Who Holds Climate Risk? Evidence from Housing Prices and Ownership Sorting *

Paymon Khorrami† Jung Sakong‡

March 26, 2026

Abstract

When did climate risk begin to affect housing markets, how large were the valuation effects, and who ended up holding the risky properties? We answer these questions using property-level measures of expected climate losses, matched to housing transactions, ownership status, and mortgage characteristics. Our empirical design exploits within-neighborhood variation in climate exposure and a repeat-sales-style framework with property fixed effects, allowing us to distinguish climate-risk pricing from time-varying local demand. We find that climate risk began to be priced into housing markets as early as 2006, well before the timing emphasized in most of the existing literature. Quantitatively, a 1% increase in expected annual climate loss is associated with a permanent 2–4% decline in house prices. At the same time, riskier properties became less likely to be owner-occupied and more likely to be held by landlords or investors, while buyers of such properties increasingly used lower-leverage. Because our risk measure is expressed in dollar expected losses, the estimates admit a disciplined interpretation in dollar price terms and provide a framework for evaluating magnitudes. Taken together, the results suggest that climate risk affected not only valuations, but also the allocation of housing risk across households, much earlier and more strongly than previously understood.

JEL classification: G50, Q54, R21
Keywords: Climate risk, house prices, owner-occupancy, neighborhood dynamics

*We would like to thank Yuna Blajer de la Garza, Vimal Balasubramaniam (discussant), Marie-Therèse Radetzky (discussant), Victor Henry Osei (discussant), and conference participants at the 10th CEPR European workshop on household finance, the Conference on Sustainable Banking and Finance, and the 19th International Behavioural Finance Conference for helpful comments. Maria Castro, Victoria Healey, and especially Yilin Li provided excellent research assistance. These views are those of the authors and do not reflect those of the Federal Reserve Bank of Chicago or the Federal Reserve System. All errors are our own. The authors have no relevant financial or other potential conflicts of interest to disclose.

†Duke University (email: paymon.khorrami@duke.edu)
‡Federal Reserve Bank of Chicago (email: jung.sakong@chi.frb.org)

1 Introduction

Climate change can affect housing markets through both valuations and reallocation. If households begin to perceive some properties as more exposed to future climate damages, those properties may sell at a discount, and the set of agents willing to hold them may change as well. This raises several fundamental questions about how housing markets adapt to climate risk: When did climate risk begin to affect housing markets? How large are the resulting price effects? And who ultimately bears the risk associated with climate-exposed properties?

This paper answers these questions using a new property-level dataset on climate risk exposures matched to housing transactions, ownership status, and mortgage characteristics. The data measure expected climate losses in dollar terms at the property level, allowing us to study not only whether climate risk is priced, but also whether the estimated price effects are economically large relative to the losses being priced. Our empirical design exploits within-neighborhood variation in climate exposure and implements a repeat-sales-style framework with property fixed effects and geography-by-time fixed effects, which helps distinguish climate-risk pricing from changes in local housing demand.

We document four main findings. First, climate risk began to be priced into housing markets as early as 2006, substantially earlier than the timing emphasized in most of the existing literature. Second, the magnitudes are economically large: a 1% increase in expected annual climate loss is associated with a permanent 2–4% discount in house prices. Third, at the same time that this discount emerges, riskier properties become less likely to be owner-occupied and more likely to be held by landlords or investors. Fourth, the buyers of riskier properties increasingly use lower-leverage mortgages, suggesting that the incidence of climate risk shifts toward financially stronger balance sheets. Taken together, the evidence implies that climate risk affected not only valuations but also the allocation of housing risk across households.

A central advantage of our setting is that the risk measure is expressed in expected dollar losses. This allows us to compare changes in expected losses to changes in prices in economically meaningful units and therefore to discipline the magnitude of the estimates. Much of the existing literature identifies whether climate-related risks matter for housing valuations, often using geographic proxies for specific perils such as sea-level rise. By

contrast, our data allow us to study a broad set of climate perils at the property level and to frame the estimates in a way that is directly informative about economic magnitude.

A key empirical challenge is that climate risk is correlated with local amenities, neighborhood composition, and property characteristics that also affect housing demand and prices. Our empirical strategy addresses this challenge in two ways. First, because the climate data are measured at the property level, we can compare properties within highly localized geographies, including Census tracts and even Census blocks. Second, by implementing a repeat-sales-style specification with property fixed effects, we difference out time-invariant property attributes that may be correlated with climate exposure for non-climate reasons. We then add flexible time-varying controls for standard hedonic characteristics such as property size, number of bedrooms, and age, allowing us to distinguish climate-risk pricing from changes in demand for housing fundamentals.

Our results show that climate risk strongly permeates housing markets. In specifications that include property fixed effects and tract-by-time fixed effects, we find that climate-risky properties begin to be discounted significantly around 2006. Relative to pre-2006 prices, a property with a 1% higher expected annual climate loss exhibits a 2–4% lower transaction price by the end of the sample. At the same time, the ownership structure of climate-risky properties changes markedly: riskier properties become less likely to be owner-occupied and increasingly purchased by landlords or investors. These patterns coincide with changes in mortgage characteristics, with buyers of climate-risky properties increasingly using lower loan-to-value ratios and becoming less reliant on FHA or VA financing.

These patterns suggest that climate risk affected both prices and the composition of buyers. One interpretation is risk-sharing: climate risk is increasingly held by investors or landlords who may be better equipped to bear it. An alternative interpretation emphasizes preferences and financial constraints. If climate resilience is a normal good, higher-income households may move away from climate-risky properties, while financially constrained households may occupy them as renters rather than owners, implying a larger role for landlords in these markets. These mechanisms are not mutually exclusive, and both imply that climate risk may influence neighborhood composition and ownership structures in addition to property valuations.

Much of the existing literature and public discussion emphasizes climate risk pricing

emerging in the mid-2010s or later. For example, Engle et al. (2020) find, using newspaper articles, that 2009 was the year with the largest quantity of climate-relevant news, followed by the period from 2014-16. Using Yale opinion surveys since 2008, Marlon et al. (2022) finds that much of the change in opinions regarding climate change has been more recent (e.g., 2016-present). On the firm and institutional investor side, 2016 is a commonly-cited inflection point for the rise of ESG investing and changes to firms' cost of capital (Ilhan et al. 2023; Gormsen et al. 2024). And the literature on house prices and sea-level-rise (SLR) risks only finds effects, if at all, starting in the mid-2010s (Keenan et al. 2018; Bernstein et al. 2019; Baldauf et al. 2020; Murfin and Spiegel 2020; Keys and Mulder 2020; Bernstein et al. 2022). The evidence that climate risk is priced in the stock market becomes robust around 2012 (Faccini et al. 2023). Based on this literature, our estimates of an early timing could be mismeasurement: the 2006 timing could simply reflect an increasing demand for some unobservable characteristics that are correlated with our climate risk measure.

Another possibility is that climate awareness among households began rising earlier than suggested by existing empirical studies. Several prominent events occurred during the mid-2000s that may have increased attention to climate risks. The Kyoto Protocol became effective in February 2005; although the US was not part of it, this constituted international news. Hurricane Katrina, one of the most severe and costly hurricanes of all time, struck in August 2005. Katrina was a high-profile event in the US that may have shifted overall climate views (see literature cited below). In 2006, the popular documentary "An Inconvenient Truth" was released, which may have played a role in environmental awareness (Nolan 2010). Beyond these specific events, Brulle et al. (2012) and Carmichael and Brulle (2017) conduct a more systematic analysis of public opinion data for the years 2002-2010 and 2001-2013, respectively, identifying several factors driving shifts in climate perceptions. These papers find that reported climate concerns starting rising in 2005 and peaked in 2007, broadly consistent with our results.[1]

Related literature. We primarily contribute to a large existing literature relating climate change risks to housing markets. Most closely related is a literature investigating sea-

[1]Sparse evidence is available on climate-related views for years in the early 2000s and prior, partly due to data availability. This is not to say that people were unaware of climate risks in these early years. In one study, Krosnick et al. (2000) finds that the 1997 efforts by Clinton administration to build support for the Kyoto Protocol sparked a change in climate-related perceptions, especially among Democrats.

level-rise (SLR) risk (Keenan et al. 2018; Bernstein et al. 2019; Baldauf et al. 2020; Murfin and Spiegel 2020; Keys and Mulder 2020; Bernstein et al. 2022). These papers construct property-level variation in SLR risk, due to elevation, coastal distance, and other geological factors, controlling for other house characteristics, to identify the impact of flood risk on housing markets.[2] There is some disagreement on the existence and magnitude of an effect on housing, with Murfin and Spiegel (2020) notably finding zero effect of SLR risk on prices through 2017. To the extent the other papers find an effect, they all agree that it appeared later in the sample (e.g., 2013 or even later).

