

The Uneven Distribution of Climate Risks and Discounts

by John Heilbron and Kevin Zhao¹

In this brief, we document the uneven distribution of climate risk in real estate using novel data on expected losses due to climate risk at the property-level. We show that properties located in counties that are poorer, less educated, older, more rural, and that have less belief in climate change tend to have more climate risk. Next, using home sales, we document heterogeneity between counties in the size of discount per unit of climate risk. We find a smaller discount per unit of climate risk in a similar set of more exposed counties.

We summarize these findings and clarify orders of magnitude by conducting a simple estimation of the loss in housing wealth resulting from a repricing of the housing stock. When we perform this repricing according to an “empirical” benchmark obtained from the most conservative discounts observed in the data, we estimate that high-risk households stand to lose up to \$3,400, equivalent to 2.3 percentage points of their home value and over 23 percentage points of their home equity. Under an alternative “frictionless” benchmark obtained from capitalization rates in financial markets, we estimate losses of \$11,000, 6.1 percentage points of home value, and 61 percentage points of home equity. Taken together, our results reveal a novel financial stability concern stemming from climate risks in real estate and suggest that climate risk exposure may be larger than previously documented, especially in vulnerable communities.

Background and Introduction

Climate risks and the energy transition present a variety of potential threats to the financial system (for an overview, see the Financial Stability Oversight Council’s Report on Climate-Related Financial Risk, 2021²). The housing market is no exception to rising climate threats, and given its sheer size (\$45 trillion in 2022³), it may be of particular concern. One possible channel is that climate shocks may devalue real estate, increasing the default risk on mortgage loans and mortgage-backed securities (Brunetti et al. (2021)). However elevated default risk does not require the occurrence of an adverse climate event. Growing concerns about future climate episodes may devalue home prices today, wiping out homeowner equity and raising default likelihoods. In this brief, we study the distribution of climate risks in the U.S. housing market and the pricing of these risks to better assess the magnitude of homeowner losses as these risks become fully capitalized into home prices.

In our analysis of climate risk and risk pricing, we use CoreLogic Climate Risk Analytics.⁴ These data provide broad and granular coverage of natural hazard risks in real estate in the United States. Most studies have focused on the effects of flood and sea-level rise risk (e.g., Gourevitch et al. (2023)), with much less work being done on wildfire risk and even less on convective storm and earthquake risk.⁵ Furthermore, analysis has thus far been largely constrained to smaller geographic regions (see Giglio et al. (2021) for a review). By contrast, CoreLogic models average annual loss from climate risk along multiple dimensions at the property-level for the entire continental United States.

In the data, we examine the distribution of climate risk geographically and across different population segments. We find that climate risk is disproportionately distributed in regions that are generally less financially resilient (for example, poorer and less-educated counties). We then estimate the sensitivity of house prices to climate risk separately by county and state while controlling for a host of other housing characteristics, such as property-level elevation and distance to the coast. On average, all else equal, homes with higher climate risk sell at a discount when priced. In addition, there is a large degree of heterogeneity across states and across counties in the size of the

discount. Notably, the regions with smaller discounts per unit of climate risk tend also to be more financially vulnerable.

Our findings suggest that accurately appraising threats to financial stability requires careful analysis of the cross-sectional incidence of climate risks, climate risk pricing, and household financial vulnerability. We document a previously unidentified financial stability concern stemming from climate risks in real estate: vulnerable households live in the areas with the highest climate risk and have also overpaid for their homes relative to their level of climate risk. These households are the least resilient to financial shocks and have the highest loan-to-value ratio on their mortgages.

In the final part of the brief, we introduce a simple framework to estimate the average amount of housing wealth that could be lost in a climate repricing event. Under a somewhat conservative “empirical” benchmark, we find that the average homeowner stands to lose around \$2,200. Importantly, losses are concentrated among the high-vulnerability group. This group is least resilient to financial shocks and most likely to pass along losses as credit events through mortgage defaults. The average loss in the high-vulnerability group is \$3,400, equivalent to 2.3 percentage points of their home value and over 23 percentage points of their home equity. By contrast, under an alternative “frictionless” benchmark, the average homeowner stands to lose \$11,000. Because losses are more dramatic across all groups in this scenario, they are somewhat less concentrated among the high vulnerability group but still dramatic. In this case, the average loss among the high-vulnerability group is \$8,800, 6.1 percentage points of home value, or 61 percentage points of home equity.

The rest of the brief is structured as follows: Section 2 introduces the climate risk data and documents an uneven distribution in climate risk levels, both geographically and when regions are sorted based on demographic characteristics; Section 3 examines the pricing of climate risk and finds that discounts associated with climate risks in real estate are also unevenly distributed; Section 4 discusses a novel financial stability concern stemming from the findings in Sections 2 and 3 and provides back of the envelope calculations for potential losses resulting from the

underpricing of climate risk in real estate; and Section 5 concludes the brief.

Climate Risk Data

Analytic Sample

We use three datasets from CoreLogic in our main analysis. CoreLogic Climate Risk Analytics provides projections of physical damages due to natural disasters at the individual property-level. CoreLogic Tax and Deeds includes information on housing units and transactions. The CoreLogic Mortgage Basic dataset contains information on the origination characteristics of mortgage loans. The three datasets can be merged at the property-level using identifiers internal to CoreLogic.

CoreLogic Climate Risk Analytics provides a one-time snapshot of property-level annual average loss (AAL) of structure value due to climate risk in the contiguous United States. The data are projections based on quantitative climate risk modeling performed by CoreLogic in December 2021. CoreLogic provides AALs for different climate perils such as earthquakes, floods, non-flood weather, and a composite of all climate events. In addition to present AALs, CoreLogic also reports AALs for future time horizons under different climate scenarios. Section 2.2 details the geographic distribution of AALs in the data and highlights a few interesting areas.

These data provide granularity in the measure of climate risk as well as the variety of perils, scenarios, and time-horizons considered. The property-level nature of the data allows us to control for home characteristics that can interact with climate risk (for example, distance-to-coast⁶ and elevation⁷), as well as amenities that may vary at the county-level. The data also covers the entire continental United States, allowing us to make statements about heterogeneity in climate risks and climate risk pricing for a more comprehensive sample than what is normally studied in the literature. Finally, because our measure of climate risk incorporates a variety of perils, it provides a more holistic view of climate risk and climate risk pricing compared to studies that focus on a specific risk.

To construct our analytic sample, we merge the three datasets from CoreLogic at the property-level. By doing so, we construct records with measures of climate risk measures, home prices, and financing conditions. The former two fields permit the measurement of climate risk pricing, while the latter allows us to assess the amount of home equity at home purchase that stands to be lost. We augment this merge with the supplementary data on housing location to better control for amenities and the demographic information to better assess heterogeneity in the cross-section.

