Sea Level Rise and Home Prices: Evidence from Long Island

Justin Tyndall

jtyndall@hawaii.edu
University of Hawai'i Economic Research Organization
and University of Hawai'i at Mānoa Department of Economics
2424 Maile Way, Saunders Hall 540, Honolulu, HI, USA, 96822

August 30, 2021

Abstract

Global sea level rise is a known consequence of climate change. As predictions of sea level rise have grown in magnitude and certainty, coastal real estate assets face an increasing climate risk. I use a complete data set of repeated home sales from Long Island in New York State to estimate the appreciation discount caused by the threat of sea level rise. The repeat sale methodology allows for time-invariant, unobserved property characteristics to be controlled for. Between 2000 and 2017, I find that residential properties that were exposed to future sea level rise experienced an annual price appreciation rate of roughly one percentage point below unexposed properties. I provide numerous robustness checks to confirm this result. I also find evidence of demand spillovers by estimating an appreciation premium for properties that are near the coast but are relatively safe from sea level rise.

Housing markets; Climate change; Sea level; Flooding
G10; R30; Q54

1 Introduction

Predictions of global sea level rise have become more dire (Chen et al., 2017). Current estimates suggest that by 2100 global sea level will be between 0.3 and 1.0 meters above 2000 levels (Church et al., 2013; Nicholls and Cazenave, 2010). The threat of sea level rise is compounded by the predicted increase in climate change induced extreme weather events that could cause water to surge into coastal areas (Seneviratne et al., 2012). Many real estate assets on the coast will experience climate change induced land erosion, flood events, or total inundation, all of which could decrease or eliminate the asset's value. I contribute new estimates of a climate change price effect in the US housing market.

I estimate the extent to which climate risk has been priced into real estate transactions for the Long Island housing market in New York State. Using repeated property sales, I compare the rate of price appreciation among properties at risk from sea level rise to other properties that are not at risk. The use of property level fixed effects allows for the price effect of coastal proximity to be controlled for and the use of detailed spatial coastline data allows for the control of unique appreciation trends for coastal properties. I calculate the sea level rise exposure of properties by combining transaction data with detailed elevation and flood map data. Results indicate that climate risk has reduced the value of exposed residential properties on Long Island.

How the risks of sea level rise influence real estate has been studied by a recent and growing literature. A starting point of the theoretical literature is the assumption that property markets are populated by forward-looking agents with information on future climate risk. Bunten and Kahn (2017) provided important theoretical work on the effect of climate change risk on property development and investment. If climate risk is included in the property valuation of buyers and sellers, financial incentives will shift investment decisions and partially insulate against sudden climate shocks. Kahn

(2016) provided an important overview of how individuals adapt to climate change and Kahn (2014) as well as Desmet and Rossi-Hansberg (2015) provided further discussion regarding adaptation in spatial location decisions. Severen et al. (2018) argued that much of the future costs of climate change in land markets are already accounted for due to pervasive information and beliefs about the probable effects of future climate change.

Some empirical work has been undertaken to estimate the effect of exposure to sea level rise on home prices. Bernstein et al. (2019) relied on Zillow home price data from across the US and found that homes exposed to sea level rise sell at a 7% discount relative to other homes that are similar, based on observable characteristics. Contrastingly, Murfin and Spiegel (2020) studied housing transaction data from US coastal states and found no evidence of climate risk being priced into home sales. The study made use of the fact that sea level rise affects different coastal areas differently due to local subsidence or uplift of continental landmasses partially offsetting or exacerbating sea level rise.

Buyers who personally believe climate change risk to be large are more likely to incorporate climate change risk into their purchase decisions. Some studies have found spatial heterogeneity in the US regarding to what extent climate change risk is priced into real estate assets, driven by spatial heterogeneity in personal beliefs about climate change risk. McNamara and Keeler (2013) supplied a model to study coastal climate change risk in the US Northeast, noting that heterogeneity in agent beliefs regarding climate change risk influence the overall support for climate mitigation measures. Bakkensen and Barrage (2018) as well as Baldauf et al. (2020) further extend theory and analysis of how belief heterogeneity affects the pricing of climate risk into coastal property assets. Both studies find that households that are skeptical regarding the risks posed by climate change are willing to pay relatively more for coastal assets that are in areas of significant risk, and this behavior leads to a difference in home prices across local markets.

Informational issues are also addressed in studies of markets that have been exposed to an extreme weather or flood event. Results have shown that individual weather events can generate sudden shifts in how future climate risk is priced into real estate assets. Gibson and Mullins (2020) investigated price declines associated with climate risk in New York City. Results show significant price reductions for properties directly exposed to information shocks, including flooding from Hurricane Sandy, changes to federal flood insurance, and the updating of FEMA floodplain maps. Ortega and Taspinar (2018) looked specifically at the role of Hurricane Sandy on climate risk discounting in New York City. The authors found a significant price discount that increased over time in areas at risk of flooding. McCoy and Zhao (2018) analyzed a similar process in New York City, finding that homeowners' property investment decisions are influenced by perceived flood risk. The effect of Hurricane Sandy on the behavior of sophisticated real estate investors in New York City was examined in Eichholtz et al. (2019), providing evidence that the implied future value of properties in storm damaged areas fell compared to properties in comparison cities. Bin and Landry (2013) found that the price of homes in North Carolina reflected flood risk much more strongly after a nearby hurricane, despite long-term risk remaining constant. McKenzie and Levendis (2010) uncovered a similar informational effect for properties in New Orleans after Hurricane Katrina. For non-coastal flooding, Yi and Choi (2019) examined the effects of a flood event in Iowa, and Zhang and Leonard (2019) examined a flood event in the Fargo, North Dakota metropolitan area, both found that homes in flooded areas began discounting the future value of their homes more heavily after the flood.

The US federal government takes an active role in regulating flood insurance (Sklarz and Miller, 2018). Homes in high flood risk areas are required to purchase insurance against the risk of floods (Kriesel and Landry, 2004). The presence of a flood insurance requirement will have two notable effects on prices in the coastal market. First, the insurance requirement represents a cost to home owners, putting downward pressure on the price of coastal real estate assets. Second, holding insurance will reduce the downside risk of a flood event, which provides value to the homeowner. There is some evidence that flood insurance is underpriced, particularly in very high risk areas (Kousky and Shabman, 2014). Discounted flood insurance would put upward pressure on prices as the property owners are able to offload downside risk to insurance programs, while maintaining the upside risk of price appreciation. Kousky et al. (2020) provided a thorough discussion of the role of flood insurance on the single family housing market in the US, including the behavioral responses of homeowners.

Overall, empirical studies on sea level rise and real estate values have found mixed and conflicting results regarding the existence and magnitude of a price effect. I provide analysis from an important setting where future sea level rise will affect a large share of properties. By using a complete data set of housing transactions from a large coastal market over an 18 year period I am able to provide precise identification of effects. I also identify demand spillovers wherein properties in areas of high climate risk experience an appreciation penalty while properties close to the coast but with lower risk actually experience an appreciation premium as buyer demand for coastal properties shifts within the coastal housing market.

The paper will proceed as follows. The next section discusses details of the Long Island property market. The third section introduces data sources. The fourth section provides the methodology. The fifth section provides primary results and numerous robustness tests and the final section concludes.

2 The Long Island Property Market

The study area includes Nassau and Suffolk counties in New York State. These counties comprise all of Long Island, outside of New York City. Figure 1 provides a map indicating the study area. The numerous coastal communities on Long Island provide an excellent opportunity to study the impact of sea level rise on a real estate market with significant coastal exposure. A significant portion of real estate assets on Long Island are within a few meters of current sea levels, and coastal areas are at risk of damage due to storm events, particularly hurricanes.

