. $s^*_{a,2} > s^*_{a,1}$, were $s^*_{a,2}$ is the equilibrium ratio of H types in neighborhood A after flooding in period 2, and $s^*_{a,1}$ is the pre-flood s in equilibrium.*

Derivation: Let $t = 1$ be the pre-flood period, and $t = 2$ the post-flood period. The equilibrium condition 5 must hold in equilibrium in both periods, given the values of s^*_a and Z_a at $t = 1, 2$.

$$\begin{aligned} P^*_{a,t} - P^*_{b,t} &= f_h(s^*_{a,t}) + u_h(Z_{a,t}) - f_h(s^*_{b,t}) - u_h(Z_{b,t}) \\ &= f_l(s^*_{a,t}) + u_l(Z_{a,t}) - f_l(s^*_{b,t}) - u_l(Z_{b,t}) \end{aligned} \tag{5}$$

Rearranging terms, and using the fact that $s_b = 1 - s_a$:

$$f_h(s^*_{a,t}) - f_h(1 - s^*_{a,t}) - \left[f_l(s^*_{a,t}) - f_l(1 - s^*_{a,t})\right] = u_l(Z_{a,t}) - u_l(Z_{b,t}) - [u_h(Z_{a,t}) - u_h(Z_{b,t})] \tag{6}$$

The left hand side in 6 is strictly decreasing in s_a, as its derivative with respect to s_a is equal to:

$$\frac{\partial f_h(s_{a,t})}{\partial s_{a,t}} - \frac{\partial f_h(1 - s_{a,t})}{\partial s_{a,t}} - [\frac{\partial f_l(s_{a,t})}{\partial s_{a,t}} - \frac{\partial f_l(1 - s_{a,t})}{\partial s_{a,t}}] = 2 \cdot f'_h(s_{a,t}) - 2 \cdot f'_l(s_{a,t})$$

and $f'_l(\cdot) > f'_h(\cdot)$.[66]

The change in the right hand side in 6 in the post-flood period with respect to pre-flood is equal to:

$$u_l(Z_{a,2}) - u_l(Z_{b,2}) - [u_h(Z_{a,2}) - u_h(Z_{b,2})] - [u_l(Z_{a,1}) - u_l(Z_{b,1})] + [u_h(Z_{a,1}) - u_h(Z_{b,1})]$$

Rearranging terms:

$$u_l(Z_{a,2}) - u_l(Z_{a,1}) - [u_l(Z_{b,2}) - u_l(Z_{b,1})] - [u_h(Z_{a,2}) - u_h(Z_{a,1})] + [u_h(Z_{b,2}) - u_h(Z_{b,1})]$$

Because of the model assumption (2), that is, that utility loss after the flood for a given resident type is proportional to pre-flood residential prices, this is equivalent to:

$$\underbrace{\left[1 - \frac{1}{\alpha}\right]}_{>0} \cdot \underbrace{[u_l(Z_{a,2}) - u_l(Z_{a,1})]}_{①} - \underbrace{\left[1 - \frac{1}{\alpha}\right]}_{>0} \cdot \underbrace{[u_h(Z_{a,2}) - u_h(Z_{a,1})]}_{②}$$

Flooding is a disamenity for both types, so both ① and ② are negative in the equation above. Because of the model assumption (1), that is, that the utility loss after the flood is larger for L types than H types, it must be that the absolute value of ① is larger than ②.

[66]As noted on section ??, a required condition for the existence of an interior partial segregation solution is that the slope of the willingness to pay for the preferred neighbor is higher for L-types than for H-types.

Then, we conclude that the change in the right hand side in 6 in the post-flood period with respect to pre-flood is negative. That is, the right hand side in 6 is smaller after the flood

As the right hand side of condition 6 decreases after the flood, and the right hand side is strictly decreasing in s_a, it must be that the equilibrium level of s_a^* which satisfies condition 5 in $t = 2$ is larger than before the flood.