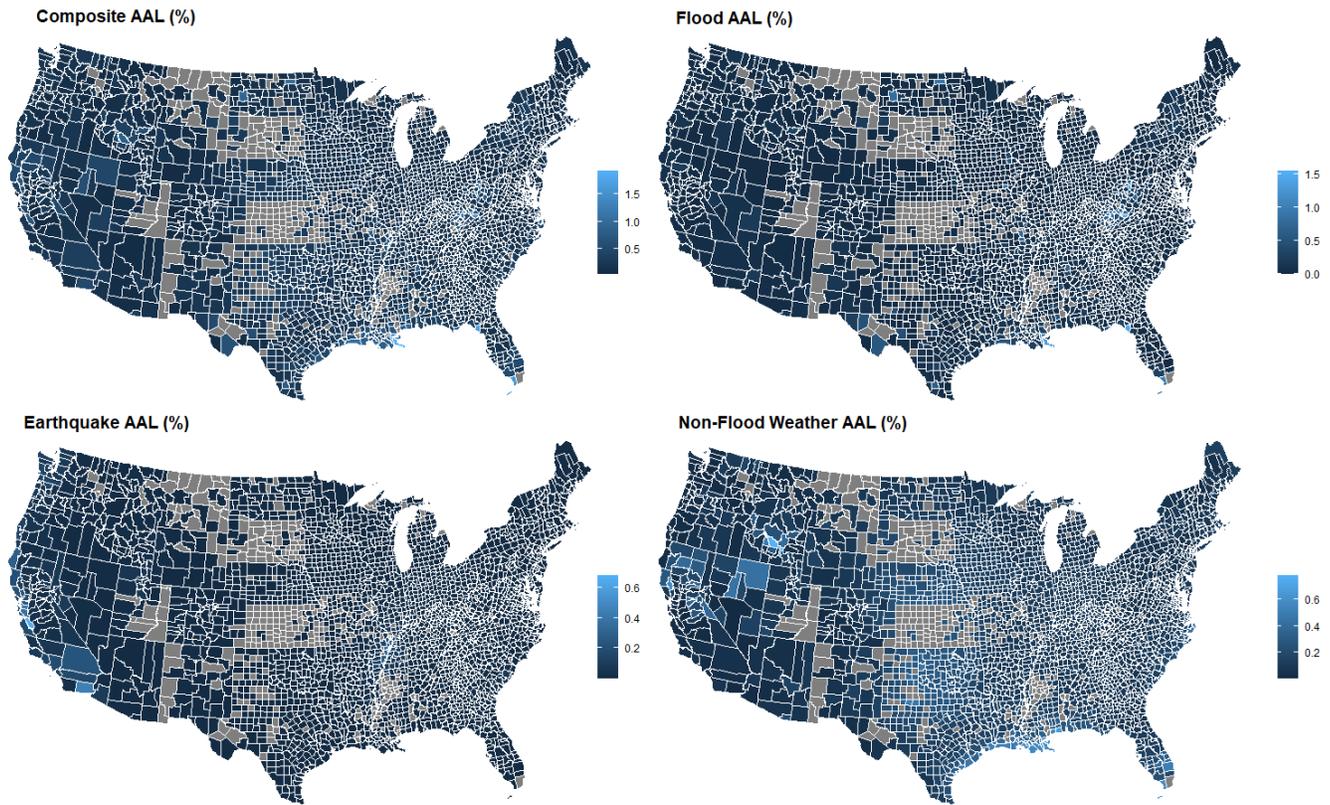
Our final sample consists of around 10 million arms-length housing transactions with a sales price of at least \$1,000 between 2020 and 2022. **Figure 1** shows summary statistics for the sample of transactions we study. The sample is representative of typical home sales in the United States: the median transaction is on a 1,700 square foot, 3-bedroom home that is 40 years old. The sample has a large degree of heterogeneity in climate risk, and the distribution is positively skewed.

Figure 1. Summary Statistics for Our Sample of Transactions Occurring Between 2020 and 2022

	Mean	Median	S.D.	Obs.
Sales price (\$ thousands)	414.02	285.00	2,049.80	10,245,937
Composite AAL (percent)	0.18	0.11	0.33	10,245,937
Flood AAL (percent)	0.05	0.00	0.30	10,245,937
Earthquake AAL (percent)	0.02	0.00	0.06	10,245,937
Non-flood weather AAL (percent)	0.11	0.09	0.11	10,245,937
Land sq. ft. (1000 ft.)	67.21	10.02	3,661.66	10,235,007
Building sq. ft. (1000 ft.)	1.94	1.72	2.08	9,317,936
# bedrooms	3.25	3	1.38	6,866,787
Structure age (years)	44.40	40.00	31.32	9,126,495
Distance to coast (km)	272.09	141.59	349.29	10,186,740
Elevation (m)	276.74	191.74	376.09	10,114,957

Sources: CoreLogic, Inc., USGS, NOAA, Authors' Analysis

Figure 2. Average AAL by County Separately for Composite Risk, Flood Risk, Earthquake Risk, and Non-flood Weather Risk



Note: Gray counties are either not covered by the data or did not have a transaction between 2020-2022.

Sources: CoreLogic, Inc., Authors' Analysis

The mean composite AAL is 0.18%, with a standard deviation of 0.33%, compared to a median of 0.11%. Home sales prices in the sample are also positively skewed: the median sales price is \$285,000, while the mean of the distribution is \$414,020.

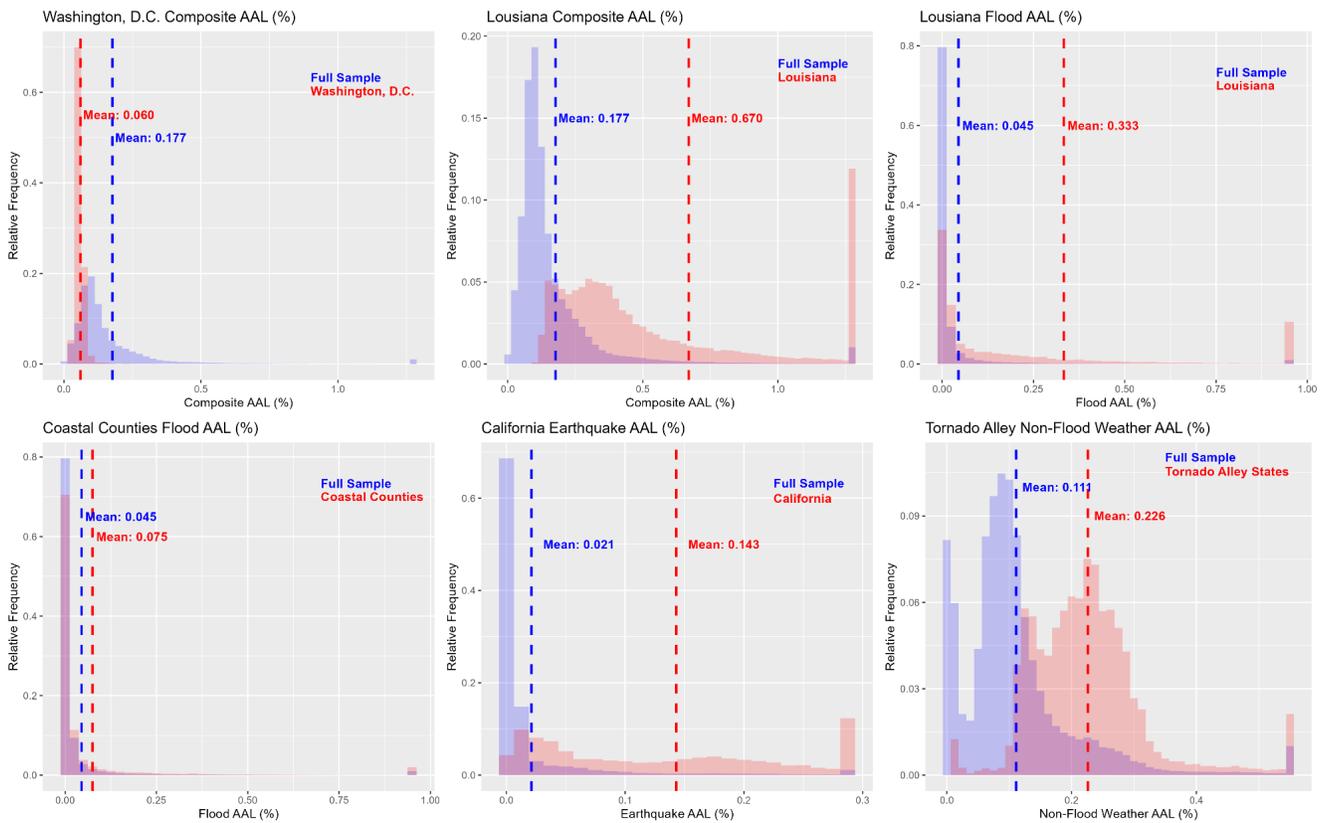
Geographic Heterogeneity in Climate Risks

Climate risks in real estate are not evenly distributed across the continental United States. **Figure 2** shows the spatial distribution of average AALs by county, separately for composite, flood, earthquake, and non-flood weather. The states with the highest composite risk are concentrated in the Southeast (Louisiana, Florida, and Mississippi), where flood and non-flood weather risks are both elevated. Within those states, the counties with the highest risk in the sample are coastal counties in Louisiana and Florida with elevated flood risk, such as Plaquemines, Louisiana

and Dixie, Florida. In contrast, the counties with the lowest composite risk are mainly in the inland regions of the Mountain West states, such as Montana, Idaho, and Utah.

