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

Davide Castellani†

Job Market Paper

October 16, 2024

Abstract

Despite tropical cyclones causing around $1.3 trillion in damages from 1980, surpassing all other recorded weather disasters in U.S. history, there remains a lack of empirical evidence regarding their long-term effects on the housing market. Here I investigate the impact of tropical cyclones on the U.S. housing market from 1988 to 2021, utilizing a novel dataset. Employing HURDAT2 data from the National Hurricane Center and a radial hurricane winds model, I develop a new measure of wind speed and I offer the first insight on the overall impact of all recordings cyclones at the county-yearly level. I use a doubly robust estimator that combines propensity score weighting and two-way fixed effect model, to compare counties with similar cyclone occurrence probabilities and estimates dynamic effects of cyclones on the housing value. Despite an initial increase in housing value following cyclones, findings reveal a significant 15% decline in the housing market over a 9-year period. The observed effects are not driven by factors such as employment, income, and migration. Instead, it appears that the driving force behind the effect on housing lies in the increased prevalence and elevated costs of insurance policies. This particular finding serves as a compelling impetus for delving further into the exploration of cyclones' impacts on risk perception, as opposed to their direct effects on the economy.

JEL-Codes: Q54, R31, G22, D81

*I am grateful to Giovanni Mastrobuoni and Elena Esposito 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, Solomon Hsiang, Enrico Moretti, Giorgio Chiovelli, Scott Cunningham, Joseph S. Shapiro, Jeffrey Wooldridge, Òscar Jordà, Kirill Borusyak, James Sallee, Christoph Albert, Daniele Paserman, Joel Ferguson, people of the Workshop on Ibeo 2023, of the ECCD Workshop at GSSI 2024, of the Global Policy Lab (Berkeley), of the Agriculture and Resource Economics dept. (Berkeley) seminar. Any remaining errors are solely my responsibility.

†PhD Candidate, Department of Economics and Statistics, University of Turin and Collegio Carlo Alberto. Email: davide.castellani@carloalberto.org

1 Introduction

While the impact of climate change on the frequency of tropical cyclones is still debated, there is now evidence that these storms are becoming more intense and dangerous. Rising global temperatures elevate sea surface temperatures, which in turn supply tropical cyclones with increased energy from the ocean (Intergovernmental Panel on Climate Change (IPCC), 2021). Furthermore, warmer air retains more moisture, leading to heavier rainfall. Cyclones are also stalling more frequently, and their slower translation speed increases the risk of severe flooding, storm surges, and prolonged exposure to high winds (Trepanier et al., 2024). Sea level rise is projected to further intensify storm surges, increasing the potential for coastal damage (Collins et al., 2019).

According to National Oceanic and Atmospheric Administration (NOAA), from 1980 the cost of damages from weather and climate disaster exceeds $2.6 trillion. Of all recorded weather disasters in U.S. history, cyclones have caused the most damages: over $1.3 trillion 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, contingent upon the severity and frequency of storms. For example, in 2020, the United States experienced a particularly active cyclones season, with a record-breaking 30 named storms, including 13 hurricanes. According to FEMA, the federal government spent over $26 billion in response to these storms, including funds for emergency measures such as search and rescue operations, temporary housing, and debris removal, as well as longer-term recovery and rebuilding efforts. Notably, predictions indicate that the most substantial economic losses stemming from climate change in the coastal area are anticipated to arise from the intensification of cyclones and sea level rise (Hsiang et al., 2017).

Living in coastal areas offers amenities, but it also entails a significant risk, particularly in light of climate change. Yet, the risk of rising sea levels hasn't influenced the real estate market as of now (Murfin and Spiegel, 2020) and the overall effects of cyclones on the housing market are not clear at first glance. A significant disruption in housing supply can have notable effects on the real estate market, but predicting the overall impact is challenging. On one hand, a shortage of housing is likely to cause a temporary increase in prices. Initially, housing markets experience a quick rise in prices due to a post-storm shortage, followed by a later adjustment as the supply returns to normal. However, cyclones can also disrupt local economic activities, affecting income and influencing housing demand. Moreover, there is a potential for increased out-migration as individuals seek areas less susceptible to cyclone-related risks. Indeed, after a hurricane, if people adjust their perception of the likelihood of future cyclones, it may reduce the attractiveness of hurricane-prone areas and lead to a lasting decrease in housing demand.

Utilizing a novel dataset sourced, HURDAT2, from National Hurricane Center and model for radial

profile of hurricane winds (Holland et al., 2010), I have conducted a comprehensive analysis of the impact of cyclones on the continental United States housing market spanning the period from 1988 to 2021. For the first time, HURDAT2 documented distinctive data characteristics at diverse distances from the hurricane's center. I then aggregate the wind speed into broader county-by-year units to align it with housing data. Indeed, for the housing market analysis I utilized the Housing Price Index from the Federal Housing Finance Agency and the Zillow Housing Value Index, specifically at the county level.

Despite having high-quality wind speed data, significant empirical challenges remain. Estimating the causal effect of cyclones requires identifying a credible comparison group that represents what would have occurred in affected counties had they not experienced cyclones. To leverage the exogenous variability of cyclone events, I ensure that both the treatment and control groups share comparable probabilities of exposure to a cyclone. To achieve this, I restricted the sample to Atlantic coastal counties and employed propensity score weighting in combination with a Two-Way Fixed Effects (TWFE) model . This method allows for a robust comparison of counties with similar cyclone exposure probabilities and enabled me to estimate the dynamic effects of cyclones on housing values. The main result show that experiencing a new cyclone leads to a significant 15% decline in housing market values over a 9-year period compared to unaffected counties. This decline reveals a far greater persistence compared to other natural disaster like wildfires (McCoy and Walsh, 2018). This finding also contrasts with earlier research on climate-related shocks, which often highlighted positive effects on housing values (Graff Zivin et al., 2023; Murphy and Strobl, 2010).

In the context of analyzing the effects of natural disasters using Difference-in-Differences (DiD), a key issue highlighted by recent literature (Sun and Abraham, 2020) is that some comparisons may be 'unclean.' This occurs when previously cyclone-affected areas are used as control groups for regions newly impacted by cyclones. Such 'unclean' comparisons can lead to biased estimates because the control group may already have experienced changes due to earlier cyclone exposure, which can confound the observed effects for newly impacted areas. Here, I show that the long-run effects of cyclones on the housing market are robust to the new local projection estimator applied to event study (Jordà et al., 2023) and to the last Difference-in-Differences Estimator for intertemporal treatment effects (de Chaisemartin and D'Haultfœuille, 2024).

Regarding heterogeneous effects, a spatial econometric model show that neighboring counties not directly impacted by the cyclones still exhibit negative effects, albeit to a lesser extent.

This paper makes two substantial contributions. Firstly, it offers 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. While existing literature has focused on this aspect (Nordhaus, 2010; Hsiang, 2016; Bakkensen et al., 2018; Hsiang et al., 2017), the current study builds on the measures developed by Hsiang and Jina (2014) and Deryugina (2017). Unlike Deryugina (2017), which utilized the "maximum wind speed radius" to identify affected counties, this study employs a higher resolution approach that incorporates multiple radii from the cyclone. This refined metric facilitates a more nuanced analysis of cyclone impacts, capturing greater variability and including a wider range of affected counties. By adopting this method, I was able to minimize measurement errors, as counties that were further away may have been inaccurately categorized as controls while actually experiencing impacts, and vice versa. Moreover, compared to Deryugina (2017) and Graff Zivin et al. (2023), this study include all cyclones occurring between 1988 and 2020, without restricting the analysis solely to hurricanes (defined as cyclones with a minimum wind speed of at least 33 meters per second). This is crucial to account for events like Hurricane Sandy. When Hurricane Sandy reached New York, it was no longer classified as a hurricane but rather as a post-tropical cyclone. This reclassification occurred because Sandy had merged with another weather system and lost many of its tropical characteristics, although it still had strong winds and caused significant damage (Ortega and Taşpınar, 2018).

The second noteworthy contribution focuses on elucidating the aggregate spatial effects of extreme natural events on the the housing market. Despite a multitude of single-event studies investigating the influence of hurricane strikes on housing prices, the empirical evidence has thus far produced inconclusive results. For instance, Biny and Polasky (2004) observed a reduction in the market value of houses following Hurricane Floyd or, another example, Ortega and Taşpınar (2018) after Hurricane Sandy. On the other hand, Graham et al. (2007), Beracha and Prati (2008) found no discernible effect on house prices when examining larger areas. In contrast, Deryugina et al. (2018) found that Hurricane Katrina led to a rise in house prices in directly affected areas. Graff Zivin et al. (2023) investigated the mixed results regarding the impact of hurricanes on the housing market by examining data from the Florida housing market between 2000 and 2016. They discovered that housing prices experienced a significant increase for three years following a hurricane. However, their findings did not provide evidence for other states and for all tropical cyclones. Moreover, recent literature in the difference-in-differences estimators suggest that the observed short-term positive effects of cyclones on housing demand may be misleading.

