The Long-Run Effects of Tropical Cyclones on the U.S. Housing Market*

Davide Castellani†

March 25, 2026

Abstract

Housing is a major component of U.S. wealth, yet the long-term impact of tropical cyclones on housing markets is poorly understood. This paper presents a systematic analysis of all U.S. coastal counties exposed to tropical cyclones from 1992–2021. Housing values decline by 15% over eight years following cyclone exposure, with the most persistent losses in less-experienced areas. Rising insurance uptake post-disaster drives increased living costs, partly explaining price declines, while higher-risk areas eventually rebound due to greater insurance coverage.

Keywords: natural disasters, housing market, insurance

JEL-Codes: Q54, R31, G22

*I am grateful to Giovanni Mastrobuoni, Elena Esposito and Solomon Hsiang for invaluable guidance and support. Special appreciation is extended to Tatyana Derugyna for a collaboration which greatly enriched this work. I also want to acknowledge the insightful feedback and engaging discussions with Paolo D'Odorico, Enrico Moretti, Giorgio Chiovelli, Joseph S. Shapiro, Jeffrey Wooldridge, Joshua Graff Zivin, Matthew Gordon, Daniele Girardi, Òscar Jordà, Guglielmo Zappalà, Juan Vargas, Yagan Hazard, Joel Ferguson, Aleksandar Andonov, people of the Ibeo Workshop 2023, of the ECCD-GSSI Workshop 2024, of the Global Policy Lab (Berkeley), of the ARE dept. (Berkeley) seminar, of the EIEE (Milan). Any remaining errors are solely my responsibility.

†Amsterdam Business School, University of Amsterdam. Email: d.castellani@uva.nl

1 Introduction

According to the National Oceanic and Atmospheric Administration (NOAA), since 1980 the cost of damage from weather and climate disasters has exceeded $2.6 trillion. Of all recorded weather disasters in US history, cyclones have caused the most damage: more than $1.3 trillion in total, with an average cost of $22.8 billion per event. The expenditure attributed to cyclones in the United States can fluctuate significantly from year to year, depending on the severity and frequency of the storms. For example, in 2020, the United States experienced a particularly active cyclone season, with a record-breaking 30 named storms, including 13 hurricanes.[1] Beyond their immediate destruction, such events can reshape local economies through large-scale damage to housing and infrastructure, population displacement, and reconstruction activity. These dynamics raise important questions about how tropical cyclones influence housing markets in the long run, a central focus of this paper.

The largest economic losses from climate change in coastal areas are expected to arise from the intensification of cyclones and sea-level rise (Hsiang et al., 2017). A substantial share of these damages stems from residential property losses, which are particularly important given the central role of housing in household wealth and the U.S. economy: in 2023, the real estate sector accounted for 12.6% of U.S. GDP. [2] According to the U.S. Census, about 18.5% of the U.S. population resides in counties along the Atlantic and Gulf coasts, regions highly exposed to tropical cyclones. While coastal locations offer significant amenities, they also entail increasing climate risk. Existing evidence on how such risks affect housing markets is mixed, both for sea-level rise (Hino and Burke, 2021; Bakkensen and Barrage, 2021; Murfin and Spiegel, 2020; Bernstein et al., 2019; Giglio et al., 2021b) and for individual hurricane (Biny and Polasky, 2004; Ortega and Taşpınar, 2018; Deryugina et al., 2018). The most comprehensive studies to date document positive effects of hurricanes on local housing markets. Graff Zivin et al. (2023) document short-run increases in Florida housing values following hurricanes, while Roth Tran and Wilson (2025) provide evidence that FEMA-declared natural disasters boost local income, wages, and home prices over horizons up to eight years across U.S. counties, consistent with federal aid enabling recovery beyond pre-disaster levels (Hsiang and Jina, 2014). However, their analysis does not limit attention to counties with a comparable ex ante likelihood of experiencing a tropical cyclone, predominantly coastal counties, even though coastal counties have been subject to very different trends over the last 20 years. This papers shows the importance of framing the identification as if otherwise identical counties are randomly assigned to either treatment or control weather conditions.

This paper quantifies the effects of tropical cyclones on housing values across all U.S. coastal counties and demonstrates how methodological choices drive divergent findings. I focus the main analysis

[1] Hurricanes are intense tropical cyclones characterized by sustained wind speeds of at least 74 mph

[2] In 2023, the U.S. GDP attributable to the housing market was $3.447 trillion out of a total GDP of $27.36 trillion. Source: Data retrieved from the Federal Reserve Economic Data (FRED) at `https://fred.stlouisfed.org/series/USREALNGSP`

on coastal counties, defined as those bordering the Atlantic Ocean or the Gulf of Mexico, not only because they experience more frequent cyclone events, but also because the damage tends to be more severe. Coastal areas are particularly vulnerable to storm surge, where rising sea levels driven by strong winds, low pressure, and storm motion can flood coastal zones (Marsooli et al., 2019). Storm surge amplifies cyclone impacts by combining extreme winds with rising water levels, leading to extensive flooding and damage to infrastructure and ecosystems. Given this heightened vulnerability, emphasizing coastal regions is critical when assessing the economic impacts of cyclones (Booth et al., 2021). Because storm surge is fundamentally driven by the wind field, maximum wind speed is a key determinant of coastal flooding intensity and provides a consistent proxy for surge-related impacts (Islam et al., 2021). Using HURDAT2 data from the National Hurricane Center in combination with a parametric model of the radial profile of hurricane winds (Holland et al., 2010), I construct a comprehensive dataset covering all recorded cyclones from 1992 to 2021. The cyclone dataset is merged with county-level housing price and quantity series to trace housing demand and supply dynamics before and after cyclone exposure. This captures both short-run market adjustments and longer-run persistence. The analysis then incorporates insurance and labor-market variables to explore the underlying mechanisms. Building on the empirical evidence, the paper develops a theoretical framework that highlights the dual role of insurance markets in disaster-prone areas: raising the cost of living while enhancing resilience by facilitating post-disaster recovery.

The overall effects of cyclones on housing markets are theoretically ambiguous. Cyclones simultaneously affect housing supply and demand. On the supply side, they can destroy or damage homes, leading to shortages and potentially driving up prices in the short term. On the demand side, cyclones disrupt local economic activity, affect income levels, and can reduce the attractiveness of affected areas, particularly if individuals revise their perception of future cyclone risks. In addition, outmigration from hurricane-prone areas can exacerbate the decline in housing prices. The overall impact of these competing forces is ultimately determined by empirical analysis. Empirically estimating cyclone impacts, however, poses several challenges. First, measuring exposure is complex, as wind speeds vary with distance from the cyclone center. To address this, I develop a new wind speed measure that accounts for multiple radii, mitigating measurement error. Second, identification is difficult because U.S. counties face heterogeneous cyclone risks, complicating causal inference for housing markets. The research design tackles this by exploiting plausibly exogenous variation in exposure, ensuring treated and control counties have comparable ex-ante cyclone probabilities. Specifically, I restrict the sample to coastal counties and applies propensity score weighting. A third challenge involves "unclean" comparisons, where previously treated counties may serve as controls for later events, biasing standard event-study estimators (Sun and Abraham, 2021). To resolve this contamination, the analysis employs local projections in a difference-in-differences event-study framework (Dube et al., 2025), which enforces a clean-control condition by excluding observations potentially influenced by prior treatments.

The main finding shows that exposure to cyclones leads to a substantial 15% decline in housing values over eight years. This long-term decline is far more persistent than the effects of other natural disasters, such as wildfires (McCoy and Walsh, 2018). Considering the mean value of houses in coastal counties, $254,079, counties experiencing cyclones witness an average decrease of around $38,000 in property value after 8 years. In the rental market, the negative effects are more transitory. From the supply side, the total housing quantity decrease in the aftermath of cyclones, explaining an initial short-term rise in prices.
The paper finds that tropical cyclones do not generate significant spillover effects on housing values in nearby "nearly missed" counties, suggesting that neither migration nor perceived risk materially affect neighboring markets, in contrast to prior single-county evidence (Hallstrom and Smith, 2005).

A second key finding concerns the adaptation capacity of high-risk areas and the mechanisms driving post-cyclone housing dynamics. Within coastal regions, residents increasingly sort into historically low-exposure counties, expanding local housing markets and pushing up prices there; yet these same areas experience persistent price losses following severe events. High-risk counties, by contrast, display a rebound in values after the initial decline, suggesting greater adaptive capacity. To identify the mechanisms behind these patterns, the same empirical strategy examines labor, population, and insurance markets. Labor-market adjustments are economically meaningful, a roughly 3 percent decline in household income after four years and a 10 percent employment, but is only indirectly connected to housing values, offering limited explanatory power for either the magnitude of the price declines or the rebound in high-risk areas. Instead, flood insurance emerges as a more direct housing-related mechanism. Following major cyclones, insurance uptake rises by about 10 percent and average premia increase by roughly 5 percent, directly raising the cost of living in risky areas. This insurance response helps explain both the price declines (via higher insurance premia) and the rebound in high-risk counties (via greater baseline coverage enabling reconstruction). I create a model of location, housing, and insurance choices to rationalize these patterns.

This paper makes substantial contributions. The main contribution focuses on elucidating the effects of physical climate risk on financial assets (Giglio et al., 2021a), particularly on the housing market. There is a substantial body of literature on the economic impacts of extreme weather events, reviewed by Ferreira (2024). For example, Strobl (2011) finds that, in U.S. coastal counties, an average hurricane reduces annual economic growth by 0.45 percentage points in the year of impact. When it comes to the housing market specifically, however, the evidence is more fragmented. Despite a multitude of single-event studies exploring how hurricanes affect housing prices, the findings have so far been inconclusive, suggesting a need for more comprehensive and systematic approaches. For example, reductions in the market value of homes were observed by Biny and Polasky (2004) after Hurricane Floyd, by Ortega and Taşpınar (2018) and Addoum et al. (2024) after Hurricane Sandy. On the other hand, Deryugina et al. (2018) and Vigdor (2008) found that Hurricane Katrina led to a rise in house prices in directly affected areas. Graham et al. (2007), Beracha and Prati (2008) found no discernible effect on house prices when examining larger areas. Graff Zivin et al. (2023)

examined the mixed findings on the impact of hurricanes on the housing market using data from the Florida housing market from 2000 to 2016. They found that housing prices increased significantly for three years after a hurricane. However, their results did not extend to other states or all tropical cyclones, particularly in the long-run. In related work, Roth Tran and Wilson (2025) use nearly four decades of U.S. county data to estimate local projection impulse responses to FEMA-declared natural disasters and show that, on average, personal income per capita, wages, and house prices are higher up to eight years after a disaster, while employment and population remain roughly unchanged. However, their analysis does not explicitly separate disaster effects from underlying differences in long-run house price trends, even though many disaster-exposed counties are coastal areas where amenity-driven demand has been pushing prices up over time. By contrast, this paper restricts the sample to coastal counties and uses an inverse-probability weighting approach. Under this more comparable coastal control group, I find negative long-run effects on housing values, suggesting that once one compares like-for-like coastal markets, tropical cyclones reduce long-run housing wealth.

By selecting a set of comparable control units and aggregating across events, I show that the effects of cyclones on housing are systematically underestimated in single-event study designs. This finding contributes to a growing literature arguing that conventional regressions of local climate shocks on local outcomes fail to account for spatial spillovers, general equilibrium adjustments, and dynamic amplification mechanisms, thereby biasing estimated damages downward (Bilal and Känzig, 2026). More broadly, recent evidence suggests that climate damages may be an order of magnitude larger once nonlinearities, long-run dynamics, and macroeconomic linkages are incorporated. For example, Burke et al. (2015) show that temperature increases generate large nonlinear productivity losses at the macro level. Similarly, Carleton et al. (2022) estimate temperature-related mortality damages that substantially exceed earlier benchmark calculations. In the context of extreme events, Young and Hsiang (2024) demonstrates that incorporating longer-run mortality effects of hurricanes dramatically revises damage estimates upward relative to short-run assessments. Consistent with this perspective, benchmarking a 15% decline after 8 years in housing values is challenging because my estimates incorporate longer time horizons and broader equilibrium responses than most existing work. Taken together, these findings suggest that localized, short-run empirical designs may substantially understate the true economic costs of natural disasters.

This paper contributes also to the extensive literature on climate change adaptation to natural disasters (Carleton et al., 2024; Hsiang and Narita, 2012a). Adaptation research distinguishes between infrastructure solutions (Benetton et al., 2022), behavioral responses, such as migration and insurance uptake (Blonz et al., 2025; Gallagher, 2014), and institutional measures, including building codes and risk disclosure (Hino and Burke, 2021; Wang, 2014). Focusing on the mechanisms behind housing-market responses, I show that cyclone exposure induces a large and persistent increase in flood insurance take-up, while effects on other local economic outcomes, such as income and employment, are comparatively modest, consistent with Gallagher (2014). This pattern points

to cost capitalization rather than local economic contraction as the primary channel. Prior studies document that insurance coverage is a key conduit through which climate risk enters housing valuations, with changes in flood insurance premia capitalized into home prices (Keys, 2025; Indaco et al., 2019; Speyrer and Ragas, 1991; Georgic and Klaiber, 2022). This paper contributes to the literature by showing that, while expanded insurance coverage increases the recurring cost of living in high-risk areas through higher premiums, it also enhances household and community resilience. I develop a theoretical model of location choice and housing prices that emphasizes the dual role of insurance: it facilitates post-disaster recovery and reconstruction while simultaneously shaping housing valuations through mechanisms such as adverse selection and delayed uptake in lower-risk areas Wagner (2022). In high-exposure counties, households hold subjective beliefs that are closer to the true risk and maintain higher insurance coverage, which enables rebuilding at similar or higher quality after a cyclone. Consistent with this mechanism, housing prices in high-risk areas tend to rebound following exposure, whereas lower-probability (less-exposed) counties experience larger and more persistent declines.

Moreover, I offer a methodological advancement in the context of climate change impact assessment, particularly in the construction of a damage function to associate economic costs to cyclones. Building on the exposure measures developed by Hsiang and Jina (2014) and Deryugina (2017), this study refines a literature that has focused on constructing cyclone damage functions (Nordhaus, 2010; Hsiang, 2016; Bakkensen et al., 2018; Hsiang et al., 2017). Unlike Deryugina (2017), which used the "maximum wind speed radius" to identify affected counties, this study employs a higher resolution approach that incorporates multiple radii from the cyclone. While some wind observations are available, maximum wind speeds are not directly observed at the county level with sufficient spatial and temporal resolution. Observational stations are often sparse, unevenly distributed, particularly for counties located far from measurement stations. To address this, I estimate wind speeds using an interpolation of the cyclone's radial wind profile (Holland et al., 2010), incorporating central pressure, pressure gradients, cyclone intensity, and environmental factors. This method produces a refined, time-specific spatial distribution of wind speeds across counties, reducing misclassification of affected versus control areas and minimizing measurement error. This refinement facilitates a more nuanced analysis of cyclone impacts, capturing greater variability and including a broader set of affected counties. Compared to prior studies (Deryugina, 2017; Graff Zivin et al., 2023), this study also includes all cyclones between 1992 and 2020, rather than restricting the analysis to hurricanes (cyclones with minimum wind speeds of 33 m/s). This is crucial to account for events like Hurricane Sandy, which had transitioned to a post-tropical cyclone upon reaching New York yet caused substantial damage (Ortega and Taşpınar, 2018).

The rest of the paper proceed as follow. Section 2 describes data and the new methodology for constructing the cyclone trajectories. Section 3 introduces the research design to identify the impact of cyclones on housing and other economic variables. Section 4 presents the results. Section 5 provides the conclusions.

2 Background and Data

2.1 Economics of natural disasters

The quality of previous estimates is affected by the endogenous nature of their independent variables: self-reported disaster counts and losses that are usually from the Emergency Events Database (EM-DAT). The quality and completeness of these self-reported measures are known to depend heavily on the economic and political conditions in a country (Kahn, 2005; **?**; Kellenberg and Mobarak, 2008; Noy, 2009; Hsiang and Narita, 2012b), factors that also affect growth and therefore can confound these results. Similarly, issues are observed in data from the Federal Emerging Management Agency (FEMA). I overcome the challenges of omitted variables bias and endogenous disaster reporting by developing novel data that describe year-to-year variation in the physical exposure of each country to disaster. To do this, I focus on tropical cyclones, the class of natural disasters that includes hurricanes and tropical storms with at least 17 m/s wind speed, and I reconstruct every storm observed in the contiguous US from 1992 to 2020. Unlike the self-reported statistics contained in EM-DAT or FEMA, my objective measures of wind speed exposure are fully exogenous, constructed using physical parameters and meteorological observations, so they are unlikely to be influenced by economic behavior or political actions within each state or county.

2.2 A new measure of wind speed for cyclones

Tropical cyclones that affect the United States form over warm waters of the Atlantic Ocean, where the sea surface temperature is at least 26.5°C to provide the necessary heat and moisture. For a cyclone to develop, the atmosphere requires warm, moist air near the surface and cooler air aloft, creating atmospheric instability. The development of a cyclone progresses through several stages. It begins as a tropical disturbance, which is a cluster of thunderstorms. Various factors, such as the convergence of wind patterns, can initiate these disturbances. As the system becomes more organized and shows signs of circulation, it is classified as a tropical depression, with sustained wind speeds reaching up to 38 mph (17 km/h). When wind speeds increase to 39-73 mph (18-33 m/s), the system becomes a tropical storm and receives a name. Once the wind speeds exceed 74 mph (>33 m/s), the storm is classified as a hurricane. Here, I consider all the observations with at least 17 m/s of maximum sustained wind speed.

Tropical cyclones are difficult to predict beyond a few days. This inherent randomness is crucial for my analysis, as it allows me to identify the causal effects of cyclones based on the unpredictable year-to-year variation in their intensity, timing, and duration at specific locations (Deschênes and Greenstone, 2007).

To accurately pinpoint the areas affected by cyclones, I introduce a novel measure of wind speed. [3] I combine data from the second generation Hurricane Data dataset (HURDAT2) from the National Hurricane Center [4] with the parametric modeling approach of Holland et al. (2010). Compared to previous measures used by Deryugina (2017) and Hsiang and Jina (2014), this measure accounts for variations in wind intensity at different radii from the cyclone center, improving spatial precision. The HURDAT2 dataset provides six-hour data for each North Atlantic cyclone, including information on the storm's central location, maximum wind speed, radius of maximum wind speed, and the radial extents of specific wind thresholds (34, 50, and 64 knots) across the four quadrants of the cyclone. This comprehensive dataset forms the basis for estimating wind speeds. To refine the wind speed estimates at varying distances from the center of the cyclone, I adopt the methodology of Holland et al. (2010), which describes the radial wind profile of a hurricane using the following parametric equation:

$$v_s = v_{ms} \left(\left(\frac{r_{v_{ms}}}{r} \right)^{b_s} \exp \left[1 - \left(\frac{r_{v_{ms}}}{r} \right)^{b_s} \right] \right)^{x} \tag{1}$$

Here:

- v_s is the wind speed at a distance r from the center of the cyclone.
- v_{ms} is the maximum wind speed and $r_{v_{ms}}$ is the radius of maximum wind speed, both reported in HURDAT2.
- The parameter b_s determines the rate of wind decay with distance and x adjusts wind speeds at larger distances using the wind thresholds provided by HURDAT2.