Figure 1: Long Island Study Area

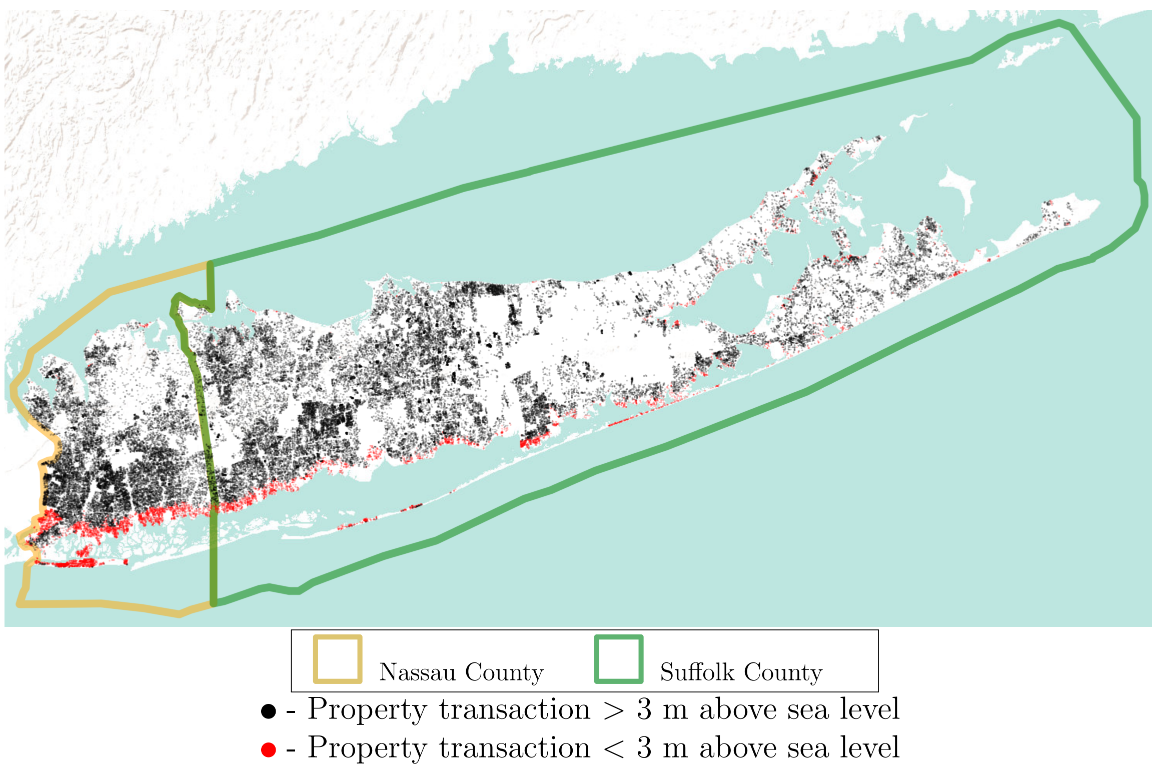

Each of the 73,346 unique properties in the repeat sales transaction data are represented as a point on the map. Red dots indicate properties that are within three meters of sea level. The large majority of properties that are close to sea level are on the southern coast of Long Island.

Table 1 provides demographic summary statistics for the study area and compares the demographics of Long Island to the US as a whole. Nassau and Suffolk counties have higher income and education levels than the US average. The racial composition of the counties are relatively representative of the national population. Nassau and Suffolk have a high rate of homeownership, with 81.3% of residents being owner-occupiers, compared to the national rate of 66.1%.

Table 1: Demographic Characteristics of Study Area

	Nassau and Suffolk Counties	**USA**
Population	2,820,124	306,603,772
Median household income	89,947	52,762
College education rate (%)	36.6	28.2
Median Age	39.8	37.0
White (%)	78.2	74.1
Black (%)	9.1	12.5
Asian (%)	5.4	4.7
Hispanic (%)	15.1	16.1
Owner-occupancy rate (%)	81.3	66.1

Data from the 2007-2011 American Community Survey.

Coastal erosion and flooding are long-term issues faced by coastal landowners on Long Island. An important event for real estate assets in Long Island was Hurricane Sandy, which struck the area in October 2012. The storm event caused significant property damage on Long Island, in New York City, and neighboring states. In addition to the direct damage caused by the storm, Hurricane Sandy may have impacted the perceptions of local homeowners as to the risk of climate change related property damage (McCoy and Zhao, 2018; Ortega and Taspinar, 2018).

3 Data

I use housing data provided by the New York State Department of Taxation and Finance (NYSDTF), Office of Real Property Tax Services. The data covers all real estate transactions within the state of New York between 2000 and mid-2017. I trim the data to only Suffolk and Nassau counties.

I drop a number of observations to focus analysis on reliable data and to enable the repeat sales methodology. I drop any observation that sold for less than \$50,000, as such transactions are unlikely to relate to legitimate, competitive sales. I keep only sales that are classified by the NYSDTF as "arms-length." I also drop observations with incomplete street address information. I exclude sales of individual condominium units. In part, I exclude condominium units because the unit number of condominiums are not reliably recorded in all cases, preventing the matching of specific units through time. I further limit the sample to the following property types as defined by the NYSDTF: One Family Year-Round Residences, Two Family Year-Round Residences, Three Family Year-Round Residences, Apartments, Multiple Residences and Residential Vacant Land. These property types cover 97.3% of property sales in Suffolk and Nassau counties. The excluded property types include commercial, industrial and agricultural land uses. I will provide robustness checks regarding the effects of these sample limitations.

I conduct analysis only on properties that sold at least twice during the study period. Limiting the sample to repeat sales will be important to the methodology wherein I introduce fixed effects at the property level, which partials out the effect of all time-invariant housing characteristics. I use the recorded address of the transacted property to identify unique properties that sold multiple times. I normalize the street names to account for the use of abbreviations, for example matching *Rd* to *Road*. The NYSDTF flags observations that have undergone significant renovations between sales. Homes that underwent significant renovations can not be reasonably assumed to represent the same asset so I consider the post-renovation property as a unique property in the repeat sales method. 0.4% of the properties in the sample underwent a significant renovation between sales.

The final repeat sales data set contains 164,026 transactions spanning 73,346 unique properties. A total of 59,132 properties sold exactly twice, 11,668 properties sold three times, and the remainder sold more than three times. The most transacted property in the data set was sold eleven times. Fully 95.0% of transactions in the final sample are classified as "One Family Year-Round Residences." The median sale price among the repeat sales properties is $443,985. Figure 2 graphs the change in median sale price over time within the repeat sales sample. Home prices on Long Island climbed significantly along with the national US housing market during the 2000-2007 period, with the median property price rising by 70.8%. The 2007-2017 period corresponds to a significant decline in home prices on Long Island, with the median home value dropping by 27.2%. The average property within the sample appreciated at a rate of 0.5% annually in real terms (2.7% in nominal terms) across the 2000-2017 study period. Figure 2 also displays annual median sales prices for the full sample of properties (N=451,237), including those that sold only once over the study period. The two series track very closely to one another. The correlation coefficient between the two series is 0.992, suggesting that the repeat sales data is representative of the overall market.

Figure 2: Trend in Repeat Sale Median Property Price

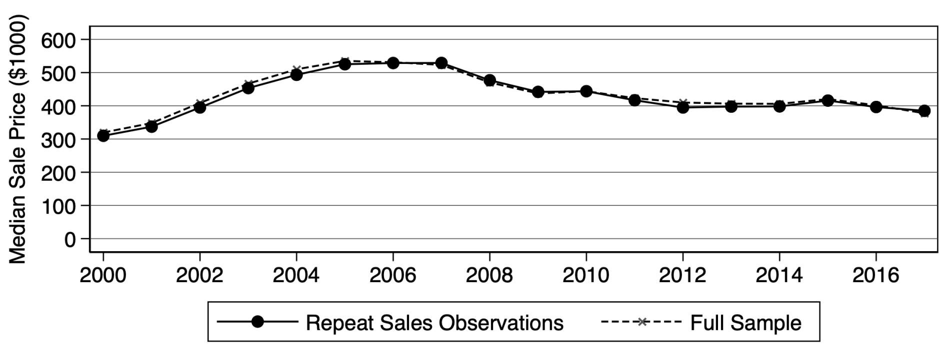

The median sale price of a property on Long Island rose between 2000 and 2007. Prices declined beginning with the Great Recession. Between 2012 and 2017 prices were relatively stable. All prices are in 2017 USD.

Figure 3 shows the distribution of sale prices across the sample of repeat sale observations. Panel A shows the sale prices in 2017 USD. I use the log of the sale price in analysis. Panel B shows the distribution of logged sale prices, which provides an approximately normally distributed dependent variable for analysis.

Figure 3: Sale Price Distribution

A. Sale Prices

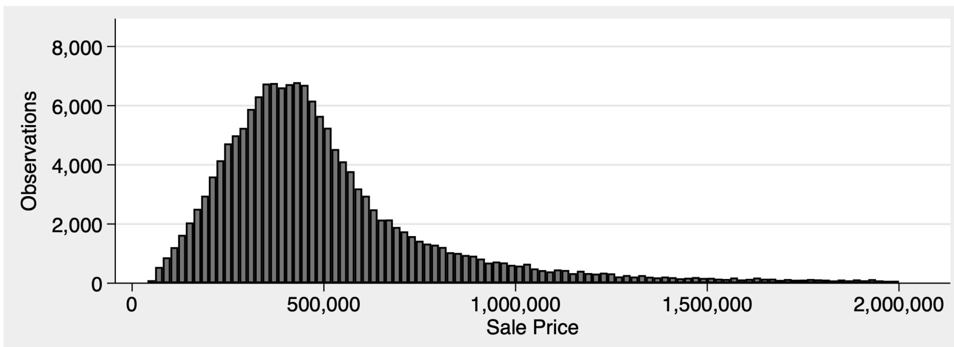

B. Logged Sale Prices

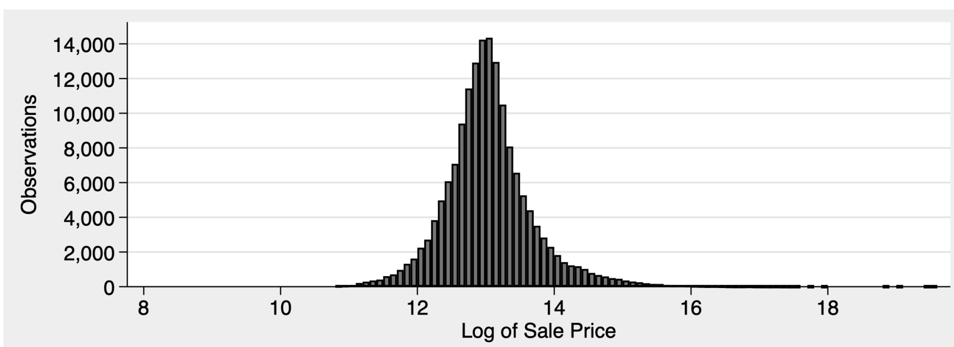

Panel A shows the distribution of sale prices, across all years in the sample. Panel A truncates the data at $2,000,000. Panel B shows the distribution of logged sale prices, which will be the variable used in analysis. All prices are in 2017 USD.