Model prediction 2: *In the new equilibrium after flooding, the premium for living in A with respect to B could increase, stay constant, or decrease with respect to pre-flood levels, depending on the relative drops in utility derived from the flood and the slopes of $f_j(\cdot)$ for both types of residents.*

Derivation: The steps below show the derivation of condition 4 in the main text, which guarantees that the premium for living in A with respect to B stays constant after a flood. As above, let $t = 1$ be the pre-flood period, and $t = 2$ the post-flood period. The equilibrium condition 7 must hold in equilibrium in both periods, given the values of s_a^* and Z_a at $t = 1, 2$.

$$\begin{aligned} P_{a,t}^* - P_{b,t}^* &= f_h(s_{a,t}^*) + u_h(Z_{a,t}) - f_h(s_{b,t}^*) - u_h(Z_{b,t}) \\ &= f_l(s_{a,t}^*) + u_l(Z_{a,t}) - f_l(s_{b,t}^*) - u_l(Z_{b,t}) \end{aligned} \tag{7}$$

For the premium for living in A to remain constant, it must be that $P_{a,1}^* - P_{b,1}^* = P_{a,2}^* - P_{b,2}^*$. Then, the following must hold for H-types in particular:

$$f_h(s_{a,1}^*) + u_h(Z_{a,1}) - f_h(s_{b,1}^*) - u_h(Z_{b,1}) = f_h(s_{a,2}^*) + u_h(Z_{a,2}) - f_h(s_{b,2}^*) - u_h(Z_{b,2}) \tag{8}$$

For notation simplicity, let δ_j be the difference between utility derived from amenities in A and B before and after the flood for resident type $j = H, L$:

$$\delta_j = u_j(Z_{a,1}) - u_j(Z_{b,1}) - \left[u_j(Z_{a,2}) - u_j(Z_{b,2})\right]$$

and Δs_a the absolute value change in the equilibrium values of s_a^* between the two periods:

$$\Delta s_a = s_{a,2}^* - s_{a,1}^*$$

Reorganizing terms in 8, making use of the fact that $s_b = 1 - s_a$, and the notation definitions above yields:

$$[f_h(s_a^* + \Delta s_a) - f_h(1 - s_a^* - \Delta s_a)] - [f_h(s_a^*) - f_h(1 - s_a^*)] = \delta_h \tag{9}$$

Diving both sides in 9 by Δs_a^* yields:

$$\frac{[f_h(s_a^* + \Delta s_a) - f_h(1 - s_a^* - \Delta s_a)] - [f_h(s_a^*) - f_h(1 - s_a^*)]}{\Delta s_a} = \frac{\delta_h}{\Delta s_a} \tag{10}$$

The left hand side in 10 is by definition of the derivative function equal to:

$$\frac{\partial\left[f_h(s) - f_h(1 - s)\right]}{\partial s} \tag{11}$$

which is arithmetically equal to $2 \cdot f_h'(s)$.

Then, 10 could be rewritten as:

$$2 \cdot f_h'(s) = \frac{\delta_h}{\Delta s_a} \tag{12}$$

Repeating the same steps for the L-types would yield an equivalent condition:

$$2 \cdot f_l'(s) = \frac{\delta_l}{\Delta s_a} \tag{13}$$

Considering 12, 13, and the fact that both H- and L-types experience the same Δs_a yields the required condition for the premium to live in A with respect to B constant after the flood (as in equation 4 in the main text):

$$\frac{f_h'(s)}{f_l'(s)} = \frac{\delta_h}{\delta_l} \tag{14}$$

Appendix E Data: further details on sources, processing, and key statistics

Property sales and buyers' demographics

Sample I only use property sales records with complete data on property location, sale amount and date. Consistently with the literature, I drop records with sale amounts larger than the 99th percentile, and smaller the 5th percentile. Dropping observations with lower sale values also ensures I only consider arm's length transactions, as transactions among kins are usually recorded with symbolic low values, e.g. \$1, \$100. I also exclude properties that have been sold 3 times or more in one calendar year.

Matching process with mortgage lending transactions Mortgage lending transactions include data on loan amount, code of lending institution, census tract, and transaction year. I linked the code of each lending institution with its name scraping the Federal Financial Institutions Examination Council webpage. These four variables, present also in the property sale records, are used for matching a property transaction record with race, ethnicity and income of the buyer.

Consumer price index I deflate property values to \$2010 using the CPI-Housing index from the Federal Reserve Bank of Saint Louis (FRED St Louis, 2019). Buyer's incomes have been deflated to \$2010 current values using the CPI - Northeast urban from the Bureau of Labor Statistics (BLS, 2019).