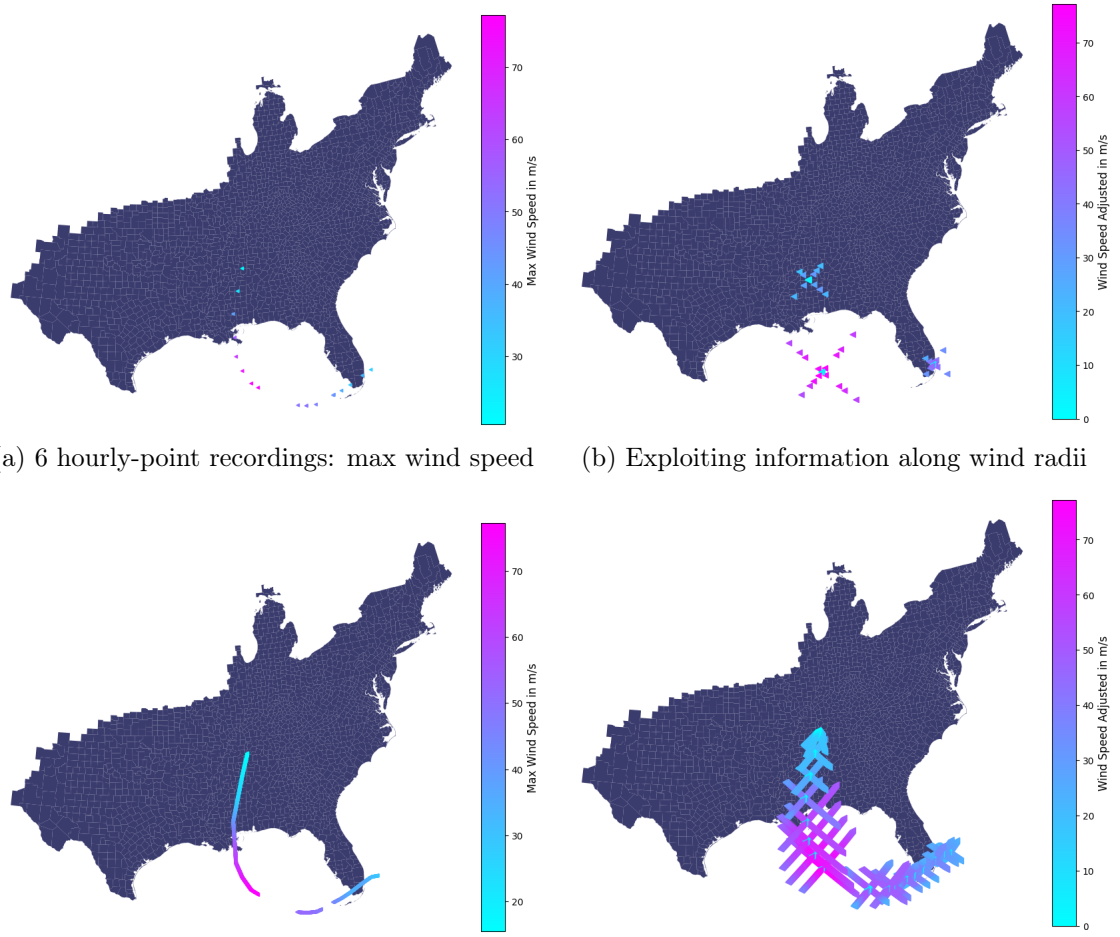
For regions beyond $r_{v_{ms}}$, I adjust the x parameter using the wind thresholds (34, 50, and 64 knots) and iteratively solve Equation 1 to match the observed wind radii. Specifically, I derive x by solving the equation for distances corresponding to the known wind radii (e.g., r_{34}, r_{50}, r_{64}). For detailed derivations of b_s, x, and their calibration, refer to Appendix A.

Having described Holland's (2010) parametric model, I demonstrate its application using Hurricane Katrina, the most damaging hurricane in U.S. history. To visually depict the reconstructed wind speed profile of Hurricane Katrina, I present various interpolations in Figure 1.

[3] This measure was developed as part of a short-term collaboration between myself and Professor Deryugina, with the mutual understanding that each of us could subsequently use the results independently in our respective research projects

[4] Available at `https://www.nhc.noaa.gov/data/`

Figure 1: Reconstruction of wind speed for Hurricane Katrina



(a) 6 hourly-point recordings: max wind speed

(b) Exploiting information along wind radii

(c) Interpolations along the trajectory

(d) Interpolations along the wind radii

Notes: Panel 1a displays the 6-hourly recording points for maximum wind speed, v_{ms}, as defined in Equation 1, excluding observations located far from the coast or with wind speeds below 17 m/s. Panel 1b presents three examples of recording points, along with the northeastern, northwestern, southwestern, and southeastern radii, where wind speeds reach at least 17.49, 25.72, or 32.9 $\frac{m}{s}$, respectively. Unlike in Panel 1a, where the focus is on maximum wind speed, this panel considers wind speed variation across quadrants. Specifically, I plot the "wind speed adjusted" v_s, as described in Equation 1. In particular, v_s is calculated at different distances from the cyclone center: at the cyclone center itself (where wind speed is lowest), at the radius of maximum wind speed (where wind speed peaks), and at the outermost extents of the 34, 50, and 64 kt wind radii. To account for gaps between the 6-hourly recording points and variations in the maximum wind radii, I apply linear interpolation—along the cyclone's main trajectory (Panel 1c) and across the different wind extents (Panel 1d).

Panel 1a shows the 6-hourly recording points for the maximum wind speed, v_{ms} in Equation 1, excluding all the observation that are far away from the coast or with a wind speed below 17 m/s. These recording points represent the center of the cyclone and they come also with other data, like the minimum pressure p_{min} and the radius of maximum wind speed $r_{v_{ms}}$, which are essential for reconstructing the true velocity at a certain distance r from the recording point. Indeed, the true

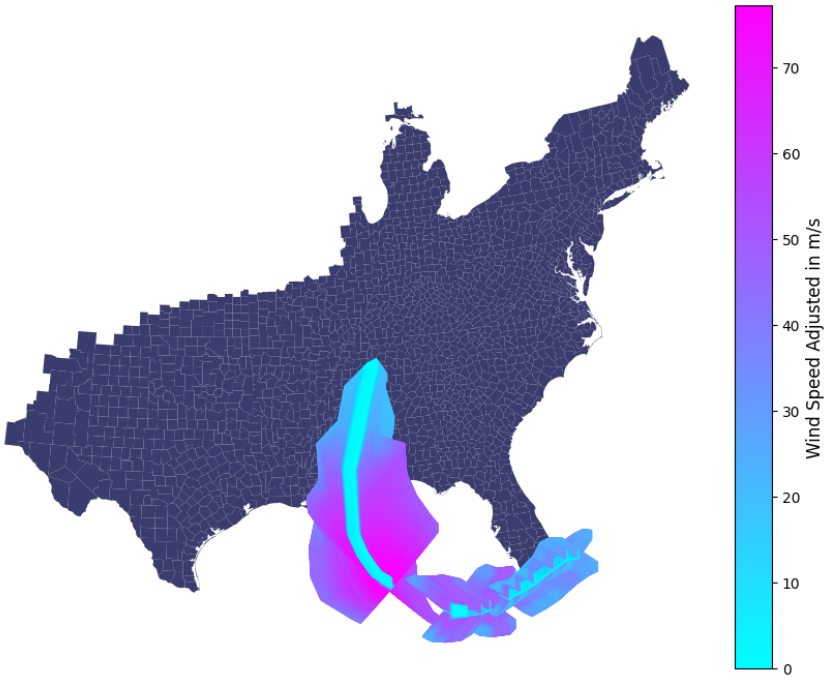
wind speed at the center of the cyclone would be close to zero: as first noted by Leonardo Da Vinci, all intense vortices are characterized by inward exponentially increasing winds to a maximum (close to the radius of maximum wind speed $r_{v_{ms}}$), then a rapid drop to a calm center. Additionally, HURDAT2 contains the maximum extents of the 34, 50 and 64 kt wind radii in the northeastern, northwestern, southwestern, and southeastern quadrants. These values indicate the distances in each of the four directions where the wind speed is at least 17.49, 25.72, or 32.9 $\frac{m}{s}$, respectively. This data is crucial for calibrating parameter x in equation 1.

In Panel 1b, I plot this information for three recording points that are distant from each other to avoid overlapping observations and to provide a clear picture of the main information contained in each point. For each quadrant, I plot the "wind speed adjusted" v_s as described in Equation 1. These plots include v_s at the center of the cyclone (where wind speed is low), at the radius of maximum wind speed (where the wind speed reaches its peak), and at the three maximum extents of the 34, 50, and 64 kt wind radii.

To fill the gaps between the 6-hourly recording points and between the different recordings of the maximum wind radii, I use linear interpolation respectively along the main trajectory (Panel 1c) and along the different wind extents (Panel 1d). The number of new points depends on the distance between two consecutive recording points: I create 20 points per kilometer.

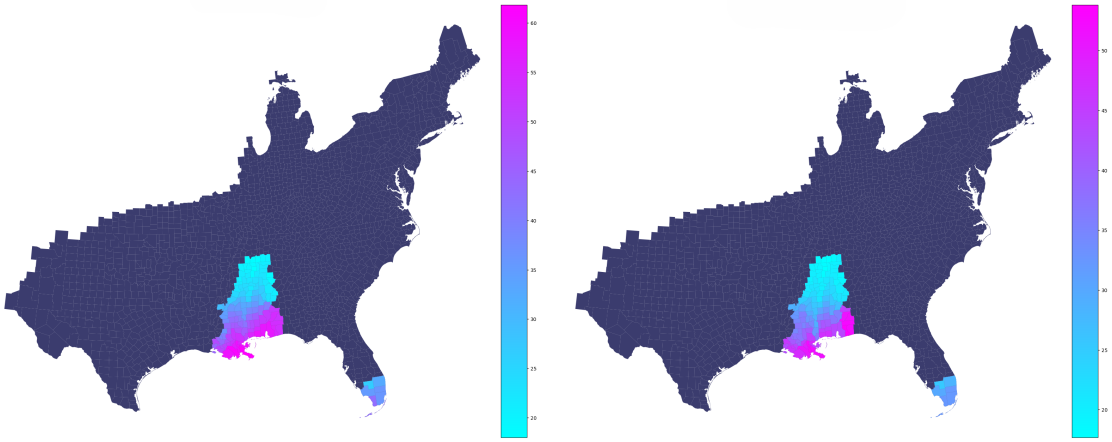
Finally, the wind speed for Hurricane Katrina is reconstructed by combining the two interpolations shown in Panels 1c and 1d.

Figure 2: Final reconstruction of Hurricane Katrina



Each location may exhibit different overlapping wind points due to the dynamic trajectory of the cyclone. In order to estimate the wind speed at the county level, I have taken into account the maximum wind speed experienced by each county.

Figure 3: Aggregate the wind speed at county level

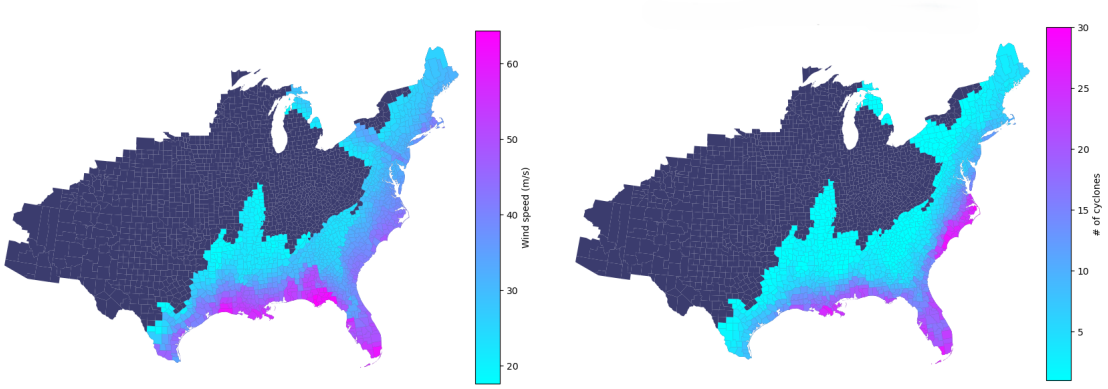


(a) Maximum wind speed of Katrina by county (b) Average wind speed of Katrina by county

Hurricane Katrina stands out as one of the most devastating natural disasters in US history, resulting in 1,833 fatalities and causing approximately $108 billion in damages. Despite their significant impact, there has been a lack of comprehensive analysis of the numerous cyclones that have occurred since 1992. In this study, I collect information on all 126 cyclonic events in the US during this period. The wind speed in each location for each event has been reconstructed following the procedure described for Hurricane Katrina.

In Figure 4, Panel a provides a visual representation of the maximum wind speed experienced by a county throughout the entire period from 1992 to 2021. Panel b illustrates the frequency of cyclones experienced in each location. For a further visualization of variability by each year, see Appendix B.1.

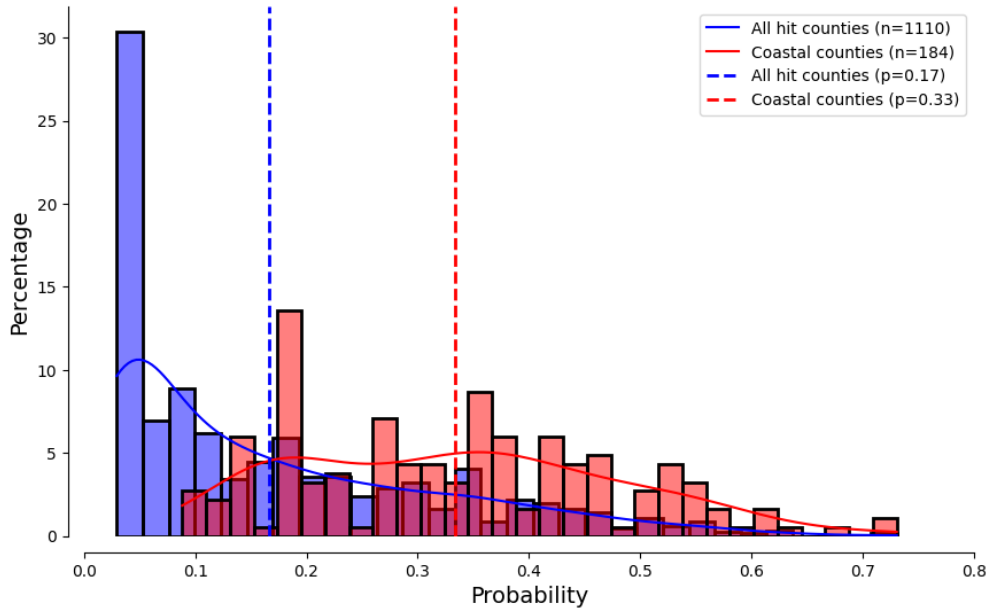
Figure 4: Aggregate the wind speed at county level



(a) Maximum wind speed of cyclones by location (b) Frequency of cyclones by location

To estimate the likelihood of cyclone exposure, I plot two distributions of the probability of experiencing at least one cyclone between 1992 and 2021: one for all counties that have recorded at least one cyclone (defined as events with wind speeds exceeding 17 m/s, as shown in Figure 4), and another restricted to coastal counties. Figure 5 displays these distributions.

Figure 5: Probability of receiving cyclones from 1992 to 2021



Note: The blue distribution represents the probability of experiencing at least one cyclone for all counties that have ever been hit, while the red distribution is limited to coastal counties. Dashed lines indicate the average probability for each group. For each county, the probability p is calculated as the number of years in which a cyclone occurred divided by 33, the total number of years in the panel (1992–2021). Median values are close to mean value for coastal area.

The sample of coastal counties is smaller (184 counties) compared to the full set of hit counties (1,110 counties). The distribution for coastal counties is more tightly clustered around the mean of 0.33, indicating a relatively homogeneous exposure pattern. In contrast, the distribution for all counties is left-skewed with a lower mean of 0.17, reflecting a more irregular pattern of cyclone exposure. This means that coastal counties, on average, experience cyclones more frequently and consistently than other counties. For the heterogeneous analysis based on event frequency in the Results section, I use a threshold of 0.27 instead of 0.33. This lower cutoff reflects the median value derived from the full distribution of probabilities, which accounts for variation in the availability of housing price data over time across counties. [5] Adopting this threshold allows me to include a larger number of counties in the high-hit group, as areas with fewer events tend to have longer housing time series and thus more observations.

I focus the main analysis on coastal counties, defined as those bordering the Atlantic Ocean or the Gulf of Mexico, not only because they experience more frequent cyclone events, but also because the damage tends to be more severe. Coastal areas are particularly vulnerable to storm surge, where rising sea levels driven by strong winds, low pressure, and storm motion can flood coastal zones (Marsooli et al., 2019). Storm surge amplifies cyclone impacts by combining extreme winds with rising water levels, leading to extensive flooding and damage to infrastructure and ecosystems. Given this heightened vulnerability, emphasizing coastal regions is critical when assessing the economic impacts of cyclones (Booth et al., 2021).

Note also that among the 184 coastal counties, there are no "never-treated" units: every county experiences at least one cyclone during the sample period. A treatment probability of 0.33 for coastal counties implies that, on average, they are hit by a cyclone roughly once every ten years between 1992 and 2021. Moreover, many counties are exposed to multiple events, which poses additional challenges for identification due to overlapping treatments. I discuss these issues in detail in the research design.

2.3 Housing data

The Federal Housing Finance Agency Housing Price Index (FHFA HPI) is the primary measure that I use to assess the impact of cyclones on the housing market.[6] The FHFA HPI operates as a weighted repeat sales index at the county level and captures average price fluctuations in properties through repeated sales or refinancing, keeping a focus on the same properties over time. The data originate from an extensive review of repeat mortgage transactions on single-family properties whose mortgages have been acquired or securitized by Fannie Mae or Freddie Mac since January 1975. Crucially, the repeat sales approach mitigates the potential influence of shifts in sample composition,

[5] Median and mean value for coastal area are very closed since distribution is normally distributed.

[6] Available at `https://www.fhfa.gov/DataTools/Downloads/Pages/House-Price-Index.aspx`

ensuring that changes in the estimated index primarily reflect market dynamics rather than shifts in the types of property included. By considering values of the same physical units at different points in time, the FHFA HPI effectively controls for variations in housing attributes across the sample, without necessitating a direct estimate of their individual contributions to the overall index. The Housing Price Index (HPI) starts at a value of 100 for the initial period when a county enters the dataset. Consequently, I focus the analysis on the year-over-year change in prices, denoted as ΔHPI, which captures variations relative to the previous year. I considered data from 1992 to 2021 to match it with cyclone data.

I include the Zillow Value Housing Index (ZVHI) [7] as an additional metric. The ZVHI is designed to reflect the value of a standard property on a national or neighborhood scale. Unlike indices focused solely on sold homes, the ZVHI provides a comprehensive perspective by considering data from the entire spectrum of homes within a specific region. The Zestimate, a key component of the Zillow Value Housing Index (ZVHI), is derived using a set of sophisticated machine learning models. This estimation process integrates various data sources, including public datasets, user-generated input, and real estate data from direct feeds or multiple listing services. The comprehensive nature of the Zestimate extends beyond mere numerical evaluations, as it takes into account specific home attributes and geographic locations. By incorporating a wide array of information, Zillow's Zestimate aims to provide a nuanced and accurate reflection of property values, contributing to a more holistic understanding of the real estate landscape. The ZVHI is at the monthly and county level and gives a value in dollars. The ZHVI is available from 2000, and I have matched it with cyclone events through 2021.

To complement the analysis of housing prices, I examine the Zillow Rent Index (ZRI) , which offers a timely and demand-sensitive measure of local housing market conditions. Rents are often considered a more immediate reflection of market dynamics than home values, as they adjust more quickly to changes in supply and demand and are less influenced by long-term financial arrangements or transaction frictions. The ZRI captures the median estimated monthly rent for all housing units—regardless of whether they are currently listed—across property types, including apartments, single-family homes, and condominiums. Reported at the county and monthly level in nominal dollars, the ZRI spans from 2015 to 2021. While coverage is limited to a smaller subset of coastal counties (only 122 instead of 180)— reducing the statistical power of the analysis relative to the housing value data—the index still offers valuable insights into short-term, month-to-month rental market responses following cyclone events.

In addition to housing prices and rents, I analyze two datasets to capture the quantity and dynamics of the housing stock. First, I use annual county-level estimates of the total number of housing units from the U.S. Census Bureau.[8] This measure includes all separate living quarters intended

[7] Available at `https://www.zillow.com/research/data/`

[8] Data obtained from the U.S. Census Bureau, available at `https://www.census.gov/programs-surveys/popest.html`

for residential use—such as single-family homes, apartments, and mobile homes—and serves as a comprehensive snapshot of the existing housing stock. However, these data are available only from 2000 onward, resulting in a more limited time frame and smaller sample compared to the housing price series. Second, to assess more immediate construction activity and local housing market responses, I incorporate annual county-level data on final residential building permits from the U.S. Department of Housing and Urban Development (HUD). [9] This dataset, which spans from 1980 to the most recent reporting year, captures official authorizations for new residential construction.

2.4 Other variables

To examine the impacts of cyclones on other variables, I use data sourced from different organizations. Employment, income, and population statistics are provided by the U.S. Bureau of Economic Analysis (BEA), which offers detailed economic insights at the county level. Employment data highlight labor market dynamics and regional economic productivity, income metrics reflect individual and household economic well-being, and population figures provide a foundation for understanding demographic trends and regional composition.