I use the street address of each property in the NYSDTF data in order to calculate the latitude and longitude coordinates of the property's centroid. I rely on the geocoding web service HERE to convert street addresses into precise latitude and longitude coordinates.

Two separate data sources will be used to evaluate the extent to which a property is exposed to the risk of sea level rise. First, I make use of the Federal Emergency Management Agency (FEMA) flood maps. In particular, I use the FEMA National Flood Hazard Layer for Nassau and Suffolk counties. The maps classify all land into zones of various flood risk. I use FEMA defined 100-year flood zones, which correspond to areas that have at least a 1% chance of flooding during a given year. I assign each property transaction as being either inside or outside of a 100-year flood zone through spatial mapping software.

The second definition I use for sea level rise exposure is the elevation at the centroid of the property. Using the latitude and longitude coordinates of each property's centroid, I make use of the United States Geological Survey (USGS) online Elevation Point Query Service. The web service can return the elevation of any set of latitude and longitude points for the US. I generate an individual query for all 73,348 unique properties, creating precise elevation estimates. The elevation estimates from USGS are interpolated from the 1/3 arc-second 3D Elevation Program DEM dataset. For Long Island, USGS elevation points are measured approximately every 8 meters and interpolated between these points, providing highly accurate elevation estimates.

Figure 4A provides an elevation map of Long Island. Most areas that are within three meters of sea level are located along the southern coast of Long Island. Figure 1 provides the locations of all the transactions in the data set that are within three meters of sea level. The elevation profile of Long Island is such that coastal areas are the only locations that have an elevation that is close to sea level. There are no inland areas of low elevation that do not extend to the coast. This fact is important to justify using elevation as a proxy for exposure to sea level rise. Figure 4B maps the location of FEMA defined 100-year flood zones. The location of FEMA flood zones are highly

Figure 4

A. Elevation Map of the Study Area

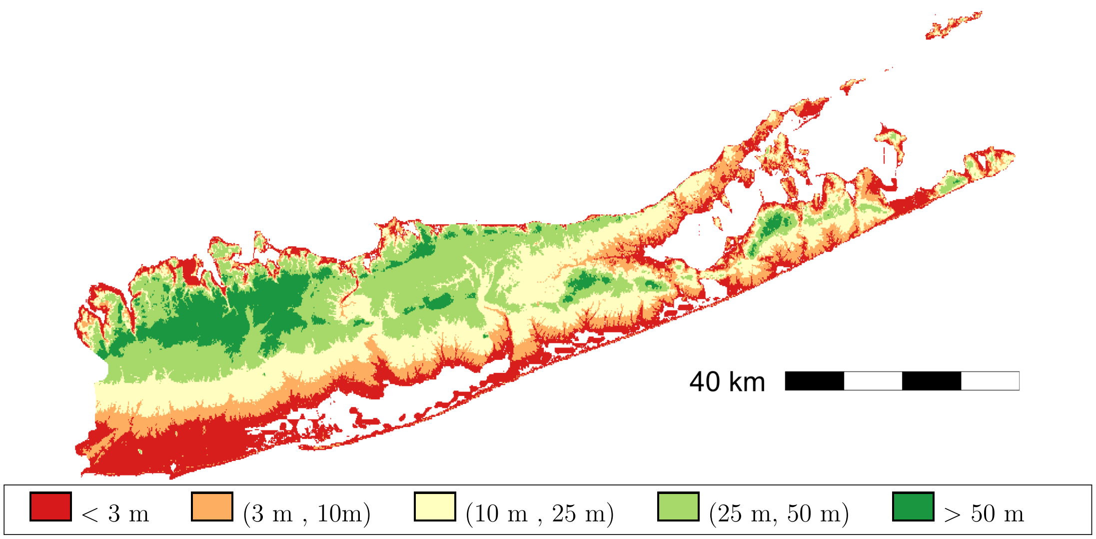

B. FEMA Flood Map

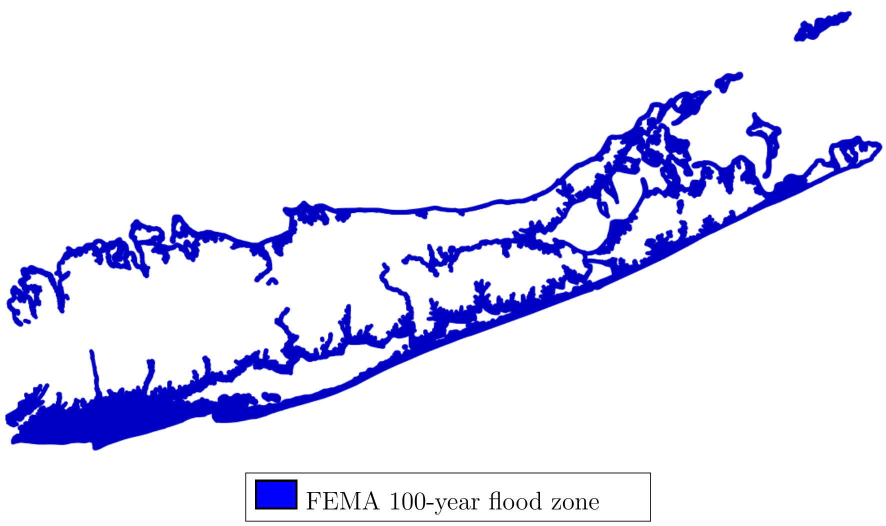

Panel A displays the elevation across Nassau and Suffolk counties. Most areas close to sea level are on the south coast. Panel B displays all areas that are classified by FEMA as being within a 100-year flood zone.

correlated with areas of low elevation.

Figure 5 provides a histogram of the elevation of properties in the repeat sales

Figure 5: Elevation of Repeat Sale Properties

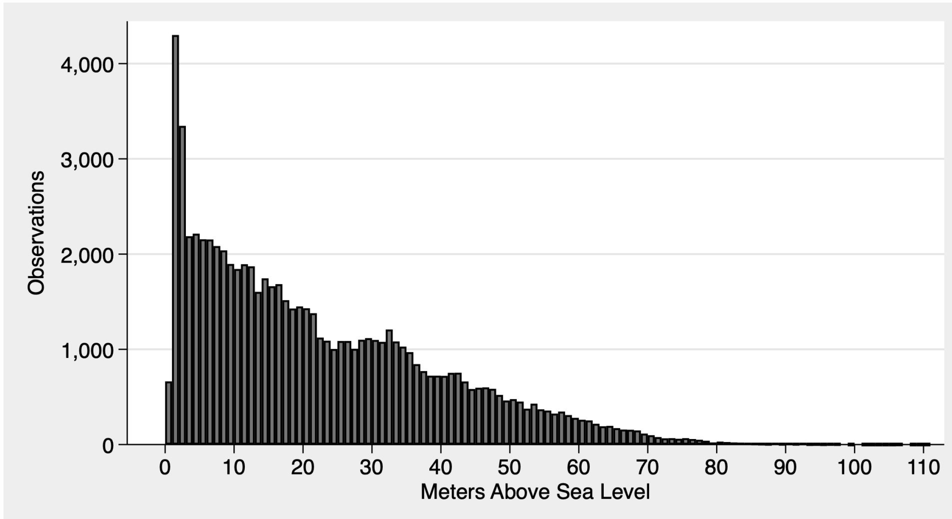

The plot shows the frequency of properties at each elevation level, using one meter bins. For example, there are 4,300 properties in the data set that are at an elevation of between one and two meters.

sample. A large share of properties are close to sea level. Across all properties, 11.4% are within three meters of current sea level and 6.8% are within two meters.

I measure the distance of each property to the coast. I use the publicly available New York State Civil Boundaries Shoreline shapefile. I calculate the shortest straight line distance, in meters, from every property observation to the nearest coastline. I make use of these measurements to control for potentially differential price appreciation for properties with coastal proximity. By combining data on elevation with data on coastal proximity, the proposed methodology will be able to separately identify the price appreciation effects of coastal proximity from the appreciation effects of sea level risk exposure.

Finally, I use the FEMA National Flood Insurance Act Redacted Claims Data Set. The data provides census tract level information on insurance payments made by FEMA. I assign each property to a census tract using the US Census TIGER shapefile.