The rest of the paper proceed as follow. Section 2 shows 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 disaster

The quality of prior estimates are 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; Strömberg, 2007; Kellenberg and Mobarak, 2008; Noy, 2009; Hsiang and Narita, 2012), factors which also affect growth and thus might confound these results. Similarly also data from the Federal Emerging Management Agency (FEMA). I overcome the challenges of omitted variables bias and endogenous disaster reporting by developing a novel data file describing year-to-year variation in each country's physical exposure to disaster. To do this, I focus on tropical cyclones, the class of natural disaster that includes hurricanes, cyclones and tropical storms with at least 17 m/s wind speed and reconstruct every storm observed on the contiguous US 1988-2020. 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 must support the formation of thunderstorms, which 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 have 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 to 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 construct a novel measure of wind speed which

constitutes the first contribution of this paper [1]. Here, I combine the Hurricane Data second generation (HURDAT2) dataset from the National Hurricane Center [2] with a parametric approach to modeling the radial profile of winds in hurricane, Holland et al. (2010). Indeed, the HURDAT2 database contains six-hourly information for each North Atlantic cyclone, detailing the storm's location and characteristics. Then, I have adopted the methodology outlined by Holland et al. (2010) to effectively integrate this information and accurately predict wind speed at specific locations. 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.

Compared to the previous measures used by Deryugina (2017) and Hsiang and Jina (2014), this new measure of wind speed provides variations in wind intensity across different radii. Indeed, 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:

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

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):

$$b_s = -4.4 \times 10^{-}5\Delta p_s + 0.03\frac{\partial p_{cs}}{\partial t} - 0.014\varphi + 0.15 v_t^{x_1} + 1 \tag{2}$$

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

[1] The measure has been developed on the insights of professor Tatyana Deryugina

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

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[3], and $x_1 = 0.6 \times (1 - (\Delta_{ps}/215))$. It is important to note that the wind speed v_s is positively influenced 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:

$$x_d = \frac{\ln\left(\frac{v_s}{v_d}\right)}{\ln\left(\left(\frac{r_{v_d}}{r_d}\right)^{b_s} e^{\left[1-\left(\frac{r_{v_d}}{r_d}\right)^{b_s}\right]}\right)} \tag{3}$$

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.

Having described Holland's (2010) parametric model, I will now demonstrate its application using the HURDAT2 data. To illustrate, I will use Hurricane Katrina, the most damaging hurricane in U.S. history. To visually depict the reconstructed wind speed profile of Hurricane Katrina, I have presented various interpolations in Figure 1. 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

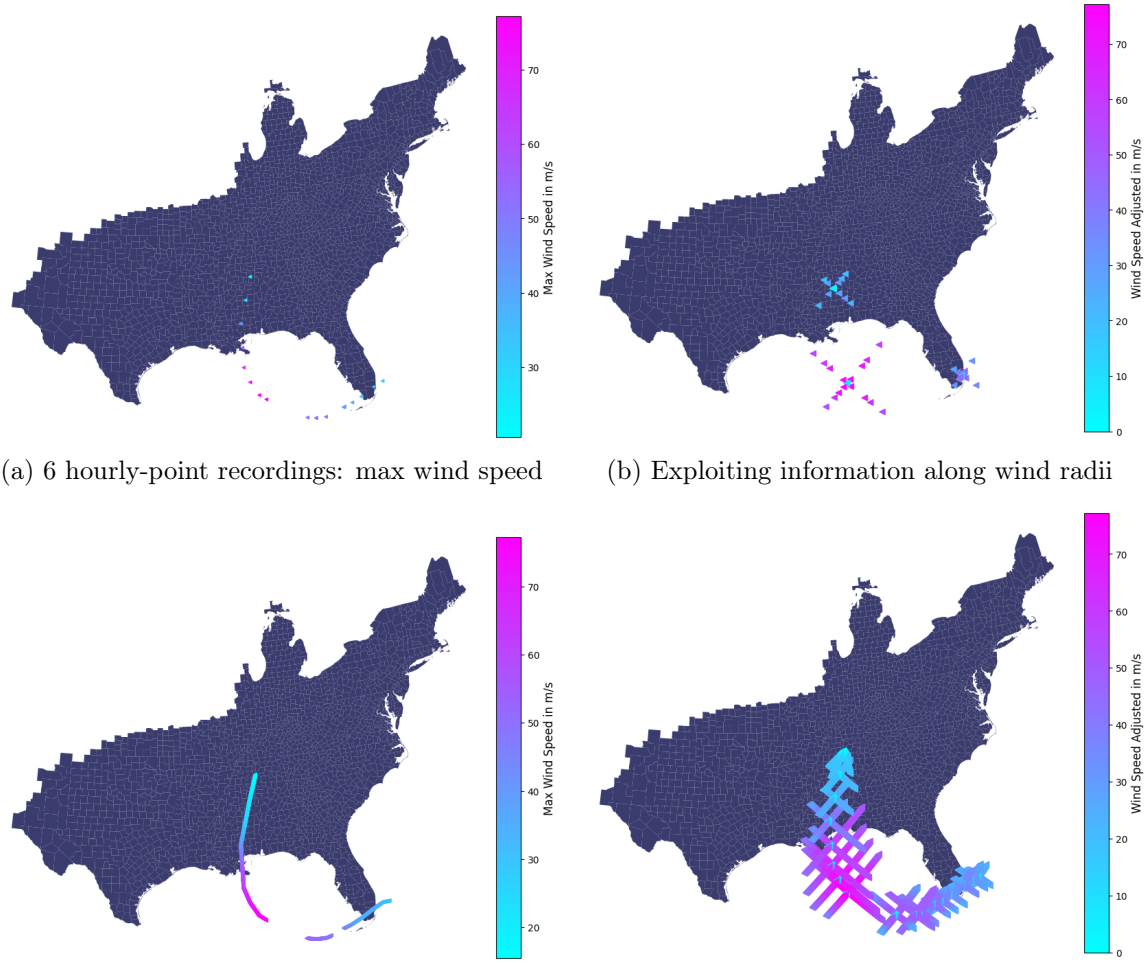
[3]I calculate the speed dividing the distance between two subsequent recording points by the time interval of 6 hours

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 (Equation 1 take into account this to construct the "wind speed adjusted" v_s close to $r = 0$). 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 3. 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 (where the wind speed is adjusted based on the parameter x_d in Equation 3). The extent of wind radii varies from observation to observation, and capturing this variation is the primary contribution of this new measure of wind speed v_s compared to methods used by Deryugina (2017) and Hsiang & Jina (2014).

To fill the gaps between the 6-hourly recording points and between the different recordings of the maximum wind radii, I use linear interpolation. This is done 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 observations for every kilometer, generating an "expansion" that covers the distance between two consecutive points and 20 new interpolated points.

To obtain all the new values for Equation 1, I apply linear interpolation, adding a proportional amount of the difference (delta) between the original points. When I interpolate along the trajectory (Panel 1c), new observations are also implicitly created along the wind radii, similar to those shown in Panel 1b. The remaining interpolation is performed along the different wind extents. Panel 1d illustrates this for the starting 6-hourly recording points, but this process is repeated for all interpolated points along the main trajectory as well.

Figure 1: Reconstruction of wind speed for Hurricane Katrina



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

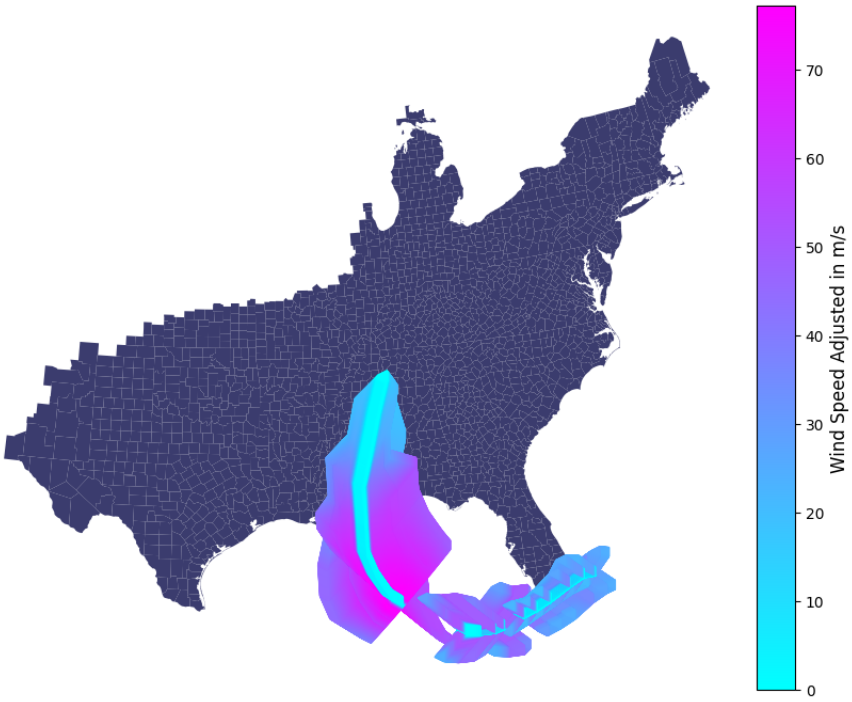
(b) Exploiting information along wind radii