Our study differs in several respects. Most importantly, our methodology differs, as discussed in the introduction. Second, we go beyond flood risk and examine all climate-related risks (e.g., wildfires, tornadoes, earthquakes, and more), allowing us to study a larger geographic sample of housing markets beyond the coastal areas. Third, our risk measure is conveniently reported in percentage units, allowing us to quantify an economically-meaningful elasticity. Finally, our results differ. Not only do we find a large impact of climate risk on house prices, we also shed light on the timing of the effect. And with the exception of Bernstein et al. (2022), who study partisan sorting, we also examine transactions and residential sorting.

Another literature has measured the impact of disasters on local house prices (McKenzie and Levendis 2010; Ortega and Taspınar 2018; Pollack and Kaufmann 2022; Ellen and Meltzer 2024).[3] These results are often highly localized, for instance the impact of Florida hurricanes on Florida house prices, and struggle to separate the direct effect of the disaster from the effect of future climate risks.[4] A notable exception is Addoum et al. (2024), which demonstrates a persistent decline in Boston house prices following New York's Hurricane Sandy.

A commonly-held view is that climate risk perceptions shift after a climate disaster. For instance, a survey-based literature shows that personal experiences modulate climate risk perceptions (Egan and Mullin 2012; Akerlof et al. 2013; Deryugina 2013; Shao 2017; Bergquist

[2]Note that SLR risk may affect household choices in other dimensions, like their stock market investments, as well (Ilhan 2020).

[3]For surveys on the housing market effects of climate change perceptions, see Schuetz (2024) and Contat et al. (2024).

[4]Some research also examines the ex-post effect of a disaster like Hurricane Katrina on households directly harmed, e.g., Gallagher and Hartley (2017) and Deryugina et al. (2018). But again, this differs from our goals. We wish to obtain the impact of *climate risk*, which requires looking at households not directly harmed and identifying effects ex-ante.

and Warshaw 2019). Relatedly, but in observational data, disaster insurance premia may increase following a disaster realization, and these higher premia can be impounded in house prices (Nyce et al. 2015; Keys and Mulder 2024). Given the coincidence of our effects and the arrival of Hurricane Katrina, such a belief-updating view may be at work. Similar to Addoum et al. (2024), we show our results are not mechanically due to Katrina in a localized sense—in fact, excluding states directly impacted only modestly changes the magnitude of our estimates. But Katrina was a high-profile event, which may have been a broader signal to consumers' attention to climate damages, inducing a belief update.

2 Data

We combine two data sources both from CoreLogic: one on climate risks at the property level (single snapshot) and another on housing-market transactions and property characteristics. We take only the properties that have a unique property identifier: CoreLogic Integrated Property (CLIP), which is used to merge the two data.

2.1 Climate Data

Data on climate risks come from CoreLogic's Climate Risk Data, which are a one-time snapshot of structure-level measures of exposures to various climate risks, computed using CoreLogic's proprietary algorithms.

The actual data are at a structure level. This can be more disaggregated than properties for properties that contain multiple structures, for example with a main home and a stand-alone garage buildings. CoreLogic computes an importance ranking, which we use to only take the climate risk scores for the primary structure, e.g., only the climate risk for the main home in our example. The resulting data we use is therefore unique at the property level.

The main measure of climate risk exposure we use is the Average Annual Loss (AAL), which is the fraction of the property value that is expected to be lost due to climate events from CoreLogic's simulations in a given year.

The main composite risk is the union of nine individual climate events, or "perils." These

include three separate earthquake-related perils (fire, shake and sprinkler leakage, and tsunami), wildfire, inland flood, severe convective storm, winter storm, hurricane storm surge, and hurricane wind.[5]

CoreLogic provides AAL for the current condition, as well as projections for future years of 2030, 2040 and 2050. For each of the future years, they provide projections under varying climate scenarios following the Representative Concentration Pathways (RCP). RCPs are hypothetical scenarios of varying intensity of climate change described by the the expected changes in radiative forcing values by year 2100. CoreLogic provides projected AAL under three of the four original pathways: Under RCP 2.6, CO_2 concentration eventually falls to and stays at around 400 parts per million (ppm); RCP 4.5 is the intermediate scenario; and under RCP 8.5, emissions continue to rise.

Importantly, AAL captures not only the climate change risk but the exposure of a particular property to such climate change. For example, flood risk can increase for an entire neighborhood due to rising temperatures, but the damage to a property can be higher if that property has a lower first floor height or if it has a basement. For another example, wildfire risk can increase for a larger geographic area, but the risk can accompany a higher expected loss for a property that is built from more combustible building material.

2.2 Housing Market Data

Data on housing-market prices and characteristics come from CoreLogic's Real Estate Data, and in particular their Property Basic data based on assessment records and Owner Transfer data based on transaction deed records.

We take all arms-length, resale transactions available in the Owner Transfer data for 2000-2024, with property ID variable (CLIP), transaction price, transaction date, Census block in which the property is located, and an indicator for whether the purchase was for owner occupation. Using the CoreLogic ID variable CLIP, we obtain property characteristics including square footage from the Property Basic and Historical Property Basic data, also from CoreLogic. There are 74 million unique properties that are resold within 2000 and 2024.

[5] In what follows, we drop the earthquake perils related to leakage and tsunami, due to very sparse coverage.

TABLE 1: Summary Statistics

AAL	count	mean	p1	p5	p10	p25	p50	p75	p90	p95	p99
Wildfire	61,712,911	0.0245	0.0000	0.0000	0.0000	0.0000	0.0002	0.0041	0.0375	0.1122	0.5025
Hurricane Storm Surge	69,055,194	0.0578	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0079	0.1500	1.6579
Hurricane Wind	69,055,194	0.1058	0.0000	0.0004	0.0013	0.0055	0.0223	0.1105	0.2842	0.5101	1.0648
Earthquake Fire Following	141,337,976	0.0004	0.0000	0.0000	0.0000	0.0000	0.0000	0.0001	0.0006	0.0018	0.0070
Inland Flood	141,337,976	0.0497	0.0000	0.0000	0.0000	0.0000	0.0000	0.0074	0.0478	0.1357	0.9241
Severe Convective Storm	141,337,976	0.0725	0.0001	0.0008	0.0017	0.0157	0.0523	0.1001	0.1544	0.2406	0.3401
Winter Storm	141,337,976	0.0166	0.0000	0.0000	0.0000	0.0011	0.0134	0.0241	0.0378	0.0501	0.0898
Composite	141,337,976	0.2522	0.0165	0.0400	0.0620	0.0936	0.1404	0.2318	0.3990	0.6840	2.5231

Notes. Summary statistics for the Average Annual Loss (AAL) by peril, with the composite at the end. The AAL is for the year 2050 under a moderate scenario (RCP 4.5), multiplied by 100 (i.e., units are percentage losses per year).

2.3 Summary Statistics

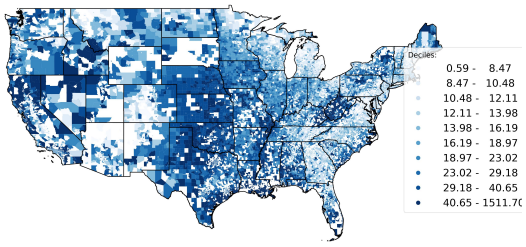
Table 1 shows the cross-sectional distribution of AALs for various perils, as well as a "Composite AAL" that sums across perils. Based on the cross-sectional means of AAL, the largest perils are severe convective storm, inland flood, hurricane wind, and hurricane storm surge (the two hurricane peril AALs are missing for approximately half the properties, which we believe corresponds to a negligible risk for those properties). The smallest risk by far is earthquake (fire following), followed by winter storm and wildfire. The mean Composite AAL is 0.25% per year. For all the perils in Table 1, AAL is significantly right-skewed, with a significant fraction of very low-loss properties and small fraction of much higher-loss properties. For instance, the median inland flood AAL is 0% per year, but its 99th percentile is 0.92% per year. A similar right-skewed pattern holds for the Composite AAL.

The perils are geographically distributed as one might expect. Figure 1 displays maps for the geographic distribution of AALs, by peril and for the composite version. For instance, wildfire risk is concentrated in the west of the US; severe convective storms—which involve winds, tornadoes, and hail—primarily hit the mid-west; winter storms mainly in the north and northeast; hurricane-related damages are on the eastern and southeastern coasts. Interestingly, the Composite AAL (panel A) has a very even distribution throughout the US.