I calculate the total value of FEMA insurance claims that occurred in each census tract in each year. Figure 6 shows the total value of claims made across the study period. Claims totaling \$2.5 billion were made in 2012, as a result of Hurricane Sandy. Across the study period, 2012 accounted for 92% of all claims by value. The average value of claims made in the other years was \$13.4 million. I calculate the cumulative value of claims that occurred in each tract through time. I use this variable as an indication of past exposure to storm damage. For two consecutive sales of the same property, the difference in this measure will be the total amount of local storm damage that occurred between the two sale dates. I make use of this variable to test whether past storm damage affects climate risk discounting.

Figure 6: Total Value of FEMA Claims on Long Island

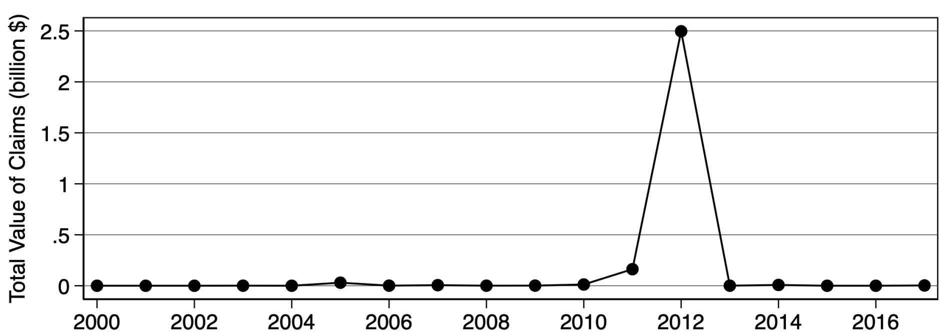

In 2012, Hurricane Sandy caused a large increase in FEMA insurance claims on Long Island.

4 Methodology

A major empirical challenge to identifying the price effect of coastal climate risk is untangling the effect of exposure to sea level rise with the amenity effect of coastal proximity. Homes closer to the beach are likely to sell at a premium to other homes, independent of climate risk (Atreya and Czajkowski, 2019; Conroy and Milosch, 2011).

I perform a repeat sales analysis and study how property appreciation differs during a period of increasing projections of sea level rise. By focusing only on homes that sold multiple times I am able to control for all time-invariant property characteristics, including nearness to the beach as well as unobserved home quality that may be correlated with coastal proximity. Any changes to the characteristics of a property that occurred between sales will not be controlled for by the property fixed effects. For example, an owner might invest in household improvements between purchasing and selling a property, such improvements could affect the value of the house but would not be controlled for. I am able to observe significant renovation events in the data, and treat homes that underwent renovations as unique homes with a unique fixed effect. However, minor renovations are not recorded and may introduce bias to estimates. Billings (2015) argues that omitted variable bias from endogenous home improvements in repeat sales analysis is likely to be small.

While I can control for the level effect of coastal proximity through fixed effects, it may be the case that homes closer to the coast are appreciating at different rates for reasons unrelated to climate risk. For example, home buyer preferences for coastal proximity may vary through time. To deal with this potential source of bias I control for the time trend of coastal proximity's relationship to price, absorbing potential changes in coastal preference. I test for the effect of adding different sets of coastal proximity trends, including linear distance and vectors of dummy variables for various coastal proximity distances. I find that these controls are important to results and that the price premium for coastal proximity appears to be increasing over the study period. Controlling for the effect of coastal proximity means that I am effectively comparing homes that are at similar distances to the coast, but face different climate risk due to their elevation or location within flood zones. Results are robust to various alternative methods to control for coastal proximity time trends.

Equation 1 represents the main estimation equation. The dependent variable $\log(P_i)$ is the log of the sale price, where i indexes a property transaction. R is a dummy variable indicating whether a property is located in an area at risk to sea level rise. Y indicates the year of sale. Y is a continuous variable generated from the date of sale. For example, a sale occurring exactly halfway through 2008 takes a value of 2008.5. U is a property level fixed effect. M is a fixed effect for year-month of sale. W^d is a dummy variable that takes a value of one if the property's centroid is within d meters of the coast. C is the county where the property is located. In alternative specifications, I will replace the county time trends with other levels of geography.

$$\log(P_i) = \beta_0 + \beta_1(R_i \times Y_i) + \Phi U_i + \Psi M_i + \chi(W_i^d \times Y_i) + \kappa(C_i \times Y_i) + \varepsilon_i \qquad (1)$$

In all specifications I use two-way clustered standard errors. Spatially, clusters are at the zip code level. Where zip codes span both at risk ($R = 1$) and not at risk observations ($R = 0$), I split the cluster in two, with one cluster containing $R = 0$ observations and the other including $R = 1$ observations. Splitting the clusters in this way allows for the possibility of error correlation within local areas that share treatment status. Temporally, I cluster standard errors at the year-month level.

The use of property fixed effects absorbs the average price of a particular property across sales, allowing coefficients to capture price changes over time. The coefficient of interest (β_1) corresponds to the average difference in annual price appreciation between properties that are located in areas at risk from sea level rise relative to properties that are not at risk. I will test multiple definitions of "at risk." Identifying the time trend of coastal exposure (β_1) accounts for the time elapsed between sales so that a home with consecutive sales two years apart would experience twice the total appreciation effect on the log price as a property that had consecutive sales one year apart. In turn, I apply treatment definitions that include whether the elevation at the property's centroid is within two, three, or four meters of sea level and whether the centroid is within a FEMA defined flood zone. While land more than two meters above sea level is generally above the reach of sea level rise projections, elevation might vary across the property, meaning the property may contain areas directly exposed to sea level rise. Additionally, storm induced flood events are a threat to properties even several meters above sea level. While 6.8% of transactions occur among properties within two meters of sea level, only 0.9% occur among properties within one meter. I therefore use two meters as the lowest cutoff as there are not a sufficient number of properties at lower elevations to precisely identify an effect. The estimate of β_1 will account for the combined effect of rising exposure to flood events in the near future, as well as long term exposure to sea level rise that may affect the property through land erosion or inundation.

Local idiosyncrasies such as tides, coastal materials, and man made structures will influence the true level of exposure of an individual property. In the absence of data and strong priors regarding their effect, I use the elevation of the property's centroid as a strong proxy for true exposure.

In the main specification, $W_i^d \times Y_i$ will include a vector of 19 unique time trends. The vector includes an interaction between a dummy variable for being within 50 meters of the coast, and time. I include similar controls at 50 meter intervals ranging from 50 to 1,000 meters. This vector of time trends controls for the potentially non linear relationship between coastal proximity and price appreciation. Note that, a property within 50 m of the coast would take a positive value for all of these coastal time trends. I cannot directly identify which properties in the data have coastal frontage, however, the inclusion of the 50 meter trend should serve as a strong proxy for any unique trends among properties that have coastal frontage. Because the coastal proximity controls are highly correlated, they may suffer from multicollinearity issues. Identifying the unique effects of the distance to coast trends is not important to the identification of the main effect.

Homes with coastal exposure are more likely to have experienced damage from past storm events. Additionally, considering past research on storm damage and risk perceptions (McKenzie and Levendis, 2010; Ortega and Taspinar, 2018), areas that have experienced past damage may price future climate risk into current prices more aggressively. The presence of storm damage over the study period may influence the above estimates through two distinct processes. First, some properties may have sustained real damage to their structure which was not repaired before resale and therefore resulted in a diminished sales price. Owners may also have responded by investing in flood mitigation measures, which affect the home's price. Second, the storm events may have increased the awareness of potential buyers and sellers regarding local flood risk, and caused them to reduce their valuation of the asset.

I test for a differential price appreciation effect between properties located in areas that experienced different degrees of storm damage over the study period. The estimation strategy follows a triple difference model setup, captured in Equation 2. I add a term to the model that captures the interaction between a transacted property being at risk (R), a measure of storm damage insurance claims made in the property's census tract (F), and time (Y). The variable F is measured using FEMA insurance claim data. In separate regressions I will calculate F as either total storm damage occurring over the study period or cumulative storm damage occurring prior to the particular sale. Because F is highly skewed, I convert F from a dollar value into a binary variable that takes a value of one if the property is in a tract above the 95$^{\text{th}}$ percentile in terms of damage. I control for F directly, as storm damage may have a direct effect on prices.

When F is calculated as total storm damage across the study period, F is absorbed by the property fixed effects. I also control for the possibility that high-damage areas appreciated at a different rate in general, regardless of coastal exposure, with a local storm damage (F) by time (Y) interaction term. The form and notation of Equation 2 is otherwise consistent with Equation 1. A negative coefficient on the triple interaction term (γ_2) would indicate that properties in areas with greater local storm damage suffered a larger price penalty from sea level rise exposure.

$$\begin{aligned}\log(P_i) = \beta_0 + \gamma_1(R_i \times Y_i) + \gamma_2(R_i \times F_i \times Y_i) + \gamma_3(F_i \times Y_i) + \gamma_4 F_i +\\ \Phi U_i + \Psi M_i + \chi(W_i^d \times Y_i) + \kappa(C_i \times Y_i) + \varepsilon_i\end{aligned} \tag{2}$$

5 Results

Table 2 presents estimates of the effect of being exposed to coastal climate risk on property appreciation. The estimation strategy corresponds to Equation 1. The first definition of exposure is whether a home is within two meters of current sea level, shown in column 1. I find that, conditional on distance to the coast, properties within two meters of sea level experienced average annual price appreciation that was 1.4 percentage points below properties at higher elevations. I find a slightly reduced estimate of 1.2 percentage points when increasing the treatment definition to three meters (column 2), and a further reduced but still highly significant effect of 0.9 percentage points when the definition is increased to four meters (column 3). The declining estimate is consistent with a more severe effect for properties at higher risk. Column 4 considers a property at risk to sea level rise if it is located within a FEMA defined 100-year flood risk zone. I estimate a similar and significant effect of an annual reduction in price appreciation of 1.0 percentage point. The estimated effects of sea level rise exposure across Table 2 are all precisely estimated and highly significant.