Flood insurance policy data are sourced from the Federal Emergency Management Agency (FEMA) through the National Flood Insurance Program (NFIP) [10]. This variable captures the number of flood insurance policies issued, proxying for a county's flood vulnerability and residents' preparedness. Notably, cyclones often generate heavy rainfall alongside storm surge and winds, heightening inland flooding risks beyond coastal zones and underscoring NFIP relevance. The National Flood Insurance Program (NFIP) is a federal initiative established by the U.S. Congress in 1968 to improve the availability and affordability of flood insurance in the United States. In principle, homeowners with federally backed mortgages in high-risk areas (as designated by Flood Insurance Rate Maps) must obtain flood insurance. However, Gallagher (2014) estimates that 97% of NFIP policyholders buy flood insurance voluntarily rather than under mandate.

2.5 Descriptive statistics

This subsection provides descriptive evidence on the characteristics of counties exposed to cyclones, with a particular focus on differences between coastal areas and the degree of cyclone risk they face. Table 1 summarizes key housing market indicators, such as housing price inflation, housing values, rental prices, housing units, and new permits, for all counties that have experienced at least one cyclone, for coastal counties, and for coastal counties with high and low historical exposure to cyclones. These statistics offer an initial overview of how housing markets vary across geographies

[9] Data obtained from the HUD State of the Cities Data System, available at `https://catalog.data.gov/dataset/residential-construction-permits-by-county`

[10] Data ara sourced from https://www.fema.gov/openfema-data-page/fima-nfip-redacted-policies-v2

with different levels of cyclone risk and help motivate the subsequent empirical analysis.

Table 1: Descriptive statistics by geography and cyclone exposure

	All Counties with cyclones	**Coastal Counties**	**High Exposure Coastal (p > 0.33)**	**Low Exposure Coastal (p < 0.33)**
Inflation in HPI (π%HPI)	3.10%	3.60%	3.74%	3.45%
ZHVI	\$169,995	\$254,544	\$215,320	\$335,782
ZRI	\$1,415	\$1,572	\$1,436	\$1,796
Housing Units	52,632	130,258	96,001	164,515
New Housing Permits	651	1,455	1,400	1,508
New Flood insurances	2,279	9,963	12,402	5,475
Obs. (HPI)	30,659	5,797	2,805	2,945
Obs. (Permits)	29,646	5,569	2,775	2,794
Obs. (Housing Units)	22,200	3,680	1,840	1,840
Obs. (ZHVI)	19,198	3,673	2,477	1,196
Obs. (ZRI)	2,341	824	513	311
# of Counties	1,040	184	117	67

Notes: This table presents housing value and quantity yearly averages for counties experiencing at least one cyclone, for coastal counties (those Adjacent to Atlantic Ocean or Gulf of Mexico) and for those with different cyclone exposure within coastal counties between 1992 and 2021. ZHVI is the Zillow Home Value Index, and ZRI is the Zillow Rent Index. High cyclone exposure is defined as a probability of landfall greater than 0.33. Observations reflect non-missing values.

Table 1 reports summary statistics for housing market indicators in U.S. counties affected by cyclones, with a particular focus on coastal areas. On average, annual housing price inflation (π%HPI) is modestly higher in coastal counties (3.6%) than in the broader set of cyclone-affected counties (3.1%). Moreover, housing values—as measured by the Zillow Home Value Index (ZHVI)—are significantly higher in coastal areas, averaging \$254,079, compared to \$169,811 across all counties that have experienced at least one cyclone.

Coastal counties are further classified into high and low exposure groups based on whether their probability of experiencing a cyclone, computed as described in Section 5, is above or below the coastal mean from 1992 to 2021. Interestingly, the average Zillow Home Value Index (ZHVI) is substantially higher in low-exposure coastal counties (\$335,858) than in their high-exposure counterparts (\$215,295), which may reflect a general preference for areas perceived to be at lower risk from cyclones. A similar pattern is observed for rental prices, with average rents (ZRI) higher in low-exposure coastal areas, suggesting that both homeowners and renters may favor locations perceived as safer from cyclone-related risks. This pattern extends to the housing stock as well: low-exposure coastal counties have, on average, a significantly larger number of housing units than high-exposure

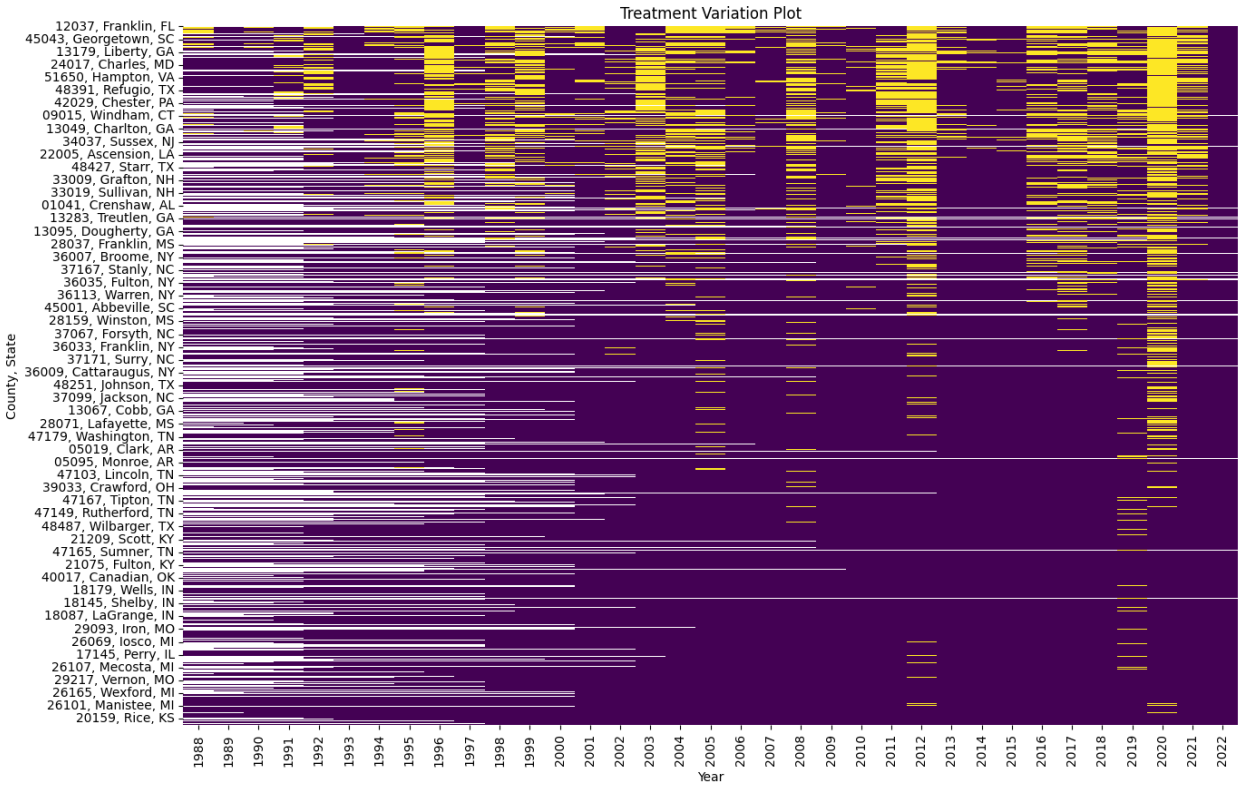
ones, indicating that both market values and residential development may concentrate in less risky areas. Insurance behavior moves in the opposite direction. When scaling the number of new flood insurance policies by the housing stock, high-exposure coastal counties exhibit a markedly higher insurance uptake than low-exposure counties. This suggests that households in riskier locations are more likely to insure against cyclone-related damages, consistent with greater exposure and risk awareness.

3 Research design

To investigate the relationship between the housing market and cyclones, I leverage a natural experiment. The quasi-random trajectory of cyclones provides a setting where the occurrence of a cyclone in a specific county and year can be treated as exogenous.

However, it is essential to acknowledge that not all U.S. counties share the same probabilities of being hit. In Figure 6, the yellow shading represents counties affected by cyclones in a specific year, while the purple shading indicates counties that are not exposed to the treatment. White areas denote regions with housing missing data.

Figure 6: Treatment variation plot



Note: Heatmap depicts cyclone impact on counties; Counties are arranged in ascending order based on their distance from the Atlantic coast, from top to bottom; White: housing market data missing, Violet: cyclone-free years, Yellow: cyclonic activity years. Because many counties have missing housing values before 1992, the analysis period begins in 1992.

The counties are listed in ascending order, from top to bottom, based on their distance to the Atlantic coast, with those at the top experiencing more cyclones than those at the bottom. In line with the definitions used by Deryugina (2017) and Strobl (2011), and as outlined in the data section, I restrict the main analysis on coastal counties, those bordering the Atlantic Ocean or Gulf of Mexico. Because many counties have missing housing values before 1992, the analysis period begins in 1992.

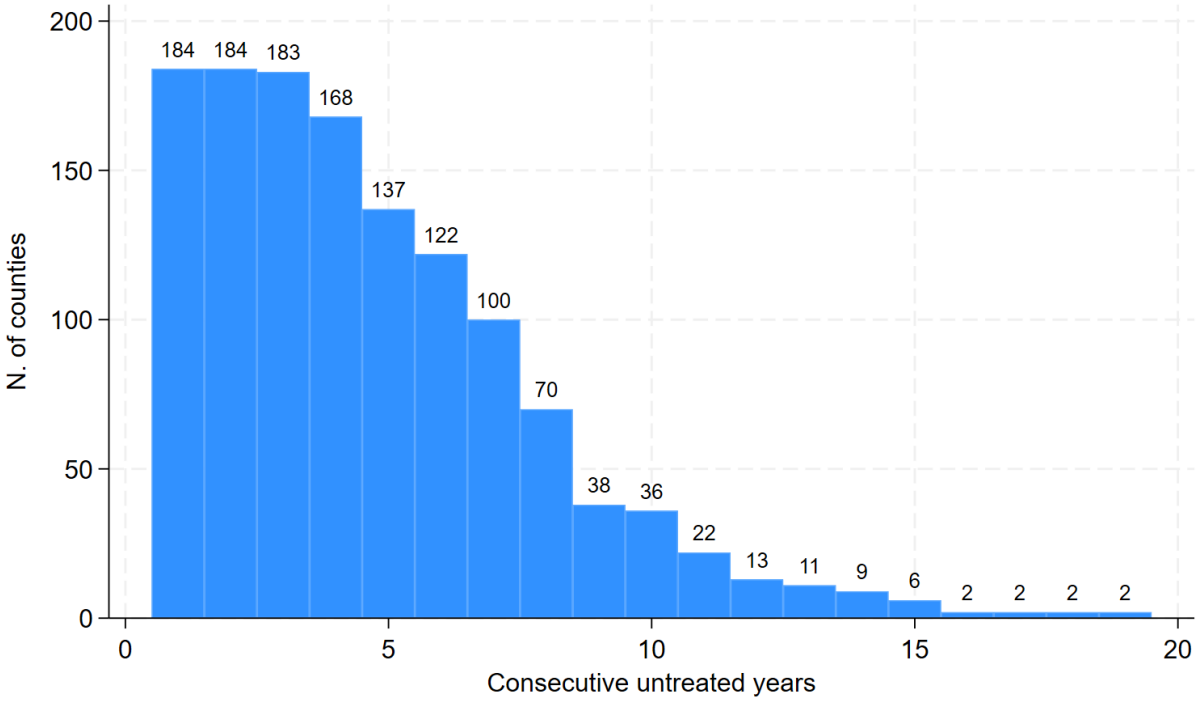
To address the varying probabilities of cyclones within the coastal area, I introduce inverse probability weighting into the regression analyses. This approach generates control observations with more comparable probabilities of experiencing a cyclone compared to the treatment observations. I compute the probability of a county being hit by a cyclone, p, considering the mean of the treatment variable for each county, given the long period available in the dataset. Then, following the Inverse Propensity Score Weighting method, I apply the following weights:

- $1/p$ for the treatment units

- $1/(1-p)$ for the control units

In this context, counties transition in and out of treatment status, and there are no "never-treated" coastal counties, as all experience at least one cyclone at some point between 1992 and 2021. To construct a credible control group within coastal areas, I focus on counties that remain cyclone-free for several consecutive years. This intuition is formalized by the "clean control condition" in Dube et al. (2025), as specified in equation 2. The condition ensures that control observations are free from recent storm exposure, thereby preventing contamination of the comparison group and yielding a more credible counterfactual. Figure 7 illustrates the distribution of these cyclone-free periods across coastal counties.

Figure 7: Consecutive cyclone-free years used for control group selection



Note: The figure displays the distribution of coastal counties by the number of consecutive years without a cyclone, where x ranges from 1 to 19. Because no counties are "never treated," the "cleanest" control group consists of the eight counties that go at least 15 consecutive years without a cyclone. As a broader control group, I use the 70 counties that experience no cyclones for at least 8 consecutive years.

This figure highlights a key methodological trade-off in constructing an appropriate control group. Longer cyclone-free streaks produce a "cleaner" control group, counties that resemble never-treated units as closely as possible. However, these long-streak counties are few in number, limiting statistical power. Using shorter cyclone-free thresholds expands the control group and increases precision, but at the cost of including counties that may be exposed more frequently.

To estimate the average treatment effect on the treated (ATT) of cyclone exposure on housing prices, I adopt the Local Projection Difference-in-Differences (LP-DiD) estimator proposed by Dube et al. (2025). This approach is well-suited for settings like mine, where treatment is non-absorbing, effects are dynamic, and all units receive treatment.

$$\Delta^h y_{i,t} = \delta_t^h + \beta_h^{LP-DID} \Delta D_{i,t} + e_{i,t}^h \tag{2}$$

In Equation 2, $\Delta^h y_{i,t}$ represents the log difference in housing prices between $t-1$ and $t+h$, and $\Delta D_{i,t}$ is the change in treatment status from $t-1$ to t. Here, $D_{i,t} = 1$ if county i experiences a cyclone in year t, so that $\Delta D_{i,t} = 1$ captures the occurrence of new cyclones in that period. The coefficient β_h^{LP-DID} captures the dynamic effect of cyclone exposure h years after the event.

To avoid "contamination effects" caused by prior exposures (Sun and Abraham, 2021) and ensure valid estimation of β_h^{LP-DID}, I restrict the sample following the criteria in Dube et al. (2025):

- 'clean control condition': Control units must be untreated throughout the interval from $t-K$ to $t+h$.
- Treated units must not have experienced flooding in the K years prior to treatment and must remain treated for the h years following it.

The first sample restriction ensures that the estimated effect reflects new cyclone exposure, free from confounding by prior storms. The LP-DiD estimator relies on the assumption that dynamic treatment effects stabilize after a fixed number of periods K^*. In the main specification, I set $K^* = 8$, meaning that a county must be untreated for at least eight consecutive years prior to being classified as newly treated. I also verify the robustness of the results by extending this window, $8 < K^* \leq 15$, allowing for even longer untreated periods, and find that the estimates remain consistent, suggesting the control group is sufficiently clean and the results are not sensitive to the choice of K^*.

One challenge with the second condition is that treated units may experience overlapping flood events within the h-year window following a new cyclone event. To address this, I isolate the subset of treated counties that experienced only a single flooding event. I compare these counties to those that either never experience a cyclone within 200 km of the coast (since there are not never-treated coastal counties) or have not yet been exposed . Moreover, I conduct a heterogeneity analysis in which I split the treated group based on the number of cyclones they receive, distinguishing counties below versus above the mean exposure level of 0.33 reported in Table 5.

As an additional robustness check, I employ the estimator proposed by de de Chaisemartin et al. (2024). However, their method only considers counties that have not yet received treatment as

controls, omitting those that do not experience a cyclone for a long period in the middle of the panel. This significantly reduces the number of observations. I also report Two-Way Fixed Effect (TWFE) and Distributed Lag Model estimates in the Appendix C, which provide useful benchmarks.

Another important assumption to identify the causal impact of cyclones on housing market is the Stable Unit Treatment Value Assumption (SUTVA) defined by Rubin (1990). SUTVA implies that the potential outcomes for each county i are not related to the treatment status of the other counties. I analyze spillover effects by taking into account both one's own treatment status and the treatment status of neighboring counties. To implement this spatial model, I enhance equation 2 by integrating the effects of cyclones on counties situated within a specified distance from the cyclone's path:

$$\Delta_h y_{i,t} = \delta_t^h + \sum_{L=-h}^{H} \beta_L^{LP} \Delta D_{i,t-L} + \sum_{L=-h}^{H} \theta_L^{LP} \Delta N_{i,t-L} + e_{i,t}^h \tag{3}$$

where the new term $N_{i,t-L}$ is a dummy variable equal one if county i in year t is at a maximum of 50km from cyclone (we also consider different thresholds in the appendix as a robustness check).

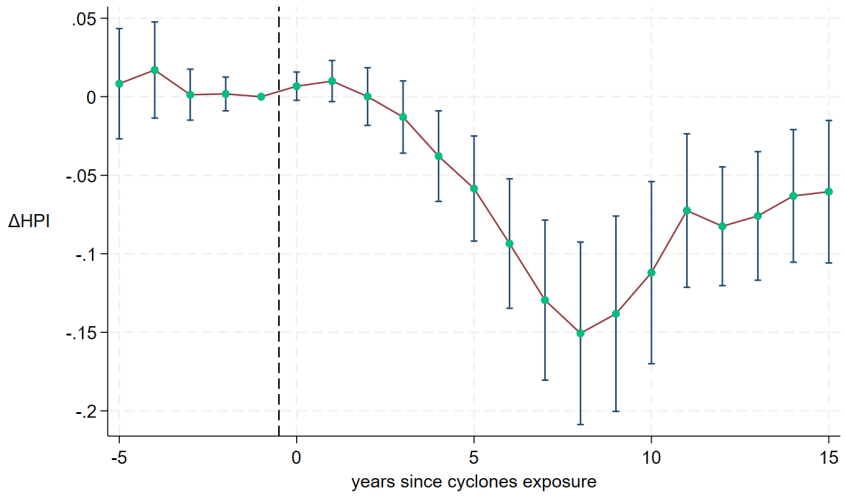
4 Results

This section presents empirical evidence on the adverse effects of cyclones on the housing market. To ensure the robustness of the main findings, I conduct several checks, demonstrating that the results hold under more traditional model specifications and when using alternative data sources. I show that the effects dissipate for counties farther inland and test the optimal number of lags to study long-run effects. I investigate heterogeneous effects, examining how the impacts differ between counties where cyclones are less expected versus those with a higher probability of occurrence, and comparing the effects of stronger cyclones to weaker ones.

4.1 Changes in housing prices following cyclones

Figure 8 presents my main result: the effects of tropical cyclones on the value of housing relative to eight years without being exposed to cyclones. The plot shows β_h^{LP-DID} for each time horizon h between -2 and and 9 after Equation 2 is estimated for coastal counties. The leads, representing periods before the cyclone exposure, act as placebo tests to verify the validity of the analysis. These tests confirm the absence of significant differences between the control and treatment counties prior to cyclones, indicating no anticipatory behavior or systematic pre-treatment divergence. In this analysis, I assume that the effects of cyclone exposure stabilize after eight years. Therefore, the control group is defined as any coastal county that has not experienced a cyclone for eight consecutive years (Figure 7).

Figure 8: Impact of cyclones on the US housing market using LP-DiD



Note: This model employs a Local Projection for Difference-in-Differences (LP-DiD) framework to estimate cumulative impulse responses to cyclone exposure over time. The sample includes 184 coastal counties. Control observations are defined as county-years in which no cyclone occurred for at least eight consecutive years. ΔHPI is the log difference in housing prices between $t-1$ and $t+h$

After a county has gone at least eight years without cyclone exposure, a new exposure tends to generate an insignificant short-lived increase in housing values during the exposure year and the year immediately following it. Starting in the second year after exposure, however, housing values begin to decline and continue falling for roughly eight years relative to a counterfactual path without exposure. The market then shows signs of recovery around year 9, with housing values stabilizing by approximately year 11. Importantly, these dynamics reflect the typical adjustment pattern following cyclone exposure, not the effect of any single cyclone. To assess whether eight cyclone-free years is sufficient for estimating long-run effects, I compare the results using broader "clean-condition" windows, $8 < K^* \leq 15$ (Appendix D). The consistency of the findings across these windows suggests that the estimates between years 9 and 11 reflect actual post-exposure recovery, rather than contamination from control counties experiencing new cyclone exposure after eight clean years. [11]

Being exposed to cyclones induces a negative impact on housing price change, amounting to approximately 15% after 8 years compared to a counterfactual trajectory in the absence of the event. Referencing Table 1, where the mean value of houses in coastal counties is \$254,079, counties experiencing cyclones witness an average decrease of around \$38,000 in property value after 8 years. I also examine the effect using wind speed as a continuous measure of exposure. Figure D.3b shows that a one–standard deviation increase in wind speed (10 m/s) leads to about a 4% decrease in housing values after nine years. For context, Hurricane Katrina, with maximum sustained winds of

[11] The β_h^{LP-DID} are cumulative coefficients, so the simultaneous estimates between 9 to 11 are positive.