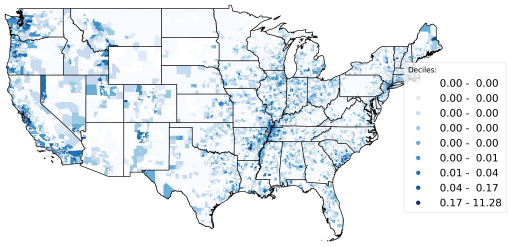
To quantify the determinants of AAL variation, Table 2a shows the R^2 of AAL in various fixed effects specifications. About 60% of the variation in AALs is driven by Census blocks and peril types. As mentioned, the residual variation within small local geographies (e.g., Census tract or block) is what we hope to exploit. Note also that the addition of scenario

FIGURE 1: Tract-level Map of Peril AAL

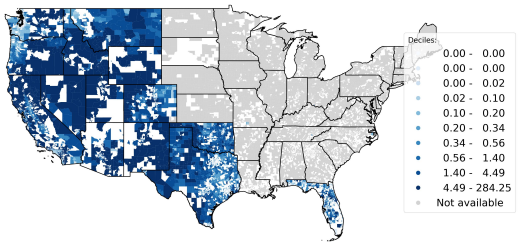
(a) Composite All Perils



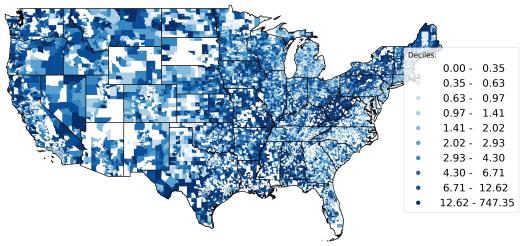
(b) Earthquake (Fire Following)



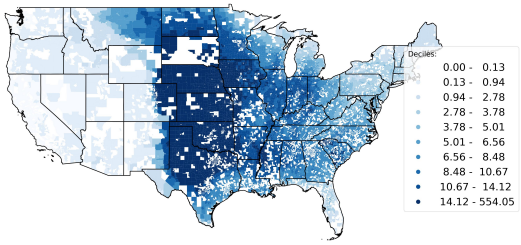
(c) Wildfire (Fire + Smoke)



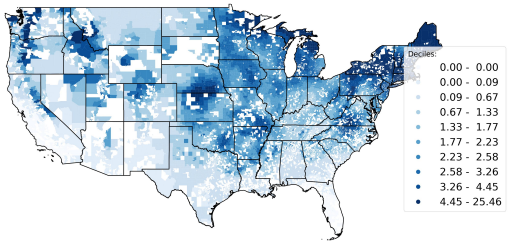
(d) Inland Flood



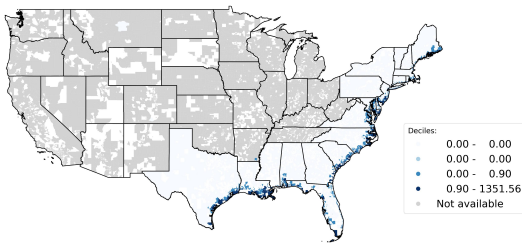
(e) Severe Convective Storm



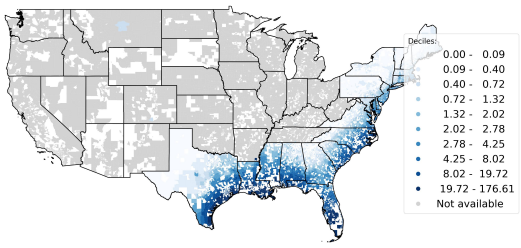
(f) Winter Storm



(g) Hurricane Storm Surge



(h) Hurricane Wind



Notes. Each map shows the Census tract average AAL in basis points, i.e. AAL × 10,000. Tracts are divided into deciles unless we cannot form distinct deciles because more than 10% of the tracts have 0 or N/A.

and timeframe fixed effects interacted with block and peril only increases the explanatory power by 1% or so. This reveals that, roughly speaking, the CoreLogic climate model scales a baseline property-level AAL by some common factor to obtain AALs at different scenarios and timeframes. For this reason, most of our results simply fix the scenario at its intermediate value (the RCP 4.5 scenario) and the timeframe at its long-horizon value (the 2050 timeframe). By contrast, significant variation is captured by heterogeneity across perils: the R^2 rises from about 11% to about 60% when including an interaction of peril with block, versus using block fixed effects alone. Comparing Composite AAL to transaction prices, Table 2b shows that prices are much better explained by geography (tract or block fixed effects) than AAL.

Table 2b also suggests that, for the question of explaining house prices with climate risk, using tract-level controls may be preferable to block-level controls. Indeed, when refining fixed effects from tract to block, the explained variation in AAL increases much moreso than for transaction prices. We ideally want an empirical design that preserves as much variation in AAL while simultaneously purging as much non-climate-related variation in housing demand. Based on these R^2 results, employing tract fixed effects seems to better manages this trade-off between preserving AAL variation and purging demand variation, justifying tract-level controls as our baseline. That being said, we will examine block-level fixed effects as a robustness exercise.

Going beyond their means, Table 2c shows the variance-covariance matrix of AALs across properties for the 8 perils for which there are future projections. (In this case, to compute a full covariance matrix, we set any missing values of peril-specific AALs to zero.) Using variances as a proxy for importance, by far the two largest perils are inland flood and hurricane storm surge, followed by wildfire, convective storm, and hurricane wind. Relative to the variances of these perils, the others pale in comparison. Additionally, the covariances between perils tend to be very small, with the exception of the two hurricane-related perils.

What drives the peril-specific risks beyond geography? In the Online Appendix, Tables A.1-A.7 display results from regressing peril-specific AAL (as well as other moments of loss risk beyond the average) onto block fixed effects and property characteristics, such as first-floor height, ground elevation, and dummies for the inclusion of various materials in the property structure (e.g., steel, timber, concrete, etc.). Judging by the R^2, Census blocks

TABLE 2: Climate Risk Exposure (AAL) Variation

(a) By Block, Peril, Scenario, Time Frame

Fixed Effects	R^2
Block	0.1312
Block X Scenario	0.1315
Block X Timeframe	0.1319
Block x Peril	0.6760
Block x Scenario X Timeframe X Peril	0.5774

(b) R^2 for Block and Tract FEs

Dependent variable	Fixed Effect	R^2
Composite AAL	Tract	0.370
Composite AAL	Block	0.613
Log Sale Amount	Tract X Year	0.840
Log Sale Amount	Block X Year	0.912

(c) Variance-Covariance by Peril ($\times 10^{-9}$)

	Earthquake Fire	Earthquake Tsunami	Wildfire	Inland Flood	Convective Storm	Winter Storm	Hurricane Storm Surge	Hurricane Wind
Earthquake Fire	0.42							
Earthquake Tsunami	0.02	0.72						
Wildfire	-0.30	-0.09	2410.0					
Inland Flood	-1.22	0.31	-14.40	27,900.0				
Convective Storm	-1.34	-0.54	-75.50	66.20	2,370.0			
Winter Storm	-0.12	-0.12	-18.20	-3.31	-0.70	37.30		
Hurricane Storm Surge	-0.04	-0.17	-17.90	73.20	-90.60	-21.80	6,610.0	
Hurricane Wind	-0.79	-0.33	38.20	-6.28	-120.00	-51.30	1,230.00	1,980.00

Notes. Panel (a) presents unadjusted R^2 from various fixed-effects regressions of AAL at the property × peril × scenario × timeframe level. Samples 1 and 2 are 5% random samples of properties in the CoreLogic Climate Risk data. Panel (b) presents R^2 from fixed effect specifications of AAL (2050 timeframe, intermediate scenario) and log sale price onto various geographic fixed effects. Panel (c) presents the variance-covariance matrix for property-level AAL for the 8 perils, for the 2050 timeframe under intermediate scenario (RCP 4.5). The actual covariances can be obtained by multiplying the numbers in panel (b) by 10^{-9}. Variances along the diagonal are in bold. For perils with limited geographic coverage (Earthquake: Tsunami, Wildfire, Hurricane Storm Surge, and Hurricane Wind), missing AAL has been assigned a value of 0.

explain the majority, and often much more, of the variation in climate-related losses. But important variation remains within Census blocks, and this variation appears to co-move in sensible ways with structure characteristics. For instance, Table A.2 shows that, controlling for block fixed effects, flood-related AAL is negatively related to first-floor height, ground elevation, and the use of "solid" materials like steel, concrete, masonry, and timber.

These summary statistics support the following interpretation of the data. Each property is primarily exposed to a single climate peril, if any, with floods, hurricanes, and convective storms being the most important economically. The climate risks to houses in a local region are largely similar, but the property-level exposures can differ significantly, due to variation in construction materials, property design, and upkeep. Ultimately, we seek to isolate these sources of variation from the localized climate-related risks.

Finally, we explore the correlations between standard hedonic variables and Composite AAL in Table 3.[6] The results show that high-AAL properties tend to be smaller (in terms of both square-footage and number of bedrooms) and older. As an example, Table 3 shows that, within Census block, a property that is 10% larger in square-footage has on average a 0.00278% lower Composite AAL—meaning that modest differences in AAL can potentially be associated to large price differences, through significant differences in property size. Given that larger properties and newer constructions tend to be more valuable on average, our "repeat-sales" style procedure (comparing a house to its own previous transaction) can help to purge these time-invariant valuation issues and identify climate discounts. Furthermore, to correct for potentially "time-varying demand" related to these basic hedonics, but having nothing to do with climate risks, we also control for them directly in some specifications below.

[6]Table A.8 in the Online Appendix also displays the raw pairwise correlations between all these variables.