(c) Interpolations along the trajectory

(d) Interpolations along the wind radii

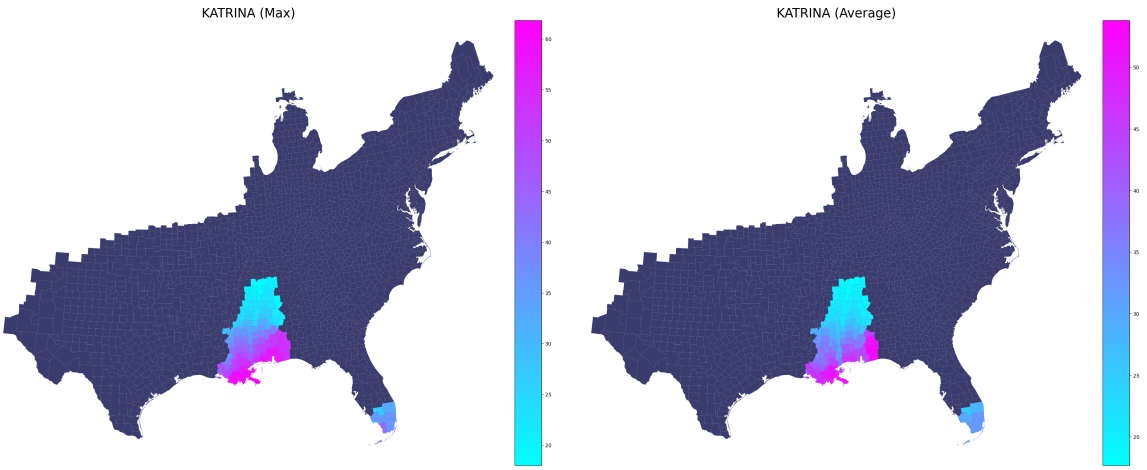
By combining the two interpolations shown in Panels 1c and 1d, I achieve the final reconstruction of wind speed for all observations of Hurricane Katrina. This comprehensive approach ensures that both the main trajectory and the variations in wind extents are accurately represented, providing a detailed and continuous depiction of the hurricane's wind speed dynamics.

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. As a robustness check, I have also utilized the average wind speed, as detailed in the Appendix. The corresponding visual representation is presented in the figure 3.

Figure 3: Aggregate the wind speed at county level



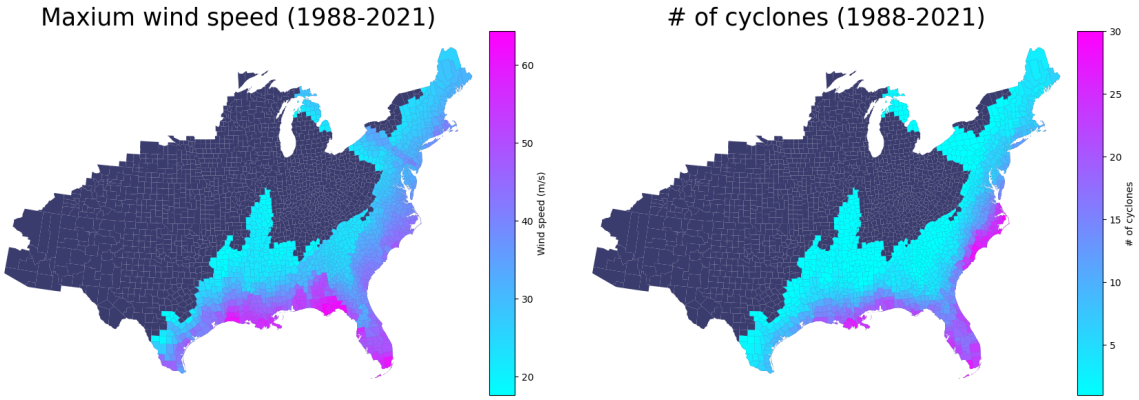
(a) Maximum wind speed of Katrina by county (b) Average wind speed of Katrina by county

Hurricane Katrina stands out as one of the most devastating events in US history, resulting in 1,833 fatalities and causing approximately $108 billion in damages. Despite its significant impact, there has been a lack of comprehensive analysis for the numerous cyclones that have occurred since 1988. In this study, I have gathered information on all 126 cyclonic events in the US during this period. The

wind speed in each location for each event has been reconstructed followed 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 span from 1988 to 2021. Panel b illustrates the frequency of cyclones experienced in each location. For a further visualization of variability by each year see Appendix A.2.

Figure 4: Aggregate the wind speed at county level



(a) Maximum wind speed of cyclones by location

(b) Frequency of cyclones by location

In constructing the panel dataset at the county-year level, I have opted to use the maximum wind speed for the main specifications. Moreover, to fortify the robustness of our analyses and account for potential heterogeneity, I incorporated the number of cyclones and the mean value of wind speed as supplementary checks in Appendix. The following table report the summary statistics for the characteristics of the new measure of wind speed aggregated at county level. Note that the sample comprises counties that have encountered at least one cyclone within the last 30 years.

Table 1: Summary statistics for Mean Wind, Max Wind, and Number of Cyclones

Variable	Mean	Variability	sd	Min	Max	Observations
Mean Wind	3.969922	Overall	10.00872	0	62.53913	N = 38819
		Between	3.853023	0.3808099	21.90429	n = 1110
		Within	9.2775	-17.93437	58.50084	T = 35
Max Wind	4.103458	Overall	10.42781	0	64.32609	N = 38819
		Between	4.058712	0.3808099	22.75997	n = 1110
		Within	9.645841	-18.65651	58.49398	T = 35
# of Cyclones	0.1755841	Overall	0.4674257	0	5	N = 38819
		Between	0.1770901	0.0208333	0.962963	n = 1110
		Within	0.4339346	-0.7873788	4.632106	T = 35

Notes: n = countries; T = # of average years per county; N = observations; Within variability = temporal variability in a given country; Between variability = variability across counties; sd = standard deviation.

The variability in the number of cyclones within a county ranges from 0 to 5. Upon examining the geographical map, it becomes evident that counties exhibiting a higher frequency of cyclones also coincide with regions characterized by more intense occurrences of these extreme weather events.

2.3 Housing value

The key measure employed to evaluate the impact of cyclones on the housing market is the Housing Price Index from the Federal Housing Finance Agency (FHFA HPI)[4]. Providing a detailed perspective, the FHFA HPI utilizes extensive data to offer insights into house price dynamics at the county level. The FHFA HPI operates as a weighted, repeat-sales index. This means it captures average price fluctuations in properties through repeated sales or refinancing, maintaining a focus on the same properties over time. The data originates 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 the sample composition, ensuring that changes in the estimated index primarily reflect genuine market dynamics rather than shifts in the types of properties 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 estimation of their individual contributions to the overall index. The Housing Price Index (HPI) initiates at a value of 100 for the initial period when a county enters the dataset. Consequently, our analysis focuses on the year-over-year change in prices, denoted as ΔHPI, which captures variations relative to the preceding year.

I incorporated the Zillow Value Housing Index (ZVHI) [5] as an additional metric in this study. The ZVHI is crafted 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 through an ensemble of sophisticated machine learning models. This estimation process integrates diverse data sources, encompassing public datasets, user-generated inputs, 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 consideration 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 monthly and county level and give a value in dollars. However, this variable is not stationary and I have also used the ΔZVHI in the following analyses.

In the following table I present the summary statistics for both HPI and ZVHI.

The housing inflation in the period

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

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

Table 2: Summary statistics for housing value

Counties that have experienced at least one cyclone

Variable	Mean	Variability	sd	Min	Max	Observations
HPI	235	Overall	160	54	1828.07	N = 36562
		Between	100	74	683	n = 1040
		Within	118.6	-348.6	1381	T = 35
π%HPI	3.9	Overall	6.27	-45.36	75.97	N = 35452
		Between	1.015343	0.88	11.02	n = 1040
		Within	6.2	-47.11	75.5	T = 35
ΔHPI	0.04	Overall	0.06	-0.6	0.56	N = 35452
		Between	0.01	0.007	0.11	n = 1040
		Within	0.06	-0.6	0.6	T = 35
ZVHI	169811	Overall	113273	29939	2287031	N = 224980
		Between	98283	33168	1347789	n = 1079
		Within	42622	-483078	1109053	T = 208
ΔZVHI	0.003	Overall	0.007	-0.222	0.1	N = 223847
		Between	0.001	-0.003	0.02	n = 1079
		Within	0.007322	-0.2	0.1	T = 207

Coastal counties (Adjacent to Atlantic Ocean or Gulf of Mexico)

ZVHI	254079	Overall	163690	42990	2287031	N = 43351
		Between	146853	66432	1347789	n = 183
		Within	67833	-398811	1193321	T= 237
ΔZVHI	0.004	Overall	0.01	-0.2	0.13	N = 43156
		Between	0.001	-0.002	0.011	n = 183
		Within	0.01	-0.2	0.12	T = 236

Notes: π% is the inflation in the HPI; $\Delta HPI = \ln HPI_t - \ln HPI_{t-1}$; $\Delta ZVHI = \ln ZVHI_t - \ln ZVHI_{t-1}$.