Figure 3 highlights a few geographic areas with elevated climate risk. AALs are winsorized at the 99% level for each risk. In general, areas of elevated climate risk in our sample align with common perceptions. For example, counties on the coast have elevated flood risk.⁸ In addition, the composite AAL in Louisiana is the highest of any state and is four times that of the full sample. By contrast, the composite AAL in Washington, D.C. is the nation's lowest. Most of Louisiana's climate risk comes from flood risk, as properties there have a mean flood risk AAL over seven times as high as the whole sample. Earthquake risk is generally higher in the western United States, especially California. California's mean earthquake risk AAL is almost seven times higher than the sample mean. Non-flood weather

Figure 3. Distribution of AALs for the Full Sample and Selected Regions



Note: AALs are winsorized at the 99th percentile.
Sources: CoreLogic, Inc., Authors’ Analysis

risk is around twice as high in states along “tornado alley”⁹ compared to states that see less tornado activity.

Demographic Heterogeneity in Climate Risks

Next, we explore whether counties with certain characteristics have higher climate risk on average. We merge in county-level demographic metrics, including average age, education level, poverty level, income level, and how urban or rural the county is.¹⁰ In addition, we use a measure of belief in climate change from the Yale Climate Opinion Map.¹¹ These demographics allow us to identify regions more “vulnerable” and less resilient to climate disasters. For example, poorer and less educated households are less likely to have the financial means to rebuild properly after being hit by a natural disaster. **Figure 4** shows summary statistics for each of these measures at the county-level and the transaction-level. Relative to the average county,

the transactions in our sample occur in counties with younger residents, are more well-educated, have higher income, are more urban, and are located on the coast. This is consistent with a larger share of homes in the United States being concentrated in more developed, urban, and coastal areas.

Figure 5 shows the distributions of composite AALs split by the county-level metrics, and their means. We split the sample of transactions into two equally sized bins. Although the differences in distributions are not as striking as they were for some geographic splits, we find that, in general, counties with more vulnerable residents have higher climate risk on average. For example, counties with older and less-educated residents, on average, have higher composite climate risk. In addition, properties located in more rural areas are exposed to higher risk.

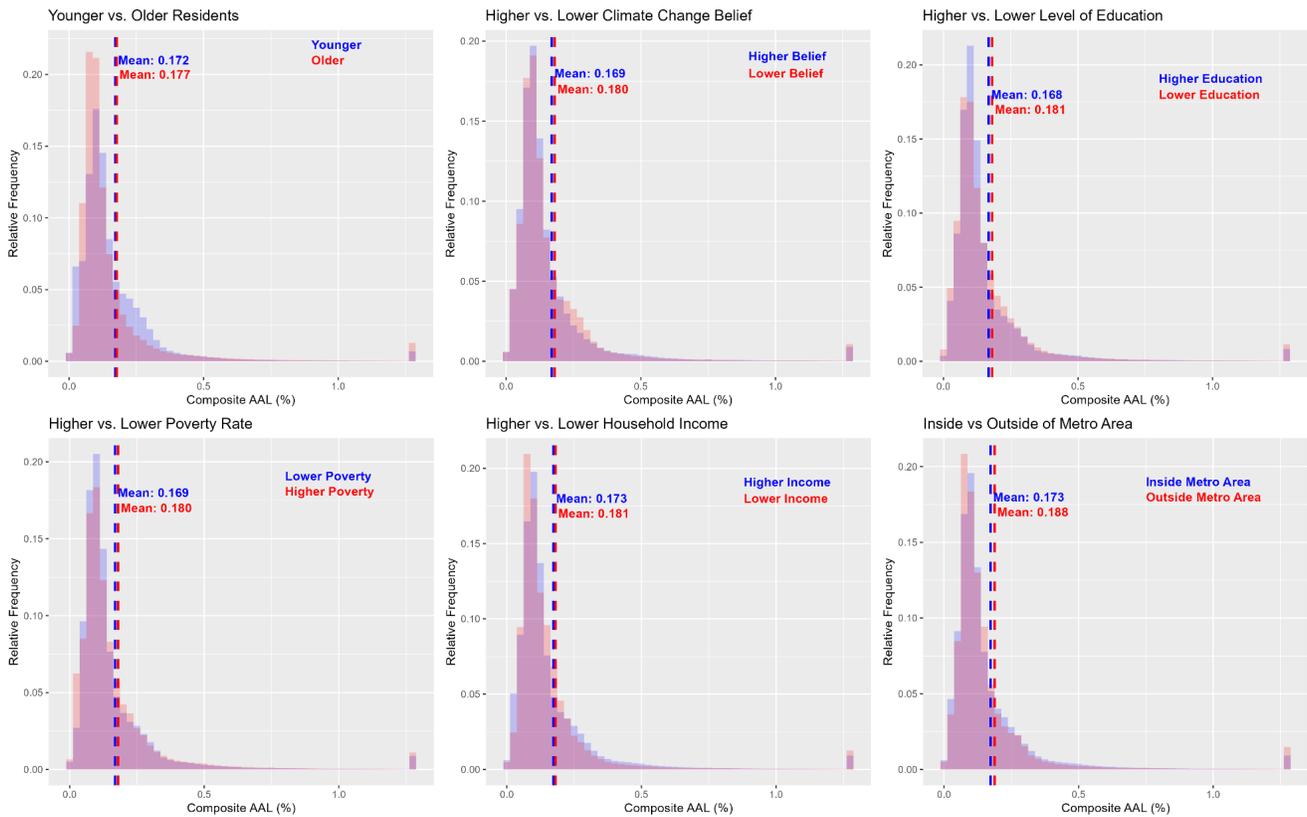
Figure 5 is also consistent with the existing climate risk literature on sorting by certain demographics into

Figure 4. Summary Statistics for Characteristics at the County-level and the Transaction-level

	County Level (N≈3,000)			Transaction-level (N≈10,000,000)		
	Mean	Median	S.D.	Mean	Median	S.D.
Belief (percent)	65.27	64.58	6.16	71.03	71.33	6.23
Age (years)	41.67	41.40	5.31	39.29	38.50	4.74
College graduate (percent)	23.19	20.86	9.88	32.86	32.78	10.82
Poverty level (percent)	13.70	12.80	5.39	11.65	11.40	3.97
HH income (\$ thousands)	57.62	55.31	14.59	69.19	65.79	17.49
Metro county	0.38	1	0.49	0.87	1	0.34
Coastal county	0.08	0	0.27	0.23	0	0.42

Sources: CoreLogic, Inc., USDA, U.S. Census Bureau, Yale Climate Opinion Map, Authors' Analysis