64 m/s, is associated with roughly a 25% reduction in cumulative housing price growth.

While the magnitude is substantial, the dynamic response is also informative: housing prices exhibit a small and statistically insignificant increase in the year of the cyclone and the year immediately following, before beginning a sustained decline from the second year onward and reaching a trough after about eight years. This delayed adjustment is economically plausible given the presence of frictions in housing markets. In the short run, reconstruction activity, insurance claim payouts, and federal disaster aid can temporarily support demand and stabilize prices, offsetting immediate negative effects, while storm-related destruction may constrain effective supply. Over the medium term, however, property reassessments, updated information about flood and cyclone risk, and gradual changes in household expectations lead to a downward revision in housing demand. These effects may be compounded by repeated cyclone exposure within treated counties, amplifying the cumulative impact on prices over time. Evidence from rental markets, presented in Appendix E, supports this interpretation: rents, typically more flexible, decline modestly (by up to 2%) following a cyclone but recover within roughly a year and a half, indicating that short-run disruptions are temporary. In contrast, housing supply responses, also shown in Appendix E, are more persistent, with new construction permits falling sharply after cyclones and remaining depressed for several years, alongside a decline in total housing units. Taken together, these patterns suggest that while short-run market adjustments and policy responses may initially mask the impact of cyclones, longer-run changes in risk perceptions, investment, and local economic conditions drive the gradual but pronounced decline in housing values over time.

This finding represents the paper's key contribution: unlike prior studies documenting positive housing market effects, either within two years post-cyclone in Florida (Graff Zivin et al., 2023) or over eight-year horizons across all U.S. counties (Roth Tran and Wilson, 2025), this paper documents housing price declines in affected coastal counties. Expanding the sample to include inland counties transforms the negative coastal estimates into positive results for the full sample, as shown in Figure D.9 of Appendix D, and replicates the positive 1% effect over eight years found by Roth Tran and Wilson (2025). This progression from negative coastal estimates to positive full-sample results may come from three factors. First, differential pre-trends introduce compositional bias: coastal counties exhibit stronger housing-price appreciation over time, attributable to demand for coastal amenities. Second, heterogeneous damage severity: coastal regions suffer compounded losses from both hurricane winds and storm surges, whereas inland effects are milder. Expanding the sample thus averages down the treatment intensity. Third, differences in exposure probability between coastal counties and inland. Prioritizing coastal counties therefore provides a more credible identification of cyclone damages. Methodological choices, like choosing clean controls far away from the coast versus having controls similar in probability of receiving cyclone, drive divergent findings.

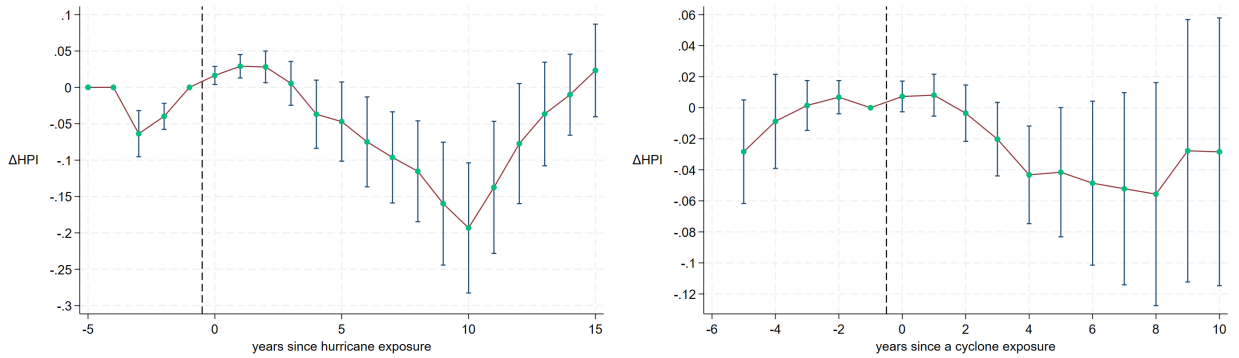
To ensure the robustness and credibility of the main finding, I conduct an extensive battery of sensitivity analyses, detailed in Appendix D. These include alternative model specification, such as

a comparison between the Local Projection Difference-in-Differences (LP-DiD) estimator and the Distributed Lag Model (DLM). I also replicate the analysis using an alternative housing market dataset (ZHVI), finding similar results. Then, I explore sample-selection robustness by varying the geographic distance from the Atlantic coast used to define the sample. Including more inland counties turns the estimates positive, with the extreme case of including all US counties. Another useful analysis is examining the impact of a single cyclone relative to no exposure. To do this, the sample must extend at least 200 km from the coast. Although both the treatment and control groups differ from the main specification, the findings reveal a 10% drop in housing prices eight years after a single cyclone struck. In addition, I assess the sensitivity of results to the choice of lag length and conduct cross-validation to guide the selection of the optimal lag structure. To guard against model misspecification or spurious correlations, I implement multiple forms of randomization tests. Lastly, I explicitly address concerns about spatial spillovers by estimating a Spatial Lag Model and examining the outcomes in neighboring but untreated counties. The findings show no significant impact on nearby municipalities, confirming the SUTVA assumption and contrasting prior single-county studies (Hallstrom and Smith, 2005).

4.2 Heterogeneous impacts of cyclone by intensity and frequency

I examine how the effects of cyclones on coastal counties vary with both storm intensity and frequency. Using the continuous wind-speed exposure measure, I can assess heterogeneous impacts according to cyclone strength. In particular, I examine the effects of exposure to hurricanes, storms with wind speeds of at least 33 meters per second, after eight years without any cyclone activity. I then contrast these with the effects of weaker storms, limiting the sample to counties that were never exposed to a hurricane. The subsample for weaker cyclones is considerably smaller: only 40 of the 184 coastal counties were never hit by a hurricane. This reduction in sample size explain the larger standard errors and the shorter post-event lags observed in the estimates.

Figure 9: Impact of strong and weak cyclones on housing

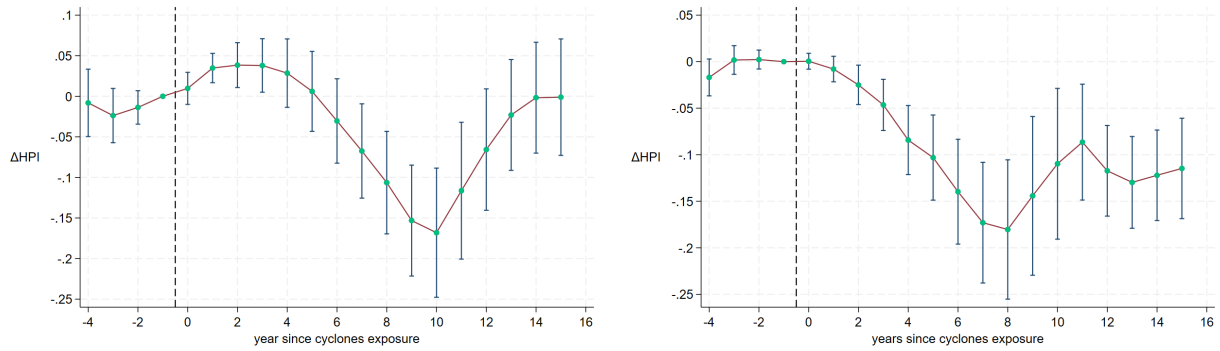


(a) Impact of hurricanes (wind speed $>$ 33 m/s) (b) Impact of weaker cyclones (wind speed $<$ 33 m/s)

The results show that hurricanes generate both a sharper short-term increase and more severe long-term decline in housing values compared to weaker cyclones, around 20% after 10 years. However, even when excluding hurricanes and focusing exclusively on storms with wind speeds below the 33 m/s threshold, the housing market still suffers significant adverse effects (around 4% after 4 years), underscoring the broader relevance of cyclone activity for local economic outcomes.

Next, I compare the impacts of cyclones based on the frequency of their occurrence. Using a threshold probability of 0.27, described in the data section, I divide the sample into two groups: counties with a higher probability of experiencing cyclones (above 0.27) and those with a lower probability (below 0.27).

Figure 10: Impact of cyclones on housing according to different probabilities of treatment



(a) High-probability areas ($p > 0.27$)

(b) Low-probability areas ($p < 0.27$)

Counties with frequent cyclone exposure show a more gradual response: housing values rise for up to three years, then decline, reaching a 15% drop by year ten before recovering. In contrast, low-exposure counties experience an immediate post-cyclone decline with no recovery, implying more permanent losses. This pattern aligns with the single event study analysis in Appendix D, which shows that counties with only one event experience a 15% decline without rebound. This suggests that areas with frequent cyclone exposure may have stronger adaptation mechanisms, such as resilient infrastructure, stricter building codes, and broader insurance coverage, that help mitigate long-term losses. Moreover, the positive correlation (0.44) between cyclone frequency and intensity helps explain why the hurricane estimates in Figure 9a show a recovery phase beginning around year 11. This interpretation is consistent with (Hsiang and Narita, 2012a), who find that cyclone impacts are more severe in places with limited historical exposure.

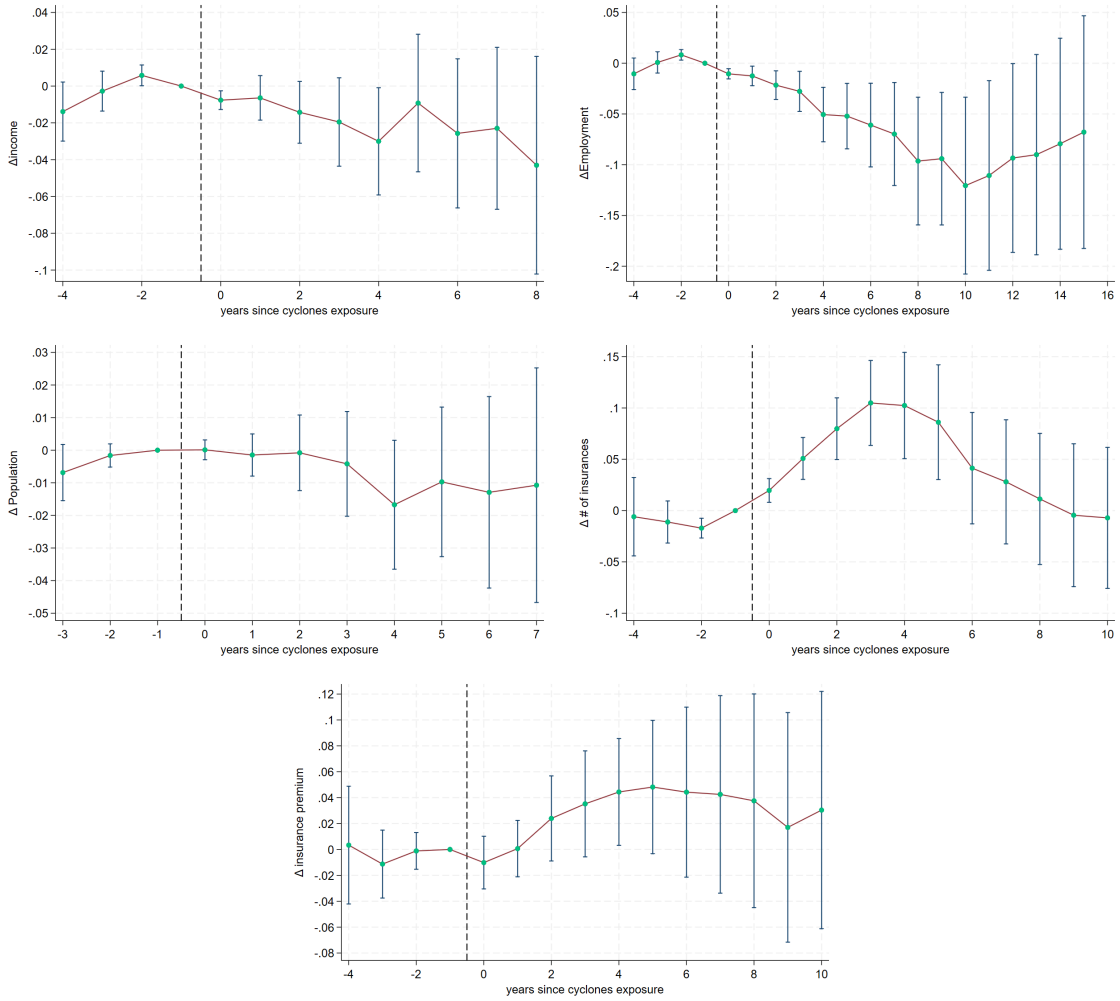
To characterize which types of counties are located in safer areas, those with a lower probability of experiencing cyclones and lower cyclone intensity, I examine their observable characteristics. As shown in Table 1, these counties tend to have higher housing units and prices. However, despite being in safer zones, these areas exhibit more persistent negative effects in response to cyclone shocks, as documented in Figure 10b. To verify whether richer counties indeed experience deeper

impacts, in Appendix F I present a heterogeneous analysis based on county income levels, as well as a quantile regression exploring differences across the distribution of housing values. Both exercises consistently show that wealthier and more developed counties suffer more pronounced negative effects.

4.3 Potential mechanism

To shed light on the mechanisms behind the long-run decline in housing values, I examine the response of other local economic outcomes, income, employment, population, and flood insurance uptake, using the same event-study specification as in Equation 2. Figure 11 reports the dynamic effects.

Figure 11: Impact of cyclones on other variables



The results show that cyclones are followed by a deterioration of local labor-market conditions: four

years later the average income is around 3 percent lower and ten years after exposure, employment is about 10 percent lower than in non-exposed counties. These effects are economically meaningful, but they are smaller than the response of housing values and are not accompanied by a significant decline in population, suggesting that most households remain in place rather than relocating, consistent with Blonz et al. (2025).

By contrast, the insurance response is both rapid and sizable. The number of flood insurance policies increases by 10 percent in the first four years after a cyclone, together with a 5% increase in insurance premium. This surge in insurance uptake implies a substantial and persistent increase in housing-related expenditures in affected counties. Taken together, the evidence points to a combination of channels: higher insurance premia and other housing costs in riskier areas, alongside weaker but non-trivial labor-market declines, jointly contribute to the long-run reduction in housing values. The strong insurance response is in line with previous work on flood events documenting increases in flood coverage following large shocks (Gallagher, 2014).

The limited out-migration following a disaster indicates that many residents choose to remain in the affected area. This can help explain the initial, insignificant increase in housing values: with fewer homes available for sale, the remaining properties may command higher prices due to constrained supply.

These empirical patterns motivate the model in the following section, where households choose location, housing, and insurance in response to evolving risk perceptions and higher housing-related costs in cyclone-prone areas.

5 Model

This section presents a simple dynamic model of location, insurance, and housing choices that rationalizes the empirical patterns in housing values and insurance uptake following cyclones. The framework is related to spatial equilibrium models with climate risk (Wrenn, 2024; Hsiang and Jina, 2014), to models of land use and development in disaster-prone areas (Wang, 2014), and to work on adaptation and insurance in markets for natural disaster risk (Gallagher, 2014; Wagner, 2022; Keys, 2025). Detailed derivations of the optimality conditions and insurance rule are provided in Appendix G.

5.1 Environment

Time is discrete and indexed by $t = 0, 1, 2, \ldots$. Counties are indexed by i and differ in exogenous cyclone risk type $\theta_i \in \{\text{low}, \text{high}\}$. Each county i faces a true annual cyclone probability p_i, with p_i

higher in high-risk counties. In year t, a cyclone occurs in county i with indicator

$$D_{it} \in \{0, 1\}, \quad \Pr(D_{it} = 1) = p_i.$$

Households in county i hold subjective beliefs $\tilde{p}_{it}$ about cyclone risk, which update via constant-gain learning (Evans and Honkapohja, 2001), a tractable alternative to full Bayesian updating that avoids prior elicitation:

$$\tilde{p}_{i,t+1} = (1 - \lambda)\tilde{p}_{it} + \lambda D_{it},$$

where $\lambda \in (0, 1]$ is an experience parameter capturing the strength of belief updating. Initial beliefs $\tilde{p}_{i0}$ reflect baseline experience by county type $\theta_i \in \{\text{low}, \text{high}\}$: high-risk counties start near true risk ($\tilde{p}_{i0} \approx p_i$), while low-risk counties substantially underestimate risk ($\tilde{p}_{i0} \ll p_i$).

The belief updating rule in equation is a constant-gain learning algorithm that approximates Bayesian updating when agents discount distant past data or suspect that the underlying risk p_i may be drifting. For small λ, beliefs evolve similarly to a Bayesian learner with a diffuse prior; for larger λ, recent events are overweighted, consistent with the empirically observed sensitivity of flood insurance take-up to recent disasters in both this paper's results. A relatively highλ is consistent with a model where rare, high-salience events trigger large revisions, as in "learning about an infrequent event (Gallagher, 2014). This is not necessarily irrational: given sparse personal experience and limited information about tail risks, households may rationally discount low-probability events ex ante and then update sharply when they occur.

The housing stock in county i at time t is H_{it}, with average quality q_{it}. A representative dwelling provides housing services

$$h_{it} = q_{it}.$$

5.2 Households

A representative household in county i maximizes expected lifetime utility

$$E_t \sum_{s=t}^{\infty} \beta^{s-t} \left[u(c_{is}) + \eta \ln(h_{is}) - \alpha \tilde{p}_{is} \right],$$

where c_{is} is non-housing consumption, h_{is} housing services, $\beta \in (0, 1)$ the discount factor, $\eta > 0$ the weight on housing, and $\alpha > 0$ captures the disutility of living under perceived cyclone risk. For simplicity, I assume a standard concave utility function $u(\cdot)$ without explicitly modeling heterogeneous risk aversion or wealth effects, as these are beyond the scope of the current analysis.

Each period, the household chooses consumption c_{it}, housing services h_{it}, and whether to purchase

insurance $I_{it} \in \{0, 1\}$. The budget constraint is

$$c_{it} + P_{it} h_{it} + \pi_{it} I_{it} = y_{it} + I_{it}\, \varphi L_{it} D_{it},$$

where P_{it} is the housing price (or rent per unit of services), π_{it} the insurance premium, y_{it} exogenous income, L_{it} the loss if a cyclone hits, and $\varphi \in (0, 1]$ the indemnity rate.

Given $(P_{it}, \pi_{it}, \tilde{p}_{it})$, the household chooses (c_{it}, h_{it}, I_{it}) to maximize expected utility subject to the budget constraint. The first-order condition for housing equates the expected marginal utility cost of an additional unit of housing services to its marginal utility benefit; the expression and algebraic steps are provided in Appendix G.

Insurance is purchased when the expected benefit of coverage exceeds the premium cost. Under standard assumptions (concave $u(\cdot)$ and modest feedback of I_{it} on h_{it}), the condition $EU(1) - EU(0) \geq 0$ can be approximated by the simple cutoff rule

$$\tilde{p}_{it}\, \varphi L_{it} \;\gtrsim\; \pi_{it}, \tag{4}$$

derived formally in Appendix G. Equation (4) shows that insurance take-up is more likely when perceived risk $\tilde{p}_{it}$ and potential losses L_{it} are large relative to the premium π_{it} and when coverage φ is generous. This inequality rationalizes higher insurance uptake in counties that perceive themselves to be at greater risk or that have just experienced a cyclone shock.

In the baseline model, income y_{it} is treated as exogenous. This assumption keeps the household problem tractable and allows the analysis to focus on how evolving risk perceptions and insurance premia shape housing demand and prices. The empirical results, however, show that cyclones are also followed by declines in local income and employment. A natural extension is therefore to allow y_{it} to respond to cyclone exposure, further weakening housing demand in affected counties. Formally, one can write income dynamics as

$$y_{i,t+1} = y_{it}\big(1 - \delta_y D_{it}\big) + \varepsilon_{it}, \tag{5}$$

where $\delta_y > 0$ captures the average proportional decline in income after a cyclone, and ε_{it} is an idiosyncratic shock.