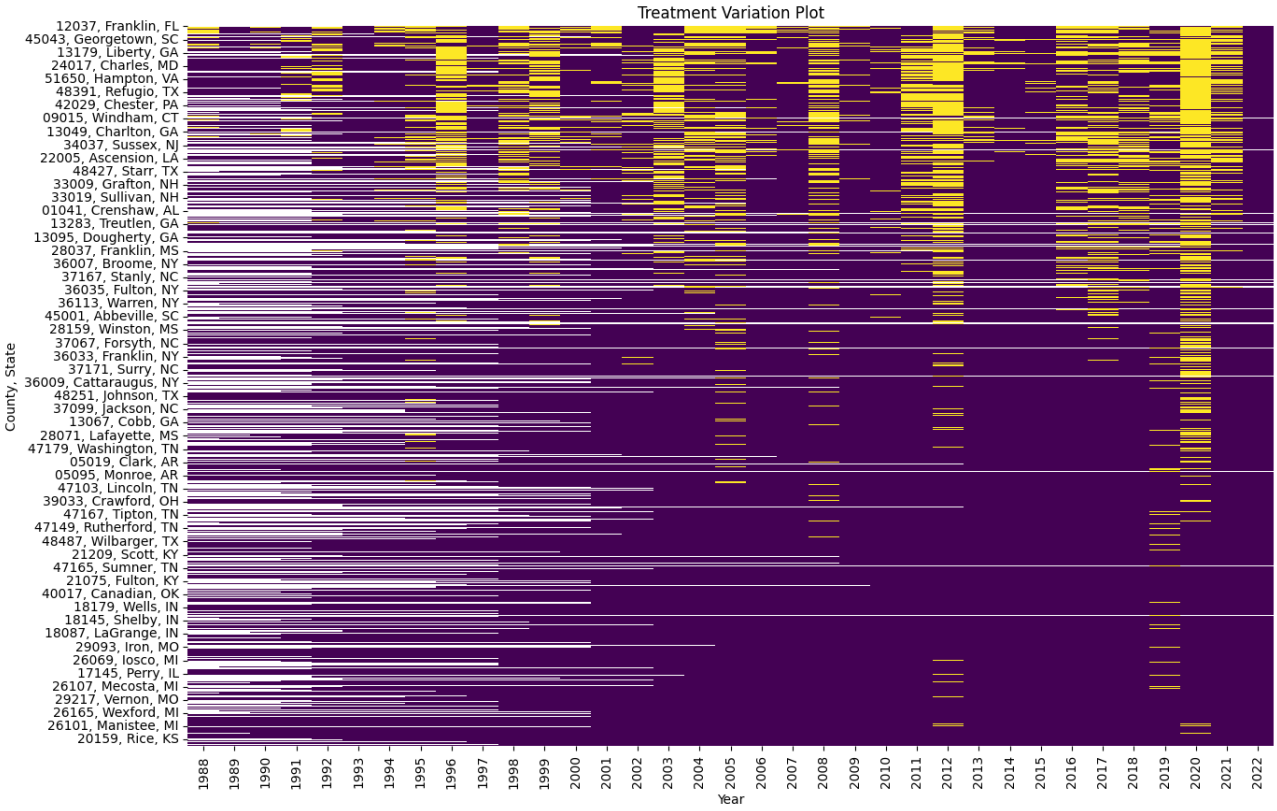
2.4 Other variables

3 Research design

3.1 Identification strategy

To investigate the relationship between the housing market and cyclones, I leveraged 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's essential to acknowledge that not all U.S. counties share the same probabilities of being hit. In Figure 5 the yellow shading indicates instances when counties are affected by cyclones in a specific year. At this stage, my treatment variable (wind speed) is binary, where it equals 1 if a cyclone hits the county and 0 otherwise.

Figure 5: 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.

Notably, the counties are ordered in ascending order, from top to bottom, based on distance to the Atlantic coast. Consequently, those at the top experience more cyclones than those at the bottom. In addition, I conducted a spatial analysis examining the varying intensities and frequencies of cyclones by year (see Appendix A.2). Proximity to the coast correlates with both more frequent and more intense cyclones. Consequently, I bifurcated my analysis into two distinct samples: one comprising exclusively coastal counties, i.e., those bordering the Atlantic Ocean or Gulf of Mexico (in accordance

with the definitions of Derugyna (2017) and Strobl (2011)); the other including all counties within 1000 km from the coast (because at this distance at least one county has experienced a cyclone). To address the variability in the probabilities of cyclone impact across different counties, I introduced inverse probability weighting into in all subsequent regression analyses. This approach generates control observations with more comparable probabilities of experiencing a cyclone compared to the treatment observations. I computed the probability p by 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 have applied the following weights:

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

Given the restricted sample and the inverse probability weighting, it's now reasonable to consider the following assumption:

Assumption 1. *The occurrence of a cyclone in a specific county and year can be treated as exogenous*

To estimate the average treatment effect in the hit counties (ATT), I've run an analysis utilizing a two-way fixed-effect (TWFE) model within a traditional event study framework.

$$\ln y_{i,t} = \alpha_i + \delta_t + \sum_{l=-Q}^{L} \gamma_l^{TWFE} D_{i,t}^l + e_{i,t} \tag{4}$$

where $\ln y_{i,t}$ is the logarithm of housing price in county i and year t; α_i is the county fixed effect; δ_t is the year fixed effect; $D_{i,t}$ is a dummy=1 if the county c experiences a cyclone in year t; γ_l^{TWFE} is the effect of receiving a cyclone l years before (leads) and after (lags) the cyclone and the contemporaneous effect (when $l = 0$).

The TWFE estimator γ_l^{TWFE} alone does not represent a design-based, nonparametric estimation strategy for causal inference (Imai & Kim, 2021). More specifically, the weights are proportional to group sizes (i.e. how many counties enter the treatment in each year) and treatment indicator's variation, where units in the middle of the panel have the highest weight [6]. We do need functional-form assumption for the TWFE to work (Imai & Kim, 2021):

Assumption 2. ***Linear additive assumption:*** *the effect of each predictor variable on the outcome variable is additive and independent of other predictors*

When there are multiple variables (e.g. leads and lags in equation 4), their total influence on the outcome is best stated by combining their effects together. To evaluate the cumulative effects over

[6] Note that I've refined the weights using the Inverse Propensity Score Weighting method. This adjustment prioritizes units that are more similar in terms of their probability of experiencing a cyclone, thereby enhancing the reliability of the analysis

time, I have used the Distributed Lag Model (DLM). Drawing from the insights of Schmidheiny & Siegloch (2023), I can establish the equivalence between the TWFE and the DLM methodologies by aggregating the effects identified within the TWFE framework:

$$\beta_k = \sum_{L=0}^{k} \gamma_L \tag{5}$$

where now the parameter β_k represents the cumulative impact of experiencing a cyclone up to the k year. It's noteworthy to highlight that this equivalence extends beyond just the two-way fixed effect model; it also holds true for the first difference model. From a macroeconomic perspective, housing values are a non-stationary time series. Therefore, I should consider using the first difference. In the results section, I leverage both equivalences to present distinct estimates for both the fixed effect and first difference models. The Dynamic Two-Way First Difference Model is:

$$\Delta y_{i,t} = \sum_{l=-Q}^{L} [\gamma_l^{TWFD} \times \Delta W_{i,t-l}] + \delta_t + \Delta e_{i,t} \tag{6}$$

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.

Despite I have adjusted the weights in equation 4 using Inverse Propensity Score Weighting to create a better match between control and treatment units, results from the recent DiD literature imply that a TWFE regression (equation 4) can suffer from bias. Since previously treated units might still be experiencing a delayed dynamic response to treatment, these treatment effect dynamics are effectively subtracted from the static TWFE treatment effect estimate (Goodman-Bacon, 2021). The problem is that some of the comparisons are 'unclean' in which previously treated units are used as controls for newly-treated units. These 'unclean comparisons' are the source of the 'negative weights' bias of TWFE. With heterogeneous treatment effects across cohorts, also including leads and lags (equation 4) is biased (Sun and Abraham, 2020). Sun and Abraham (2020) show that the relative-period coefficients (ie, the coefficients on leads and lags of treatment) can be contaminated by effects from other periods. In Equation 6, the "contamination effects" can be expressed, for example using $L = 0$, with $\Delta W_{it} = 0$ but $\Delta W_{i,t-j} \neq 0$ for some $j \geq 1$.