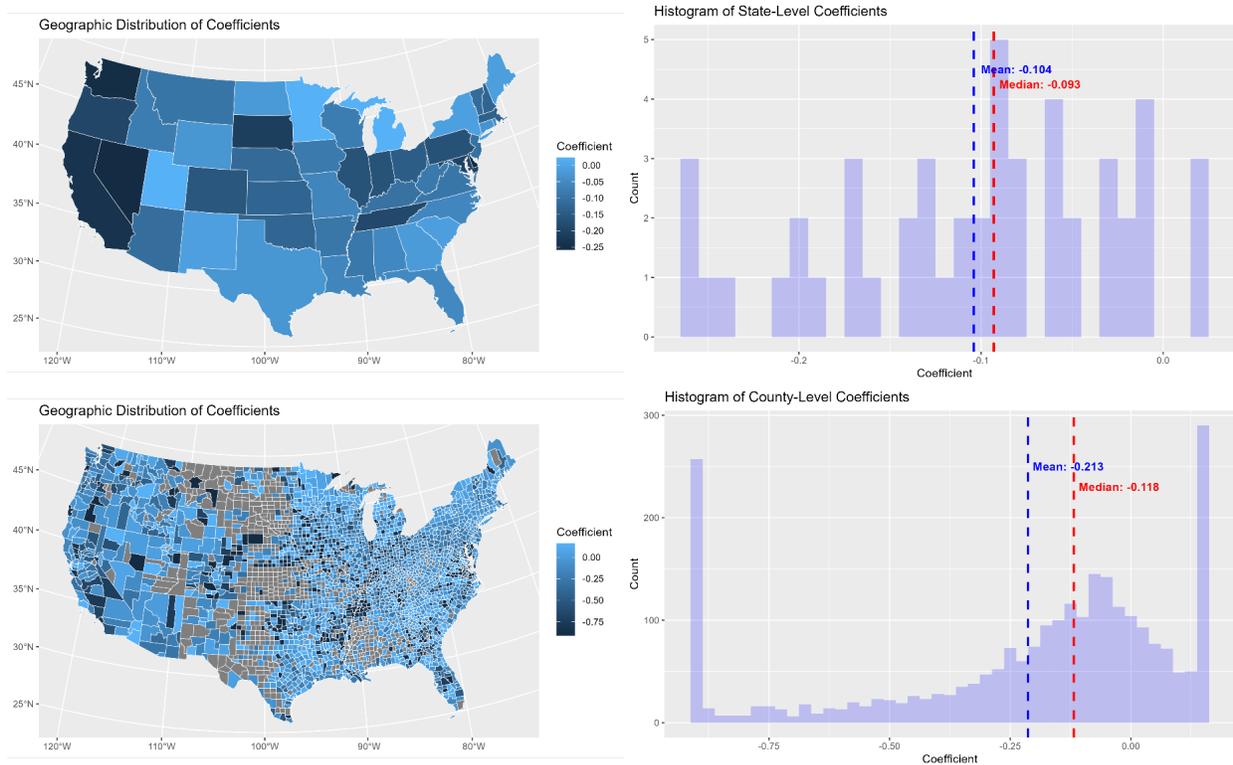
Figure 5. Distribution of AALs Split by County Characteristics



Note: AALs are winsorized at the 99th percentile.

Sources: CoreLogic, Inc., Authors' Analysis

Figure 6. Distributions of β 's on Climate Risk from State- and County-level Regressions



Note: β 's winsorized at the 5th and 95th percentile for state-level regressions and at the 10th and 90th percentile for county-level regressions. Gray counties are either not covered by the data or did not have at least 100 qualifying transactions between 2020-2022. Sources: CoreLogic, Inc., USGS, NOAA, Authors' Analysis

high- and low-climate risk areas. Although the existing literature has mostly focused on flood and sea level rise risk, our findings suggest that the results may generalize to a broader set of climate risks. For example, Bakkensen and Ma (2020) find that low-income residents are more likely to sort into high flood risk areas. Similarly, in our data, counties with a higher poverty rate and lower average income have a higher composite climate risk. This is a financial stability concern because the segment of the population that are least financially resilient to potential shocks bears the largest share of climate risk. Therefore, the expected damages in dollar terms may understate the strain caused on the financial system by climate shocks. Section 4 contains a longer discussion on this point.

Pricing of Climate Risk

To test for the discount placed on climate risk in each geographic region, we run hedonic regressions

separately for each state and county. To do so, we use the following regression specification:

$$\log(\text{Sales Price})_i = \alpha + \beta * (\text{Composite AAL})_i + \text{controls}_i + FE + \varepsilon_i$$

We use a restricted set of controls for the main analysis to preserve observations in the sample with certain unpopulated fields. Results remain qualitatively similar after adding more extensive controls and fixed effects.¹²

For each geographic region, we recover the coefficient on climate risk. β can be interpreted as a measure of the sensitivity of home prices in a geographic area to climate risk (“climate risk sensitivity”). A negative β coefficient indicates that homes in the region with higher climate risk sell at lower prices. A more negative β means that the size of the discount per unit of climate risk is larger.

Geographic Heterogeneity

Figure 6 shows the distribution of climate risk sensitivities at the state and county levels. In the average state and county, we find that $\beta < 0$, indicating that homes with higher climate risk tend to sell at lower prices. The level of discount per unit of climate risk is economically meaningful. For the median (mean) state, a within-state 1 standard deviation increase in AAL is associated with a 2.26% (2.85%) decrease in transaction price. For the median (mean) county, a within-county 1 standard deviation increase in AAL is associated with a 2.49% (5.95%) decrease in transaction price. 46 out of the 49 states and 1,836 out of 2,529 counties in the sample have a negative β .¹³

Our analysis demonstrates that, on average, composite climate risk is priced across the entire continental United States. In contrast, prior literature has usually focused on a specific type of climate risk (especially sea-level rise or wildfire risks) in a smaller geographic region (usually a single state or a subset of coastal counties). However, it is important to note that while a negative β indicates that climate risk is priced, by itself, it does not provide any insight into whether the size of the discount is higher or lower than what we would expect by looking at AALs. Section 4 contains a longer discussion on this topic, and we find that relative to a simple benchmark, the median β s that we estimate are too small in absolute magnitude. That is, climate risk may not be fully reflected in home sales prices in the median state and county.

Demographic Heterogeneity

We merge in county-level demographic metrics to test whether counties with certain characteristics are more likely to have a larger discount for climate risk. Along this line, recent papers have shown that flood risk is priced differently by people based on their income and race (Gourevitch et al. (2023)), their belief in climate change (Baldauf et al. (2020)), and their political affiliation (Bernstein et al. (2021)). Our data allows for analysis of a more comprehensive measure of climate risk and explores heterogeneity in pricing along a more extensive set of dimensions. An additional difference is that the climate risk we study in this paper is a measure of the *current level* of risk. Interpreting the level of discount per unit of current climate risk can be