Place characteristics

Data sources As table B.18 shows, data come from several publicly available sources. Many of the variables come from Chetty et al. (2014) and Chetty et al. (2018) [67]. The benefit of using these datasets is that they contain not only harmonized data from publicly available sources such as the US Census, the FBI Uniform Crime Reporting Program; but also aggregate summary variables derived from non-publicly available federal income tax records (such as fraction middle class, and top 1% income share.) In order to gain more spatial granularity, I get directly from the US Census[68] data that it is publicly available (e.g. racial shares, income per capita, share of occupied housing units at the census tract level.) Other data come from the National Center for Education Statistics (NCES, 2019), the Stanford Education Data Archive (Stanford CEPA, 2019), the County Business Patterns at the US Census (US Census, 2019), NASA's Socioeconomic Data and Applications Center at Columbia University (Van Donkelaar et al., 2018; Geddes et al., 2017), PRISM Climate Group at Oregon State University (PRISM, 2019), , and Open Street Map (Open Street Map, 2019).

Time of measurement: transport All data have been measured prior to 2011, so their levels are not possibly affected by the 2012 flooding. The only exception is for data on transport

[67] Data available at the website of the Opportunity Insights group at Harvard University
[68] Specifically, from the 5-year 2010 American Community Survey and the 2010 Census.

infrastructure location (motorway crossings and subway, train, and bus stations), that has been downloaded from the Open Street Map website as of May 2018. It should also be noted that the transport accessibility variables do not take into account temporary closures for repairs after Sandy.[69] Even if there is no evidence that most of transport infrastructure has changed significantly[70] and property sales decisions might not be overly influenced by temporal transit disruptions, the coefficients on the transport variables should be interpreted with caution.

Matching process I assign each property sale record to the relevant place variables using different techniques according to how these variables were spatially defined. Many of the variables are defined within administrative boundaries (county or census tract.) I use the properties' longitude and latitude coordinates to determine which 2010 census boundaries they correspond to, which it is then used to merge to these variables. For variables measured in a grid, i.e. those on the Environment and Weather categories in table B.18, properties were assigned the value of the cell their longitude and latitude coordinates fell into. Distance to transport infrastructure was measured as the linear distance between the property's longitude and latitude coordinates and those of the closest infrastructure.

Flood risk maps

I obtain data on flood hazard designation from Flood Insurance Rate Map (FIRM) digitized maps from FEMA Flood Map Service (FEMA, 2019b). In particular, I obtain spatial delimitation of the Special Flood Hazard Area (SFHA.) The SFHA marks areas with an annual probability of flood equal to 1% or higher. Owners of properties in the SFHA are required by law to purchase flood insurance if receiving a federally backed mortgage (Horn and Webel, 2019).

These data come with two caveats. First, digitized maps are not available throughout the area of study. Notably, Atlantic county in New Jersey is missing a map. Records in these areas are hence not included in models which control for being in the SFHA. Second, the maps available are current as of October 2018. Some communities have seen their flood maps updated since Sandy's flooding in October 2012. This will introduce some error in my SFHA classification, as some properties that I assign as inside the SFHA were actually not in the flood hazard area in 2012 (with false positive type of errors, i.e. properties not classified as SFHA that were actually in the flood hazard area in 2012, also possible but less likely.)[71]

[69] For instance, nine New York City subway tunnels were affected by Sandy flooding, whose repairs are still ongoing on 2019 (MTA, 2016). .

[70] A major exception to this would be the planned 15-month closure of the L subway line in NYC, connecting Brooklyn and Manhattan (MTA, 2019). Even if it was not finally carried, the expectation of a complete shutdown of the line might have had an impact on property prices locally. Further, four subway stations opened in NYC since Sandy. As these changes are affecting NYC, with an already existing dense transport network, I don't expect the coefficients on variables measuring distance to transport infrastructure to be overly affected

[71] Note that the decision to update flood hazard maps is endogenous to the community, and oftentimes involves a legal contentious process (Ramey, 2015). Preliminary flood maps are made available to the general public during the updating process, so community residents could be taking them into account in decision making before they are legally made effective.

Flood insurance claims and policies

Finally, I use two datasets from FEMA's National Flood Insurance Program (NFIP) to evaluate changes in flood insurance before and after Sandy (FEMA, 2019d,c). The NFIP is the largest provider of residential flood insurance in the United States. In 2016, 90% of all flood insurance premiums were written for the NFIP, while only 10% were for private flood insurance (Insurance Information Institute, 2018). Data from the NFIP would then draw an informative picture on how residents affected by Sandy engaged with flood insurance.