To address the potential bias associated with the TWFE model and establish a more robust control, I've adopted a new estimator derived from the Local Projection (LP) Model, designed for the Difference-in-Difference event studies (Jordà et. al, 2024). Before delving deeper into the application of the LP method within the DiD framework, I want to highlight the necessity of making strong assumptions with the DiD strategy, particularly concerning two fundamental assumptions: no anticipation and parallel

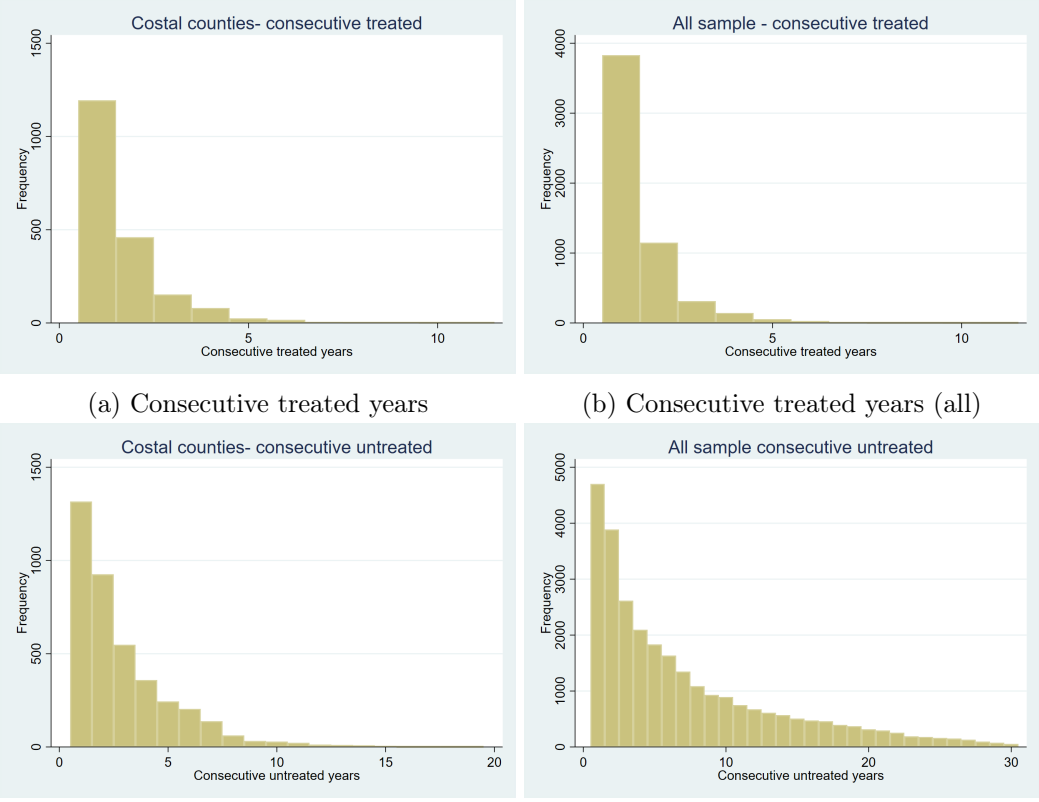
trend.

Assumption 3. ***No Anticipation:*** $E[y_{it}(p) - y_{it}(0)] = 0$, *for all* p *and* t *such that* $t < p$

Assumption 4. ***Parallel trends:*** $E[y_{it}(0) - y_{i1}(0)|p_i = p] = E[y_{it}(0) - y_{i1}(0)]$, *for all* $t \in \{2, ..., T\}$ *and for all* $p \in 1, .., T, \infty$

By restricting the analysis to Atlantic coastal counties and employing an inverse probability weighting scheme, I treated the cyclone trajectory as "randomly assigned" (assumption 1). If treatment timing is as good as randomly assigned, then parallel trends will hold and assumption 2 becomes redundant. Nevertheless, conducting placebo tests to examine housing market dynamics in the years preceding cyclones in the event study further helps validate these assumptions.

Secondly, in the "canonical" DiD approach the treatment status D_{it} represent an absorbing state and I cannot use some of the new DiD estimator in the literature (e.g. Callaway & Sant'Anna, 2021): once a unit is treated they remain treated for the remainder of the panel. Indeed in my case counties switch in and out of treatment and only a few coastal counties experience cyclones in consecutive years, especially for duration exceeding four years consecutively (as depicted in Figure 6a).



(a) Consecutive treated years

(b) Consecutive treated years (all)

(c) Consecutive untreated years

(d) Consecutive untreated years (all)

Figure 6: Consecutive treated and untreated years

on the other hand, there are no "never-treated" units among the coastal counties. This complex scenario highlights another often-overlooked issue in the literature: the trade-off between forming a "clean control" group (counties that have not experienced cyclones for many consecutive years, as

shown in Figure 6c) and creating a control group with a similar probability of experiencing cyclones compared to the treatment group. A more flexible TWFE estimator enables the retention of "transient treatment" effects, whereas the Difference-in-Differences (DiD) approach imposes strong assumptions on the duration of effects—an aspect I aim to estimate rather than assume. Nonetheless, DiD can be beneficial in revealing potential biases in TWFE, as evidenced by the comparative analysis of different estimators in the results section.

After outlining the limitations of the DiD approach, I can now introduce the application of Local Projections within the DiD strategy as proposed by Jordà et al. (2024):

$$\Delta^h y_{i,t} = \delta_t^h + \beta_h^{LP-DID} \Delta D_{i,t} + e_{i,t}^h \tag{7}$$

where $\Delta^h y_{i,t}$ is the difference $\ln(HPI_{i,t+h}) - \ln(HPI_{i,t-1})$; $\Delta D_{i,t}$ is the difference between current treatment $D_{i,t}$ and the year before $D_{i,t-1}$; β_h^{LP-DID} is the estimator. The LP model involves running separate regressions for each time horizon h to examine whether there persists an effect of the cyclone on the housing market h years later. Jordà et al. (2024) particularly fit in my case because they extend the setting to "non-absorbing" treatment. However, the authors make an additional assumptions to deal with such cases:

Assumption 5. ***Dynamic effects stabilize after*** K^* ***periods:*** $\tau(K^*) = \tau(K^* + k)$ *for* $k \geq 0$

Assumption 4 implies that, for any h and for any $j \geq K$, we have $\tau(h + j) = \tau(j - 1)$. To create a clean control, Jordà et al. (2024) restricts the sample to observations that are either:

$$\begin{cases} \text{new treatment} \quad (D_{i,t+j} = 1 \quad \text{for} \quad 0 \leq j \leq h) \& (D_{i,t-j} = 0 \quad \text{for} \quad 1 \leq j \leq K) \\ \text{clean controls} \quad D_{i,t-j} = 0 \quad \text{for} - h \leq j \leq K \end{cases} \tag{8}$$

To isolate the effects of each cyclone and mitigate potential "contamination effects" from previous cyclones, we ensure that those in the treatment group remain treated for h period and untreated for K period before the cyclones. On the other hand, counties in the control group must not experience a cyclone from $t - k$ to $t + h$. Under those conditions, I obtain a convex weighted avarage treatment effect of treated (ATT) for the effect of entering treatment and staying treated.

To add: de Chaisemartin and D'Haultfœuille (2023a): "Those issues are also absent in the careful extension of local projection regressions to panel data with binary and staggered treatments recently proposed by Dube et al. (2023), in work posterior to ours. (I HAVE TRANSFORMED MY TREATMENT IN BINARY STAGGERED TO APPLY Dube et al. (2023))

When the other variables in the regression are treatments rather than covariates (see below for the difference between a treatment and a covariate), one can show that the parallel trends condition in

Theorem S4 implicitly assumes that the effect of the other treatments is constant, which is why the contamination term disappears.

Another important assumption to identify the causal impact of cyclones on housing market is the Stable Unit Treatment Value Assumption (SUTVA) defined by Angrist, et. al., 1996:

Assumption 6. ***SUTVA: a.*** *If* $Z_i = Z_i'$*, then* $D_i(\boldsymbol{Z}) = D(\boldsymbol{Z}')$
b. *If* $Z_i = Z_i'$ *and* $D_i = D_i'$*, then* $Y_i(\boldsymbol{Z}, \boldsymbol{D}) = Y_i(\boldsymbol{Z}', \boldsymbol{D}')$

SUTVA implies that potential outcomes for each county i are unrelated to the treatment status of the other counties. But what if potential outcomes are affected by treatment status of other counties? In the next subsection I have analyzed spillover effects by using Spatial Lag Model, taking into account both one's own treatment status and the treatment status of neighboring counties.

3.2 Accounting for spillover effects

The spatial lag model stands out for its simplicity and can be easily implemented using a least squares estimation (Ruttenauer, 2024). This modeling approach has found precedent in prior studies, such as those conducted by Hoffmann et al.(2022) and Ogunbonde et al. (2019), where it was effectively employed to analyze the repercussions of extreme events in close proximity to affected areas. In the context of my study, this model enables the assessment of the influence exerted by proximity to cyclones, even in cases where direct impact is absent. To achieve this, I enhance the equation 7 by integrating the effects of cyclones on counties situated within a specified distance from the cyclone's path:

$$\Delta_h y_{i,t} = \delta_t^h + \sum_{L=-h}^{H} \beta_L^{LP} \Delta D_{i,t-L} + \sum_{L=-h}^{H} \theta_L^{LP} \Delta N_{i,t-L} + e_{i,t}^h \tag{9}$$

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 100km from cyclone (we also consider different thresholds in the appendix as a robustness check).

4 Results

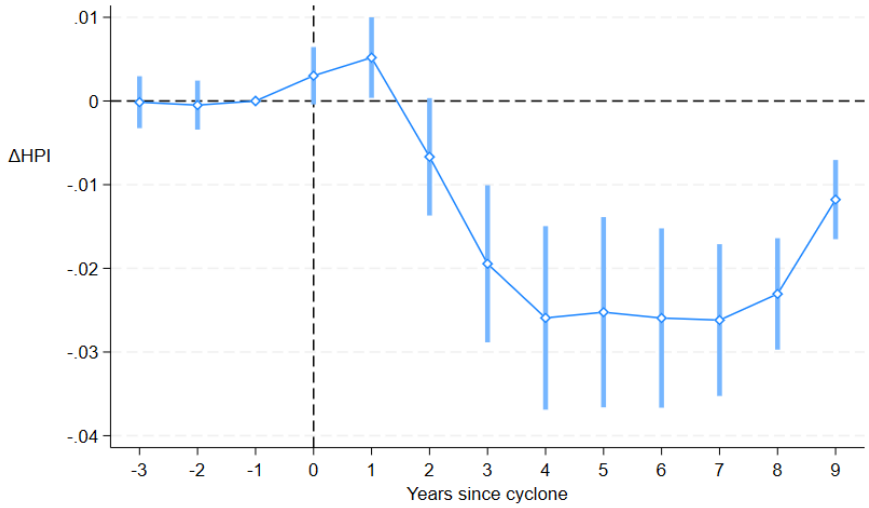
In this study, I provide empirical evidence highlighting the adverse effects of cyclones on the housing market, both in directly affected counties and neighboring regions that narrowly escaped the cyclone's path.