TABLE 3: Property Characteristics AAL Composite

	AAL 2050 Composite RCP 4.5 ×10,000			
Log Sq Ft	-3.030			-2.779
	(0.099)			(0.111)
Age		0.034		0.028
		(0.001)		(0.001)
Bedrooms=1			9.758	7.680
			(0.432)	(0.414)
Bedrooms=2			3.841	2.832
			(0.103)	(0.093)
Bedrooms=4			-0.829	0.049
			(0.034)	(0.039)
Bedrooms=5			-1.252	0.124
			(0.075)	(0.077)
Bedrooms=6			-0.931	0.835
			(0.151)	(0.154)
N	1.05e+08	8.82e+07	7.59e+07	6.47e+07
R^2	0.469	0.483	0.520	0.533
R^2 within	0.002	0.001	0.004	0.005

Notes. The dependent variable is Composite AAL for the 2050 timeframe under intermediate scenario (RCP 4.5), in units of basis points, i.e. AAL × 10,000. Standard errors, in parentheses, are computed by clustering at the Census block level.

3 Empirical Framework

CoreLogic's property-level climate risk exposure measure permits us to explore the consequences of granular variation in climate riskiness between properties. Broadly speaking, we are interested in how various property-transaction-level outcomes change over time due either to evolving climate riskiness or climate-change awareness.

Our overarching empirical framework is as follows. We index properties by i and transaction dates by t. Let y_{it} denote a generic property-transaction-level outcome—we investigate transaction price, owner-occupancy, the buyer's loan-to-value ratio, and some non-standard mortgage characteristics (e.g., ARMs, FHA loans, VA loans). Let A_i denote the property-level climate risk proxy—unless mentioned otherwise, our results employ the Composite AAL variable (which adds losses across all perils) projected for the 2050 horizon under the intermediate scenario (RCP 4.5). We are interested in the time-varying sensitivity

of y_{it} to A_i, controlling for various important factors:

$$y_{it} = \sum_s \phi_s \mathbf{1}_{\{t=s\}} \times A_i + \text{controls}_{it} + \nu_{it}. \tag{1}$$

How we implement the controls, to be discussed shortly, is particularly critical to the interpretations one can draw. We are interested in time-series pattern in the AAL coefficients $\{\phi_t\}$.

How should one interpret the pattern in the $\{\phi_t\}$ coefficients? Our climate risk measure A_i is a property-level measure estimated from a single cross-sectional snapshot in 2020. This is not to say that climate risk is fixed over time at the property-level, and our results do not rely on this being true or false. In general, one can attach two possible interpretations to a time-series pattern in $\{\phi_t\}$: this can reflect either (i) changes in the *quantity of climate risk* and/or (ii) changes in *prices of risk*.

Given our data, risk-quantity and risk-price explanations look observationally equivalent. On the one hand, it is possible that climate risk itself is changing at the property level, whereas our data only see it measured in 2020. To the extent that changes in actual climate risk through our sample (2000-2025) are correlated with estimated levels as of 2020, the pattern in $\{\phi_t\}$ could reflect these underlying changes in risk quantity. For instance, a natural conjecture is that the highest-risk properties experienced the largest increases in risk over the sample. On the other hand, it is possible that true climate risk was relatively stable over this period, but awareness or fear of these risks has changed over time. In that case, the pattern in $\{\phi_t\}$ reflects changes in the required compensation for bearing climate risks, due to the change in awareness or fear. For instance, a natural guess is that climate risk concerns have risen broadly over time, which would then affect the highest-risk properties the most.

Let us now turn to some details of our implementation. We first estimate (1) with controls for geography-by-year fixed effects and hedonic property characteristics (such as square-footage, number of bedrooms, or age). In other words, we estimate

$$y_{it} = \alpha_{gt} + \sum_s \beta_s \mathbf{1}_{\{t=s\}} \times X_i + \sum_s \phi_s \mathbf{1}_{\{t=s\}} \times A_i + \nu_{it} \tag{2}$$

where α_{gt} is a geography-by-time fixed effect and X_i are hedonic controls (such as square-footage, number of bedrooms, or age). This specification effectively compares properties

within a given geography g and within hedonic categories captured by X, and therefore we refer to it as a "cross-sectional" approach. Despite the terminology "hedonics," the reader should note that our inclusion of characteristics X_i in (2) differs slightly from the standard approach, because we interact it with time dummies $\mathbb{1}_{\{t=s\}}$, which allows flexibly for time-varying impacts of the same characteristic (e.g., if y_{it} were house prices and X_i were house size, the β_s coefficients would measure the time-varying valuation of size).

One difficulty in interpreting the year-to-year variation in $\{\phi_t\}$ from (2) is that the set of properties being transacted can be different in each year, as opposed to the climate-risk discount rising over time. The problem is that, by construction, certain value-relevant characteristics are driving our AAL variation (e.g., the presence of basements, carpet versus hardwood floors, and other such characteristics), and the prevalence of these characteristics in housing transactions can be changing over time for reasons unrelated to climate perceptions. Suppose, for example, that houses with basements have higher AAL but are also generally less valued in the market, in a time-invariant way, simply because basements are undesirable. If the propensity of homes with basements to transact is time-invariant, then there would be a constant downward bias in all of the $\{\phi_t\}$ coefficients measured from (2). This problem is minor, in the sense that the time-series pattern of the $\{\phi_t\}$ remains preserved. However, if the propensity of homes with basements to transact is rising over time, then even the time-series pattern will be affected: the (negative) bias in the estimates of $\{\phi_t\}$ would be rising over time, a pattern which mimics that of a growing climate risk discount. We attempt to correct for this issue below.

What cannot be definitively addressed in the specifications that follow is the time-varying bias that arises if the *demand* for basements is time-varying. However, we do try to address time-varying demand, to some extent, via some important and standard hedonic controls (the X_i variables), similarly to the benchmark cross-sectional approach in (2). Crucially, by interacting X_i with time dummies, we can purge time-varying demand related to those specific hedonic factors.

To address some of these concerns about time-varying compositional effects and some types of time-varying demand, our core method implements a "repeat-sales" style approach,

augmented by flexible time-varying loadings on hedonic characteristics. We estimate:

$$y_{it} = \gamma_i + \alpha_{gt} + \sum_s \beta_s \mathbf{1}_{\{t=s\}} \times X_i + \sum_s \phi_s \mathbf{1}_{\{t=s\}} \times A_i + \nu_{it}, \tag{3}$$

where γ_i is a property fixed effect. This specification, similar to a difference-in-differences, effectively compares the change over time in the difference between two properties within a geography and hedonic categories.[7]

We refer to specification (3) as a "repeat-sales" analysis. To understand this, compare it to the simpler regression

$$y_{it} = \gamma_i + \alpha_t + \nu_{it}. \tag{4}$$

If y_{it} is property i's transaction price in year t, then the estimate of α_t in (4) delivers a repeat-sales house price index. By including a property fixed effect in (4), we are effectively considering how housing valuations change over time within a given property. Analogously, by putting the property fixed effect in (3), we are effectively considering how the impact of climate risk exposures changes over time within a given property. Even though a given property may not transact very often, by putting the linear structure on how properties are impacted by AAL, a large enough dataset allows us to estimate ϕ_s reliably using this within-property variation. Akin to the standard logic in favor of repeat-sales indexes, our ability to look within property is helpful to control for unobserved cross-sectional quality differences that are correlated with climate risk exposures. The connection between the two-way fixed-effects framework in (3)—property and year fixed effects—and the traditional repeat-sales methodology is noted for example in Gao and Wang (2007).

[7] As an alternative for robustness, we also implement the analysis in two stages. Consider first residualizing Composite AAL by geography fixed effects:

$$A_i = \alpha_g + \varepsilon_i$$

and then estimating

$$y_{it} = \gamma_i + \alpha_t + \sum_s \beta_s \mathbf{1}_{\{t=s\}} \times X_i + \sum_s \phi_s \mathbf{1}_{\{t=s\}} \times \varepsilon_i + \nu_{it}$$

where γ_i is a property fixed effect, α_t is a time fixed effect, and ε_i are the residuals from the first-stage regression. The estimates of $\{\phi_t\}$ from (3) and the two-stage approach described above should mechanically be identical in a balanced panel; but if some properties transact more than others, the results could potentially differ.