One of the datasets contains data on NFIP claims, including insurance payment amounts and the date the water entered the building. The second dataset lists NFIP policies in force, including beginning and termination date, premiums, coverage amounts, and census tract of the property. Both datasets contain data since January 1st, 2009.

The unit of observation in both datasets is the individual claim or policy, respectively. Data have been redacted to ensure anonymity, so the smallest geography indicator is the census tract. Hence, the analysis of flood insurance claims and policies will be done at the census tract level, as it is not possible to assign a claim or policy to a particular property sale. In order to create a sample of census tracts that is plausibly comparable to the set of properties located less than 500m from the shoreline, I select census tracts that have at least one property sale record on the 500m buffer from the shoreline.

Appendix F Principal Component Analysis

I summarize data on place variables and buyers' demographics using Principal Component Analysis (PCA). PCA is a suitable technique to obtain a low-dimensional representation of the data that explains as much as possible of the original variation.[72] Formally, PCA finds the linear projection of the data that minimizes the distance to the original data points.[73] PCA yields a series of unit vectors – eigenvectors, or principal components (PCs) – which are linear combinations of the original variables, and their associated eigenvalues, which determine the proportion of variance explained by each PC. Hence, the eigenvectors with the largest associated eigenvalues identify the directions in which the original data exhibits the largest variation (Woods and Edwards, 2011; James et al., 2015).

With respect to buyers' demographics, I construct a summary variable of buyers' race and income, whereby buyers who rank high on this composite variable have higher income and are more likely to be white. This variable allows me to evaluate empirically predictions of the model described on 2, which predicts join sorting along these dimensions.

With respect to place characteristics, the main benefit of combining indicators into principal components in this setting is to avoid spuriously capturing heterogeneity alongside single place variables. Out of random chance alone, 1 out 20 of the place variables identified would yield results with a p-value smaller than 0.05. Single indicators within each of the categories are multicollinear. The first principal component captures the data signal within each of the overarching place categories, and hence it is less noisy than a single indicator. Hence, PCA increases confidence on not measuring overfitted models compared with the use of a single indicator to represent an overarching category, while it does not sacrifice the original goal of describing as many observable place characteristics as possible. One drawback of PC variables, as linear combinations of variables, is the loss of tractability. As a robustness check, I replicate main results using single variables instead of a PC combination for key variables in section 7.

Figure A.6 summarizes the results of the principal component analysis for the place variables. Specifically, it shows the scree plots for each of the twelve categories, that is, the eigenvalue or proportion of variance explained by each of the PCs. I determine how many PCs to include in the analysis for each category by visually inspecting where the proportion of variance explained significantly drops.[74] This results in picking only the first PC in most of the categories, and the second one as well in the segregation and social capital categories.

Figure A.6 also shows the loadings of each PC variable, that is, the coefficients assigned to each of the original variables in the linear combinations that yield the PCs. The original variables were normalized to have mean zero and standard deviation equal to one before performing PCA, to avoid the analysis being affected by variable scaling. Considering these loadings, I also narratively describe each of the PCs in figure A.6.

[72] Some examples in the economics literature which use a low-dimension representation of urban amenities are Collinson and Ganong (2018) and Diamond (2016). In particular, Diamond (2016) also uses PCA in her approach.

[73] This is equivalent to finding which linear combination of the original variables maximizes variance.

[74] This visual test is known as the "scree" test (James et al., 2015)

Appendix G Machine learning procedure for estimating heterogeneous effects

G.1 Sketch of procedure in Chernozhukov et al. (2018)

For completeness, I sketch below the steps taken to measure heterogeneity following Chernozhukov et al. (2018). For formal proofs, I refer to the original paper:

1. The data is randomly split into two samples, main and auxiliary, 100 times.
2. For each of the 100 iterations, the auxiliary sample is used to train two models to estimate proxies for the baseline and treatment effects using a machine learning method, i.e. the model which optimally estimates the outcomes based on covariates is chosen separately for flooded and non-flooded properties. The baseline and treatment proxies are then obtained by applying the chosen models to the main sample.[75] Following Chernozhukov et al. (2018) , I compute results with four different ML methods: random forests, elastic net, boosted tree, and neural networks.[76]
3. With the baseline treatment proxies in hand, consistent key features of the treatment effects are estimated using the main sample data as follows:
   (a) The best linear predictor of the treatment effect is obtained by running the following weighted linear regression:

$$Y = \alpha_0 + \alpha_1 B(Z) + \alpha_2 S(Z) + \beta_1 [D - p(Z)] + \beta_2 [(D - p(Z)][S(Z) - \mathbb{E}(S(Z))] + \epsilon \tag{15}$$

   with

$$\mathbb{E}[\omega(Z)\epsilon X] = 0$$

   and where Y is the outcome variable, i.e. the log of the property sale price; D is a variable indicating treatment status, i.e. equal to one for properties on Sandy's floodplain, and to zero for those outside; Z is a vector of place covariates, as described in section 3; $B(Z)$ and $S(Z)$ are, respectively, the proxy baseline and treatment effects obtained on the previous step; $p(Z)$ is the probability of assignment; $\omega(Z)$ is equal to $\{p(Z)(1 - p(Z))\}^{-1}$; and X is $[1, B(Z), S(Z), D - p(Z), (D - p(Z))(S(Z) - \mathbb{E}(S(Z))]$. Standard errors are two-way clustered at the census tract and quarter of year levels, to maintain equivalency with the main specification of the paper.
   Chernozhukov et al. (2018) prove that the coefficients β_1 and β_2 estimated in 16, pin down the best linear predictor of the true treatment effect given the proxy treatment effect, i.e. solve the following maximization problem:

[75] These proxies for treatment effects might be noisy and biased predictors. Hence, they need to be further post-processed to generate consistent estimates.

[76] Main results presented are derived from the preferred random forest method, which performs better according to criteria on Chernozhukov et al. (2018). Results from other methods available from the author upon request.

$$(\beta_1, \beta_2)' = \underset{b_1, b_2}{\operatorname{argmin}} \mathbb{E}[s_0(Z) - b_1 - b_2 S(Z)]^2 \tag{16}$$

where $s_0(Z)$ is the true treatment effect given covariates.
In particular: $\beta_1 = \mathbb{E}[s_0(Z)]$ and $\beta_2 = \text{Cov}[s_0(Z), S(Z)]/\text{Var}[S(Z)]$. Rejecting the hypothesis that β_2 – which Chernozhukov et al. (2018) refer to as the "heterogeneity loading" – is equal to zero means that (1) there is heterogeneity in the treatment effect according to covariates, and (2) the proxy $S(Z)$ is a relevant predictor.

(b) Average effect by groups, defining groups as the quintiles according to the predicted treatment effect, are obtained by running the following weighted linear regression:

$$Y = \alpha_0 + \alpha_1 B(Z) + \alpha_2 S(Z) + \sum_{k=1} K\gamma_k \cdot [D - p(Z)] \cdot \mathbb{1}(G_k) + \nu \tag{17}$$

with

$$\mathbb{E}[\omega(Z)\nu W] = 0$$

and where G_k, with $k = 1, ..., 5$, representing one of the quintiles which divide the support of $S(Z)$ into five non-overlapping regions; and W is equal to $[1, B(Z), S(Z), D - p(Z), \sum_{k=1}^{K} \mathbb{1}(G_k)]'$. As above, standard errors are two-way clustered at the census tract and quarter of year levels.
Chernozhukov et al. (2018) prove that the coefficients γ_k estimated in equation 17 are the group average treatment effects, i.e.:

$$\gamma_k = \mathbb{E}[s_0(Z)|G_k] \qquad \text{with} \quad k = 1, ..., 5$$

(c) Average characteristics of the properties on the top and bottom quintiles are obtained by averaging the place characteristics, $g(Y, Z)$ of the properties in each of the quintiles:

$$\delta_i = \mathbb{E}[g(Y, Z)|G_i] \qquad \text{with} \quad i = 1, 5$$

4. Point estimates of key features are computed as the median of the estimates obtained in the 100 splits.
5. The final step is to compute confidence intervals and p-values that take into account uncertainty coming both from the estimation as well as on the data splitting. Chernozhukov et al. (2018) prove that this could be achieved by taking the medians of confidence interval bounds and p-values obtained in the 100 steps and adjusting its nominal level (from a $1 - \alpha$ confidence level to a $1 - 2\alpha$ level) to allow for splitting uncertainty.
6. Additionally, Chernozhukov et al. (2018) compute metrics which allow comparing the performance of the different ML methods. In particular, the paper suggests that the best ML method could be chosen as the one that maximizes the following metric:

$$\Lambda := |\beta_2|^2 \text{Var}(S(Z))$$

which is equivalent to maximizing the correlation between the true treatment value $s_0(Z)$ and the computed proxy predictor $S(Z)$.