4.1 Main result: changes in housing prices following cyclones

Figure 7 presents my main result: the effects of tropical cyclones on the housing value relative to a county's pre-disaster baseline trend. The plot depicts $\gamma_l \in [-4, 17]$ after Equation 6 is estimated for

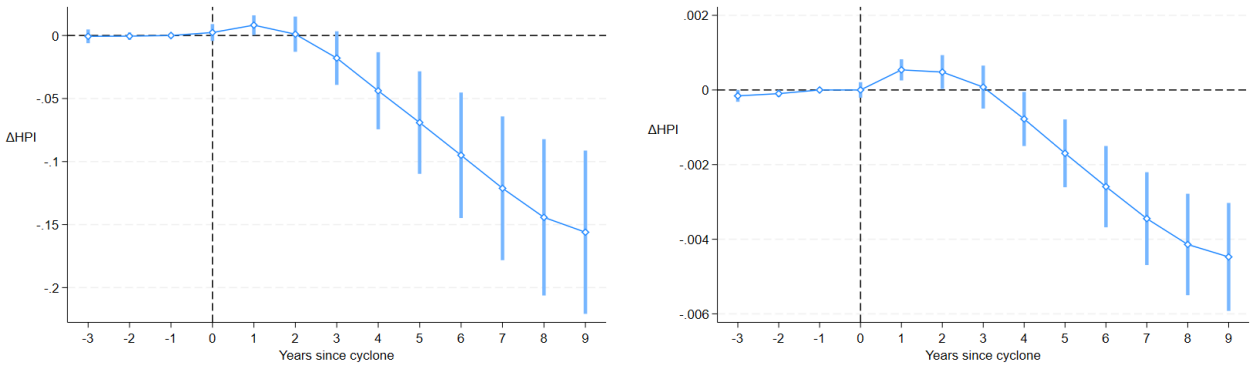
the coastal counties.It's important to note that the leads serve as placebo tests, demonstrating that there are no significant differences between the counties in the control and treatment groups before the cyclone. Indeed, the random trajectory of cyclones strengthens the assumption of exogeneity, allowing me to isolate and identify the effects starting from the year of the event.

Figure 7: Impact of cyclones on the US counties using Two-Way First Difference 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 5). Figure 8a illustrates the cumulative impact of cyclones over time, taking into account both the binary cyclone variable and the continuous measure of wind speed.

Figure 8: 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 8a, 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 2, where the mean value of houses in coastal counties is $254,079, the counties ex-

periencing 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. In the figure, an increase of 1 meter per second in wind speed corresponds to a housing inflation reduction of 0.008. Referring to Table 1 in the summary statistics, 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 9 years. The maximum wind speed observed (64 m/s during Hurricane Katrina) implies a substantial 25% reduction in housing price growth.

4.2 Robustness of the main result

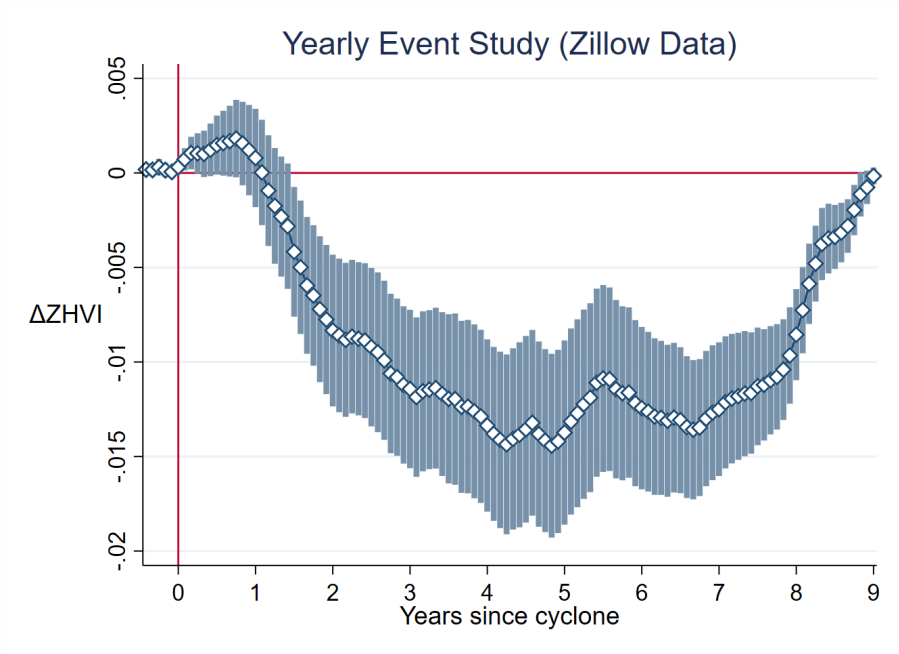
I check the robustness of this result by using alternative data sets, alternative specifications, randomization tests, subsampling of data, and spatial lag models.

Data selection

I replicate the main finding using the ZHVI, our alternative measure of housing price, and the average wind speed instead of maximum, as alternative measure of cyclone exposure.

Figure 9 reveals a prolonged impact over nine years. However, it's crucial to note that incorporating numerous lags can lead to overfitting, diminishing the predictive accuracy of the model. To address this concern, I opted to utilize only the yearly dataset from the Federal Housing Finance Agency, allowing for a focused examination of the long-term effects on the housing market.

Figure 9: Yearly Event study using Zillow Value Housing Index



Both datasets highlight a surge in the housing market within the areas directly impacted by cyclones during the year of the storm, reaching its peak in the subsequent year. However, the dynamics shift in the aftermath, with demand-side factors beginning to outweigh supply-side influences. This transition

leads to a negative impact on the housing market starting from the second year post-cyclone.In Figure 7, the repercussions of a new cyclone become evident, resulting in a roughly 15% decrease in housing values within affected counties after four years.

Sample selection

While Strobl (2011) focused exclusively on coastal counties, other studies, such as Derugyna (2017), included all counties affected by hurricanes in their analysis. To assess the varying impact of cyclones by including inland counties, I have analyzed samples at increasing distances from the coast.

Figure 10: Cyclone Impact Based on Proximity to the Atlantic Coast

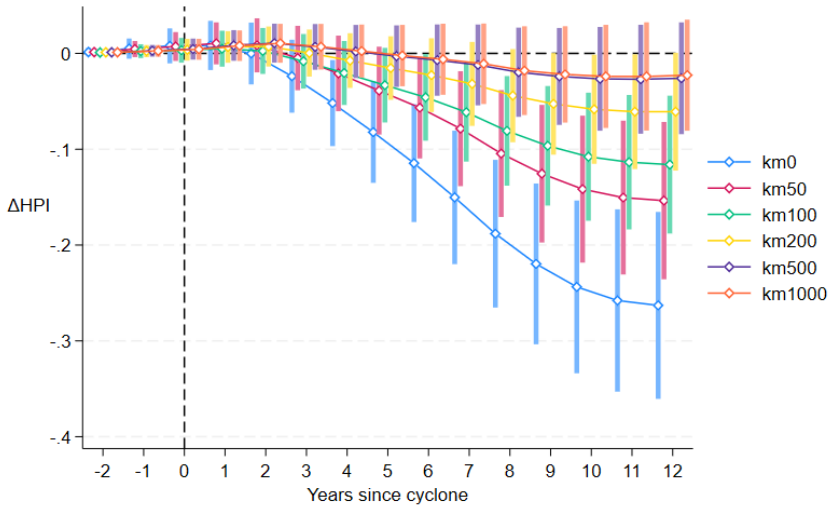
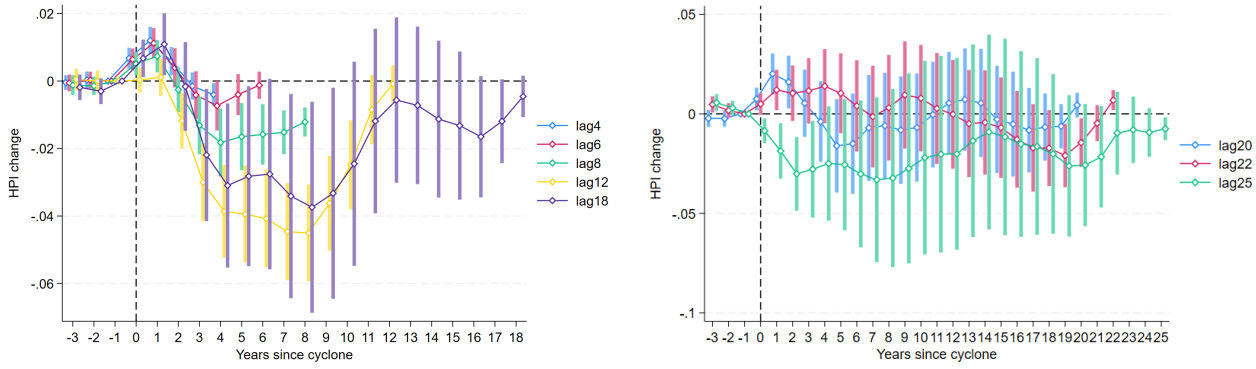


Figure 13 shows that as we include counties farther away from the Atlantic coast, the effects of cyclones disappear. Indeed, the strongest impact of the cyclones are received by the coastal counties becuase , together with the storm wind, they are affected also by storm surge. Because major devastating effect are focus in the coast, I focus the main analysis in this part.