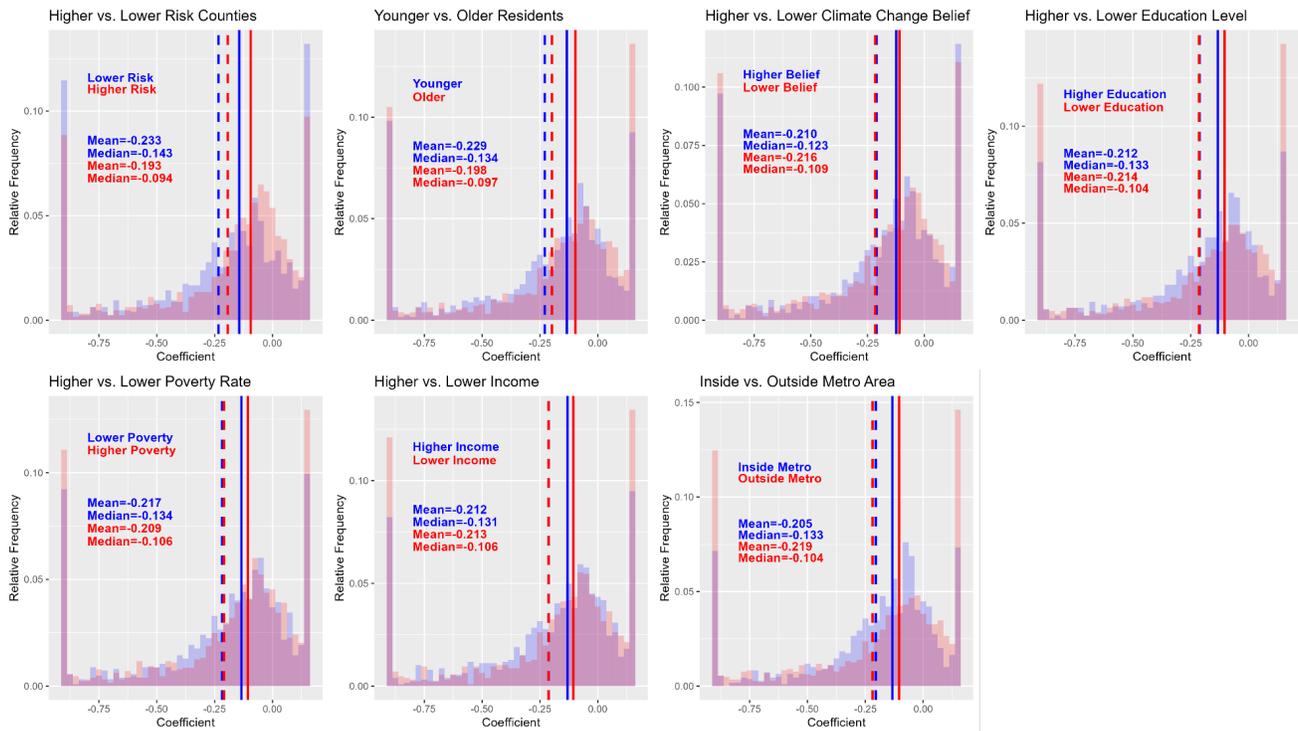
different than if we were studying future projections of risk, as is done in the literature examining pricing of sea level rise risk (Baldauf et al. (2020); Bernstein et al. (2021)). By using current climate risk, we can examine the impacts of characteristics correlated with financial literacy and attention to climate risk, such as education or household income. This channel is related to but distinct from studying heterogeneity between how people with different climate change *beliefs* price climate risk *projections*.

Figure 7 shows the county-level climate risk sensitivities distribution derived from regression specification (1), split by county metrics. We split the sample of counties into two equally sized bins. The top left chart in each panel shows that there is a difference between higher and lower risk areas in their discount on climate risk: the β and *t-statistic* in lower risk areas (in blue) is more negative, indicating that lower risk areas tend to have a higher discount per unit of climate risk.

The rest of **Figure 7** shows that there is also a difference in climate risk sensitivities when split by demographic characteristics, as in Section 2.3. The counties with demographics correlated with higher climate risk, as shown in Section 2.3, are plotted in red, while counties with demographics correlated with lower climate risk are plotted in blue. As evident from the charts, the red distributions tend to be shifted towards the right, indicating a less negative β corresponding to a *smaller* discount per unit of climate risk.

Counties with a higher level of climate risk and lower discount per unit of climate risk (in red) are those that are typically perceived as more vulnerable. For example, counties with older, less educated, and lower-income residents have less of a discount on climate risk. In addition, rural counties and those with a smaller share of residents believing in climate change also have less discount per unit of climate risk. These results suggest a further financial stability concern that compounds with the findings in Section 2. The more financially vulnerable segments of the population are living in more risky areas while receiving a smaller discount per unit of climate risk when compared to less vulnerable segments. This coincidence between risks and risk mispricing means that climate risks in real estate can be understated if either of the components are examined in isolation.

Figure 7. Distributions of β 's on climate risk from county-level regressions split by county characteristics



Note: β 's are winsorized at the 10th and 90th percentiles

Sources: CoreLogic, Inc., Authors' Analysis

Discussion

Benchmarks for Observed Climate Risk Sensitivity

In the data, we measure that climate risk sensitivities $\{\beta_i\}$ are negative on average, but vary across different geographies, i . To interpret these coefficients and assess potential financial stability implications of the pricing of climate risk, we compare the observed β_i to two benchmarks, β^* . We refer to these as the “frictionless” and “empirical” benchmarks. Each benchmark represents an appropriate climate risk sensitivity under certain assumptions, which are discussed below. Although the benchmarks impose strong assumptions, they allow us to make progress in sizing up the welfare consequences and vulnerabilities that climate risks present.

The “frictionless” benchmark considers how households might reasonably respond to climate risks under certain clarifying, though potentially counterfactual,

conditions. Suppose at home purchase, the household anticipates the expenses it may incur due to damages resulting from climate risks and pays less accordingly. Their climate risk sensitivity precisely measures the size of the home discount for a given level of anticipated damages. Suppose climate damages are idiosyncratic (i.e., unrelated to other sources of household risk), fully insurable at fair premiums, and that households are fully attentive to them. We would expect households’ climate risk sensitivity to be roughly the inverse of the risk-free rate, $\beta^* = -1/r^f$.¹⁴

The “empirical” benchmark considers how certain households are already responding to climate risks. As discussed, different households in the data have different climate risk sensitivities. We use those households already exhibiting the greatest degree of sensitivity as our benchmark for the other households in the sample, $\beta^* = \min\{\beta_i\}$. Put differently, this benchmark presumes that the largest price discounts are the most accurate and that smaller discounts are supported only by misperceptions or market frictions that will not survive in the long run.

In our later analysis, we use these benchmarks to consider how house prices will change if, in some repricing event, households become more sensitive to climate risks. We do not take a stand on the appropriate benchmark, but we note that each has certain advantages. For example, it is possible that even the households that are the most sensitive to climate risk at present nevertheless understate these risks, and so a repricing event will have more dramatic consequences than their behavior suggests. By contrast, some factors keeping households' climate risk sensitivity low in magnitude, like limited access to credit at home purchase, may not be changed in a repricing of climate risk. In this case, the results from using the “frictionless” benchmark may be overstated.

In the population of home purchasers, we measure a median climate risk sensitivity, $med\{\beta_i\}$, of about -0.1 . By comparison, with a risk-free rate during our sample of about 3%, the “frictionless” benchmark instead implies a climate risk sensitivity of $\frac{-1}{0.03} * \frac{1}{100} = -0.33$.¹⁵ Finally, in our analysis of a repricing event below, we measure a minimum climate risk sensitivity, $min\{\beta_i\}$, of about -0.15 ,¹⁶ which serves as the “empirical” benchmark. Considered together, we have that households tend to be sensitive to climate risk, but the empirical and frictionless benchmark suggest, respectively, greater degrees of sensitivity, $-1/r^f < min\{\beta_i\} < med\{\beta_i\} < 0$.

To conclude our analysis of benchmarks, we provide additional intuition about different economic forces that might affect the observed climate risk sensitivities. These forces might cause the observed climate risk sensitivities to deviate from the “frictionless” benchmark, $med\{\beta_i\} \neq -1/r^f$. They may also contribute to the observed heterogeneity in sensitivities, which itself causes the “empirical” benchmark to differ from the population tendency, $min\{\beta_i\} < med\{\beta_i\}$. In what follows, we note the direction of these economic forces, and whether it is consistent with our observations. However, we do not take a stance on which forces are operative in the data. Several could be in play simultaneously, and identifying the separate influence of each is beyond the scope of this note.