Table 2: Effect of Sea Level Exposure on Price Appreciation

	(1)	(2)	(3)	(4)
At risk × Year	-0.014**	-0.012**	-0.009*	-0.010**
	(0.004)	(0.004)	(0.004)	(0.004)
Time fixed effects	Y	Y	Y	Y
Property fixed effects	Y	Y	Y	Y
Distance to coast time trends	Y	Y	Y	Y
County time trends	Y	Y	Y	Y
At risk definition	2 m	3 m	4 m	FEMA Zone
R^2	0.888	0.888	0.888	0.888
Adjusted R^2	0.797	0.797	0.797	0.797
N	164,026	164,026	164,026	164,026

Significance levels: $*$: 5% $**$: 1%. Two-way clustered standard errors in parenthesis. The coefficient estimates for At risk × Year correspond to the difference in annual price appreciation attributable to sea level rise exposure.

Compared to the typical price appreciation on Long Island over the study period, the estimated effects of sea level rise exposure are economically sizable. For example, a home purchased on Long Island for $500,000 in nominal dollars in 2000 appreciated, on average, to $730,000 by 2017. However, if the same home was within three meters of sea level it would have only appreciated to $628,000 by 2017, on average.

An important component of the identification strategy is to control for the possibility that there had been differential appreciation of properties close to the coast for reasons unrelated to climate change. For example, home buyer tastes may have changed over the study period. I address this in the main specification by including a series of controls that capture differential time trends for coastal properties (Equation 1). The distances from the coast I select to construct control variables are somewhat arbitrary. In Table 3, I evaluate the robustness of the main result to different sets of coastal proximity time trends. I repeat the Table 2, column 2 estimate, which used the three meter definition of exposure, in Table 3 but alter the coastal time trend controls.

Table 3: **Effect of Coastal Proximity Time Trend Controls on Results**

	(1)	(2)	(3)	(4)	(5)
At risk × Year	0.004	-0.007	-0.012**	-0.012**	-0.013**
	(0.003)	(0.004)	(0.004)	(0.004)	(0.004)
Distance to coast × Year		-0.027**			-0.019**
		(0.006)			(0.006)
Within 100 m of coast × Year				0.001	0.001
				(0.003)	(0.003)
Within 200 m of coast × Year			0.004	0.004	0.004
			(0.002)	(0.002)	(0.002)
Within 300 m of coast × Year				-0.000	-0.000
				(0.002)	(0.002)
Within 400 m of coast × Year			0.001	0.001	0.001
			(0.002)	(0.002)	(0.002)
Within 500 m of coast × Year				0.001	0.000
				(0.003)	(0.003)
Within 600 m of coast × Year			0.004	0.004	0.004
			(0.002)	(0.003)	(0.003)
Within 700 m of coast × Year				-0.000	-0.001
				(0.003)	(0.003)
Within 800 m of coast × Year			0.002	0.002	0.002
			(0.002)	(0.002)	(0.002)
Within 900 m of coast × Year				0.001	0.001
				(0.002)	(0.002)
Within 1000 m of coast × Year			0.012**	0.012**	0.004
			(0.003)	(0.003)	(0.002)
Time fixed effects	Y	Y	Y	Y	Y
Property fixed effects	Y	Y	Y	Y	Y
County time trends	Y	Y	Y	Y	Y
At risk definition	3 m	3 m	3 m	3 m	3 m
R^2	0.887	0.888	0.888	0.888	0.889
Adjusted R^2	0.794	0.797	0.797	0.797	0.798
N	164,026	164,026	164,026	164,026	164,026

Significance levels: $*$: 5% $**$: 1%. Two-way clustered standard errors in parenthesis. The coefficient estimates for At risk × Year correspond to the difference in annual price appreciation attributable to sea level rise exposure. The variable "Distance to coast" is in units of 10 kilometers.

In Table 3 column 1, I remove all coastal proximity time trends. In this specification, I estimate a small, insignificant positive effect of sea level exposure on price appreciation. The positive effect would suggest that the price appreciation trend of properties close to the coast outperformed properties further from the coast, and the exposure variable is proxying for this amenity value of coastal access. In columns 2-5 I add time trend controls that capture differential appreciation for properties with coastal proximity. Column 2 adds only a linear time trend of time interacted with distance to the coast, column 3 adds a unique time trends for properties within 200, 400, 600, 800 and 1,000 meters of the coast, column 4 increases the number of time trends by reducing the interval between controls to 100 meters and column 5 includes both the linear distance time trend and the 100 meter interval time trends. The main estimate of interest changes little between columns 3-5, suggesting the main model is able to control for changing market preferences for coastal proximity and isolate the partial effect of climate risk.

I display coefficients for all coastal proximity control variables in Table 3 for transparency, though these coefficient estimates may be unreliable due to multicollinearity. The coefficient on the linear time trend of distance to the coast (columns 2 and 5) is estimated as significant and negative, suggesting that properties close to the coast appreciated at a faster rate, conditional on sea level rise exposure.

In addition to coastal proximity time trends, I include time trends at the county level. In Table 4 I test for the sensitivity of results to time trends implemented at alternative geographic units. Column 1 shows the estimate when county time trends are removed, column 2 repeats the main specification, column 3 includes time trends unique to local school districts and column 4 includes zip code level time trends. There are 125 school districts and 172 zip codes with at least one repeated sale property. The main estimate is essentially unchanged in columns 1-3. With the inclusion of 172 unique time trend controls by zip code in column 4, the estimated effect of sea level exposure becomes statistically insignificant. The inclusion of so many time trends in column 4 is absorbing much of the identifiable statistical variation and attenuating estimates towards zero.

Table 4: Effect of Local Time Trend Controls on Results

	(1)	(2)	(3)	(4)
At risk × Year	-0.011**	-0.012**	-0.011*	-0.001
	(0.004)	(0.004)	(0.004)	(0.001)
Time fixed effects	Y	Y	Y	Y
Property fixed effects	Y	Y	Y	Y
Distance to coast time trends	Y	Y	Y	Y
At risk definition	3 m	3 m	3 m	3 m
Time trends	none	County	School Dist.	Zip Code
R^2	0.888	0.888	0.889	0.899
Adjusted R^2	0.797	0.797	0.799	0.816
N	164,026	164,026	164,026	164,026

Significance levels: $*$: 5% $**$: 1%. Two-way clustered standard errors in parenthesis. The coefficient estimates for At risk × Year correspond to the difference in annual price appreciation attributable to sea level rise exposure.

In the above analysis, I include all repeat sales occurring within Nassau and Suffolk counties. However, properties far from the coast may represent a different submarket and may therefore be relatively poor control observations. Coastal homes may share unique characteristics that are experiencing unique demand trends over the study period. While I tightly control for coastal proximity, limiting control observations to properties that are somewhat close to the coast may provide a cleaner estimate. As a robustness check, I repeat the analysis while limiting the observations to include only those properties within two kilometers of the coast. This strategy substantially reduces the number of transactions in the sample from 164,026 to 62,477. I provide a map of the areas of Long Island that are within two kilometers of the coast in Appendix A. If demand for properties close to the coast had a differential trend over the study period for uncontrolled for reasons this could potentially be a source of bias for the full sample specification, but this bias should be reduced in the coastal sample analysis.

Coastal sample results are provided in Table 5. I find results that are very similar to the main estimates when I limit the sample to properties within two kilometers of the coast. The result demonstrates that coefficient estimates are not driven by observations far from the coast.

Table 5: Effect of Sea Level Exposure on Price Appreciation, Coastal Sample

	(1)	(2)	(3)	(4)
At risk × Year	-0.013**	-0.011**	-0.008*	-0.007
	(0.004)	(0.004)	(0.004)	(0.004)
Time fixed effects	Y	Y	Y	Y
Property fixed effects	Y	Y	Y	Y
Distance to coast time trends	Y	Y	Y	Y
County time trends	Y	Y	Y	Y
At risk definition	2 m	3 m	4 m	FEMA Zone
R^2	0.900	0.900	0.899	0.899
Adjusted R^2	0.817	0.817	0.817	0.816
N	62,477	62,477	62,477	62,477

Significance levels: $*$: 5% $**$: 1%. Two-way clustered standard errors in parenthesis. The coefficient estimates for At risk × Year correspond to the difference in annual price appreciation attributable to sea level rise exposure.

In Table 5 estimates, I choose a cutoff of two kilometers to limit the sample. Figure 7 shows β_1 estimates (using the three meter exposure definition) but varies the distance to coast cutoff value used for filtering observations. I test alternative values ranging from one to 10 kilometers. I find the negative price effect is highly robust to truncating the sample at any of these distances.