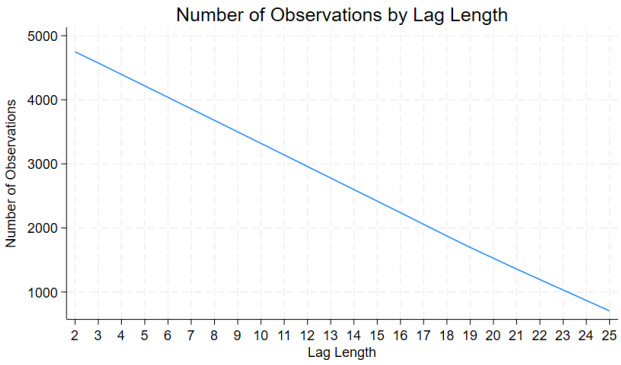
Lag Length To determine the optimal number of lags to include in Equation 6, I plot several models with varying lag lengths. For example, the model with "lag12" includes all lags from 0 to 12 in Equation 6.

Figure 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

This subset of models provides better visualization by excluding those that do not exhibit significant changes from the preceding ones. As we increase the number of lags, we generally observe larger standard errors due to a reduction in the number of observations available for each model. Additionally, in panel 11a, the placebo leads before cyclones become statistically distinguishable from zero, leading to a lack of comparable estimates between the treated and control units.

In panel 11a, most models initially show an increase in housing value during the cyclone year, followed by a decline in subsequent years. The model with 12 lags appears as an outlier, showing no increase in housing price in the two year 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 10 to 20. The cross-validation approach involves predicting housing prices for each county by training on the time series data from all other counties. For consistency, I use the number of observations corresponding to the maximum lag length as a baseline. I then compute the MSE for each lag length by comparing the predicted values to the actual values.

Cross-Validation Methodology 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.

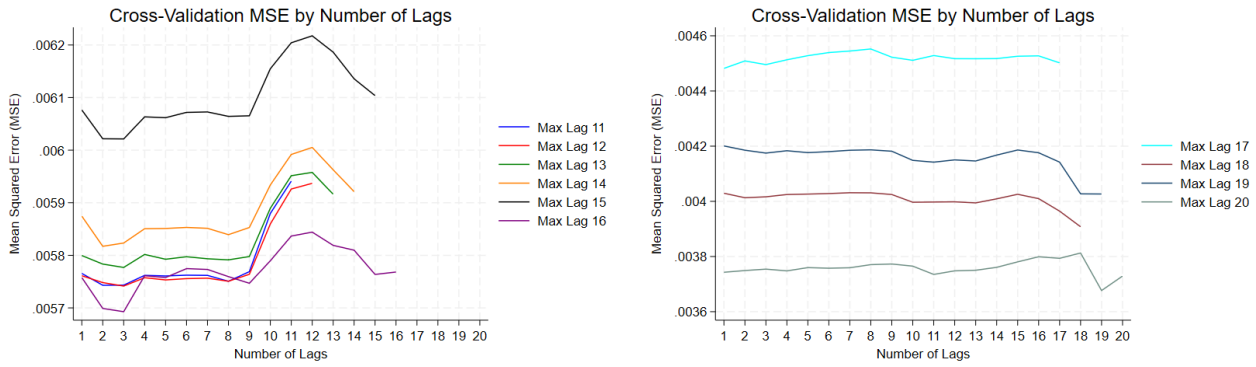
2. **Leave-One-Out Cross-Validation**: For each geographic unit i (e.g., a county), I treat it as the validation set while using all other units as the training set. The model is fit on the training set, then predictions are made for the validation set, computing the squared prediction error for that 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 12: Cross validation (MSE) by number of lags



(a) MSE from lag11 model to lag16

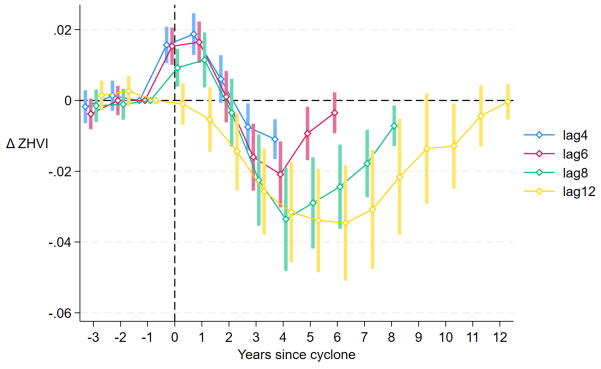
(b) MSE from lag17 to lag20

For clarity, I report results starting from lag 11, as models with shorter lag lengths exhibit relatively stable MSE. In panel 12a, MSE values rise after lag 9, indicating a decline in predictive precision across models with additional lags up to lag 16. Consequently, I retain lag 9 in the main results. Beyond lag 12, MSE values increase, reflecting greater imprecision due to the diminishing sample size and increasing standard errors.

While MSE decreases slightly beyond lag 15, the number of observations continues to drop, introducing further uncertainty. These findings underscore the importance of ample data availability, as smaller datasets may underestimate long-run impacts due to a lack of predictive power at higher lag lengths.

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

Figure 13: Lag Selection with Zillow Data



Randomization tests

To check for potential model misspecification, which could lead to spurious or biased findings, I randomize the sample to create false data and then re-estimate the model in Equation 6. I conduct this randomization in three distinct ways:

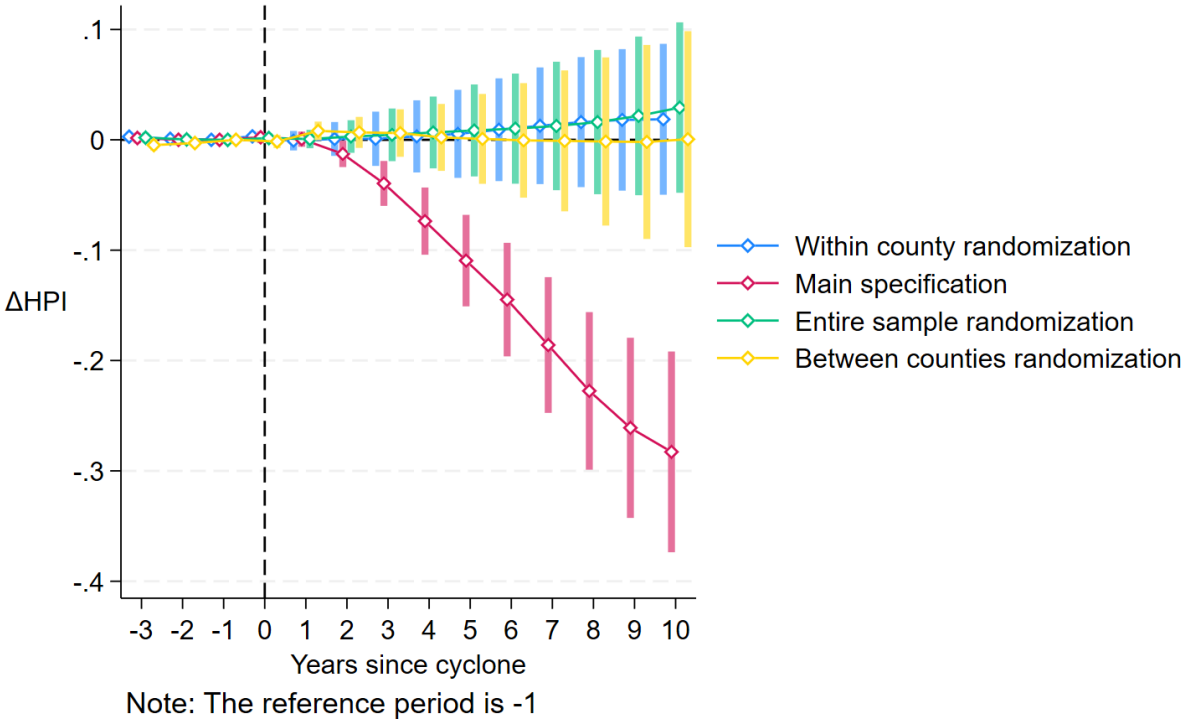
1. Entire Sample: Each cyclone observation is randomly re-assigned.

2. Between Countries: Each country's complete history of cyclone exposure is randomly re-assigned 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 Country: 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 14 demonstrates that the three types of randomization yield insignificant results, thereby reinforcing the robustness of my findings.