As noted above, we observe climate risk sensitivities that differ from the “frictionless” benchmark. Possible explanations include the following:

1. If we allow climate risks to be correlated with other household risks, we may find observed sensitivities have a larger magnitude than the frictionless benchmark, $med\{\beta_i\} < -1/r^f$. In such a case, each unit of climate risk requires further discounting beyond the risk-free rate.
2. If some homes are not able to be fully insured at a fair price, home prices should become more sensitive to climate risks, again pushing observed sensitivities below the frictionless benchmark, $med\{\beta_i\} < -1/r^f$.
3. On the other hand, if insurers are underestimating climate risk and payouts are subsidized, we may find that observed sensitivities are of lesser magnitude than the frictionless benchmark, $med\{\beta_i\} < -1/r^f$, because households are not fully financially exposed to additional units of climate risk.
4. If households are inattentive and do not fully capitalize climate risks into their home purchase decision, we may also observe moderated sensitivities, $med\{\beta_i\} < -1/r^f$.

In the data, we measure that the areas with more vulnerable households have less negative β 's compared to areas with less vulnerable households, indicating a larger discount per unit of climate risk in areas with less vulnerable households. We elaborate on potential explanations for this heterogeneity below:

1. Different levels of attention to climate risk. If certain groups are less attentive to climate risks than others, they may not discount home prices as much per unit of risk. We observe that the *less-educated* households place a smaller discount on climate risk, consistent with a lack of attention to climate risk among the less financially literate.
2. Different access to fair insurance. As noted above, if all households have access to and pay for actuarially fair insurance, the premiums should be passed through to home prices and not impact climate risk sensitivities. However, if one segment of the population has greater access to insurance or subsidized insurance, we may observe a less negative β among that group. In that case, we would not expect climate damages to be fully passed through to home prices because the homeowner is not

responsible for the full amount of the reported climate risk. For differential access to insurance to explain the heterogeneity in climate risk sensitivity we observe in the data, it needs to be the case that the more vulnerable households have better access to insurance.

3. Different levels of belief about climate risk. Certain groups may have beliefs about the level of climate risk that differ from the objective measure we employ. Consistent with this, we observe that households living in areas with lower belief in climate change place a smaller discount on climate risk.
4. Different credit constraints. More credit constrained households may be less willing to pay extra for safer homes with less climate risk compared to less credit constrained households. Consistent with this, we observe that β is less negative in poorer counties, meaning that the difference in price between a risky home and a safe home is smaller in these counties.
5. Different risk aversion. We would expect that the more vulnerable, less financially resilient households have higher risk aversion. Higher risk aversion would translate to a more negative β , meaning more risk-averse households discount climate risk more than less risk-averse households. However, the data show the opposite: the more vulnerable households have a less negative β . Therefore, it is unlikely that differential risk aversion is the source of heterogeneity in climate risk sensitivities.

Loss Decomposition Framework for Repricing Event

With the two benchmarks climate risk sensitivities in mind, we calculate the potential losses should climate risks be repriced according to each benchmark. This provides an order of magnitude for the welfare costs of climate change that have not already been priced. The greater the financial system's exposure to the housing sector, the greater the extent to which this repricing shock represents a financial vulnerability.

Intuitively, the losses of an individual household will be the size of their climate risk exposure multiplied by the losses per unit of climate risk exposure. Additionally,

these losses may be further scaled according to the unit of analysis, e.g., whether we are measuring losses in (fractional) units of housing or dollars of housing wealth. Mathematically, we write the losses as:

$$\omega_i \cdot \Delta\beta_i \cdot AAL_i$$

In this expression, AAL_i is the extent of climate risk facing the home. The data measures this as the average annualized loss relative to the home's replacement cost. $\Delta\beta_i$ is the change in pricing of climate risk in the repricing event. This is computed as $\beta^* - \beta_p$, the difference between an asserted appropriate climate risk sensitivity benchmark, β^* , and the climate risk sensitivity under present market conditions, β_p . Therefore, is larger in magnitude for homes with a smaller discount per unit of climate risk at the time of purchase. Together, these comprise the percent loss in housing value due to the repricing. The first term, ω_p , are weights that can be chosen to reflect the units in which we wish to measure losses. We consider three different units of analysis: dollars of lost housing wealth, (fractional) units of lost housing, and (fractional) units of lost housing equity.¹⁷

We can draw valuable insights from separate examinations and comparisons between the different weighting schemes. The first scheme might be appropriate if concerned with total lost wealth. The second scheme may be useful to a social planner concerned with allocations of livable shelter. By contrast, a focus on financial stability recommends the third weighting scheme, as default risk is heightened after the exhaustion of housing equity. As will be discussed below, the choice of units matters for the assessment of the cross-sectional distribution of risks.

To consider the distribution of climate relating losses in the population, we begin by aggregating the losses faced by each household. We then disaggregate the sum in two decomposition exercises. In the first exercise, we consider the fraction of aggregate losses borne by different population segments that we delineate according to demographic characteristics. We find that more vulnerable population segments bear disproportionate losses, particularly when measured as fractions of housing equity. In the second exercise, we evaluate whether our estimates of aggregate losses are the result of high climate risk alone or whether they are exacerbated by the coincidence of high climate risk and

more extensive repricing of risk in key populations. We confirm and quantify the finding in the data that households facing more climate risk also face potentially steeper losses in climate risk repricing.

Estimation of Losses in a Repricing Event

We take a simple, back-of-the-envelope approach to estimate the loss of housing wealth in the event of climate risk repricing. We assume there is a single appropriate climate discount that applies to all households, corresponding to either our frictionless or empirical benchmark. We further suppose that any house sold at a more modest discount faces a potential repricing and fall in price to bring it in line with the

benchmark. We imagine a repricing shock in which discounts increase and housing values fall, and we consider the loss in housing wealth suffered as a result. The estimation illustrates our simple framework and conveys the insights from Section 3 on cross-sectional heterogeneity in risk pricing.

For each county, we construct a vulnerability index based on seven of the metrics studied above. Counties are classified as “vulnerable” or “not vulnerable” along each dimension separately. Vulnerable counties are those with higher average climate risk level, lower household income, higher poverty rates, lower percentage of college graduates, higher average resident age, less belief in climate change, and that are

Figure 8. Average Composite AAL, Sale Price, and Climate Discount (β) by Vulnerability Bin

Variable		Low Vulnerability		Medium Vulnerability		High Vulnerability	
Bin Size	Count (millions)	2.0		6.9		1.3	
	Frequency (percent)	20		67		13	
Within-bin Repricing Characteristic	Median sales price (\$ thousands)	412		278		145	
	Median CF price (\$ thousands)	420		283		146	
	Median LTV (percent)	82		88		90	
	Median leverage	5.6		8.9		10	
	Median down-payment (\$ thousands)	74		33		15	
	Mean AAL (percent)	0.13		0.18		0.21	
	Realized discount, β	-0.15		-0.10		-0.04	
	Repricing Benchmark	$\min\{\beta\}$	$-1/r^f$	$\min\{\beta\}$	$-1/r^f$	$\min\{\beta\}$	$-1/r^f$
"True" discount, β^*	-0.15	-0.33	-0.15	-0.33	-0.15	-0.33	
Change in discount, $\Delta\beta$	0	-0.18	-0.05	-0.23	-0.11	-0.29	
Within-bin Loss Estimate	Housing wealth (\$ thousands)	0	-9.8	-2.5	-12	-3.4	-8.9
	CF house price (percent)	0	-2.3	-0.9	-4.1	-2.3	-6.1
	Down-payment (percent)	0	-13	-7.5	-34	-23	-61

Sources: CoreLogic, Inc., USGS, NOAA, USDA, U.S. Census Bureau, Yale Climate Opinion Map, Authors' Analysis

outside of a metro area. The sample of transactions are then split such that an equal number of transactions are classified as “vulnerable” and “not vulnerable” along each metric. The transaction-level vulnerability index is the total number of metrics for which the county the property resides in was classified as vulnerable, and ranges from 0 (least vulnerable) to 7 (most vulnerable). We further split transactions into low-, medium-, and high-vulnerability bins based on their vulnerability index.¹⁸ We estimate the risk-pricing within each vulnerability bin, find the average AALs, and estimate the counterfactual zero-AAL home prices. We report these, together with the relative transaction frequency in each segment in **Figure 8**.