G.2 Machine learning methods: sample and tuning parameters

I complete the CDDF procedure with four different ML methods: random forests, elastic net, boosted tree, and neural networks using the *caret* package in *R*. Following Chernozhukov et al. (2018), I chose model tuning parameters by repeated 2-fold cross-validation[77] in the elastic net, boosted tree, and neural networks methods; and use the default tuning parameters in the random forest method.

Given the large number of observations in my dataset, allocating half of the observations to the auxiliary sample as in CDDF is too taxing computationally for some of the methods. Hence, I maximize which percentage to allocate to the auxiliary sample while maintaining computational times within reasonable limits,[78] and check that the main results are not sensitive to the specific percentage chosen.[79]

[77] In 2-fold cross-validation, the auxiliary sample is divided into two sets, or *folds*. The model is fit in one of the folds, and the mean squared error given by the model is computed in the other fold. The process is repeated twice, each time fitting the model in one of the folds. The chosen tuning parameters are the averages of both folds. The process is repeated two times.

[78] Specifically, the following percentage of observations are allocated into the auxiliary sample in each of the methods: 12.5% in random forests; 10% in elastic net and boosted trees; and 5% in neural networks.

[79] Further robustness checks in which different percentages of the data are allocated to the auxiliary samples are available from the author upon request.

Appendix H Discrete change at the flood boundary: regression discontinuity

H.1 Regression discontinuity: model

I evaluate whether there is a discrete jump in outcomes above and below the flood extent boundary using a regression discontinuity design, that is, I evaluate differential changes in outcomes for properties that were barely flooded with respect to those that were nearly missed by the flood. This approach follows the spatial boundary discontinuity design, frequently applied in the literature (e.g. Black (1999); Bayer et al. (2007). However, unlike these cases, the running variable will not be the layout distance to a boundary, but rather, the vertical distance to the maximum level reached by the flood surge nearby each property.

This analysis complements the DD model described in 4.1 by analyzing whether changes between flooded and non-flooded properties are driven partly by discrete changes alongside the flood boundary extent. The exact flood extent depends on highly localized atmospheric and physical phenomena, and hence, it would not have been possible for individuals to sort on both sides of the threshold in foresight of the flood. I would expect a discontinuity in the outcomes at the flood extent boundary if Sandy flooding sent a strong signal about the risk of flooding. On the other hand, if the risk of flooding was updated similarly for barely flooded and nearly missed properties, I would expect outcomes to change more continuously across the threshold determined by the flood extent boundary.

Besides complementing it, a regression discontinuity design allows a more relaxed identification assumption than the DD model. In the DD model, I assume that any unobservable factors affecting the outcomes change synchronously in the treatment and the control groups (i.e. flooded and non-flooded properties.) The RD model, on the other hand, can be understood as a "local randomized experiment" around the threshold (Lee and Lemieux, 2010). In particular, the RD model allows unobservable variables to have a nonlinear effect, as long as they do not change discontinuously at the flood boundary extent. The penalty for a more relaxed identification assumption is that the effect is only locally identified around the flood boundary extent.[80]

As described in section 3, I construct a metric to measure vertical distance to flood level. This metric is positive for properties whose ground stood above the flood maximum water elevation, and negative for those below. With this running variable in hand, I run the following specification to estimate whether there is a discrete jump in outcomes around the flood extent boundary using only post-Sandy data:

$$y_p = \beta \mathbb{1}\{c_p \leq 0\} + f({c_p}^+) + \mathbb{1}\{c_p \leq 0\} f({c_p}^-) + \epsilon_p \tag{18}$$

where y_p is the residualized log sale price of property p, net of census block and month fixed effects; c is the running variable, i.e. the difference in elevation between the property ground and the flood water elevation at the closest point; and $f(\cdot)$ represents a function of

[80] Lee and Lemieux (2010) suggest a broader interpretation of RD coefficients, as the weighted average treatment effect, with weights proportional to the ex ante likelihood of being close to the threshold. The more similar the weights are among observations, the closer the RD estimator would be to the average treatment effect. Given the documented sorting between coastal and non-coastal regions (Bakkensen and Barrage, 2017), I abstract from assuming similarity of weights for observations further away from the flood extent boundary, and hence prefer the localized interpretation of the RD estimator

the running variable, whose parameters are allowed to be different above and below the threshold determined by $c = 0$ (${c_p}^+$ and ${c_p}^-$, respectively). I present results for seven different functional forms of $f(\cdot)$: globally linear, quadratic, third order and four order polynomial fits; as well as local non-parametric, linear, and quadratic regressions.[81] Standard errors ϵ_p are clustered at the census tract level.