4 Empirical Results

4.1 House prices

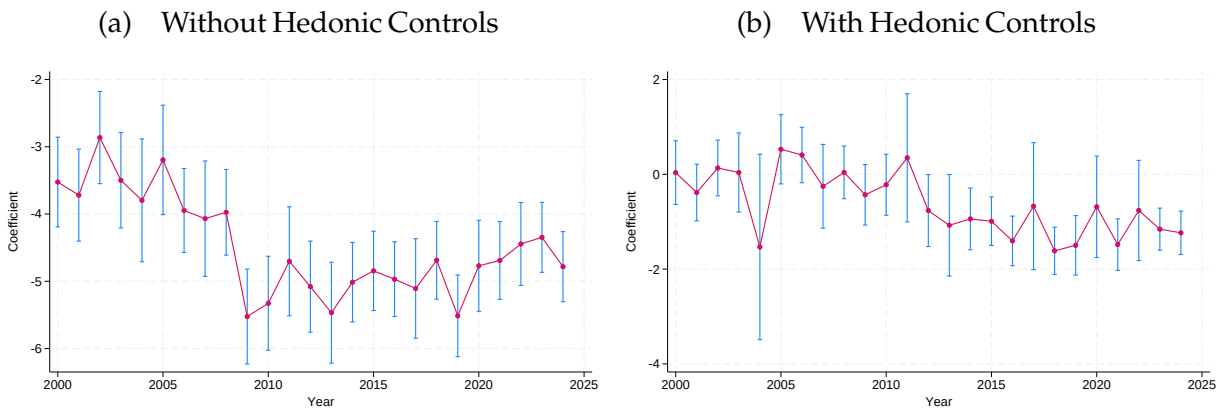
Our first analysis examines the association between climate risk and house prices. The outcome variable y_{it} is the log transaction price for property i in year t. We first estimate the coefficients $\{\phi_t\}$ from equation (2), which is the equivalent of a cross-sectional regression year by year, and then we estimate $\{\phi_t\}$ from equation (3), which is akin to a "repeat-sales" methodology.

Figure 2a plots the loadings of log prices on AAL without controlling for any hedonic characteristics. Most evident in the figure is a stark drop from 2005 to 2010 for the sensitivity of log price to AAL, which is persistently present thereafter. Quantitatively, $\phi_{2024} - \phi_{2000} \approx -1.5$, representing a 1.5% price drop per 1% of climate risk exposure. This is a preview of our main result, that climate concerns had a sizable impact on housing markets. In addition, the timing of the discount is potentially interesting as it occurs relatively early on, although this will change somewhat as we introduce different controls and methodologies. Note also that there is a negative coefficient even in the early 2000s, suggesting some validity to the type of story above that high-AAL properties have characteristics that are value-reducing.

Figure 2b repeats the same cross-sectional analysis but including square-footage, number of bedrooms, and property age as a hedonic controls, allowing all of these variables to have time-varying loadings. There are two key changes to the results. First, while the cumulative magnitude of the climate discount is similar across the two approaches (both with and without hedonic controls give about $\phi_{2024} - \phi_{2000} \approx -1.5$), the timing of the discount arrives later when hedonic controls are included. With hedonic controls, the discount appears only post-2011. Second, taking out these hedonic controls completely eliminates the non-zero coefficients in the early 2000s.[8] We view it as reassuring that, after controlling for some relatively standard hedonic variables, high-AAL and low-AAL properties have similar prices in the earlier dates when climate risk was plausibly less of a concern.

[8] In separate analysis where we include the hedonic controls one-by-one, we have found that the square-footage variable is primarily responsible for absorbing the early 2000s price differences between low-AAL and high-AAL properties. High-AAL properties tend to be smaller, which itself garners a discount, and this size discount was the primary driver of any non-zero difference in pricing in the early 2000s.

FIGURE 2: Price against within-Tract AAL, by Year: Cross-Sectional Regressions



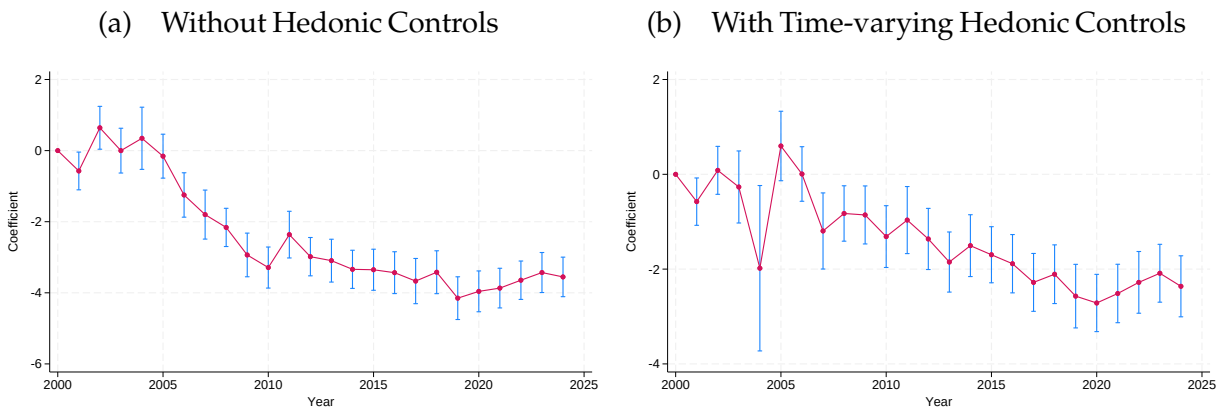
Notes. These figures plot the coefficients $\{\phi_t\}$ from equation (2) where y_{it} is the log transaction price of property i, and geography g in fixed effect α_{gt} is the Census tract. This regression is run year by year. The hedonic controls (X_i variables) used in panel b are log square footage, indicators for number of bedrooms (from 1 to 6+), and age of the property. The blue bars indicate the 95% confidence intervals on the coefficient estimates, with standard errors clustered by the Census tract.

Figure 3a plots the loadings of log prices on AAL from equation (3), which is the repeat-sales-style specification that corrects for time-varying compositional effects in the market. Relative to the cross-sectional approach in Figure 2a, the effects are substantially larger in magnitudes with $\phi_{2024} - \phi_{2000} \approx -4$. However, the timing of the climate discount remains similar: the decline in the ϕ_t coefficients begins around 2006 and continues to 2010, after which the coefficient remains permanently depressed. In addition, the pattern in the $\{\phi_t\}$ coefficients is smoother, and measured better statistically, in this repeat-sales version, suggesting that some random variation due to time-varying compositional effects may have been at play in Figures 2a-2b.[9]

Figure 3b augments our repeat-sales specification with time-varying hedonic controls for square-footage, number of bedrooms, and property age. Here, the addition of these hedonic controls is meant to capture some aspects of time-varying fundamental housing demand, which could be correlated with climate risk. For instance, larger houses are climate-safer (recall Table 3), so a growing demand for larger houses will show up as a growing climate

[9]Online Appendix Figure A.1 repeats the analysis in a two-stage regression described in footnote 7. The results are almost identical to the one-stage approach in Figure 3a, suggesting that the impact of heterogeneous transaction rates is limited (since that is the only way the two analyses could produce different results). In subsequent analyses, we therefore explore only the one-stage approach.

FIGURE 3: Price against within-Tract AAL, by Year: Repeat-Sales



Notes. These figures plot the coefficients $\{\phi_t\}$ from equation (3) where y_{it} is the log transaction price of property i, and geography g in fixed effect α_{gt} is the Census tract. The hedonic controls (X_i variables) used in panel b are log square footage, indicators for number of bedrooms (from 1 to 6+), and age of the property. The blue bars indicate the 95% confidence intervals on the coefficient estimates, with standard errors two-way clustered by the Census tract and year.

risk discount in any specification that does not control for house size. Broadly speaking, consistent with our earlier results, we continue to see a gradually-building climate risk discount, although the cumulative magnitude by 2024 is only about a -2% discount per 1% AAL (in contrast to the no-hedonic-controls baseline of -4% from Figure 3a). In addition, the timing of a climate discount is a bit more gradual, beginning in 2006 and continuing through 2020, after which it stabilizes.[10]

A key difference when including hedonic controls is that they are not available for every property, meaning the sample of properties shrinks from about 45 million to 28 million in moving from Figure 3a to the hedonic-controlled result in Figure 3b. How much of the change in our estimates can be accounted for by the different sample? To answer this, Figure A.2 in the Online Appendix shows loadings on AAL, without hedonic controls, but restricting to properties that do have hedonic controls populated in our data. The cumulative

[10]Using the estimated loadings on the hedonic controls, we can also understand what types of time-varying demand are present in the data. Figure A.3 in the Online Appendix displays these coefficients (for the hedonic controls logSqFt_i, dummies for Bedrooms_i, and Age_i). Between 2005-2010, a strongly increasing pattern is visible on the square-footage loading, which then mean-reverts—an temporary increase in demand for larger houses that mean-reverts. Given that climate-safer properties are larger (recall Table 3), this pattern in the loadings on square-footage helps explain why the analysis without hedonic controls in Figure 3a features an estimated climate discount that appears earlier.

climate discount for this sample is also about -2% per 1% AAL. In this sense, the difference in magnitudes between Figures 3a and 3b is primarily driven by a different sample.

Therefore, we conclude that the cumulative climate discount is between 2-4%, depending on the sample of transactions, and it appears starting around 2006 and continuing through the late-2010s. Relative to the existing literature, this effect arrives earlier and is larger in magnitude.