5.3 Housing stock dynamics

When a cyclone hits, dwellings suffer damage that depends on their quality q_{it} and insurance status. The aggregate housing stock evolves according to

$$H_{i,t+1} = (1 - \phi) H_{it} - D_{it}\, \kappa(q_{it})(1 - \bar{I}_{it}) H_{it} + R_{it}, \tag{6}$$

where $\phi \in (0,1)$ is the depreciation rate, $\kappa(q_{it}) \in (0,1)$ is the damage share (decreasing in quality q_{it}), and $\bar{I}_{it}$ is the average insurance coverage rate in county i. Reconstruction investment is given by

$$R_{it} = \chi \bar{I}_{it} D_{it} L_{it}, \tag{7}$$

where $\chi \in (0,1]$ captures the efficiency with which insurance payouts are translated into rebuilding and quality upgrades. Equations (6)–(7) imply that counties with greater insurance penetration experience faster recovery of H_{it} and potentially higher post-disaster quality q_{it}, consistent with a "build back better" channel discussed in the disaster and adaptation literature (Hsiang and Jina, 2014).

5.4 Housing prices

Housing prices clear the local housing market and reflect user cost and risk capitalization, in line with user-cost approaches to housing under climate risk (Keys, 2025). A simple representation is

$$P_{it} = \frac{MU_{h,it}}{UC_{it}}, \tag{8}$$

where $MU_{h,it}$ is the marginal utility value of housing services (derived from the household problem in Appendix G) and UC_{it} is the user cost of housing. The user cost is given by

$$UC_{it} = r + \delta_i + \tau_{it} + \pi_{it} - E_t\left(\frac{\Delta P_{i,t+1}}{P_{it}}\right), \tag{9}$$

where r is the risk-free rate, δ_i captures location-specific obsolescence, τ_{it} denotes local taxes or fees, and π_{it} includes the insurance premium and other risk-related costs. An equivalent asset-pricing representation and the steps leading to (8)–(9) are shown in Appendix G.

A cyclone shock $D_{it} = 1$ reduces the effective housing stock through damage, raising prices due to tighter supply. At the same time, the event raises perceived risk $\tilde{p}_{it}$ and, over time, insurance premia π_{it}, which increases user cost and gradually depresses demand as risk is capitalized into housing values.

Rising insurance premia π_{it} play a dual role: while coverage enables reconstruction and resilience (7), higher premia also increase user cost UC_{it} (9), capitalizing risk into lower housing values (Keys, 2025). This generates a dynamic feedback amplifying adaptation disparities: repeated shocks in high-risk counties cumulatively raise $\tilde{p}_{it}$, sustaining high coverage $\bar{I}_{it}$ and reconstruction R_{it}, which attracts resilient development despite elevated π_{it}. Low-probability counties, starting from low $\tilde{p}_{i0}$, experience post-shock premium spikes after coverage gaps, deterring investment and perpetuating vulnerability, a self-reinforcing divergence in climate adaptation.

In the model, the fundamental shifter of willingness to pay for housing in a county is the perceived risk $\tilde{p}_{it}$, because it enters both utility directly via $-\alpha\tilde{p}_{it}$ and indirectly via insurance premia π_{it} and user cost UC_{it}. Thus, any observed change in P_{it} after a cyclone can be interpreted as a change in the implicit shadow price of hurricane risk. The empirical event-study design with cyclone shocks is then approximately a hedonic regression where the treatment captures a shift from "underpriced" to "more correctly priced" risk for counties that were previously underestimating risk.

Proposition 1 (Model predictions). The model predicts:

(i) Destruction tightens H_{it} (6)

(ii) Risk capitalization raises housing costs UC_{it} (9), causing decline in price (8)

(iii) **Heterogeneity**:

- Low-risk ($\tilde{p}_{i0} \ll p_i$): Low baseline coverage $\bar{I}_{i0} \to$ small R_{it} ((7)), persistent declines.
- High-risk ($\tilde{p}_{i0} \approx p_i$): High baseline coverage $\bar{I}_{i0} \to$ large R_{it}, rebound.

Insurance's *dual role*: premia π_{it} drive declines; coverage enables resilience.

5.5 Link to empirical patterns

The model predictions align closely with the data for counties differing by initial risk and insurance penetration.

Low-risk counties ($\tilde{p}_{i0} \ll p_i$, low $\bar{I}_{i0}$, high effective λ): A first cyclone sharply revises beliefs ($\tilde{p}_{it} \uparrow$), spikes premia ($\pi_{it} \uparrow$), but yields limited reconstruction (R_{it} small due to low baseline coverage). Rising user costs dominate, producing persistent housing declines with only modest quality recovery—matching the 15 percent losses over eight years in low-exposure counties (Wagner, 2022).

High-risk counties ($\tilde{p}_{i0} \approx p_i$, high $\bar{I}_{i0}$, low effective λ): Shocks induce mild $\Delta\tilde{p}_{it}$ but trigger large payouts and rebuilding (R_{it} large via (7)), raising quality q_{it} and supporting rebounds after initial supply tightening.

The model thus rationalizes: (i) short-run rises from destruction ((6)), (ii) medium-run peaks then declines from risk capitalization ((9)), and (iii) divergent long-run outcomes: persistent losses in low-experience counties versus rebounds in high-risk areas, matching Figure 10.

6 Conclusion

This study provides the first plausibly causal long-run estimates of how tropical cyclones reshape U.S. housing markets, documenting a large and persistent decline in housing values of about 15 percent over an eight-year horizon. These losses extend far beyond the short-run disruptions typically emphasized after natural disasters and appear more enduring than those associated with other hazards, such as wildfires (McCoy and Walsh, 2018). Most importantly, this paper shows that methodological choices, such as restricting the analysis to coastal counties versus including all U.S. counties, can generate divergent results and help explain why previous work finds increases in housing prices (Roth Tran and Wilson, 2025).

The analysis moves beyond case studies of individual hurricanes by incorporating all tropical cyclones over more than three decades, combined with a newly constructed wind-speed exposure measure. The results show that climate shocks, and cyclones in particular, can generate deep and long-lasting economic disruptions that are often underestimated in both research and policy discussions (Carleton and Hsiang, 2016; Hsiang and Jina, 2014; Young and Hsiang, 2024).

A central contribution is to show that these aggregate effects mask substantial spatial heterogeneity. Counties with historically low cyclone exposure experience the largest and most persistent price declines, consistent with limited preparedness and weaker resilience. By contrast, more frequently exposed, high-risk counties display a rebound in housing values after an initial price increase and subsequent decline, suggesting that adaptation buffers long-run devaluation.

The paper then develops a simple dynamic model of location, housing, and insurance choices that rationalizes these patterns. In the model, households choose where to live, whether to insure, and how much housing to consume based on evolving beliefs about cyclone risk. Less-experienced coastal counties underestimate risk ex ante and have low pre-disaster insurance coverage, so a cyclone leads to sharp upward revisions in perceived risk, higher insurance premia, and only limited insurance-financed rebuilding. This combination raises long-run user costs and depresses housing prices, consistent with the persistent declines observed in the data. In contrast, high-risk counties start with more accurate risk beliefs and higher insurance uptake, so cyclone shocks trigger substantial reconstruction and quality upgrading financed by existing coverage. This "build back better" margin, together with belief revisions, helps explain the rebound in housing values in more exposed places (Hsiang and Jina, 2014).

Empirically, the study finds relatively modest effects on income, employment, and population, but large and persistent increases in insurance uptake and premiums in cyclone-prone areas. These findings suggest that insurance plays a dual role: it facilitates recovery and resilience through reconstruction while simultaneously increasing the cost of living as higher premia are capitalized into housing values. The model formalizes this dual role and shows how adverse selection and delayed

uptake in historically low-risk areas can amplify long-run price declines, aligning with recent work on climate risk and insurance markets (Wagner, 2022; Keys, 2025).

Several limitations remain. The analysis focuses on wind speed and does not separately identify other dimensions of cyclone damage, such as storm surge or inland flooding, which may matter for both physical losses and insurance pricing. Moreover, the 30-year cyclone probability weighting helps to place events in a long-run risk context, but these probabilities are themselves evolving under climate change. Lastly, the county-level approach cannot fully capture within-county heterogeneity in exposure, housing quality, or household balance sheets. Future work linking these dynamics to household-level data, mortgage markets, and firm behavior would help uncover additional mechanisms and distributional consequences. Moreover, insurance markets alone are unlikely to account for the sizable 15% decline in housing values, as households in risky areas may also face higher ongoing costs of living, such as elevated maintenance expenses.

From a policy perspective, the results imply that counties with infrequent cyclone histories are not necessarily low-risk in economic terms. Low baseline probabilities of landfall can coexist with very high realized losses when preparedness, insurance coverage, and local institutions are weak. Improving risk communication, tailoring insurance design, and investing in resilience, especially in coastal areas that have attracted development under a perception of safety, could mitigate long-run housing losses and support a more equitable adaptation to climate risk.

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.

Bakkensen, L. A. and L. Barrage (2021): "Going Underwater? Flood Risk Belief Heterogeneity and Coastal Home Price Dynamics," *The Review of Financial Studies*, 35, 3666–3709.

Bakkensen, L. A., X. Shi, and B. D. Zurita (2018): "The Impact of Disaster Data on Estimating Damage Determinants and Climate Costs," *Economics of Disasters and Climate Change*, 2, 49–71.

Benetton, M., S. Emiliozzi, E. Guglielminetti, M. Loberto, and A. Mistretta (2022): "Do House Prices Reflect Climate Change Adaptation? Evidence from the City on the Water," Occasional Paper 735, Bank of Italy, bank of Italy Occasional Paper No. 735.

Beracha, E. and R. S. Prati (2008): ""How Major Hurricanes Impact Housing Prices and Transaction Volume"," *Real Estate Issues*, 33.

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.

Bilal, A. and D. R. Känzig (2026): "The Macroeconomic Impact of Climate Change: Global Versus Local Temperature*," *The Quarterly Journal of Economics*, qjag011.

Biny, O. and S. Polasky (2004): "Effects of Flood Hazards on Property Values: Evidence Before and After Hurricane Floyd," *Land Economics*, 80.

Blonz, J., S. Bowdle, and J. A. Weill (2025): "Hurricanes and migration: New evidence from credit bureau microdata," *Journal of Environmental Economics and Management*, 133, 103180.

Booth, J. F., V. Narinesingh, K. L. Towey, and J. Jeyaratnam (2021): "Storm Surge, Blocking, and Cyclones: A Compound Hazards Analysis for the Northeast United States," *Journal of Applied Meteorology and Climatology*, 60.

Burke, M., S. M. Hsiang, and E. Miguel (2015): "Global non-linear effect of temperature on economic production," *Nature*, 527, 235–239.

Callaway, B. and P. H. Sant'Anna (2021): "Difference-in-Differences with multiple time periods," *Journal of Econometrics*, 225, 200–230, themed Issue: Treatment Effect 1.

Carleton, T., E. Duflo, B. K. Jack, and G. Zappalà (2024): "Adaptation to Climate Change," in *Handbook of the Economics of Climate Change*, ed. by L. Barrage and S. Hsiang, North-Holland, vol. 1 of *Handbook of the Economics of Climate Change*, 143–248.

Carleton, T., A. Jina, M. Delgado, M. Greenstone, T. Houser, S. Hsiang, A. Hultgren, R. E. Kopp, K. E. McCusker, I. Nath, J. Rising, A. Rode, H. K. Seo, A. Viaene, J. Yuan, and A. T. Zhang (2022): "Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits*," *The Quarterly Journal of Economics*, 137, 2037–2105.

Carleton, T. A. and S. M. Hsiang (2016): "Social and economic impacts of climate," *Science*, 353, aad9837.

de Chaisemartin, C., X. d'Haultfoeuille, and G. Vazquez-Bare (2024): "Difference-in-Differences Estimators with Continuous Treatments and No Stayers," *working paper*.

de Chaisemartin, C. and X. D'Haultfœuille (2024): "Difference-in-Differences Estimators of Intertemporal Treatment Effects," *The Review of Economics and Statistics*, 1–45.

Deryugina, T. (2017): "The Fiscal Cost of Hurricanes: Disaster Aid versus Social Insurance," *American Economic Journal: Economic Policy*, 9, 168–98.

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, 202–33.

Deschênes, O. and M. Greenstone (2007): "The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather," *American Economic Review*, 97, 354–385.

Dube, A., D. Girardi, J. Òscar, and A. M. Taylor (2025): "A Local Projections Approach to Difference-in-Differences," *Journal of Applied Econometrics*.

Evans, G. W. and S. Honkapohja (2001): *Learning and Expectations in Macroeconomics*, Princeton University Press.

Ferreira, S. (2024): "Extreme Weather Events and Climate Change: Economic Impacts and Adaptation Policies," *Annual Review of Resource Economics*, 16, 207–231.

Gallagher, J. (2014): "Learning about an Infrequent Event: Evidence from Flood Insurance Take-Up in the United States," *American Economic Journal: Applied Economics*, 6, 206–33.

Georgic, W. and H. A. Klaiber (2022): "Stocks, flows, and flood insurance: A nationwide analysis of the capitalized impact of annual premium discounts on housing values," *Journal of Environmental Economics and Management*, 111, 102567.

Giglio, S., B. Kelly, and J. Stroebel (2021a): "Climate Finance," *Annual Review of Financial Economics*, 13, 15–36.

Giglio, S., M. Maggiori, K. Rao, J. Stroebel, and A. Weber (2021b): "Climate Change and Long-Run Discount Rates: Evidence from Real Estate," *The Review of Financial Studies*, 34, 3527–3571.

Goodman-Bacon, A. (2021): "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 225, 254–277, themed Issue: Treatment Effect 1.

Graff Zivin, J., Y. Liao, and Y. Panassié (2023): "How hurricanes sweep up housing markets: Evidence from Florida," *Journal of Environmental Economics and Management*, 118, 102770.

Graham, J. E., W. W. Hall, and P. W. Schuhmann (2007): "Hurricanes, Catastrophic Risk, and Real Estate Market Recovery," *The Journal of Real Estate Portfolio Management*, 13, 179–190.

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

Hino, M. and M. Burke (2021): "The effect of information about climate risk on property values," *Proceedings of the National Academy of Sciences*, 118, e2003374118.

Holland, G. J., J. I. Belanger, and A. Fritz (2010): "A Revised Model for Radial Profiles of Hurricane Winds," *Monthly Weather Review*, 138, 4393 – 4401.

Hsiang, S. (2016): "Climate Econometrics," *Annual Review of Resource Economics*, 8, 43–75.

Hsiang, S., R. Kopp, A. Jina, J. Rising, M. Delgado, S. Mohan, D. J. Rasmussen, R. Muir-Wood, P. Wilson, M. Oppenheimer, K. Larsen, and T. Houser (2017): "Estimating economic damage from climate change in the United States," *Science*, 356, 1362–1369.

Hsiang, S. M. and A. S. Jina (2014): "The Causal Effect of Environmental Catastrophe on Long-Run Economic Growth: Evidence From 6,700 Cyclones," NBER Working Papers 20352, National Bureau of Economic Research, Inc.

Hsiang, S. M. and D. Narita (2012a): "ADAPTATION TO CYCLONE RISK: EVIDENCE FROM THE GLOBAL CROSS-SECTION," *Climate Change Economics*, 03, 1250011.

——— (2012b): "ADAPTATION TO CYCLONE RISK: EVIDENCE FROM THE GLOBAL CROSS-SECTION," *Climate Change Economics*, 3, 1–28.

Indaco, A., A. Orrego, and M. Tang (2019): "The Effects of Flood Insurance on Housing Markets," *Cityscape*, 21, 129–156.

Islam, M. R., C. Lee, K. T. Mandli, and H. Takagi (2021): "A new tropical cyclone surge index incorporating the effects of coastal geometry, bathymetry and storm information," *Scientific Reports*, 11, 16747.

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

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

Keys, B. J. (2025): "Housing, Climate Risk, and Insurance," *NBER Reporter*, 1–8, national Bureau of Economic Research.

Kolak, M. and L. Anselin (2019): "A Spatial Perspective on the Econometrics of Program Evaluation," *International Regional Science Review*, 43, 128–153, original work published 2020.

Lima, R. C. d. A. and A. V. B. Barbosa (2019): "Natural disasters, economic growth and spatial spillovers: Evidence from a flash flood in Brazil," *Papers in Regional Science*, 98, 905–924.

Marsooli, R., N. Lin, K. Emanuel, and K. Feng (2019): "Climate change exacerbates hurricane flood hazards along US Atlantic and Gulf Coasts in spatially varying patterns," *Nature Communications*, 10.

McCoy, S. J. and R. P. Walsh (2018): "Wildfire risk, salience and housing demand," *Journal of Environmental Economics and Management*, 91, 203–228.

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.

Nordhaus, W. D. (2010): "Economic aspects of global warming in a post-Copenhagen environment," *Proceedings of the National Academy of Sciences*, 107, 11721–11726.

Noy, I. (2009): "The macroeconomic consequences of disasters," *Journal of Development Economics*, 88, 221–231.

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

Roth Tran, B. and D. J. Wilson (2025): "The Local Economic Impact of Natural Disasters," *Journal of the Association of Environmental and Resource Economists*, 12, 1667–1704.

Rubin, D. B. (1990): "Formal mode of statistical inference for causal effects," *Journal of Statistical Planning and Inference*, 25, 279–292.

Schmidheiny, K. and S. Siegloch (2023): "On event studies and distributed-lags in two-way fixed effects models: Identification, equivalence, and generalization," *Journal of Applied Econometrics*, 38, 695–713.

Speyrer, J. F. and W. R. Ragas (1991): "Housing prices and flood risk: An examination using spline regression," *Journal of Real Estate Finance and Economics*, 4, 395–407.

Strobl, E. (2011): "The Economic Growth Impact of Hurricanes: Evidence from U.S. Coastal Counties," *The Review of Economics and Statistics*, 93, 575–589.

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

Vigdor, J. (2008): "The Economic Aftermath of Hurricane Katrina," *Journal of Economic Perspectives*, 22, 135–54.

Wagner, K. R. H. (2022): "Adaptation and Adverse Selection in Markets for Natural Disaster Insurance," *American Economic Journal: Economic Policy*, 14, 380–421.

Wang, C. (2014): "Regulating land development in a natural disaster-prone area: The roles of building codes," *Resource and Energy Economics*, 36, 209–228.

Wrenn, D. H. (2024): "The Effect of Natural Disasters and Extreme Weather on Household Location Choice and Economic Welfare," *Journal of the Association of Environmental and Resource Economists*, 11, 1101–1134.

Young, R. and S. Hsiang (2024): "Mortality caused by tropical cyclones in the United States," *Nature*, 635, 121–128.

A Supplementary information on the construction of cyclones wind speed

Holland et al. (2010) describe the internal dynamics of cyclones, leveraging data from the United States Air Force's reconnaissance missions. These missions involve aircraft flying directly into tropical cyclones in the Atlantic Ocean, the Caribbean Sea, the Gulf of Mexico, and the Central Pacific Ocean to measure weather data within and around the storms.
Holland et al. (2010) enhanced the capacity to account for both external and core wind observations. They introduced the following equation for calculating a tropical cyclone's wind speed at a given distance r from the cyclone center:

$$v_s = v_{ms} \left(\left(\frac{r_{v_{ms}}}{r} \right)^{b_s} \exp \left[1 - \left(\frac{r_{v_{ms}}}{r} \right)^{b_s} \right] \right)^{x} \tag{10}$$

where the subscript s refers to surface values (at a nominal height of 10 m) and v_s is the wind speed (in m/s) at radius r (in meters) - also called "wind speed adjusted". The maximum wind speed v_{ms} represents the peak wind speed in the cyclone, while the radius of maximum wind speed $r_{v_{ms}}$ defines the distance from the cyclone's center where the maximum wind speed occurs. As r approaches $r_{v_{ms}}$, v_s increases reaching it's peak at $r = r_{v_{ms}}$. Beyond $r_{v_{ms}}$, v_s decreases. The term $\left(\left(\frac{r_{v_{ms}}}{r}\right)^{b_s}\right.$ in the equation creates an exponential decay factor that modify v_s based on the distance from the radius of maximum wind speed. The decay rate is controlled by b_s, reflecting how quickly wind speeds decrease as one moves away from $r_{v_{ms}}$. Both v_{ms} and $r_{v_{ms}}$ are reported by HURDAT2. The parameter b_s can be calibrated as follows (equation 8 in Holland et al. 2010):

$$b_s = -4.4 \times 10^{-}5\Delta p_s + 0.03\frac{\partial p_{cs}}{\partial t} - 0.014\varphi + 0.15v_t^{x_1} + 1 \tag{11}$$

where $\Delta p_s = p_{ns} - p_{cs}$ is the difference in pressure between some external pressure and the pressure of the hurricane center, p_{cs}. Wind speed in cyclones is driven by the pressure gradient force, which arises from differences in atmospheric pressure. The larger the difference between the central pressure of the cyclone (p_{cs}) and the surrounding pressure (p_{ns}), the stronger the pressure gradient force, and consequently, the higher the wind speeds. Here, I take p_{ns} to be the typical atmospheric pressure, 1013.25 hPa. The value of p_{cs} is calculated as $p_{cs} = p_{min} + v_{ms}^2$ where p_{min} is the minimum pressure reported by HURDAT2. The lower central pressures typically correspond to more intense cyclones with stronger winds. $\frac{\partial p_{cs}}{\partial t}$ is then calculated as the change in pressure between two subsequent HURDAT2 points (after initial interpolation between them to reduce the time interval). The variable φ is the absolute value of latitude in degrees, which we also observe. Finally, v_t is the cyclone translation speed in $\frac{m}{s}$, calculated as the average speed between two points[12], and $x_1 = 0.6 \times (1 - (\Delta_{ps}/215))$. It is important to note that the wind speed v_s is positively influenced

[12] I calculate the speed dividing the distance between two subsequent recording points by the time interval of 6 hours

by the translation speed of the cyclone. This means that faster-moving cyclones tend to have higher wind speeds. Consequently, the persistence of wind at a given location becomes less significant compared to the overall velocity of the cyclone.