The parameter of interest is β. It estimates the change in the outcome of interest at the flood extent boundary.

As noted above, identification relies on all observable and unobservable variables affecting the outcome of interest changing smoothly at the threshold defined by the flood boundary extent. Exploiting the fact that I have data pre-Sandy, I can convincingly test this is not the case by checking whether there was a discontinuity in the outcomes of interest along the flood extent boundary using data prior to Sandy. Following the analogy of the RD model as a "local randomized experiment", this test would be equivalent to check whether randomization has been achieved by comparing outcome values on both sides of the threshold pre-treatment (Lee and Lemieux, 2010).

H.2 Regression discontinuity: results

I present suggestive evidence against a discrete jump in property prices above and below the flood maximum water elevation on average.

Table H.1 shows results from regression discontinuity models with a running variable measuring elevation with respect to the flood maximum water level. Models in the top panel use all the data, where models in the bottom panel are local regressions using only data around the threshold. Further, the models differ in the functional form assumed for the polynomial controlling for elevation with respect to the flood. Figure H.1 plots each of the polynomial approximations in panel (a) against the data, and hence provide a visual confirmation of fit.

A global linear polynomial in elevation (column (1) in the top panel) yields a significant negative effect of the flood. Properties barely flooded sold an average of 3.7 pp below nearly missed non-flooded properties. This discrete jump is also evident on the top-left panel in figure H.1. However, this result is not robust to the change in the order of the polynomial. The result is not significant for polynomials of order 2 and 4, and it even changes sign for the third order polynomial.

Following Lee and Lemieux (2010) I present two metrics to guide the selection of the optimal model, the Akaike information criterion (AIC)[82] and the p-value of the hypothesis which test the joint significance of a set of elevation bin dummies,[83] noted in the "goodness of fit" row in table H.1.

Results from these tests are inconclusive. According to the goodness-of-fit test, the linear approximation would be preferred, as it yields the higher p-value. However, with the p-

[81] To run the RD models with local regressions I use the `rdrobust` package in Stata, developed by Calonico et al. (2014). I use the package defaults to select optimal bandwidths and kernel functions to construct the estimators (the one common mean square error-optimal bandwidth selector procedure, and the triangular function, respectively.)

[82] As AIC penalizes complexity, the model with the lowest value of AIC would be the most parsimonious.

[83] The logic of this test is that, if the polynomial were a good fit, these bin dummies should all be jointly equal to zero

Table H.1: Regression Discontinuity: results

	(1)	(2)	(3)	(4)
$\text{Flooded}_{\text{p}}$	-0.037**	0.006	0.016**	0.006
	(0.015)	(0.006)	(0.008)	(0.010)
Order of Polynomial	linear	quadratic	3^{rd}	4^{th}
AIC	135534	135428	135427	135419
Goodness of fit	0.0376	0.0000	0.0000	0.0000
Observations	102571	102571	102571	102571

(a) Global polynomials

	(1)	(2)	(3)
$\text{Flooded}_{\text{p}}$	0.002	-0.009	-0.007
	(0.019)	(0.023)	(0.026)
Order of Polynomial	zero	linear	quadratic
Bandwidth	0.283	0.488	0.744
Observations - below cutoff	10956	17609	25485
Observations - above cutoff	9622	15131	20702

(b) Local polynomials

Notes: Regression discontinuity results. Dependent variable in all models is residuals of property price sale after Sandy, net of block and month level fixed effects, approximating all data with a global polynomial – top panel, linear polynomial in column (1), quadratic in (2), third order in (3) and fourth order in (4) – and using only data around the threshold in a local regression (bottom panel, non-parametric comparison of averages in column (1), linear polynomial in (2), and quadratic in (3)) The running variable measures elevation with respect to flood water level, positive for properties p above the flood, and negative for properties p below. The regression discontinuity estimate then measures whether there was a discrete jump in property prices above and below the flood maximum water elevation. Results in the bottom panel were computed using the `rdrobust` package in Stata, developed by Calonico et al. (2014). Optimal bandwidths are obtained by the one common mean square error-optimal bandwidth selector procedure. A kernel triangular function is used to construct the estimators. Standard errors clustered at the census tract level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Figure H.1: Property price sale with respect to vertical distance to flood level, polynomial fit