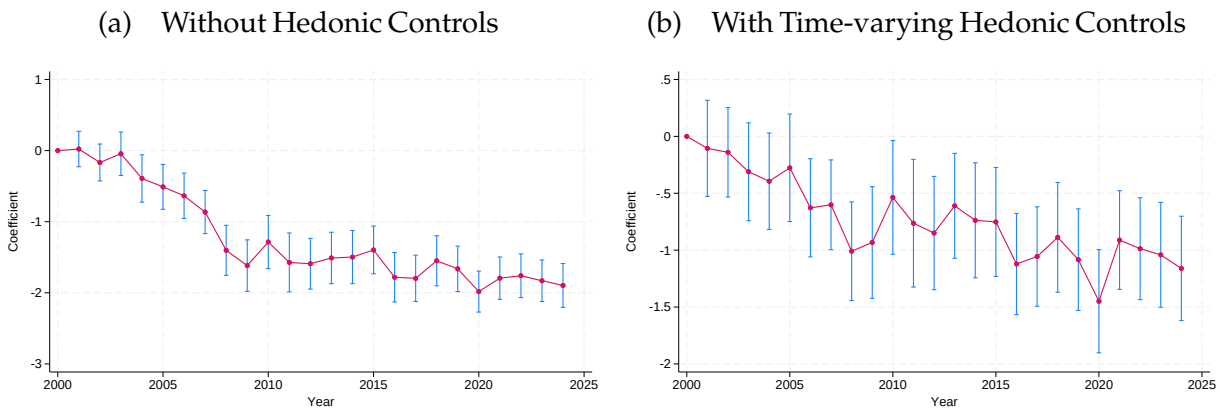
4.2 Owner-occupancy

We next examine owner-occupancy as opposed to transaction prices. We jump directly to the more rigorous repeat-sales-style specification (3), with the outcome variable y_{it} being an indicator for whether the buyer in a housing transaction is an owner-occupant. Figure 4a plots the loadings $\{\phi_t\}$ of this owner-occupancy indictor on AAL. Interestingly, although the outcome variable is different, the pattern in the coefficients largely mimics that of the coefficients of the log price outcome. In particular, high-AAL properties are increasingly likely to be non-owner-occupied starting around 2005, with the change continuing for several years. Quantitatively, a 1% higher AAL translates to a 1.5% decline in owner-occupancy by 2010. Like the price regressions, this effect remains permanently thereafter. Figure 4b repeats the same analysis, but including time-varying controls for hedonic characteristics. The results remain similar but with a slightly dampened magnitude and more gradually building up over the years. This slight dampening and more gradual pattern mimics very closely how the price coefficients changed when including hedonic controls.

4.3 Mortgage-level outcomes

Next, we examine several features of the buyer's mortgage. We are able to see the loan-to-value ratio (LTV), an indicator for whether the mortgage has an adjustable rate (ARM), and indicators for whether the mortgage is an FHA or VA loan. Figure 5 displays the results for loadings of these mortgage outcomes on AAL, in the repeat-sales-style specification (3). As before, each outcome variable is featured in two panels, one without and one with the hedonic controls (square-footage, number of bedrooms, and property age). Relative to

FIGURE 4: Owner Occupation against within-Tract AAL, by Year: Repeat-Sales



Notes. These figures plot the coefficients $\{\phi_t\}$ from equation (3) where y_{it} is an indicator that equals 1 when property i was bought by a buyer to be owner-occupied, and geography g in fixed effect α_{gt} is the Census tract. The hedonic controls (X_i variables) used in panel b are log square footage, indicators for number of bedrooms (from 1 to 6+), and age of the property. The blue bars indicate the 95% confidence intervals on the coefficient estimates, with standard errors two-way clustered by the Census tract and year.

low-AAL properties, the buyers of high-AAL properties obtain lower LTV ratios (by between 1-2% per 1% of AAL), are less likely to obtain FHA loans (by between 0.5-1% per 1% of AAL), and are slightly less likely to obtain VA loans (by between 0-0.3% per 1% of AAL). These effects all appear to materialize starting between 2006 and 2010, which roughly aligns with the timing of the evidence on prices and owner-occupation. (Note that there is little evidence that the prevalence of ARMs changes differentially for high-AAL and low-AAL properties.)

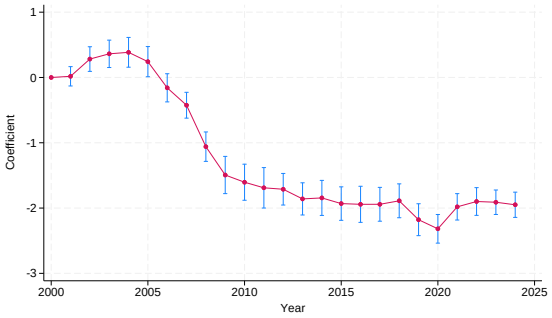
4.4 Back-of-the-envelope calculations

Some back-of-the-envelope calculations suggest quantitatively sizable effects at the aggregate level. Given the cross-sectional mean AAL of 0.25% from Table 1, we can multiply by the cumulative coefficient change ($\phi_{2024} - \phi_{2000}$) for any outcome variable to obtain an estimate of the cumulative aggregate impact due to climate risk. For prices, the cumulative discount was between (−4,−2) per unit of AAL, which implies a 0.5-1% lower level of house prices due to climate risk. For owner-occupancy, the same type of calculation implies a 0.25-0.4% lower owner-occupancy due to climate risk. Finally, for LTV ratios, this same type of calculation implies a 0.25-0.5% reduction in LTV due to climate risk. Of course, such

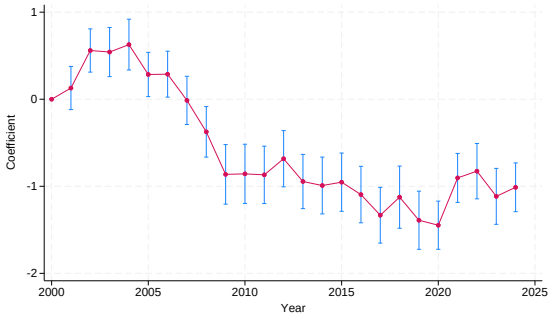
back-of-the-envelope calculations ignore reallocation and general equilibrium effects, which would imply they overstate or understate, respectively, the true aggregate effect.

FIGURE 5: Mortgage Outcome Variables

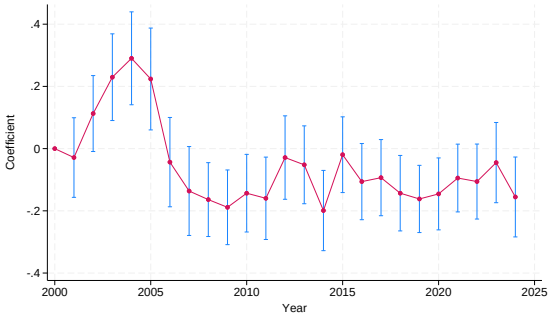
(a) LTV, without Hedonic Controls



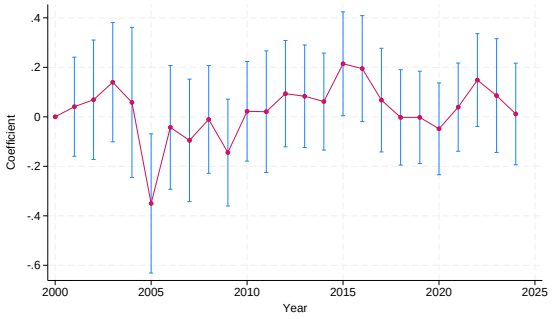
(b) LTV, with Hedonic Controls



(c) ARM, without Hedonic Controls



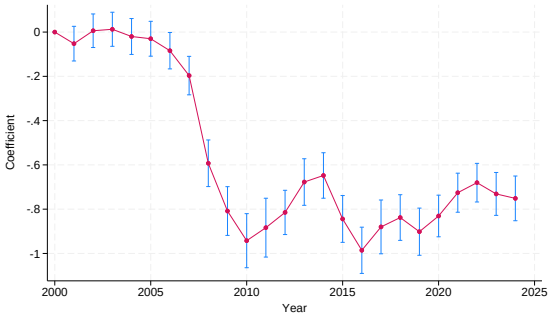
(d) ARM, with Hedonic Controls



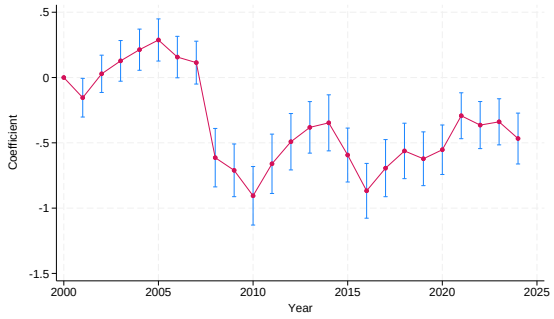
Notes. These figures plot the coefficients $\{\phi_t\}$ from equation (3) where y_{it} varies by panel, and where geography g in fixed effect α_{gt} is the Census tract. In the first row (panels a-b), y_{it} is the LTV of the buyer of property i in year t. In the second row (panels c-d), y_{it} is an indicator that equals 1 when property i was bought by a buyer with an ARM. The hedonic controls (X_i variables) used in each of the right-hand panels (b,d) are log square footage, indicators for number of bedrooms (from 1 to 6+), and age of the property. The blue bars indicate the 95% confidence intervals on the coefficient estimates, with standard errors two-way clustered by the Census tract and year.