I expect that the direct relationship between climate risk and elevation is only relevant for properties fairly close to sea level, as properties at relatively high altitudes face no risk from rising seas. However, while climate risk should clearly reduce demand for at risk properties, it may increase demand for properties that are good substitutes for at risk properties. Properties that have good coastal access, but face a low risk from sea level rise may experience an increase in demand. I test for the presence of demand spillovers by estimating unique price appreciation premiums for sets of properties in

Figure 7: Effect of Sample Selection on Main Estimate (β_1)

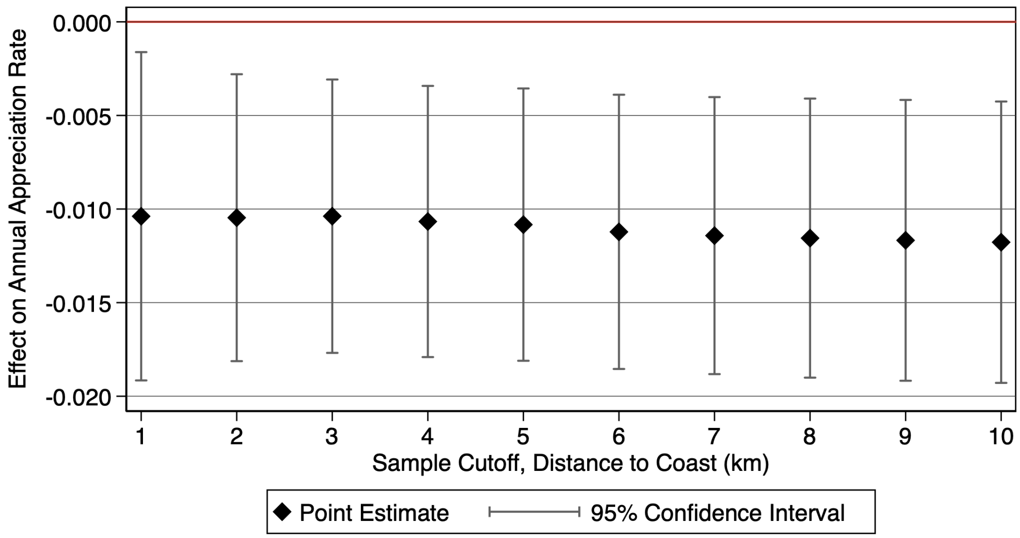

The point estimates correspond to the partial effect of sea level rise exposure on annual property appreciation (β_1). I use the three meter elevation definition of exposure. The sample cutoff is varied along the horizontal axis, with the sample size increasing from left to right. The figure displays 10 unique regression estimates. Estimates where the zero line falls outside of the confidence interval are statistically significant.

various elevation buckets. I limit this analysis to properties within two kilometers of the coast to focus on the coastal real estate market. Figure 8 shows the results of 10 unique estimates of β_1 (Equation 1) as the "at risk" definition is changed. The leftmost data point corresponds to the effect of being within two meters of sea level, which is equivalent to the Table 5, column 1 estimate. Moving rightwards, the subsequent point shows the estimated effect of being within 2-4 meters of sea level and the elevation buckets similarly increase by two thereafter.

Conditional on coastal proximity, climate risk seems to have significantly diminished the demand for low elevation properties, but increased the demand for higher elevation properties. Figure 8 shows the strong negative price appreciation effect of being within two meters of sea level. For estimates that span higher elevation buckets, I find some evidence of a price premium relative to the rest of the coastal market. For the bucket spanning 8-10 meters I find these properties had an annual price appreciation rate of 0.8% above other coastal properties. I interpret this pattern as representing demand substitution within the coastal real estate market, away from the riskiest locations and towards relatively less risky locations. Because the estimates are controlling for distance to the coast, the effect estimated is essentially comparing homes that are the same distance from the coast but face different risk exposure due to elevation.

Figure 8: Estimating Demand Spillovers in the Coastal Market

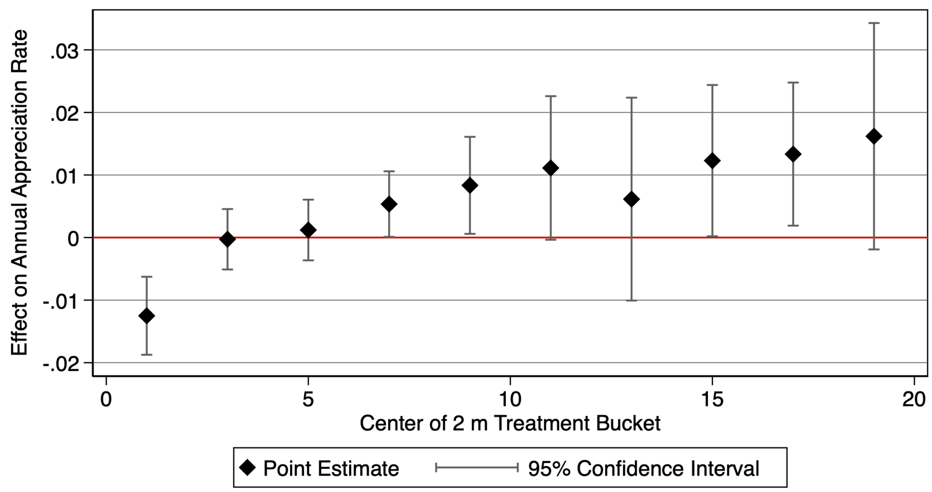

The point estimates correspond to the partial effect of a property being within a specific two meter elevation bucket. For example, the leftmost point estimate is the partial effect of being within two meters of sea level and the subsequent point estimate shows the effect of being within two to four meters of sea level. I report the β_1 estimates according to Equation 1. The figure displays 10 unique regression estimates. Estimates where the zero line falls outside of the confidence interval are statistically significant.

In addition to the robustness tests shown above, I provide several additional robustness tests in this paper's appendices. Limiting properties to only those that sold multiple times over the study period may cause the sample to be unrepresentative of the larger market if properties that sell frequently are different than those that rarely sell.

In Appendix B I provide results where I limit the sample to those properties that sold exactly twice over the study period, omitting those that sold more frequently. Filtering the data in this way reduces the sample size from 164,026 to 118,264 transactions. I find that estimates of sea level exposure are very similar when using this reduced sample, providing some evidence that results are not driven by properties that sell with an unusually high frequency. As described in Section 3, I also filter the data based on property type. In Appendix B, I provide results where I remove this sample restriction and find that results are robust.

The north and south coasts of Long Island are significantly different in their topography. Whereas the south coast has significant low-lying areas close to the shoreline, the north coast has a steeper shoreline with more cliffs (Figure 4A). The ways in which sea level rise may affect these shorelines may be significantly different. In Appendix C I show results for the north and south coasts separately. I find that the results of this paper are driven by the effect of sea level exposure on the south coast, which faces significantly higher storm risk.

The market may respond to both the immediate short term threat of a flood event, and the longer term risk of sea level overtaking the property. In Appendix D I present results for a model specification where the FEMA flood zone exposure definition is included simultaneously with the elevation based definition. I find both measures have negative coefficients. Being within a current flood zone represents potential short term flood risk, while being outside of a flood zone but near to sea level may capture longer term beliefs regarding risk caused by future sea level rise. The result provides some suggestive evidence that concerns over both types of risk are salient.

The expectations of future climate change induced property risk will be incorporated into present prices by forward-looking buyers. However, storm events that may have been induced by climate change have already affected Long Island during the study period. The most severe event being Hurricane Sandy in 2012. Past storm damage may affect home prices either through actual structural damage or from changing the expectations of buyers and sellers regarding the likelihood of future damage. I use equation 2 to test for a heterogeneous effect of coastal exposure in areas that have been subjected to high levels of past storm damage.

As an initial test, I rank all census tracts in terms of total FEMA insurance payouts made during the entire study period, and classify "high storm damage areas" as those above the 95$^{\text{th}}$ percentile. I present results of this specification in Table 6. I find that properties in high storm damage areas had stronger negative appreciation penalties from coastal exposure. For example, using the three meter exposure definition, exposure caused a statistically insignificant 0.9 percentage point annual reduction in appreciation among properties outside of the high storm damage areas, but properties with the same exposure but located in high damage areas experienced an average annual appreciation penalty of 2.6 percentage points.

Table 6: Triple Difference Estimates, Heterogeneous Effects by Local Storm Damage

	(1)	(2)	(3)	(4)
At risk × Year	-0.012*	-0.009	-0.006	-0.007
	(0.005)	(0.005)	(0.004)	(0.005)
At risk × Year × High storm damage area	0.002	-0.017*	-0.023**	0.008
	(0.007)	(0.007)	(0.007)	(0.006)
High storm damage area × Year	-0.010*	0.007	0.012*	-0.015**
	(0.004)	(0.005)	(0.006)	(0.004)
Time fixed effects	Y	Y	Y	Y
Property fixed effects	Y	Y	Y	Y
Distance to coast time trends	Y	Y	Y	Y
County time trends	Y	Y	Y	Y
At risk definition	2 m	3 m	4 m	FEMA Zone
R^2	0.888	0.888	0.888	0.888
Adjusted R^2	0.797	0.797	0.797	0.797
N	164,026	164,026	164,026	164,026

Significance levels: $*$: 5% $**$: 1%. Two-way clustered standard errors in parenthesis. "High storm damage area" corresponds to a tract above the 95$^{\text{th}}$ percentile in the total amount of storm damage insurance claims made across the study period.