The final unknown in Equation (1) is x. This term adjusts the wind speed outside the main trajectory, accounting for the fact that the wind speed does not decrease monotonically at different radii. Holland et al. (2010) suggest a number of ways that x can be estimated. I utilize the fact that HURDAT2 reports the extend of 34-knot, 50-knot, and 64-knot winds in the Southwest, Southeast, Northeast, Northwest directions. Therefore, for distances from the center that exceed the radius of maximum winds but are smaller than the extent of the corresponding wind threshold (34 knots, 50 knots or 64 knots), we derive an adjusted x value by solving Equation (1) for x and utilizing the known wind speeds at known distances from the hurricane center:

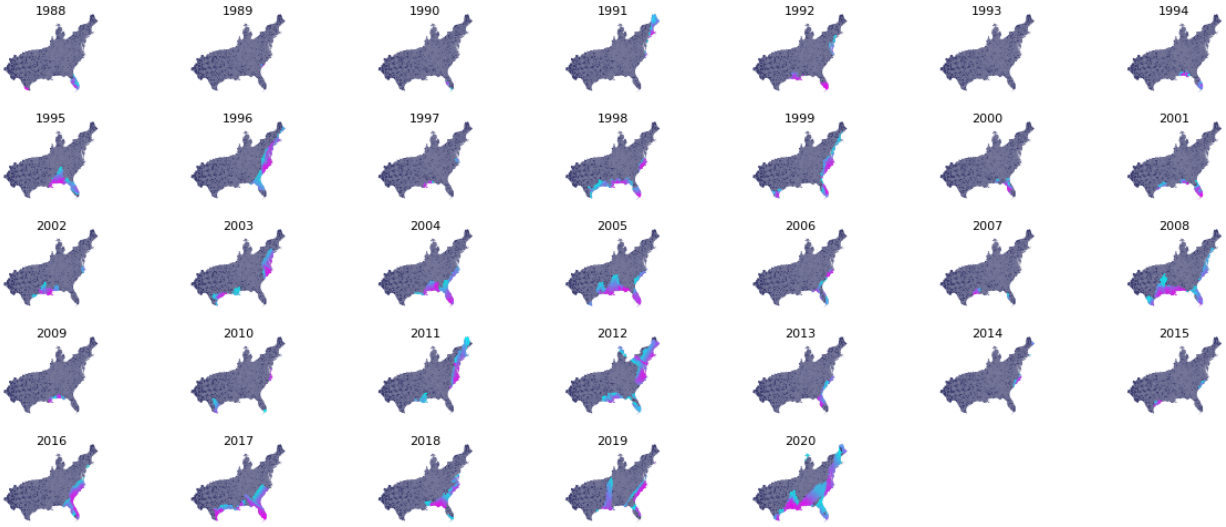
$$x_d = \frac{\ln\left(\frac{v_s}{v_d}\right)}{\ln\left(\left(\frac{r_{v_d}}{r_d}\right)^{b_s} e^{\left[1-\left(\frac{r_{v_d}}{r_d}\right)^{b_s}\right]}\right)} \tag{12}$$

where r_d is the radius of either 34-knot, 50-knot, or 64-knot wind speeds and v_d is the corresponding wind speed (17.49, 25.72, or 32.9 $\frac{m}{s}$, respectively). Finally, we use x_d along with other known values to calculate wind speed according to Equation (1). Specifically, we set $x = x_{64}$ for $R \leq r_{64}$, $x = x_{50}$ for $R \in (r_{64}, r_{50}]$ and $x = x_{34}$ for $R \in (r_{50}, r_{34}]$. Building upon the methodologies proposed by Deryugina (2017) and Hsiang and Jina (2014), my approach extends the allowance for wind speed variation based on different radii information.

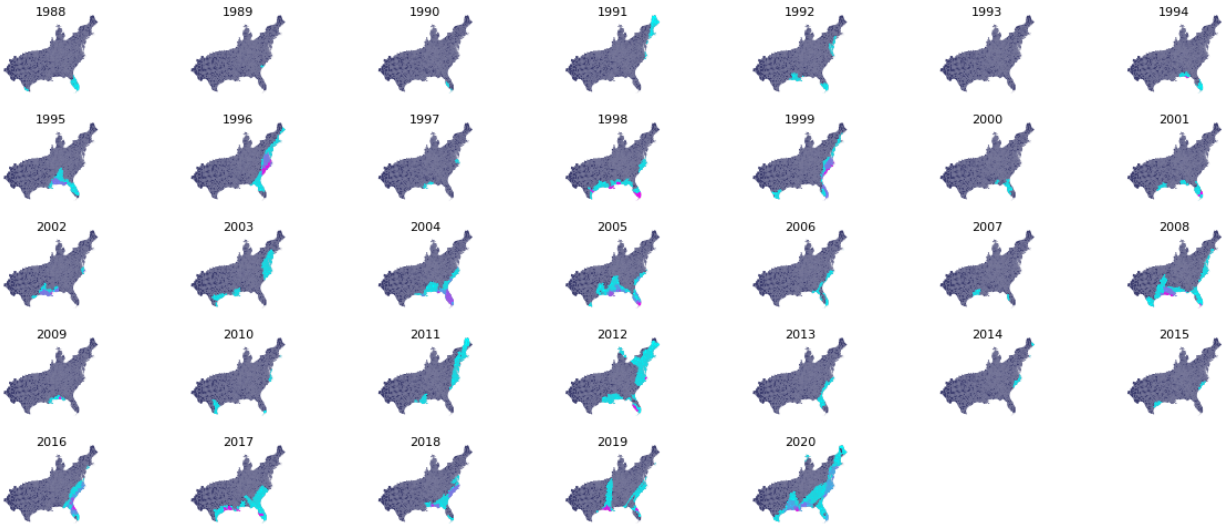
B Supplementary information on data variability

Variability of wind speed intensity and frequency by year:

Figure B.1: Variability of intensity and frequency by year



(a) Maximum wind speed of cyclones by year and location



(b) Frequency of cyclones by year and location

C Alternative Estimation Strategies: TWFE and Related Models

Two-Way Fixed Effects (TWFE) Event Study Model

To benchmark the results, I also estimate a standard event study model using a Two-Way Fixed Effects (TWFE) regression:

$$\ln y_{i,t} = \alpha_i + \delta_t + \sum_{l=-Q}^{L} \gamma_l^{TWFE} D_{i,t}^l + e_{i,t} \tag{13}$$

where $\ln y_{i,t}$ is the logarithm of housing price in county i and year t; α_i and δ_t are county and year fixed effects, respectively; $D_{i,t}^l$ is a set of relative-time dummies for l leads and lags around cyclone exposure. The coefficients γ_l^{TWFE} trace the dynamic treatment effects over time.

Distributed Lag Model (DLM)

When there are multiple variables (e.g. leads and lags in equation 13), their total influence on the outcome is best stated by combining their effects together. To evaluate the cumulative effects over time, I use the Distributed Lag Model (DLM). Drawing from the insights of Schmidheiny and Siegloch (2023), I can establish the equivalence between the TWFE and the DLM methodologies by aggregating the effects identified within the TWFE framework:

$$\beta_k = \sum_{L=0}^{k} \gamma_L \tag{14}$$

where β_k captures the total impact of a cyclone over k post-treatment years. where now the parameter β_k represents the cumulative impact of experiencing a cyclone up to k year. It's noteworthy to highlight that this equivalence extends beyond just the TWFE; it also holds true for the two-way first difference (TWFD).[13] In Appendix D, I leverage both equivalences to present distinct estimates for both the TWFE and TWFD.

Limitations of TWFE

Recent Difference-in-Differences literature (e.g., Goodman-Bacon 2021, Sun and Abraham 2021) shows that TWFE models can suffer from "negative weights" and contamination bias. Specifically, when previously treated units are used as controls, ongoing dynamic responses can bias relative-period coefficients. For example, in Equation 13, when $D_{i,t} = 0$ but $D_{i,t-j} \neq 0$ for $j \geq 1$, the comparison is "unclean."

Because cyclones occur multiple times per county and treatment is non-absorbing, TWFE models make problematic comparisons and fail to isolate marginal treatment effects. Furthermore:

[13]The Dynamic Two-Way First Difference Model is $\Delta y_{i,t} = \sum_{l=-Q}^{L} [\gamma_l^{TWFD} \times \Delta W_{i,t-l}] + \delta_t + \Delta e_{i,t}$ where $\Delta y_{i,t}$ is $\ln(HPI_{i,t}) - \ln(HPI_{i,t-1})$; $\Delta W_{i,t-l}$ is the change in dummy=1 if cyclones strike in county i in year t ; δ_t is the year fixed effects and $\Delta e_{i,t}$ is the error term

- There are no "never-treated" units in the sample (all coastal counties are treated at some point).
- Exposure is binary and irregular—not continuous like temperature shocks.
- Methods like Callaway and Sant'Anna (2021) that require absorbing treatment status are inapplicable.

Given these limitations, I report TWFE and TWFD results for transparency, but the primary specification used throughout the paper relies on the Local Projections estimator (LP-DiD), which offers a more robust design-based approach in this setting.

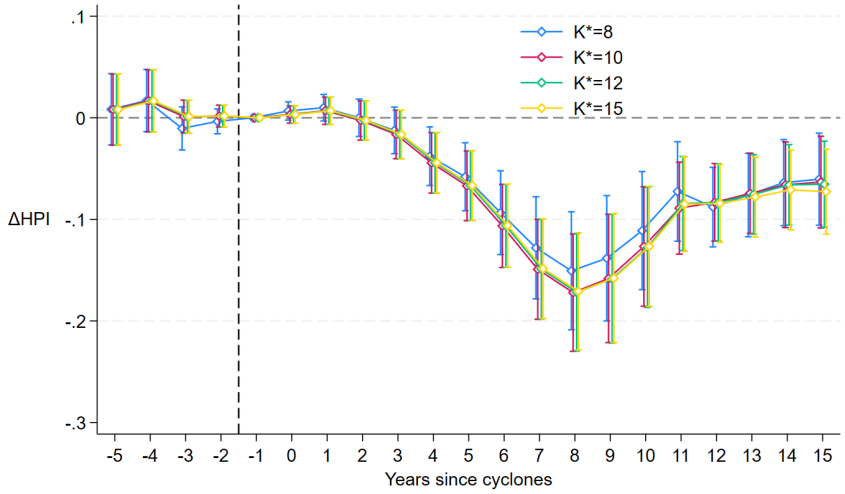
D Robustness of the main result

To ensure the robustness of the main findings, I use various strategies, including alternative model specifications, the use of different datasets, randomization tests, and subsampling. Additionally, I address potential spillover effects by implementing a Spatial Lag Model, as specified in Equation 3.

Different assumptions on the clean control condition

To assess whether the control units are sufficiently "clean", I extend the clean-control condition and examine whether the results vary for thresholds of $8 < K^* \leq 15$ consecutive cyclone-free years.

Figure D.1: Testing different clean control conditions on equation 2



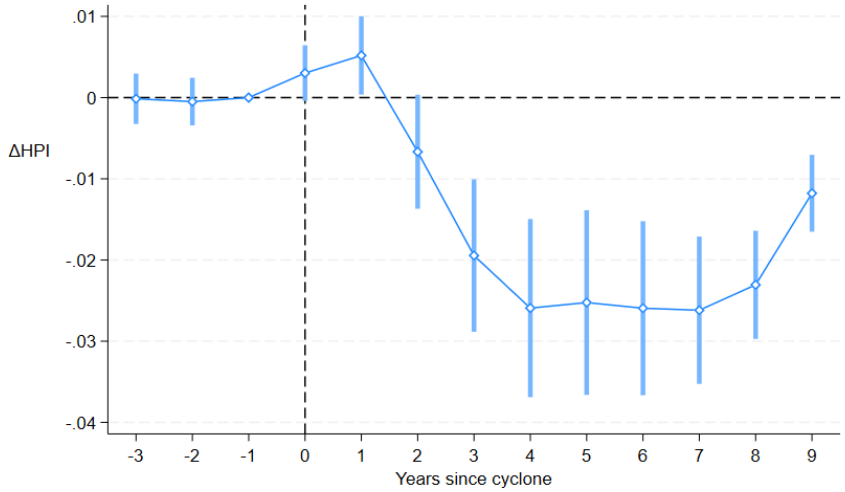
Note: This model uses a Local Projection Difference-in-Differences (LP-DiD) framework to estimate cumulative impulse responses to cyclone exposure over time. The sample includes 184 coastal counties. Control observations are defined as county-years with no cyclone for at least 8, 10, 12, or 15 consecutive years. ΔHPI denotes the log change in housing prices between $t-1$ and $t+h$.

The estimates are largely overlapping and not statistically different across the alternative thresholds, indicating that a clean-control condition of 8 consecutive cyclone-free years is sufficient. Note that already with

Results from traditional event study approach

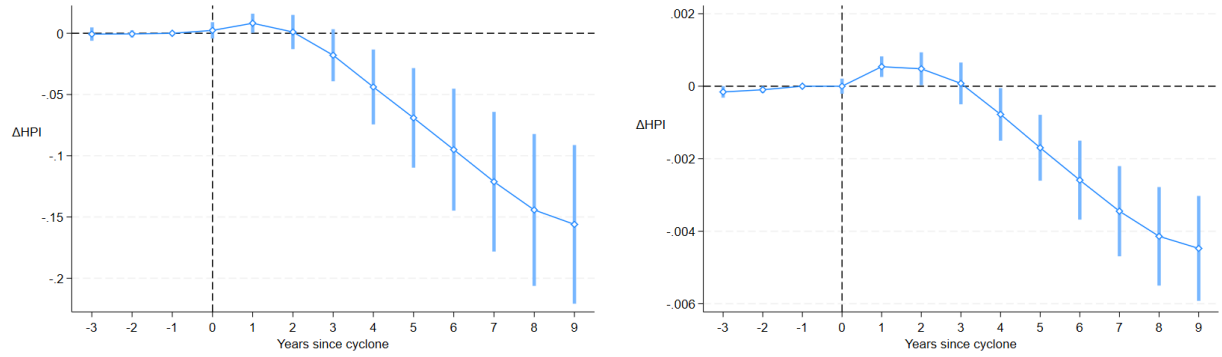
Figure D.2 presents the result of a traditional event study within Dynamic Two-Way Fixed Effect Model: The plot depicts $\gamma_l \in [-3, 9]$ after Equation 13 is estimated for the coastal counties. It's important to note that the leads serve as placebo tests, demonstrating that there are no significant differences between the counties in the control and treatment groups before the cyclone. Indeed, the random trajectory of cyclones strengthens the assumption of exogeneity, allowing me to isolate and identify the effects starting from the year of the event.

Figure D.2: Impact of cyclones on the US counties using Two-Way Fixed Effect Model



Following a cyclone event, housing values do not decline immediately; however, starting from the second year, they begin to decrease steadily over a period of roughly nine years compared to a counterfactual scenario where the event never occurred. To better visualize the magnitude of this long-term effect, I have employed a distributed lag model (Equation 14). Figure D.3a illustrates the cumulative impact of cyclones over time, taking into account both the binary cyclone variable and the continuous measure of wind speed.

Figure D.3: Cumulative effects of cyclone using Distributed Lag Model



(a) Long-run effects of a new cyclone on hosing market

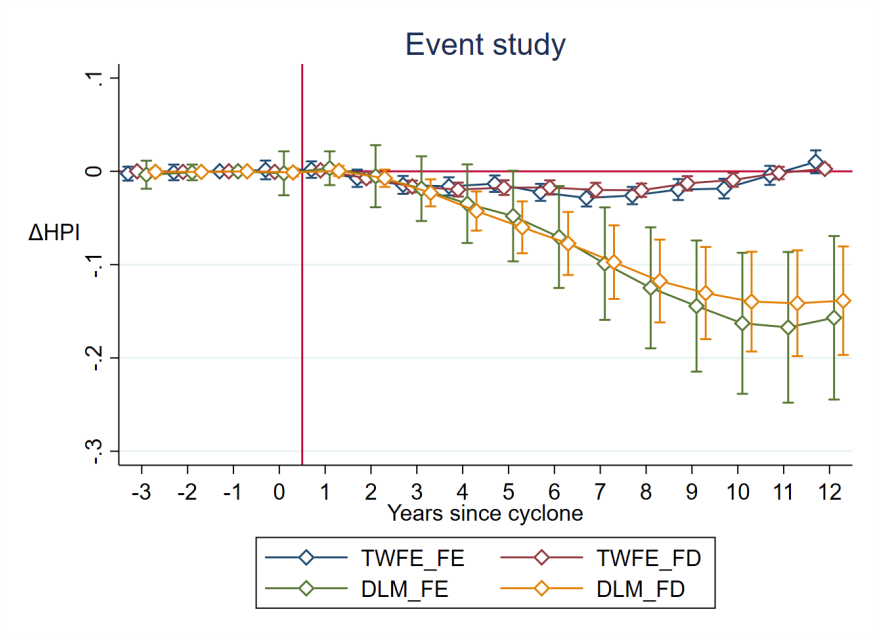
(b) Wind speed as a continuous variable

In Figure D.3a, a cyclone induces a negative impact on housing price change, amounting to approximately 15% after 9 years compared to a counterfactual trajectory in the absence of the event. Referencing Table 1, where the mean value of houses in coastal counties is $254,079, the counties experiencing cyclones witness an average decrease of $38,000 in property value after 12 years.
Additionally, we can analyze the same effect using wind speed as a continuous variable. A one

standard deviation increase in wind speed (equivalent to 10 m/s of wind exposure) results in a loss of $0.004 \times 10 = 4\%$ in housing value after 8 years. The maximum wind speed observed (64 m/s during Hurricane Katrina) implies a substantial 25% reduction in housing price growth.

To evaluate the differences between Fixed Effects (FE) and First Difference (FD) models, I applied both approaches to analyze the impact of cyclones on the housing market in directly affected counties. Here I show the outcomes of the Two-Way Fixed Effects (TWFE) model (as depicted in equation 13) and distributed lag model (DLM) (illustrated in equation 14), employing both Fixed Effect (FE) and first difference (FD) models.

Figure D.4: Comparison between fixed effect model and first difference model



The event study graph D.4 reveals no significant differences between employing a Fixed Effects (FE) or First Difference (FD) model, whether in the context of TWFE or DLM.

Model specification

To examine the potential bias of TWFE resulting from the inclusion of previously treated counties as controls, I conduct a comparison between the new Local Projection model for Difference-in-Differences (LP-DiD) and the results of the Distributed Lag Model (DLM) in equation 14 .