First, we consider the empirical benchmark, $\beta^* = \min\{\beta_i\}$. We take the risk-pricing in the low-vulnerability segment, which has the most conservative risk-pricing, to be the appropriate risk-pricing, β^* . Using these values, we compute the average loss in the aftermath of a repricing for a household within each of the three vulnerability segments.

Note that the medium-vulnerability bin comprises an outsized proportion of transactions. Higher-vulnerability households have lower priced homes, lower climate discounts, and face more climate risk. In a repricing event, low-vulnerability households – which have the largest discounts – incur no losses in housing wealth by construction. The medium- and high-vulnerability households incur losses of \$2,500 and \$3,400, respectively. The growing severity of losses is even more stark when taken as a percentage of home price. Medium- and high-vulnerability households lose 0.9 percentage points and 2.3 percentage points of the value of their homes respectively. This represents 7.5 percentage points and 23 percentage points of the equity they initially put into their home purchase.

To provide additional context, we compute the average dollar loss in housing wealth among the whole population under the empirical benchmark. On average, households lose about \$2,200 in housing wealth, 0.9 percentage points of their housing value, and 8.0 percentage points of their housing equity. We proceed to decompose these aggregates in two separate fashions, as described in the previous section. We report these results in **Figure 9a**.¹⁹

In the first decomposition, we compute the extent to which low-, medium-, and high-vulnerability households contribute to the average loss. Again, by construction, low vulnerability households incur no losses. As a baseline, high vulnerability households make up 16 percentage points of the population after excluding low vulnerability households. By contrast, these households bear 20 percentage points of the losses when measured in housing wealth, 33 percentage points when measured as a proportion of housing value, and 37 percentage points when measured relative to housing equity. Regardless of how losses are measured, the most vulnerable households bear a disproportionate share of the losses.

In the second decomposition, we assess differences between a naïve estimate of housing climate risks, measured as cross-sectional average AAL, and our more robust approach. Regardless of our measurement approach, the coincidence of high climate risks and high mispricing of risk requires a term that revises the naïve estimate upward. This revision amounts to 10 percentage points of lost wealth, 8 percentage points of lost housing, or 7 percentage points of lost housing equity. Naïve estimates need to be further adjusted to account for the relationship between loss sizes and the units of analysis. Households with large relative losses tend to live in lower priced houses, so estimates of housing wealth losses must be revised downward by 23 percentage points. By contrast, households with large relative losses tend to be more highly levered, so estimates of fractional equity losses must be further revised upward by 10 percentage points. Taking these two adjustment factors together, we find that the naïve estimate overstates losses measured as housing wealth but understates losses measured as a proportion of housing equity.

We repeat the decomposition using the frictionless benchmark, $-1/r^f$. In this case, we find that households lose, on average, \$11,000 in housing wealth, 4 percentage points of their housing value, and 33.7 percentage points of their housing equity. Again, we decompose these losses and report the results in **Figure 9b**. Under the frictionless benchmark, there are larger price discounts than under the empirical benchmark and so even the low vulnerability households incur some losses. Roughly 80-90 percentage points of losses are still attributable to the medium

Figure 9a. Decomposition of Aggregate/Average Losses, Repricing Under “Empirical” Benchmark

Benchmark: $\beta^* = \min\{\beta_i\}$			Unit of Analysis		Housing Wealth (\$ thousands)		Housing (percent)		Housing Equity (percent)	
			Welfare Weight		$\omega_i = P_i$		$\omega_i = 1$		$\omega_i = P_i / D_i$	
Aggregation	Component	Calculation	Estimate	Percent	Estimate	Percent	Estimate	Percent		
Aggregate	Average Loss	$E[\omega_i \cdot \Delta\beta_i \cdot AAL_i]$	-2.2	100	-0.9	100	-8.02	100		
Decomposition 1: Vulnerability Types	Low Vulnerability	$Pr(i = L) \cdot [\omega_i \cdot \Delta\beta_i \cdot AAL_i i = L]$	0	0	0	0	0	0		
	Medium Vulnerability	$Pr(i = M) \cdot [\omega_i \cdot \Delta\beta_i \cdot AAL_i i = M]$	-1.7	80	-0.61	67	-5.0	63		
	High Vulnerability	$Pr(i = H) \cdot [\omega_i \cdot \Delta\beta_i \cdot AAL_i i = H]$	-0.4	20	-0.29	33	-2.9	37		
Decomposition 2: Covariance Adjustments	“Naïve” Average Loss	$E[\omega] \cdot E[\Delta\beta] \cdot E[AAL]$	-2.4	113	-0.83	92	-6.7	83		
	Risk-Pricing/Risk Covariance	$E[\omega] \cdot Cov(\Delta\beta, AAL)$	-0.21	10	-0.07	8	-0.57	7		
	Loss Scaler/Loss Covariance	$Cov(\omega, \Delta\beta \cdot AAL)$	0.49	-23	0	0	-0.79	10		

Sources: CoreLogic, Inc., USGS, NOAA, USDA, U.S. Census Bureau, Yale Climate Opinion Map, Authors’ Analysis

Figure 9b. Decomposition of Aggregate/Average Losses, Repricing Under “Frictionless” Benchmark

Benchmark: $\beta^* = -1/r^f$			Unit of Analysis		Housing Wealth (\$ thousands)		Housing (percent)		Housing Equity (percent)	
			Welfare Weight		$\omega_i = P_i$		$\omega_i = 1$		$\omega_i = P_i / D_i$	
Aggregation	Component	Calculation	Estimate	Percent	Estimate	Percent	Estimate	Percent		
Aggregate	Average Loss	$E[\omega_i \cdot \Delta\beta_i \cdot AAL_i]$	-11.0	100	-4.0	100	-33.7	100		
Decomposition 1: Vulnerability Types	Low Vulnerability	$Pr(i = L) \cdot [\omega_i \cdot \Delta\beta_i \cdot AAL_i i = L]$	-1.9	18	-0.46	11	-2.6	8		
	Medium Vulnerability	$Pr(i = M) \cdot [\omega_i \cdot \Delta\beta_i \cdot AAL_i i = M]$	-7.9	72	-2.8	69	-23.3	69		
	High Vulnerability	$Pr(i = H) \cdot [\omega_i \cdot \Delta\beta_i \cdot AAL_i i = H]$	-1.1	10	-0.78	19	-7.8	23		
Decomposition 2: Covariance Adjustments	“Naïve” Average Loss	$E[\omega] \cdot E[\Delta\beta] \cdot E[AAL]$	-11.6	106	-3.96	98	-31.7	94		
	Risk-Pricing/Risk Covariance	$E[\omega] \cdot Cov(\Delta\beta, AAL)$	-0.21	2	-0.07	2	-0.57	2		
	Loss Scaler/Loss Covariance	$Cov(\omega, \Delta\beta \cdot AAL)$	0.81	-7	0	0	-1.4	4		

Sources: CoreLogic, Inc., USGS, NOAA, USDA, U.S. Census Bureau, Yale Climate Opinion Map, Authors’ Analysis

and high vulnerability households. Of these, the high vulnerability households share 10 percentage points of losses of housing wealth, 19 percentage points of housing, and 23 percentage points of housing equity.