In Table 7 I replace the time invariant definition of high damage area with a variable that captures the cumulative amount of FEMA claims that accrued within the tract up to the date of sale for an observation and create a dummy variable to represent whether the tract was above the 95$^{\text{th}}$ percentile in terms of damage accrued to that point in time. This approach captures time variation in storm damage. Because Hurricane Sandy represents almost all of the variation in damage, this specification can test for differential exposure penalties that occur because of Hurricane Sandy. I find no evidence that properties exposed to storm damage experienced an acceleration in value decline after experiencing the storm damage. However, fewer than a third of observations in the data occurred after Hurricane Sandy, so the model may simply be unable to precisely identify changing price trends over the post-Sandy period.

Table 7: Triple Difference Estimates, Heterogeneous Effects by Cumulative Local Storm Damage Occurring Prior to Sale

	(1)	(2)	(3)	(4)
At risk × Year	-0.014**	-0.011**	-0.008	-0.010*
	(0.004)	(0.004)	(0.004)	(0.004)
At risk × Year × High past storm damage	0.00003	0.00003	0.00004	0.00003
	(0.00002)	(0.00002)	(0.00002)	(0.00002)
High past storm damage × Year	0.003	0.002	0.000	0.002
	(0.004)	(0.004)	(0.004)	(0.004)
High past storm damage	-5.773	-4.940	-0.385	-3.270
	(7.443)	(7.912)	(8.017)	(7.514)
Time fixed effects	Y	Y	Y	Y
Property fixed effects	Y	Y	Y	Y
Distance to coast time trends	Y	Y	Y	Y
County time trends	Y	Y	Y	Y
At risk definition	2 m	3 m	4 m	FEMA Zone
R^2	0.888	0.888	0.888	0.888
Adjusted R^2	0.797	0.797	0.797	0.797
N	164,026	164,026	164,026	164,026

Significance levels: $*$: 5% $**$: 1%. Two-way clustered standard errors in parenthesis. "High past storm damage" corresponds to a tract above the 95$^{\text{th}}$ percentile in the cumulative amount of storm damage insurance claims made as calculated at the time of the property transaction.

Reconciling Tables 6 and 7 requires noting that the areas damaged by Hurricane Sandy were somewhat predictable, in that they were clustered on the high-risk southern coast of Long Island. Because the market already had information on this risk, the risk of possible future storms was already putting downward pressure on prices in high storm damage areas. Therefore, while high storm damage areas experienced the greatest decline in prices (Table 6), Hurricane Sandy appears to have provided limited new information to the market (Table 7). However, limitations on the time period covered by the data may be masking some of this informational effect. Expectations about the continued availability of insurance may have also caused the market to have little reaction to information on the likelihood of future storms introduced by Hurricane Sandy.

6 Conclusion

Sea level rise poses a significant and growing threat to coastal real estate. Over recent decades, climate research has increased predictions regarding the extent of future sea level rise and extreme coastal weather events. Current real estate prices will reflect not only the current utility of the asset but future monetary and utility flows. Because expectations of sea level rise became both more dire and less far into the future over the study period, properties exposed to sea level rise should sell at an increased discount if the market is populated by agents who are forward-looking, profit-maximizing and have full information on climate risk.

Implementing a repeat sales method on a complete set of housing transactions from Long Island, I find that properties exposed to the risks of sea level rise suffered a significant appreciation penalty over the 2000-2017 study period. Properties within three meters of current sea level were found to have annual price appreciation that was 1.2 percentage points per year lower than unexposed homes. I subject estimates to a battery of robustness checks to confirm the magnitude and significance of results. I also find evidence that some demand for high risk coastal properties has been diverted to lower risk coastal properties.

The declining value of high risk real estate suggests that the housing market is able to price, at least a portion, of the cost of climate risk. A gradual decline in the value of coastal real estate will help to buffer property owners from the costs of sudden climate event shocks, such as coastal floods and storm damage. However, some literature has suggested that the US housing market sets prices in ways that systematically reflect incomplete information on risk. Chivers and Flores (2002), Atreya and Ferreira (2015) as well as Hino and Burke (2020) provided empirical evidence that home buyers do not fully understand flood risk. Bakkensen and Barrage (2018) estimated that the US market for coastal homes in the 2007-2016 study period exceeded fundamental values by 10% because the market failed to fully account for flood risk. An initiative to increase information on flooding among home buyers in Finland was analyzed in Votsis and Perrels (2016), finding that the initiative was successful in closing the information gap and led to reduced prices among homes at risk of flooding. The possibility that US buyers are making real estate purchases with incomplete information about future climate risk suggests that the appreciation penalty I estimate is possibly less than what would arise in a market with perfect information. Government initiatives to ensure the availability and salience of climate risk information during the purchase process could be important in overcoming issues of incomplete information.

References

Atreya, A. and Czajkowski, J. (2019). Graduated flood risks and property prices in Galveston county. *Real Estate Economics*, 47(3):807–844.

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

Bakkensen, L. A. and Barrage, L. (2018). Flood risk belief heterogeneity and coastal home price dynamics: Going under water? Technical report, National Bureau of Economic Research.

Baldauf, M., Garlappi, L., and Yannelis, C. (2020). Does climate change affect real estate prices? Only if you believe in it. *The Review of Financial Studies*, 33(3):1256–1295.

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.

Billings, S. B. (2015). Hedonic amenity valuation and housing renovations. *Real Estate Economics*, 43(3):652–682.

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

Bunten, D. and Kahn, M. E. (2017). Optimal real estate capital durability and localized climate change disaster risk. *Journal of Housing Economics*, 36:1–7.

Chen, X., Zhang, X., Church, J. A., Watson, C. S., King, M. A., Monselesan, D., Legresy, B., and Harig, C. (2017). The increasing rate of global mean sea-level rise during 1993–2014. *Nature Climate Change*, 7(7):492–495.

Chivers, J. and Flores, N. E. (2002). Market failure in information: The national flood insurance program. *Land Economics*, 78(4):515–521.

Church, J. A., Clark, P. U., Cazenave, A., Gregory, J. M., Jevrejeva, S., Levermann, A., Merrifield, M. A., Milne, G. A., Nerem, R. S., Nunn, P. D., et al. (2013). Sea level change. Technical report, PM Cambridge University Press.

Conroy, S. J. and Milosch, J. L. (2011). An estimation of the coastal premium for residential housing prices in San Diego County. *The Journal of Real Estate Finance and Economics*, 42(2):211–228.

Desmet, K. and Rossi-Hansberg, E. (2015). On the spatial economic impact of global warming. *Journal of Urban Economics*, 88:16–37.

Eichholtz, P., Steiner, E., and Yönder, E. (2019). Where, when, and how do sophisticated investors respond to flood risk? *Available at SSRN 3206257*.

Gibson, M. and Mullins, J. (2020). Climate risk and beliefs in New York City floodplains. *Journal of the Association of Environmental and Resource Economists*.

Hino, M. and Burke, M. (2020). Does information about climate risk affect property values? Technical report, National Bureau of Economic Research.

Kahn, M. E. (2014). Climate change adaptation: Lessons from urban economics. Technical report, National Bureau of Economic Research.

Kahn, M. E. (2016). The climate change adaptation literature. *Review of Environmental Economics and Policy*, 10(1):166–178.

Kousky, C., Kunreuther, H., LaCour-Little, M., and Wachter, S. (2020). Flood risk and the us housing market. *Journal of Housing Research*, 29(sup1):S3–S24.

Kousky, C. and Shabman, L. (2014). Pricing flood insurance: How and why the nfip differs from a private insurance company. *Resources for the Future Discussion Paper*, pages 14–37.

Kriesel, W. and Landry, C. (2004). Participation in the national flood insurance program: An empirical analysis for coastal properties. *Journal of Risk and Insurance*, 71(3):405–420.

McCoy, S. J. and Zhao, X. (2018). A city under water: A geospatial analysis of storm damage, changing risk perceptions, and investment in residential housing. *Journal of the Association of Environmental and Resource Economists*, 5(2):301–330.

McKenzie, R. and Levendis, J. (2010). Flood hazards and urban housing markets: The effects of Katrina on New Orleans. *The Journal of Real Estate Finance and Economics*, 40(1):62–76.

McNamara, D. E. and Keeler, A. (2013). A coupled physical and economic model of the response of coastal real estate to climate risk. *Nature Climate Change*, 3(6):559–562.

Murfin, J. and Spiegel, M. (2020). Is the risk of sea level rise capitalized in residential real estate? *The Review of Financial Studies*, 33(3):1217–1255.

Nicholls, R. J. and Cazenave, A. (2010). Sea-level rise and its impact on coastal zones. *Science*, 328(5985):1517–1520.