Figure D.5: Comparison between DLM and LP-DiD

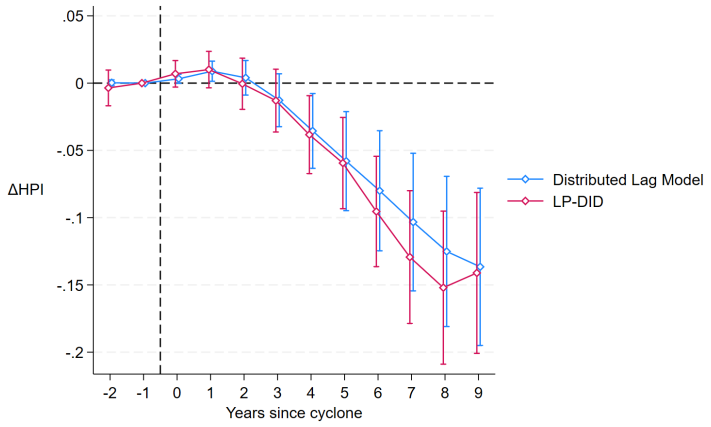
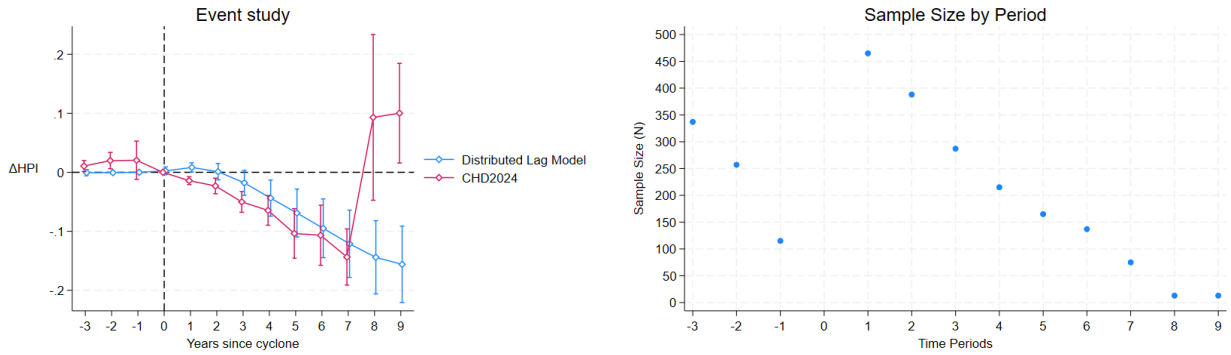


Figure D.5 demonstrates that the estimates obtained from the Distributed Lag Model (DLM) and the LP-DiD approach are statistically consistent, with no significant bias for the DLM. Moreover, Appendix D presents the results of the DLM using wind speed as a continuous variable to provide insight into the intensive margin estimates. In Appendix D, I also confirm that the difference between estimates derived from the Fixed Effects (FE) and First Difference (FD) models is not statistically significant.

I further test robustness to the estimator used by applying the method proposed by de Chaisemartin and D'Haultfœuille (2024). Figure D.6a compares the results from the Distributed Linear Model (DLM) with those from the de Chaisemartin and D'Haultfœuille (2024) estimator, while Figure D.6b illustrates the number of observations available for each lag estimate using the de Chaisemartin and D'Haultfœuille (2024) estimator. In particular, the de Chaisemartin and D'Haultfœuille (2024) estimator relies on substantially fewer observations beyond lag 7 and generally uses fewer observations overall compared to the DLM. This discrepancy arises because, unlike LP-DiD estimators, the de Chaisemartin and D'Haultfœuille (2024) estimator restricts the control group to counties that have not yet been treated. In contrast, LP-DiD estimators allow previously treated counties to serve as controls, provided they have remained cyclone-free for at least 8 consecutive years.

Figure D.6: Comparison between DLM and CHD2024



(a) Comparison of DLM and CHD2024

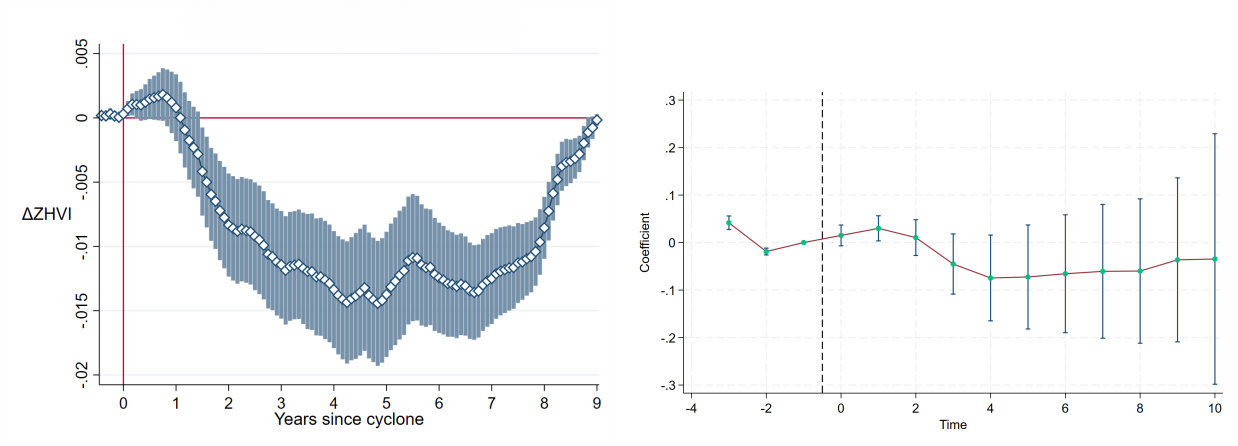
(b) Observations by each lag estimate

Similar to the LP-DiD approach, the de Chaisemartin and D'Haultfœuille (2024) estimator also identifies negative impacts of cyclones on the housing market. However, due to its reliance on fewer treatment units per period, it imposes stricter assumptions and limits the number of observations available for each lag, particularly for those beyond year 7. Despite these limitations, the consistent negative effects observed up to year 7 across both models further strengthens the conclusion that cyclones adversely affect housing market prices.

Data selection

To validate the main findings derived from the Housing Price Index (HPI) provided by the Federal Housing Finance Agency, I replicate the analysis using the Zillow Housing Value Index (ZHVI). Given that the ZHVI dataset is available at a monthly level, I first estimate a Two-Way Fixed Effects (TWFE) model to capture the month-by-month dynamics in the year following cyclone events (Figure D.7a). I then aggregate the ZHVI data to the yearly level (Figure D.7b) and I use the LP-DiD estimator to facilitate direct comparisons with the HPI.

Figure D.7: Event Study with Zillow Housing Value Index (ZHVI)



(a) Monthly Estimates

(b) Yearly (mean aggregation)

The ZHVI dataset, available from 2000 to 2021, has a narrower temporal range compared to the HPI, which spans from 1992 to 2021. This limitation contributes to larger standard errors, particularly beyond four years post-cyclone due to the smaller number of observations. While longer lags can provide valuable insights into the extended effects of cyclones, incorporating too many increases the risk of overfitting, potentially compromising the model's predictive reliability.

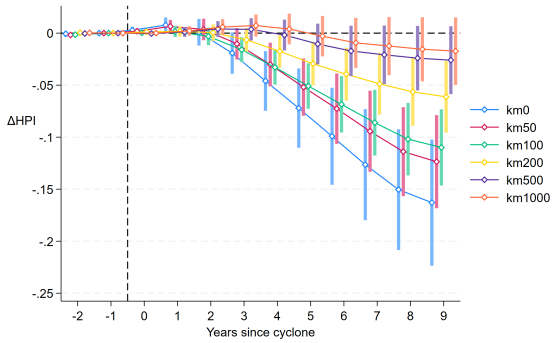
A notable difference arises when examining pre-treatment trends between control and treatment groups. I observe significant differences in the second and third years before the event— differences that are not evident in Figure D.7a, where the focus is limited to 12 months prior. This discrepancy may be attributed to the unique construction of the ZHVI, which captures a broader set of characteristics of the county housing market. Unlike the HPI, which is primarily based on repeated sales prices of the same housing units, the ZHVI incorporates a more comprehensive housing market perspective. This methodological distinction may explain the observed variations in housing market trends in the years leading up to a cyclone event.

Apart from these differences and limitations, the ZHVI results broadly align with the narrative of the housing market response to cyclone events, reinforcing the findings derived from the HPI analysis.

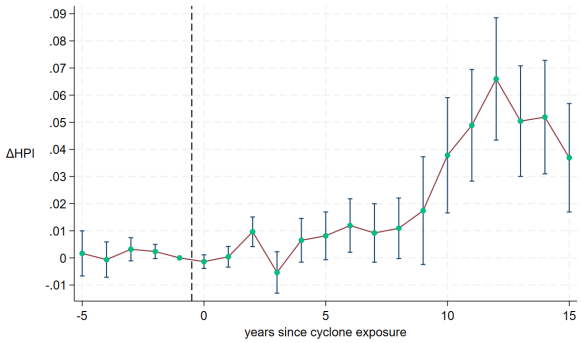
Sample selection

While Strobl (2011) restricts the analysis to US coastal counties when assessing hurricanes' effects on economic growth, other studies, such as Roth Tran and Wilson (2025), include all US counties. To examine how cyclone impacts vary with distance from the coast, I progressively incorporate inland counties using the distributed lag model, as shown for intermediate samples in Figure D.8a and the full US sample in Figure D.8b.

Figure D.8: Cyclone impacts on housing markets using different samples



(a) Cyclone impact by proximity to the Atlantic coast



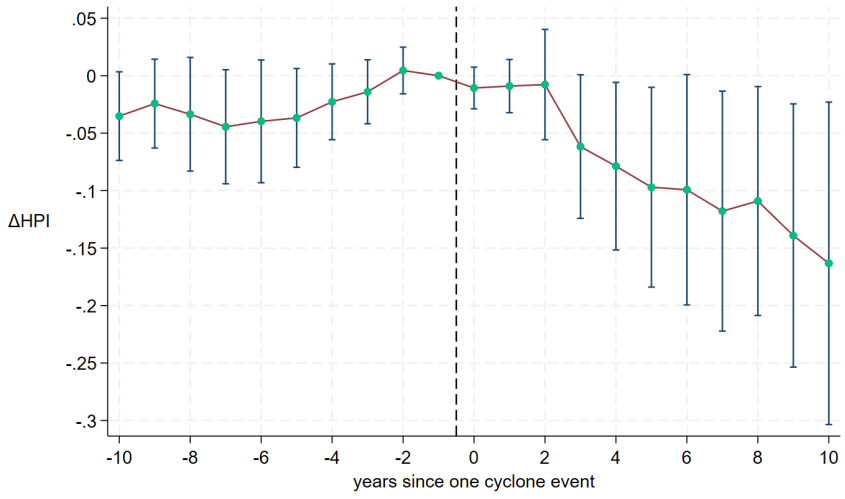
(b) Analysis across all US counties

Figure D.9: Progressive sample expansion and cyclone effects on housing prices. Panel (a) shows diminishing impacts as inland counties are included; Panel (b) reconciles the 1% housing price increase over an 8-year period reported by Roth Tran and Wilson (2025) when extending to all US counties.

As illustrated in Figure D.9, estimated cyclone effects attenuate as inland counties are incorporated into the sample (Panel a), culminating in the 1% housing-price increase over an 8-year horizon observed in the full US sample (Panel b), consistent with Roth Tran and Wilson (2025). This progression from negative coastal estimates to positive full-sample results may come from three factors. First, differential pre-trends introduce compositional bias: coastal counties exhibit stronger housing-price appreciation over time, attributable to demand for coastal amenities. Second, heterogeneous damage severity: coastal regions suffer compounded losses from both hurricane winds and storm surges, whereas inland effects are milder. Expanding the sample thus averages down the treatment intensity. Third, stark differences in exposure probability between coastal counties and inland. Prioritizing coastal counties therefore provides a more credible identification of cyclone damages.

However, to estimate the impact of a single cyclone event, it is necessary to extend the analysis to a radius of at least 200 km. This allows for comparison between counties that have experienced only one cyclone and those that have never been exposed. Within this 200 km range, only 2 counties experienced a single event, while 26 counties were never treated.

Figure D.10: The impact of one cyclone event



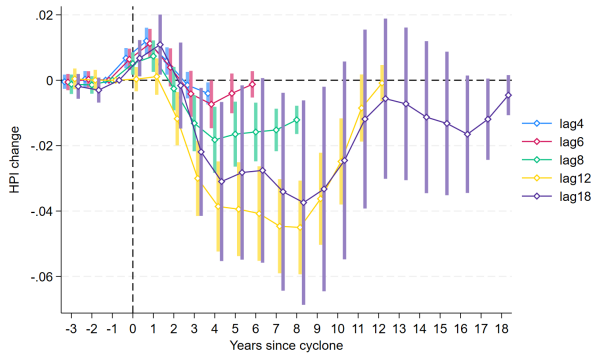
The results indicate a roughly 10% decline in housing prices following a single cyclone event. Unlike the main results shown in Figure 8, there is no evidence of subsequent recovery. This difference may reflect the role of repeated exposure: counties experiencing multiple cyclones may develop some degree of adaptation or resilience, whereas a single, unexpected event produces larger, more persistent losses. Alternatively, the absence of recovery could simply be due to the limited and different sample composition.

Selecting counties far away from the coast highlights a key trade-off: a clean control group (e.g. counties never treated away from the coast) versus one similar in treatment probability to high-exposure areas.

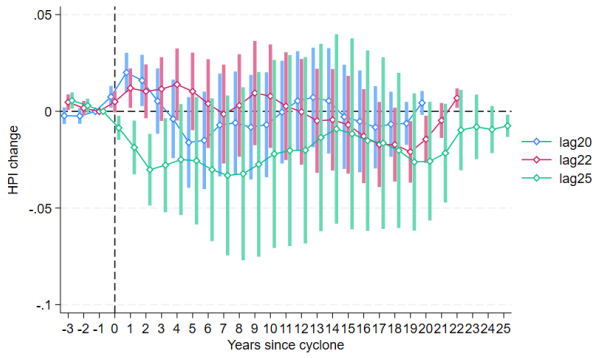
Lag length selection

To determine the optimal number of lags to include in Equation 13, I plot several models with varying lag lengths. For example, the model with "lag12" includes all lags from 0 to 12 in Equation 13.

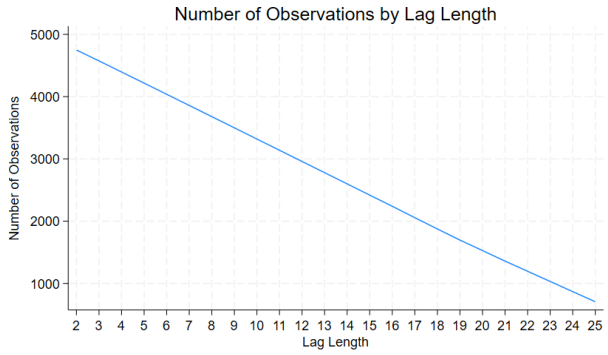
Figure D.11: Dynamic Two-Way First Difference Model with different lag length



(a) Lag selection from lag4 to lag18



(b) Lag selection from lag20 to lag25



(c) Number of observations for each lag length

As I increase the number of lags, I generally observe larger standard errors as a result of a reduction in the number of observations available for each model. Furthermore, in panel D.11b, from lag22, the placebo leads become statistically significant, leading to a lack of comparable estimates between the treated and control units.

In panel D.11a, most models initially show an increase in the housing value during the cyclone year, followed by a decline in subsequent years. The model with 12 lags appears to be an outlier, showing no increase in housing price in the two years following the cyclone.

To identify the best-fitting model, I perform a cross-validation analysis, evaluating the Mean Squared Error (MSE) for models with lag lengths from 2 to 20. The cross-validation procedure is a leave-one-out cross-validation (LOOCV) approach that computes the MSE for each lag selection model. For each lag length L, the steps are as follows:

1. **Data Restriction**: For each lag length L, I retain only observations that contain non-missing values for all lag variables up to L.

2. **Leave-One-Out Cross-Validation**: For each geographic unit i (e.g., a county), I treat it as

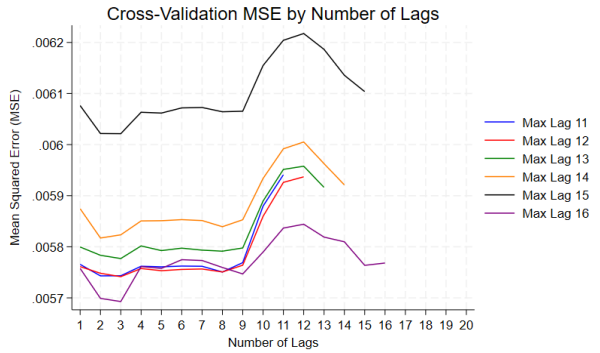
the validation set while using all other units as the training set. The model is fit on the training set, then predictions are made for the validation set, computing the squared prediction error for that unit.

3. **Mean Squared Error Computation**: This process is repeated for all units i, averaging their squared prediction errors to obtain the MSE for lag length L, which is given by:

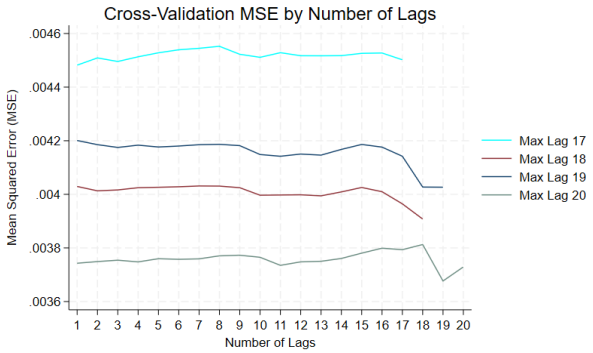
$$\text{MSE}(L) = \frac{1}{N} \sum_{i=1}^{N} \left(Y_i^t - \hat{Y}_i^t(L) \right)^2$$

where $\hat{Y}_i^t(L)$ is obtained by excluding unit i from the training data, fitting the model, and then predicting for Y_i^t. Repeating this for each lag length allows for a comparison of MSEs across models.

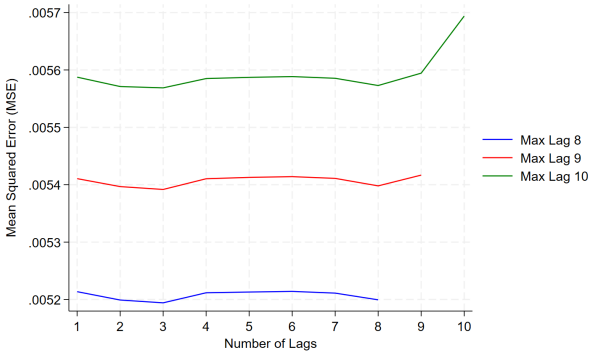
Figure D.12: Cross validation (MSE) by number of lags



(a) MSE from lag11 model to lag16



(b) MSE from lag17 to lag20



(c) MSE from lag8 to lag10

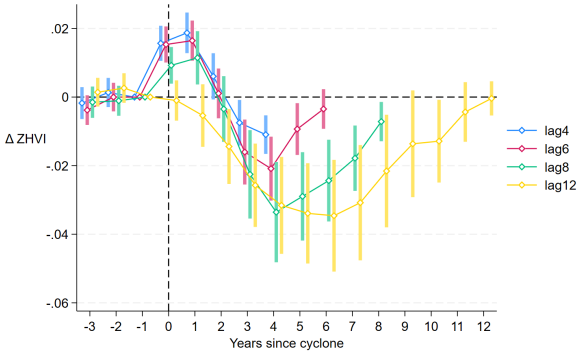
To improve clarity, I focus on results beginning from lag 11, as models with shorter lags exhibit relatively stable Mean Squared Error (MSE). As shown in panel D.12a, MSE values begin to increase after lag 8, indicating a decline in predictive accuracy across models as lag length extends up to lag 16. Consequently, I select lag 8 as the main specification. Beyond lag 12, MSE values increase further, suggesting greater imprecision driven by diminishing sample sizes and rising standard errors.

Although MSE continues to decline when lags extend beyond 15, the number of observations drops significantly, introducing greater uncertainty due to wider standard errors. These findings underscore the importance of data availability, as smaller datasets may underestimate long-run impacts due to limited predictive power at higher lag lengths.

Therefore, I concentrate on lags 8, 9, and 10. Analyzing these recursively, lag 10 yields a slightly better predictive accuracy than lag 9, which has a lower MSE than lag 10 when the maximum lag length is 10. Similarly, lag 9 exhibits a higher MSE than lag 8 when a maximum lag length of 9 is used. Lag 8, in turn, achieves a lower MSE than lag 7 when a maximum lag of 8 is applied. Based on this recursive assessment, I conclude my discussion of the long-run effects of cyclones at lag 8.

Finally, I confirmed similar patterns using data from Zillow, as shown below:

Figure D.13: Lag Selection with Zillow Data



Based on cross-validation results, I find that the long-run impacts of cyclones are best modeled with an eight-year lag in the TWFE specification. This evidence supports the assumption that the dynamic effects converge after eight years ($K^* = 8$) in the Local Projection framework (Equation 2).

Randomization tests

To check for potential model misspecification, which could lead to spurious or biased findings, I randomly select the sample to create false data and then re-estimate the model in Equation 13. I conduct this randomization in three distinct ways:

1. Entire sample: Each cyclone observation is randomly reassigned.
2. Between counties: Each county's complete history of cyclone exposure is randomly reassigned to another country while preserving the order of years. This approach maintains the temporal

structure within the data, testing whether global or regional trends might produce spurious correlations.