FIGURE 5: Mortgage Outcome Variables (Continued)

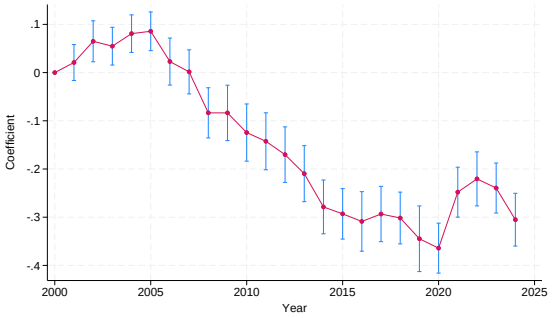
(e) FHA Loan, without Hedonic Controls



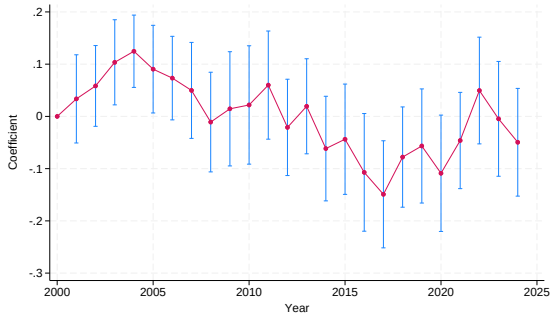
(f) FHA Loan, with Hedonic Controls



(g) VA Loan, without Hedonic Controls



(h) VA Loan, with Hedonic Controls



Notes. These figures plot the coefficients $\{\phi_t\}$ from equation (3) where y_{it} varies by panel, and where geography g in fixed effect α_{gt} is the Census tract. In the first row (panels e-f), y_{it} is an indicator that equals 1 when property i was bought by a buyer with an FHA loan. In the fourth row (panels g-h), y_{it} is an indicator that equals 1 when property i was bought by a buyer with a VA loan. The hedonic controls (X_i variables) used in each of the right-hand panels (f,h) are log square footage, indicators for number of bedrooms (from 1 to 6+), and age of the property. The blue bars indicate the 95% confidence intervals on the coefficient estimates, with standard errors two-way clustered by the Census tract and year.

5 Robustness

5.1 Block-level geographic controls

The granularity of our property-level data allows us to refine our geographic controls to the block-level. This is an important exercise that distinguishes our study from many others in the literature. In most studies, climate risk metrics are based on geographic variation, and so controlling aggressively for geography eliminates the useful variation in climate risk.

In a sense, our choice to use Census tract fixed effects as a baseline is motivated by similar concerns. As discussed earlier (recall Table 2b), using Census tract fixed effects preserves significant variation in AAL while simultaneously purging most fundamental variation in housing demand. That being said, our data is rich enough to allow finer geographic controls, since our property-level climate risk measures contain variation from small details of the construction materials, property design, and upkeep. This type of variation is unique to our study.

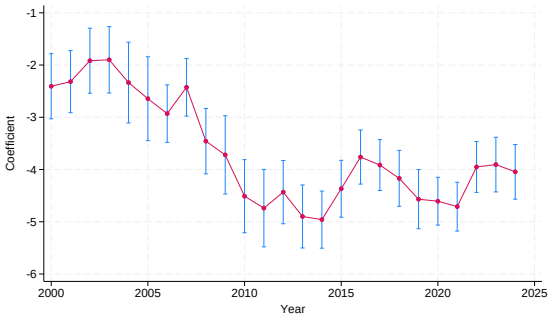
Figure 6 repeats the analyses for transaction price and owner-occupancy using only within-block variation. That is, the term α_{gt} in equations (2)-(3) is now a block-by-year fixed effect. The figure plots the loadings $\{\phi_t\}$ on AAL. Going across the various panels, we find very similar results to our baseline results with tract-level controls—both in terms of the timing of the climate risk discount and the cumulative magnitude. In particular, our preferred repeat-sales methodology with time-varying hedonic controls (Figures 6d and 6f) shows a gradually building climate discount alongside a decreasing tendency for climate-risky properties to be owner-occupied. The cumulative climate discount is $\phi_{2024} \approx -2$ per unit of AAL, while the cumulative decrease in owner-occupation is $\phi_{2024} \approx -1.5$ per unit of AAL. These gradual patterns and cumulative numbers are quite similar to our baseline results. Given the extremely refined geographic controls, we view this as evidence that the discount is caused by climate risk and not some other unobservable confound.

5.2 Single-family homes vs apartment buildings

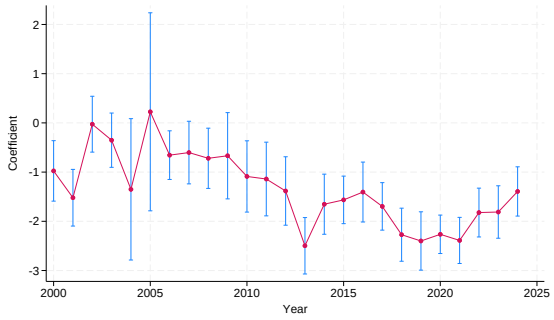
Figure A.4 in the Online Appendix repeats the analyses above after restricting the sample to single-family homes (to avoid potential confounds from apartment buildings). In these

FIGURE 6: Using Within-Block AAL

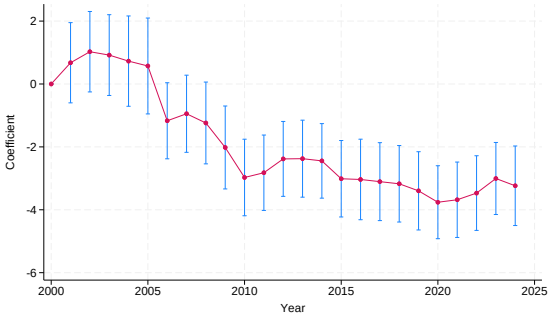
(a) Price, Raw



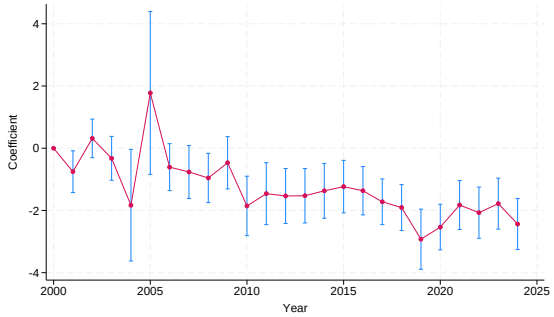
(b) Price, Raw, Hedonic Controls



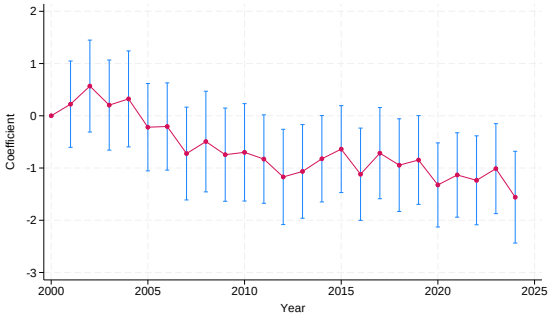
(c) Price, Repeat Sales



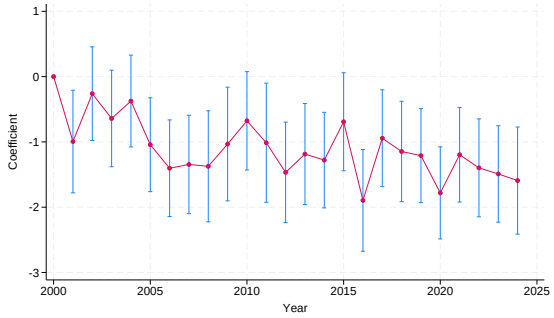
(d) Price, Repeat Sales, Hedonic Controls



(e) Owner-Occ., Repeat Sales



(f) Owner-Occ., Repeat Sales, Hedonic Controls



Notes. This figure repeats the analyses in Figures 2-3-4 but with α_{gt} representing a block-by-time fixed effect (rather than tract-by-time). The blue bars indicate the 95% confidence intervals on the coefficient estimates, with standard errors two-way clustered by the Census block and year.

results, we use the more demanding specifications with block-by-time fixed effects, as opposed to tract-by-time fixed effects. The results are very similar to the ones discussed for the whole sample, even slightly larger in magnitude.

5.3 Foreclosure Crisis

Given the main shift in estimated climate risk pricing occurs starting around 2005-06, our results could be due to the housing market crisis that triggered the Great Recession. Indeed, Figure A.5a in the Online Appendix plots the count of REO sales—a proxy for foreclosures—over time, increasing sharply starting around 2007 and only reverting back to its pre-crisis level by around 2020.