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.

Seneviratne, S., Nicholls, N., Easterling, D., Goodess, C., Kanae, S., Kossin, J., Luo, Y., Marengo, J., McInnes, K., Rahimi, M., et al. (2012). Changes in climate extremes and their impacts on the natural physical environment.

Severen, C., Costello, C., and Deschenes, O. (2018). A forward-looking ricardian approach: Do land markets capitalize climate change forecasts? *Journal of Environmental Economics and Management*, 89:235–254.

Sklarz, M. and Miller, N. (2018). The impact of waterfronts on residential home values, part 2: Considering climate change and flood risks.

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

Yi, D. and Choi, H. (2019). Housing market response to new flood risk information and the impact on poor tenant. *The Journal of Real Estate Finance and Economics*, pages 1–25.

Zhang, L. and Leonard, T. (2019). Flood hazards impact on neighborhood house prices. *The Journal of Real Estate Finance and Economics*, 58(4):656–674.

Appendix A

In Table 5 I provide results where I limit the sample to properties within two kilometers of the coast. In Figure A1 I provide an elevation map where I indicate which areas of Long Island are within two kilometers of the coast. On the south coast, much of this area is at very low elevations, whereas the north coast contains steeper terrain that extends to higher elevations.

Figure A1: Elevation Map of the Study Area

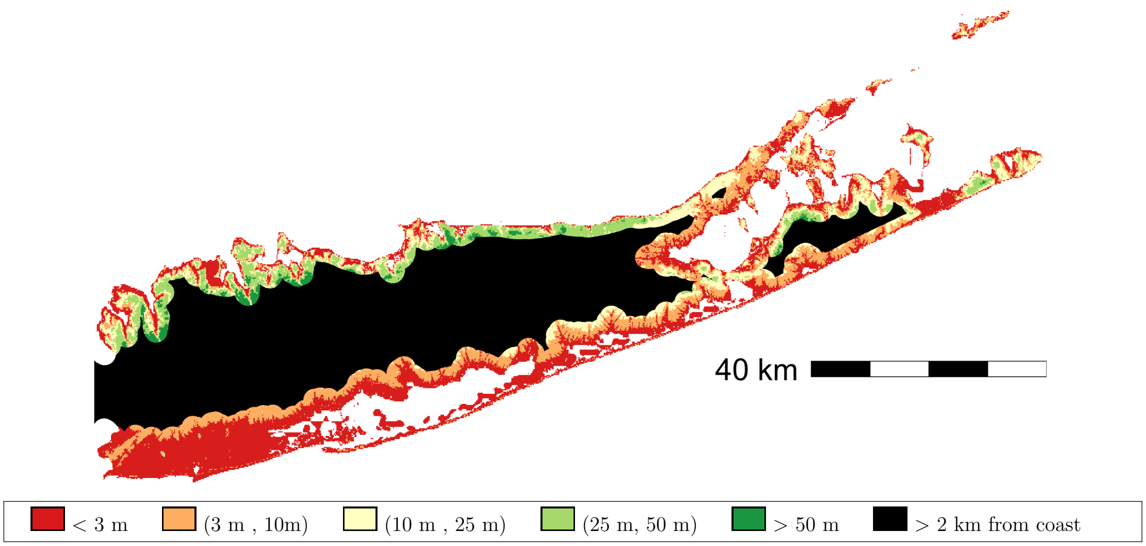

Areas of Nassau and Suffolk counties that are more than two kilometers from the coast are shown in black.

Appendix B

To generate the sample of properties used in analysis I drop a number of observations to focus the sample on observations that are most relevant to the research question and fit with the repeat sales methodology. In this Appendix I provide results using alternative samples to test the robustness of results. Overall, I find results are not generally sensitive to the specific decisions made regarding which observations are dropped.

I limit the sample to only properties that sold multiple times in order to include fixed effects at the property level. If properties that sell frequently are not representative of overall trends in the market the regression results will also not be representative of the overall market. In Table B1 I show results among properties that sold exactly twice, rather than at least two times. If the bias in the repeat sale sample is correlated with sale frequency, this subsample should be more representative of the overall market. I find results on the limited sample are robust and estimates are very close to the full repeat sale sample analysis (Table 2).

The analysis of this paper is focused on the housing market. I exclude other property types, particularly commercial and agriculturally zoned land. These non housing property types make up 4.4% of transactions. In Table B2, I repeat the main analysis but include these other property types. I find results are very similar to the main estimates of the paper.

Table B1: Effect of Sea Level Exposure on Price Appreciation, Only Properties that Sold Exactly Twice

	(1)	(2)	(3)	(4)
At risk × Year	-0.012**	-0.010**	-0.007*	-0.008*
	(0.004)	(0.003)	(0.003)	(0.003)
Time fixed effects	Y	Y	Y	Y
Property fixed effects	Y	Y	Y	Y
Distance to coast time trends	Y	Y	Y	Y
County time trends	Y	Y	Y	Y
At risk definition	2 m	3 m	4 m	FEMA Zone
R^2	0.904	0.904	0.903	0.903
Adjusted R^2	0.806	0.806	0.806	0.806
N	118,264	118,264	118,264	118,264

Significance levels: $*$: 5% $**$: 1%. Two-way clustered standard errors in parenthesis. The coefficient estimates for At risk × Year correspond to the difference in annual price appreciation attributable to sea level rise exposure.

Table B2: Effect of Sea Level Exposure on Price Appreciation, All Property Types

	(1)	(2)	(3)	(4)
At risk × Year	-0.014**	-0.011**	-0.008*	-0.009*
	(0.004)	(0.004)	(0.004)	(0.004)
Time fixed effects	Y	Y	Y	Y
Property fixed effects	Y	Y	Y	Y
Distance to coast time trends	Y	Y	Y	Y
County time trends	Y	Y	Y	Y
At risk definition	2 m	3 m	4 m	FEMA Zone
R^2	0.888	0.888	0.888	0.888
Adjusted R^2	0.797	0.797	0.797	0.797
N	171,663	171,663	171,663	171,663

Significance levels: $*$: 5% $**$: 1%. Two-way clustered standard errors in parenthesis. The coefficient estimates for At risk × Year correspond to the difference in annual price appreciation attributable to sea level rise exposure.

Appendix C

Table C1 provides split sample results where the effect of exposure is estimated separately for properties on the North and South sides of the study area. I split the sample area with a line that runs equidistant from the north and south coasts.

I find that the significant negative price effect of exposure is driven by properties on the south coast. In fact, I find a marginally significant positive price effect for exposure on the north coast (column 1). Because hurricane activity approaches from the south, the north coast is largely protected from storm surge events. The marginally positive price effect could be evidence of demand substitution towards coastal properties that are considered to be relatively low risk.

Table C1: Effect of Sea Level Exposure on Price Appreciation, North Coast vs South Coast

	(1)	(2)	(3)	(4)
At risk × Year	0.009*	-0.007	-0.018**	-0.012*
	(0.004)	(0.004)	(0.005)	(0.005)
Coast	North	North	South	South
Time fixed effects	Y	Y	Y	Y
Property fixed effects	Y	Y	Y	Y
Distance to coast time trends	Y	Y	Y	Y
County time trends	Y	Y	Y	Y
At risk definition	3 m	FEMA Zone	3 m	FEMA Zone
R^2	0.907	0.907	0.881	0.880
Adjusted R^2	0.829	0.829	0.785	0.784
N	51,011	51,011	113,013	113,013

Significance levels: $*$: 5% $**$: 1%. Two-way clustered standard errors in parenthesis. The coefficient estimates for At risk × Year correspond to the difference in annual price appreciation attributable to sea level rise exposure.

Appendix D

In Table D1 I provide a horse-race regression where I include the elevation based at risk definition simultaneously with the FEMA flood zone definition. Across the three regressions, I find that all six coefficients are negative, though not all are statistically significant. The result suggests that conditional on elevation, homes within FEMA defined flood zones appreciated less quickly. Also, the result suggests that conditional on being in a flood zone, lower elevation properties suffered a greater price penalty. The two risk indicators are highly correlated, potentially leading to multicollinearity issues. Therefore the results should be interpreted with caution.

Table D1: Effect of Sea Level Exposure on Price Appreciation, Impact of Elevation vs Flood Zone

	(1)	(2)	(3)
At risk (elevation) × Year	-0.012**	-0.010*	-0.006
	(0.004)	(0.004)	(0.004)
At risk (FEMA zone) × Year	-0.003	-0.003	-0.007*
	(0.003)	(0.003)	(0.003)
Time fixed effects	Y	Y	Y
Property fixed effects	Y	Y	Y
Distance to coast time trends	Y	Y	Y
County time trends	Y	Y	Y
At risk definition	2 m	3 m	4 m
R^2	0.888	0.888	0.888
Adjusted R^2	0.797	0.797	0.797
N	164,026	164,026	164,026

Significance levels: $*$: 5% $**$: 1%. Two-way clustered standard errors in parenthesis. The coefficient estimates for At risk × Year correspond to the difference in annual price appreciation attributable to sea level rise exposure.

Declarations

Funding: none

Conflicts of interest: none

Availability of data and material: upon request

Code availability: upon request