Turning to our second decomposition, we find that again the coincidence of large adjustments to risk pricing in areas with the greatest risk imply losses above and beyond the naïve estimates. This adjustment increases the losses by 2 percentage points, regardless of whether measuring in units of housing wealth, housing, or housing equity. Although the losses are larger under this second repricing scenario, they are uniformly larger for all types of households, and so the total effect is not driven by riskier units facing steeper price adjustments.

Conclusion

Using a novel dataset on property-level climate risk, we find that homebuyers in high climate risk areas do not sufficiently adjust purchase prices to account for climate risk properly. This phenomenon is particularly prevalent among the most vulnerable portion of population that is least financially equipped to recover following a climate shock. As a result, failing to account for the coincidence between climate risk and overpricing may understate the financial stability implications of climate risk in real estate.

References

- Bakkensen, Laura A., and Lala Ma. “Sorting over flood risk and implications for policy reform.” *Journal of Environmental Economics and Management* 104 (2020): 1023-62.
- Baldauf, Markus, Lorenzo Garlappi, and Constantine Yannelis. “Does climate change affect real estate prices? Only if you believe in it.” *The Review of Financial Studies* 33.3 (2020): 1256-1295.
- Bernstein, Asaf, Stephen B. Billings, Matthew T. Gustafson, and Ryan Lewis. “Partisan residential sorting on climate change risk.” *Journal of Financial Economics* 146, no. 3 (2022): 989-1015.
- Bernstein, Asaf, Matthew T. Gustafson, and Ryan Lewis. “Disaster on the horizon: The price effect of sea level rise.” *Journal of Financial Economics* 134.2 (2019): 253-272.
- Brunetti, Celso, Benjamin Dennis, Dylan Gates, Diana Hancock, David Ignell, Elizabeth K. Kiser, Gurubala Kotta, Anna Kovner, Richard J. Rosen, and Nicholas K. Tabor. “Climate change and financial stability.” FEDS Notes (2021).
- Financial Stability Oversight Council. “Report on Climate Related Financial Risk 2021.” <https://home.treasury.gov/system/files/261/FSOC-Climate-Report.pdf>
- Giglio, Stefano, Bryan Kelly, and Johannes Stroebel. “Climate finance.” *Annual Review of Financial Economics* 13 (2021): 15-36.
- Gourevitch, Jesse D., Carolyn Kousky, Yanjun “Penny” Liao, Christoph Nolte, Adam B. Pollack, Jeremy R. Porter & Joakim A. Weill. “Unpriced climate risk and the potential consequences of overvaluation in U.S. housing markets.” *Nature Climate Change* 13 (2023): 250-257.
- Howe, Peter D., Matto Mildenberger, Jennifer R. Marlon, and Anthony Leiserowitz. “Geographic variation in opinions on climate change at state and local scales in the USA.” *Nature Climate Change* 5.6 (2015): 596-60

Endnotes

- 1 John Heilbron, Interdisciplinary Researcher, Office of Financial Research (John.Heilbron@ofr.treasury.gov); and Kevin Zhao, Interdisciplinary Researcher, Office of Financial Research (Kevin.Zhao@ofr.treasury.gov).
- 2 [FSOC Report on Climate-Related Financial Risk \(treasury.gov\)](#)
- 3 <https://www.redfin.com/news/housing-market-loses-value-2023/>
- 4 <https://www.corelogic.com/data-solutions/property-data-solutions/climate-risk-analytics/>
- 5 CoreLogic’s composite risk aggregates flood risk, non-flood weather risk, and earthquake risk. Throughout the brief, we will use “composite risk” and “climate risk” interchangeably. In a strict sense, earthquakes may be classified as a non-climate natural disaster risk. This distinction is not crucial to our analysis and does not change any of the takeaways.
- 6 We compute the distance from the parcel to the nearest shoreline. For the shoreline map, we use the National Oceanic and Atmospheric Administration (NOAA) Medium Resolution Shoreline: <https://shoreline.noaa.gov/data/datasheets/medres.html>.
- 7 We compute the elevation of the property based on its coordinates and the National Map API from the U.S. Geological Survey (USGS): <https://apps.nationalmap.gov/epqs/>.
- 8 A list of counties located on the coast is from the U.S. Census Bureau. Coastal counties are the focus of most flood risk papers (Bernstein et al. (2019), Baldauf et al. (2020), Bernstein et al. (2022)).
- 9 Texas, Louisiana, Oklahoma, Kansas, Nebraska, and South Dakota.
- 10 Data on median age are from the U.S. Census Bureau. Data on education, poverty, income, and whether the county is in a metro area are from the U.S. Department of Agriculture.
- 11 Data on climate change beliefs at the county level are from the Yale Climate Opinion Map (Howe et al. (2015)). We measure the percentage of respondents that answered “yes” to a question about whether global warming is happening: *“Recently, you may have noticed that global warming has been getting some attention in the news. Global warming refers to the idea that the world’s average temperature has been increasing over the past 150 years, may be increasing in the future, and that the world’s climate may change as a result. What do you think: Do you think that global warming is happening?”* The Yale Program on Climate Change Communication bears no responsibility for the analyses or interpretations of the data presented here.

- 12 The minimal set of controls and fixed effects include land square footage, sale quarter, distance to coast bin, and elevation bin. For the state-level regressions, the minimal set of controls also includes a county fixed effect and sale month instead of sale quarter. Additional controls and fixed effects include building square footage, structure age, and number of bedrooms, which are only available for a subset of homes. Standard errors are clustered by sale month or sale quarter.
- 13 In terms of statistical significance, 40 out of 49 states have a significantly negative β , while 3 states have a significantly positive β . 990 out of 2,529 counties have a significantly negative β , while 255 counties have a significantly positive β .
- 14 This benchmark comes from the formula for valuing a perpetuity of cash-flows, c , at discount rate, r . The value of these cash-flows is $c * 1/r$.
- 15 In our regressions, AAL is expressed in percentage terms. To compare this benchmark to our results, it is necessary to divide by 100.
- 16 In our preceding measurement of climate risk sensitivities, we found counties that have $\beta_i < -0.15$, which may cause confusion. Our choice of empirical benchmark comes from the following analysis in which we work with more aggregated granularity bins for expositional purposes. In this latter exercise, it is the case that $\min\{\beta_i\} = -0.15$. If anything, this will serve to moderate our results and make our estimates of losses more conservative.
- 17 Setting $\omega_i = P_i$ will measure losses in dollars of housing wealth; $\omega_i = 1$ will measure losses in fractions of housing units; and $\omega_i = P_i / D_i = 1 / (1 - LTV)$ will convert losses to the fraction of housing equity.
- 18 The low, medium, and high vulnerability groups consist of properties classified as vulnerable along 0-1, 2-5, and 6-7 dimensions, respectively.
- 19 Though unnecessary for understanding the main results of the analysis, we present the formal basis of our decompositions here. The first relies on the identity relating conditional and unconditional expectations:

$$E^{(0)}[\omega_i \cdot \Delta\beta_i \cdot AAL_i] = \sum_i \Pr(i) * E^{(0)}[\omega_i \cdot \Delta\beta_i \cdot AAL_i \mid i]$$

Our second decomposition relies on the identity relating covariances and the expectation of the product of two variables. We apply this twice to obtain:

$$E^{(0)}[\omega_i \cdot \Delta\beta_i \cdot AAL_i] = E[\omega_i] \cdot E[\Delta\beta_i] \cdot E[AAL_i] + E[\omega_i] \cdot Cov(\Delta\beta_i, AAL_i) + Cov(\omega_i, \Delta\beta_i \cdot AAL_i)$$

The first term we deem the “naïve” estimate, the second term reflects correlations between risk and risk-pricing, and the third term

reflects correlations between the loss magnitudes and various units of analysis.