Figures A.5b and A.5c in the Online Appendix run raw cross-sectional and repeat-sales regressions of REO sale likelihood on within-block AAL (i.e., the outcome variable y_{it} is an indicator for an REO sale, and the geographic fixed effects α_{gt} are block-by-year fixed effects). Properties with higher climate-risk exposure within blocks experienced higher REO sale likelihoods starting around 2007 but revert back to a null relationship by around 2012. Because of this transitory relationship between climate risk exposure and foreclosures, it is unlikely that the permanent discount in prices and the permanent effect on ownership patterns can be due to foreclosures.

To directly test whether foreclosures can be responsible for our results, we directly exclude REO transactions and repeat the entire analysis, with the results in Figure A.6. The patterns and magnitudes are virtually unchanged from our baseline results.

5.4 Hurricane Katrina

Again, given the main shift in estimated climate risk pricing occurs starting around 2005-06, the change might be due to the direct impact of Hurricane Katrina, which struck southeastern US in August 2005. To explore this possibility, we re-run the main specification excluding Alabama, Louisiana and Mississippi, which are the three states that were directly affected by Katrina.

Results are shown in Figure A.7 in the Online Appendix, for both prices and owner-occupation. In these results, we use the more demanding specifications with block-by-time fixed effects, as opposed to tract-by-time fixed effects. The results are remarkably similar to our baseline results: the timing of the decline is still around 2005-06, the pattern of decline over several years remains largely the same (though slightly more drawn out for prices),

and the cumulative effects are quantitatively of similar magnitude.

While these results are meant to control for the *direct* effect of Katrina, they do not necessarily rule out an *indirect role* for Katrina. As mentioned in the introduction, a prominent hypothesis, supported by survey data and even some observational studies, is that extreme weather events may raise climate awareness and alter risk perceptions. The fact that Katrina was a large-scale, high-profile event may have caused a shift in climate risk perceptions nationwide. In that sense, our results are completely consistent with Katrina playing a role, through changing household beliefs.

5.5 Peril-specific analysis

We can repeat the same repeat-sales analysis but with peril-specific AALs rather than composite AALs. Doing so helps us explore two issues. First, do peril-specific exposures garner a discount over and above the an overall composite risk exposure? Second, do different risks become salient at similar or different times?

Figures A.8-A.14 in the Online Appendix display the results from this analysis. We re-run the repeat-sales analysis from equation (3), but with two modifications. Rather than the Composite AAL A_i, we project transaction prices onto a peril-specific AAL A_i^p (for $p =$ earthquake, hurricane, flood, etc.). We plot the time series of coefficients on A_i^p, which are conditional on property and block-by-time fixed effects. In all of these regressions, we include the hedonic controls for square-footage, number of bedrooms, and age.

While these estimates are much noisier than the Composite AAL estimates, we do see climate risk discounts for hurricane storm surge, hurricane wind, inland floods.[11]

[11]Some words of caution apply to these analyses. The estimated discount for winter storm perils seems unrealistically large, probably because it is a "small" risk—according to the diagonals in Table 2c, winter storm has a small exposure magnitude. In such cases with small risks, covariances to other perils and fundamental variables like size and age are non-trivial relative to the peril-specific variances, meaning that the estimated coefficient on the peril in question likely picks up these other exposures. The perils less subject to these potential issues are wildfire, hurricane storm surge, inland flood, and convective storm, all of which have more stable and reasonable estimates for peril-specific risk discounts.

6 Conclusion

Using a novel property-level dataset on climate risk exposures, exploiting only variation in exposures within highly-localized geographies, and employing a new repeat-sales-style methodology, we show that climate risk impounded a significant discount on house prices. Depending on the methodology, this discount began appearing between 2006 and 2010 and culminated by 2024 in a 2-4% price discount per 1% of additional climate risk. At exactly the same time, climate-risky houses began being disproportionately sold to landlords and/or investors rather than owner-occupants, the cumulative effect by 2024 being a 1% lower owner-occupancy per 1% of additional climate risk. Finally, over a similar time-frame, the buyers of climate-risky houses became less levered (lower LTV) and less likely to be FHA and VA borrowers. All these effects are permanent, consistent with an interpretation that climate risk perceptions shifted earlier and had larger impacts than the literature has appreciated.

References

Addoum, J. M., P. Eichholtz, E. Steiner, and E. Yönder (2024): "Climate change and commercial real estate: Evidence from Hurricane Sandy," *Real Estate Economics*, 52, 687–713.

Akerlof, K., E. W. Maibach, D. Fitzgerald, A. Y. Cedeno, and A. Neuman (2013): "Do people "personally experience" global warming, and if so how, and does it matter?" *Global Environmental Change*, 23, 81–91.

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

Bergquist, P. and C. Warshaw (2019): "Does global warming increase public concern about climate change?" *The Journal of Politics*, 81, 686–691.

Bernstein, A., S. B. Billings, M. T. Gustafson, and R. Lewis (2022): "Partisan residential sorting on climate change risk," *Journal of Financial Economics*, 146, 989–1015.

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

Brulle, R. J., J. Carmichael, and J. C. Jenkins (2012): "Shifting public opinion on climate change: an empirical assessment of factors influencing concern over climate change in the US, 2002–2010," *Climatic change*, 114, 169–188.

Carmichael, J. T. and R. J. Brulle (2017): "Elite cues, media coverage, and public concern: an integrated path analysis of public opinion on climate change, 2001–2013," *Environmental Politics*, 26, 232–252.

Contat, J., C. Hopkins, L. Mejia, and M. Suandi (2024): "When climate meets real estate: A survey of the literature," *Real Estate Economics*, 52, 618–659.

Deryugina, T. (2013): "How do people update? The effects of local weather fluctuations on beliefs about global warming," *Climatic Change*, 118, 397–416.

Deryugina, T., L. Kawano, and S. Levitt (2018): "The economic impact of Hurricane Katrina on its victims: Evidence from individual tax returns," *American Economic Journal: Applied Economics*, 10, 202–233.

Egan, P. J. and M. Mullin (2012): "Turning personal experience into political attitudes: The effect of local weather on Americans' perceptions about global warming," *The Journal of Politics*, 74, 796–809.

Ellen, I. G. and R. Meltzer (2024): "Heterogeneity in the recovery of local real estate markets after extreme events: The case of Hurricane Sandy," *Real Estate Economics*.

Engle, R. F., S. Giglio, B. Kelly, H. Lee, and J. Stroebel (2020): "Hedging climate change news," *The Review of Financial Studies*, 33, 1184–1216.

Faccini, R., R. Matin, and G. Skiadopoulos (2023): "Dissecting climate risks: Are they reflected in stock prices?" *Journal of Banking & Finance*, 155, 106948.

Gallagher, J. and D. Hartley (2017): "Household finance after a natural disaster: The case of hurricane Katrina," *American Economic Journal: Economic Policy*, 9, 199–228.

Gao, A. and G. Wang (2007): "Multiple transactions model: a panel data approach to estimate housing market indices," *Journal of Real Estate Research*, 29, 241–266.

Gormsen, N. J., K. Huber, and S. S. Oh (2024): "Climate capitalists," Tech. rep., National Bureau of Economic Research.

Ilhan, E. (2020): "Sea level rise and portfolio choice," Available at SSRN 3743046.

Ilhan, E., P. Krueger, Z. Sautner, and L. T. Starks (2023): "Climate risk disclosure and institutional investors," *The Review of Financial Studies*, 36, 2617–2650.

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

Keys, B. J. and P. Mulder (2020): "Neglected no more: Housing markets, mortgage lending, and sea level rise," Tech. rep., National Bureau of Economic Research.

——— (2024): "Property Insurance and Disaster Risk: New Evidence from Mortgage Escrow Data," Tech. rep., National Bureau of Economic Research.

Krosnick, J. A., A. L. Holbrook, and P. S. Visser (2000): "The impact of the fall 1997 debate about global warming onAmerican public opinion," *Public Understanding of Science*, 9, 239.

Marlon, J. R., X. Wang, P. Bergquist, P. D. Howe, A. Leiserowitz, E. Maibach, M. Mildenberger, and S. Rosenthal (2022): "Change in US state-level public opinion about climate change: 2008–2020," *Environmental Research Letters*, 17, 124046.

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

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

Nolan, J. M. (2010): ""An inconvenient truth" increases knowledge, concern, and willingness to reduce greenhouse gases," *Environment and Behavior*, 42, 643–658.

Nyce, C., R. E. Dumm, G. S. Sirmans, and G. Smersh (2015): "The capitalization of insurance premiums in house prices," *Journal of Risk and Insurance*, 82, 891–919.

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

Pollack, A. B. and R. K. Kaufmann (2022): "Increasing storm risk, structural defense, and house prices in the Florida Keys," *Ecological Economics*, 194, 107350.

Schuetz, J. (2024): "How will US households adjust their housing behaviors in response to climate change?" *Real Estate Economics*.

Shao, W. (2017): "Weather, climate, politics, or God? Determinants of American public opinions toward global warming," *Environmental Politics*, 26, 71–96.