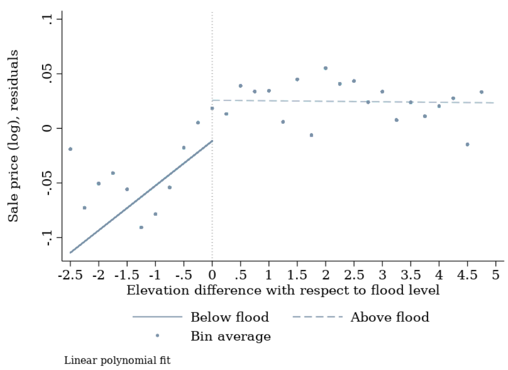

(a) Linear

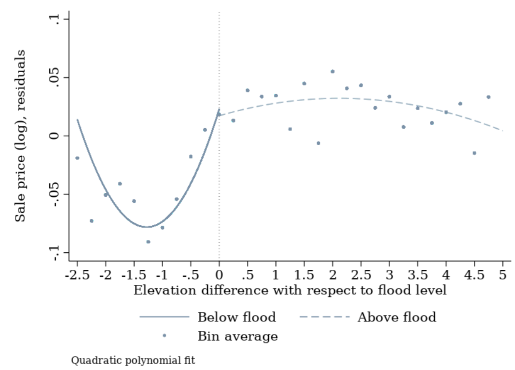

(b) Quadratic

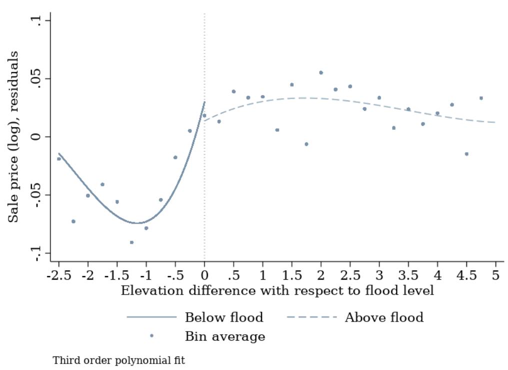

(c) Third order

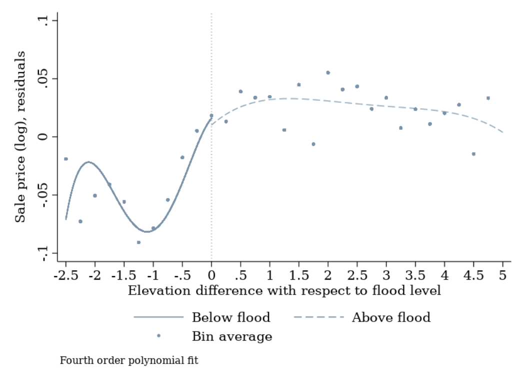

(d) Fourth order

Notes: Figure shows fit of different order polynomial approximations to the raw data, from top-left figure, clockwise: linear polynomial, 2nd, 4th and 3rd order polynomial approximations. Specifically, the y-axis represents property price sale residuals post-Sandy, net of census block and month fixed effects. The x-axis shows the difference between a property's ground level and the maximum level the flood reached closest to the property, which is negative for flooded properties, and positive for non-flooded properties. The dotted line at zero marks the maximum flood surge level. Figure shows polynomial fit, and price averages at 0.25m-wide elevation bins of the raw data.

value of 0.038, I can still rule out the hypothesis that the elevation bins are all jointly equal to zero at 95% significance level. According to AIC values, the 4$^{\text{th}}$ polynomial in column (4) would be optimal as it yields the lowest value.[84]

Results using local polynomial approximations (bottom panel (b) in table H.1) show non-significant coefficients in the non-parametric model (column (1)), and linear and quadratic models (columns (2) and (3), respectively.)

Overall, these results provide evidence against a discrete jump in prices around the maximum flood water level, and are suggestive of prices changing smoothly around the flood boundary. Barely flooded and nearly missed properties did not see a differential change in prices.

[84] In fact, the AIC including even higher order polynomials keeps decreasing (the lowest value is achieved with a 6$^{\text{th}}$ order polynomial, after which it starts increasing again – results not shown.) However, following Gelman and Imbens (2017), which argue against the approximation of high order polynomials in empirical settings, I do not consider that approximations using higher order polynomials are warranted in this case.)