3. Within county: Each country's time-series of cyclone exposure is randomly re-ordered while remaining assigned to the original country. This method only changes the temporal structure of the data, testing whether time-invariant cross-sectional patterns across countries might lead to spurious correlations.

Figure D.14: Randomization of cyclones

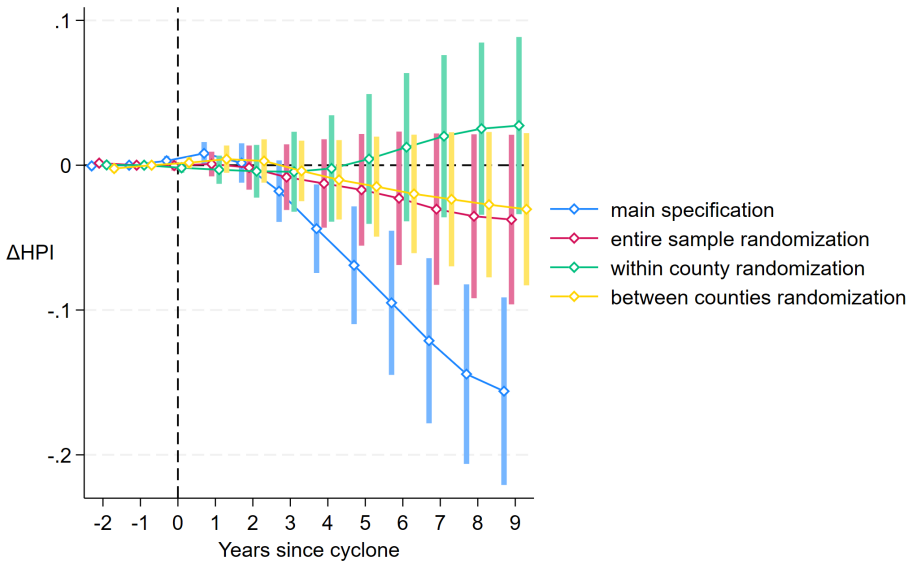


Figure D.14 shows that the three types of randomization yield insignificant results, thus reinforcing the robustness of my findings.

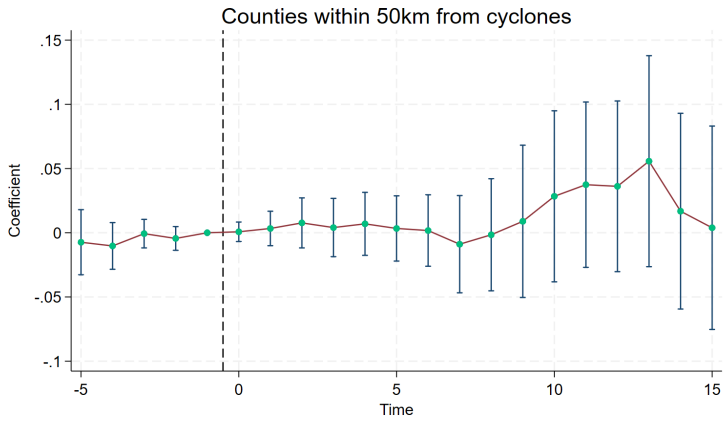
Cyclones impacts in neighboring missed counties

Regional economies are often deeply interconnected. Economic conditions in one area can influence those of its neighbors, even in the absence of direct shocks. This spatial interdependence raises concerns about potential spillover effects from natural disasters, which, if present, could violate the Stable Unit Treatment Value Assumption (SUTVA) and compromise causal identification (Kolak and Anselin, 2019). While this issue has been largely overlooked in the disaster economics literature, a few notable exceptions, such as Strobl (2011) and Lima and Barbosa (2019), explicitly incorporate spatial dependence into their econometric models.

To contribute to this limited but important body of work, I empirically test for spillover effects by examining whether counties neighboring cyclone-affected areas, specifically those located within a

50-kilometer radius, experience significant changes in housing values.

Figure D.15: Distributed Lag Model Results



Note: The red line represents the impact on counties located within a 50-kilometer radius from cyclone-affected areas.

The results indicate that these neighboring counties do not experience a statistically significant decline in housing values. This finding supports the validity of SUTVA in this context, suggesting that the treatment effect is indeed confined to directly impacted areas and that spatial spillovers, at least in terms of housing market outcomes, are minimal.

E Rent and housing quantity

Due to data limitations, the analysis of cyclone impacts on the rental market is conducted at the monthly level using data available only from 2015 onward. This restricts the temporal coverage but allows for a more granular examination of short-term market dynamics.

Figure E.1: Impact of cyclones on rent

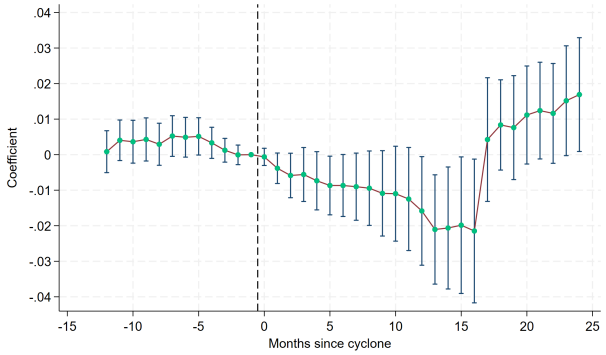
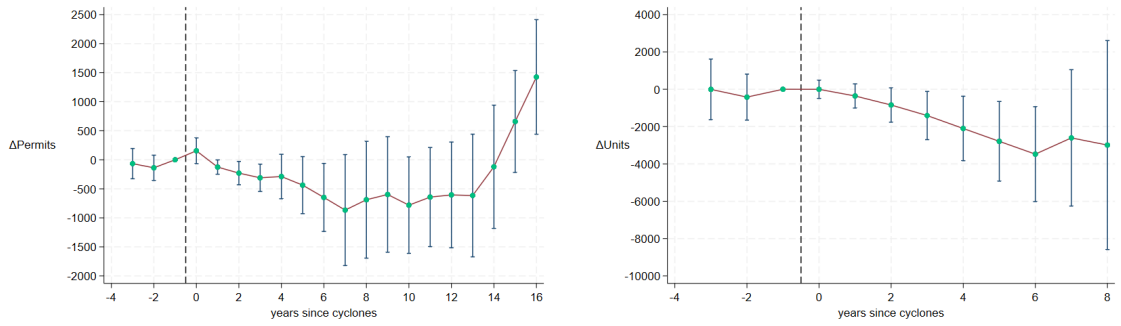


Figure E.1 presents an event study of the effect of cyclones on rent prices. The months leading up to the cyclone landfall exhibit no statistically significant changes in rent, supporting the assumption of no anticipation effects. Following a cyclone, however, rents decline noticeably, with a cumulative decrease of approximately 2% observed up to 15 months after the event. Starting around the 17th month post-cyclone, rents begin to recover and trend upward. Unfortunately, the time span of the data limits the ability to analyze longer-term recovery dynamics, as the series cannot be extended beyond the available period.

In addition to rents, I examine how cyclone events affect housing supply by analyzing both new residential construction permits and total housing units. While permits reflect forward-looking investment intentions, the estimated stock of housing units from Zillow captures realized changes in supply over time. Together, these measures offer complementary perspectives on how the housing market adjusts to climate shocks.

Figure E.2: Impact of cyclones on new housing permits and total housing units



(a) New construction permits

(b) Total housing units

Figure E.2 presents the event study estimates for both outcomes. Panel A shows that construction permits decline sharply in the year following a cyclone and remain significantly below baseline for approximately seven years—dropping by about 750 units, or 50% relative to the sample average (Table 1). From year 14 onward, permit issuance begins to rise steeply, potentially signaling delayed reconstruction or renewed development activity. Figure E.2 confirms that the total number of housing units also decreases following cyclone events, with levels persistently below the counterfactual for several years. This result indicates that reduced new construction translates into a slower-growing or shrinking housing stock, limiting supply in affected areas.

F Heterogeneous impacts of cyclones by income and housing distribution

To investigate how the impact of cyclones varies across different economic contexts, I employ two complementary approaches. First, I perform a subsample analysis based on average mean county income relative to the coastal area median (Figure F.2). Second, I use a quantile regression based on the distribution of housing values (Figure F.3). For the quantile regression, I implement a Distributed Lag model, as the Local Projection Difference-in-Differences (LP-DiD) framework does not yet support quantile estimation. As shown in Appendix D, the distributed lag model provides estimates that are not significantly biased relative to LP-DiD in the first nine years after the cyclone exposure.

Figure F.1: Impacts of cyclones by income and housing distribution

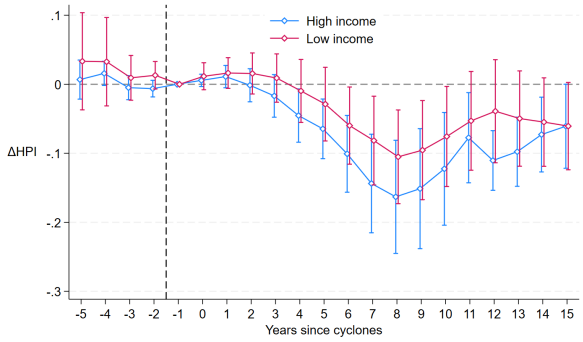


Figure F.2: By income

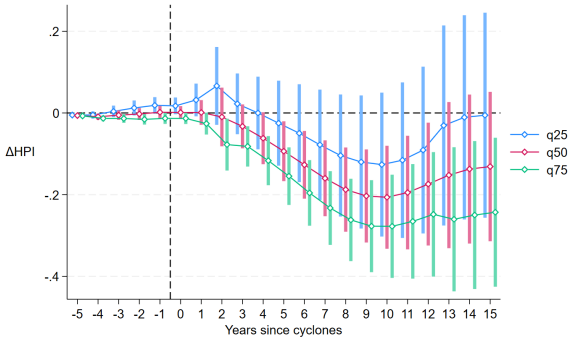


Figure F.3: By housing distribution

Note: Figure 1a is generated from a Local Projection Difference-in-Differences (LP-DiD) model, which estimates cumulative impulse responses over time. Figure 1b is generated from a standard Quantile Regression (QR) model incorporating county and year fixed effects, where the estimates represent contemporaneous effects.

Figure F.2 shows that high-income areas, defined as counties with average income above the median and typically more densely populated, experience larger declines in housing price growth following cyclone exposure. The difference is statistically insignificant (since confidence intervals overlap), but this pattern may reflect greater economic exposure in wealthier counties, such as larger disruptions to local employment and business activity, or higher perceived risk and recovery costs after extreme events. Consistently, the quantile regression in Figure F.3 confirms that counties with high-value housing markets are particularly affected by cyclones. However, in the heterogeneous analysis based on cyclone frequency (Figure 10), I find that counties with lower historical exposure suffer more persistent negative effects. Since housing values tend to be higher in areas with lower cyclone frequency (Table 1), the more persistent impact in high-income counties may just reflect the fact that cyclones are less expected and thus more disruptive in these areas.

G Derivations for the Model

This appendix derives the main optimality conditions for the household problem, the insurance decision rule, and the user-cost representation of housing prices.

G.1 Household problem

The per-period problem maximizes:

$$\begin{aligned}&(1-\tilde{p}_{it})\Big[u(y_{it}-P_{it}h_{it}-\pi_{it}I_{it})+\eta\ln(h_{it})-\alpha\tilde{p}_{it}\Big]\\&+\tilde{p}_{it}\Big[u(y_{it}-P_{it}h_{it}-\pi_{it}I_{it}-L_{it}+\varphi L_{it}I_{it})+\eta\ln(h_{it})-\alpha\tilde{p}_{it}\Big].\end{aligned}\tag{15}$$

Housing FOC: $P_{it}\left[(1-\tilde{p}_{it})u'(c_{it}^0)+\tilde{p}_{it}u'(c_{it}^1)\right]=\eta/h_{it}$.

For tractability, consider a representative household in county i facing a per-period problem in year t, taking beliefs $\tilde{p}_{it}$ and prices (P_{it},π_{it}) as given. Ignoring migration for now, the household chooses non-housing consumption c_{it}, housing services h_{it}, and insurance $I_{it}\in\{0,1\}$ to maximize expected utility

$$\max_{c_{it},h_{it},I_{it}\in\{0,1\}}\ (1-\tilde{p}_{it})\left[u(c_{it}^0)+\eta\ln(h_{it})-\alpha\tilde{p}_{it}\right]+\tilde{p}_{it}\left[u(c_{it}^1)+\eta\ln(h_{it})-\alpha\tilde{p}_{it}\right],\tag{16}$$

where superscripts 0 and 1 denote the no-cyclone and cyclone states, respectively.

The budget constraints in the two states are

$$c_{it}^0+P_{it}h_{it}+\pi_{it}I_{it}=y_{it},\tag{17}$$

$$c_{it}^1+P_{it}h_{it}+\pi_{it}I_{it}=y_{it}-L_{it}+\varphi L_{it}I_{it},\tag{18}$$

where y_{it} is income, L_{it} is the loss if a cyclone occurs, and $\varphi\in(0,1]$ is the indemnity rate.

Substituting (17)–(18) into the objective yields an expected-utility problem in terms of h_{it} and I_{it} only:

$$\begin{aligned}\max_{h_{it},I_{it}}\ &(1-\tilde{p}_{it})\left[u(y_{it}-P_{it}h_{it}-\pi_{it}I_{it})+\eta\ln(h_{it})-\alpha\tilde{p}_{it}\right]\\&+\tilde{p}_{it}\left[u(y_{it}-P_{it}h_{it}-\pi_{it}I_{it}-L_{it}+\varphi L_{it}I_{it})+\eta\ln(h_{it})-\alpha\tilde{p}_{it}\right].\end{aligned}\tag{19}$$

Treating I_{it} as given, the first-order condition for h_{it} is

$$(1-\tilde{p}_{it})u_c\big(c_{it}^0\big)\,(-P_{it}) + \tilde{p}_{it}u_c\big(c_{it}^1\big)\,(-P_{it}) + \eta\left(\frac{1-\tilde{p}_{it}+\tilde{p}_{it}}{h_{it}}\right) = 0, \tag{20}$$

which simplifies to

$$P_{it}\left[(1-\tilde{p}_{it})u_c\big(c_{it}^0\big) + \tilde{p}_{it}u_c\big(c_{it}^1\big)\right] = \frac{\eta}{h_{it}}. \tag{21}$$

Equation (22) states that the expected marginal cost of an additional unit of housing services (the left-hand side) equals the marginal utility gain from housing (the right-hand side).

Rearranging gives:

$$P_{it}\Big[(1-\tilde{p}_{it})u_c(c_{it}^0) + \tilde{p}_{it}u_c(c_{it}^1)\Big] = \frac{\eta}{h_{it}}, \tag{22}$$

or equivalently:

$$P_{it} = \frac{\eta/h_{it}}{(1-\tilde{p}_{it})u_c(c_{it}^0) + \tilde{p}_{it}u_c(c_{it}^1)}. \tag{23}$$

Definition of $MU_{h,it}$: The *marginal utility of housing services* is

$$MU_{h,it} \equiv \frac{\eta}{h_{it}}, \tag{24}$$

so (22) becomes:

$$P_{it} = \frac{MU_{h,it}}{(1-\tilde{p}_{it})u_c(c_{it}^0) + \tilde{p}_{it}u_c(c_{it}^1)}. \tag{25}$$

The denominator is the *expected marginal utility of consumption across states*, which links to user cost UC_{it} in the aggregate market-clearing condition (8)

Extension with endogenous income. If income follows (5), the household internalizes that a cyclone today lowers expected future income, $E_t[y_{i,t+s}|D_{it}=1] < E_t[y_{i,t+s}|D_{it}=0]$. This additional channel further reduces the expected marginal utility of future consumption in cyclone-prone locations, reinforcing the downward effect on housing demand described in the baseline model. For tractability, the main analysis keeps y_{it} exogenous, and treats (5) as a reduced-form extension that matches the empirically observed 5 percent long-run income loss.

G.2 Insurance decision rule

The insurance choice $I_{it} \in \{0, 1\}$ is made by comparing expected utility with and without coverage. Let $EU(I)$ denote expected utility when the household chooses insurance status $I \in \{0, 1\}$:

$$\begin{aligned} EU(I) = (1 - \tilde{p}_{it}) \left[u(y_{it} - P_{it} h_{it}(I) - \pi_{it} I) + \eta \ln(h_{it}(I)) - \alpha \tilde{p}_{it} \right] \\ + \tilde{p}_{it} \left[u(y_{it} - P_{it} h_{it}(I) - \pi_{it} I - L_{it} + \varphi L_{it} I) + \eta \ln(h_{it}(I)) - \alpha \tilde{p}_{it} \right]. \end{aligned} \tag{26}$$

The household chooses $I_{it} = 1$ if and only if

$$EU(1) - EU(0) \geq 0. \tag{27}$$

Under standard assumptions (e.g. concave $u(\cdot)$ and modest feedback of I on h_{it}), this condition can be approximated by a simple comparison between the premium and the expected utility value of coverage. A convenient sufficient condition is

$$\tilde{p}_{it}\, u_c(\bar{c}_{it})\, \varphi L_{it} \;\gtrsim\; u_c(\bar{c}_{it})\, \pi_{it}, \tag{28}$$

where $\bar{c}_{it}$ is a reference consumption level. Dividing both sides by $u_c(\bar{c}_{it})$ gives the intuitive inequality

$$\tilde{p}_{it}\, \varphi L_{it} \;\gtrsim\; \pi_{it}. \tag{29}$$

Condition (29) shows that insurance take-up is more likely when perceived risk $\tilde{p}_{it}$ and potential loss L_{it} are large relative to the premium π_{it}, which rationalizes larger post-disaster uptake in high-risk or newly shocked counties.

G.3 Housing stock and quality

The law of motion for the aggregate housing stock is

$$H_{i,t+1} = (1 - \phi) H_{it} - D_{it}\, \kappa(q_{it}) (1 - \bar{I}_{it}) H_{it} + R_{it}, \tag{30}$$

where ϕ is the depreciation rate, $\kappa(q_{it})$ is the damage share (decreasing in quality q_{it}), and $\bar{I}_{it}$ is the average insurance coverage rate. Reconstruction investment is given by

$$R_{it} = \chi \bar{I}_{it} D_{it} L_{it}, \tag{31}$$

where χ captures how efficiently insurance payouts are converted into rebuilding and quality improvements. Equations (30)–(31) imply that high-coverage counties experience faster recovery of H_{it} and potentially higher post-disaster quality q_{it}.

G.4 User cost and price

A simple asset-pricing relation for housing equates the current price P_{it} to the present value of future marginal utility flows from housing services and the continuation value of the asset:

$$P_{it} = \frac{1}{1+r} E_t \left[MU_{h,i,t+1} + (1-\delta_i) P_{i,t+1} - \Xi_{i,t+1} \right], \tag{32}$$

where $MU_{h,i,t+1}$ is the marginal utility value of housing services next period, δ_i captures physical and locational obsolescence, and $\Xi_{i,t+1}$ collects risk-related costs such as disaster losses not covered by insurance and expected insurance premia.

Rearranging (32) in user-cost form yields

$$P_{it} = \frac{MU_{h,it}}{UC_{it}}, \tag{33}$$

with user cost

$$UC_{it} = r + \delta_i + \tau_{it} + \pi_{it} - E_t \left(\frac{\Delta P_{i,t+1}}{P_{it}} \right), \tag{34}$$

where τ_{it} represents local taxes and fees and π_{it} summarizes risk-related premia. Cyclone events increase perceived risk $\tilde{p}_{it}$ and thereby raise π_{it} and τ_{it} over time, pushing up user cost and putting downward pressure on prices, even as short-run destruction and insurance-financed reconstruction can temporarily raise prices.

G.5 Implications for dynamics

Combining (22), (29), and (30)–(31) with the user-cost expression yields the qualitative dynamics observed in the data:

(i) A cyclone shock $D_{it} = 1$ reduces H_{it}, tightening supply and raising P_{it} for several years while reconstruction proceeds.

(ii) In high-risk, highly insured counties, large R_{it} and higher q_{it} generate a pronounced rebound in prices once rebuilding is underway.

(iii) In less-experienced, low-insurance counties, the shock induces a large increase in perceived risk $\tilde{p}_{it}$ and premia π_{it}, but limited reconstruction R_{it}, leading to more persistent declines in housing values as risk is increasingly capitalized into user cost.