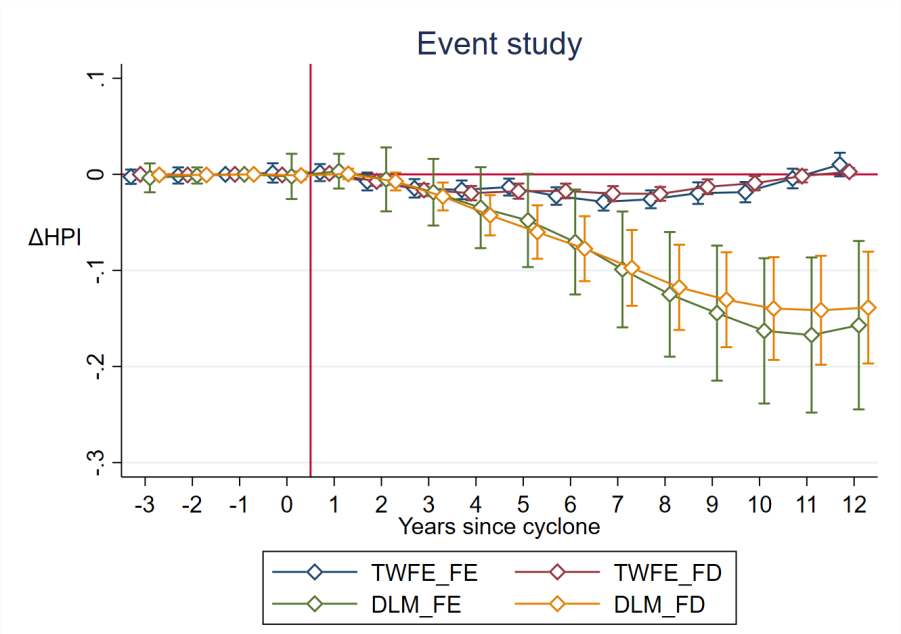
Figure 14: Randomization of cyclones



Model specification

To evaluate the changing impact of cyclones on the housing market in directly affected counties, I have I began by contrasting the outcomes of the Two-Way Fixed Effects (TWFE) model (as depicted in equation 4) and distributed lag model (DLM) (illustrated in equation 5), employing both Fixed Effect (FE) and first difference (FD) models.

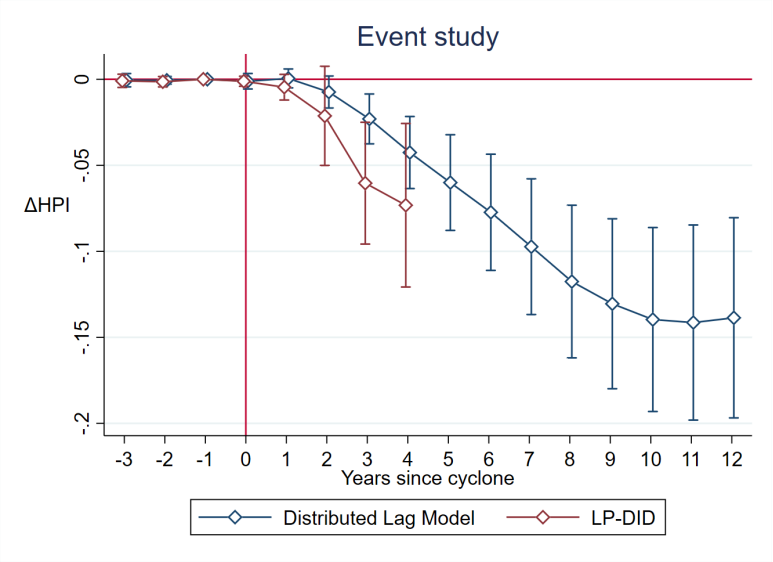
Figure 15: Comparison between fixed effect model and first difference model



The event study graph 15 reveals no significant differences between employing a Fixed Effects (FE) or First Difference (FD) model, whether in the context of TWFE or DLM. However, given the non-stationary nature of the housing price variable, it is preferable from a macroeconomic perspective to utilize the FD model. Consequently, I have adopted the Two-Way First Difference Model as the baseline model.

To mitigate the potential bias of TWFE resulting from the inclusion of previously treated counties as controls, I have conducted a comparison between the results of the Distributed Lag Model (DLM) and the new Local Projection model for Difference-in-Differences (LP-DiD), as outlined in equation 7. However, it's important to note that without enforcing the treatment variable $D_{i,t} = 1$ for a certain number of t following the shock, the LP-DiD cannot be exclusively estimated within the coastal counties sample. Nevertheless, by encompassing all counties that have experienced at least one cyclone, we can already compare the estimates of the DLM and LP-DiD for up to 4 years (since only a few counties have encountered cyclones for more than 4 consecutive years, making it impossible to estimate effects beyond this timeframe).

Figure 16: Comparison between DLM and LP-DiD keeping $D_{i,t} = 1$ only in the year of treatment

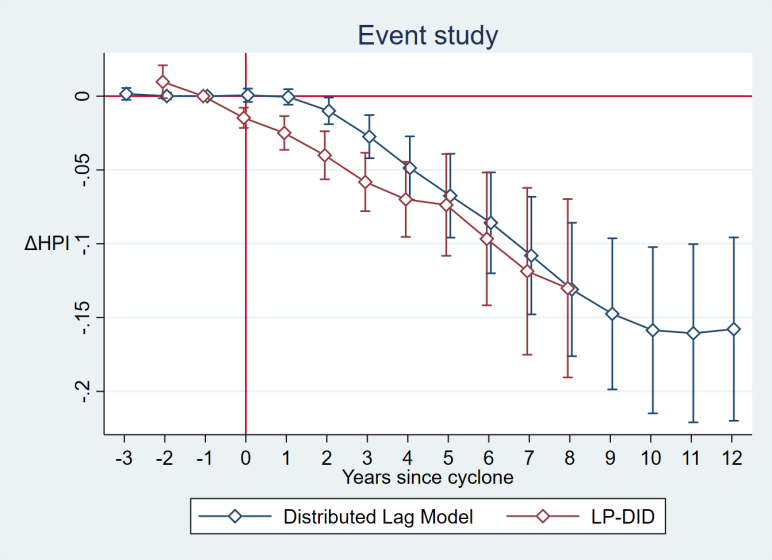


Note: The sample for the Local Projection Model here is different from the sample of the DLM

The notion of an upward bias in the distributed lag model aligns with the assumption that previously treated counties are being utilized as controls. However, this bias appears to be negligible in Figure 16. To delve deeper into the potential bias of TWFE, I have introduced an analysis in the *Appendix*. By setting $D_{i,t} = 1$ for the year following the cyclone, then for the subsequent years up to 9 years post-cyclone, we observe a noteworthy shift in the treatment structure. Interestingly, when applying this adjustment to the entire sample, the effect appears to reverse, indicating an increase in housing value for impacted counties, as also observed with TWFE. However, extending the imposition

of $D_{i,t} = 1$ for multiple consecutive years allows us to narrow the analysis back to coastal counties.

Figure 17: Comparison of DLM and LP-DiD imposing $D_{i,t} = 1$ after treatment

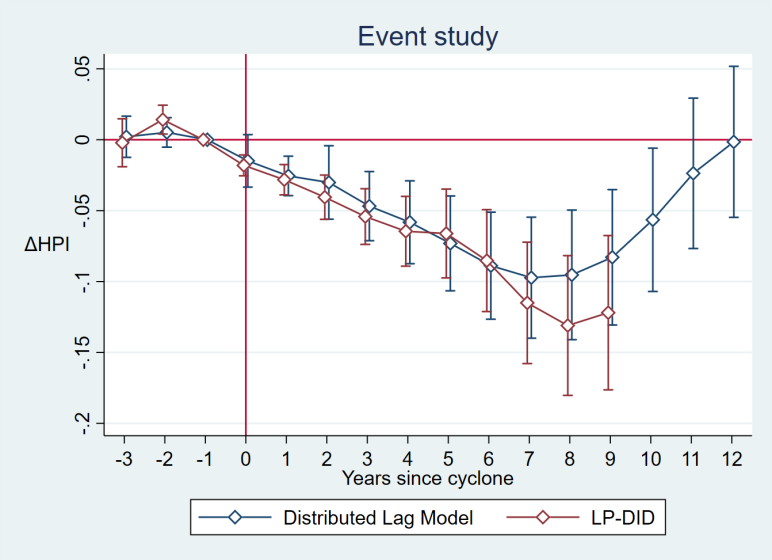


Note:

In this scenario, we can illustrate the trade-off discussed in the previous section regarding the selection of control units. Figure 17 showcases placebo tests (leads in the event study) beginning to deviate from 0. This deviation occurs because we are obligating certain potential "control years" to be treated. However, we observe a significant upward bias in TWFE close to the shock year. Furthermore, the effect of cyclones using LP-DID begins to recover after 8 years, contrasting with the 12-year recovery period indicated by the DLM.

Finally, to validate the bias of TWFE, I conducted regressions under the extreme assumption of DiD "absorbing state", where $D_{i,t}$ is set to 1 from the very first occurrence of a cyclone in a county.

Figure 18: Comparison of DLM and LP-DiD imposing $D_{i,t} = 1$ for all years after the cyclone

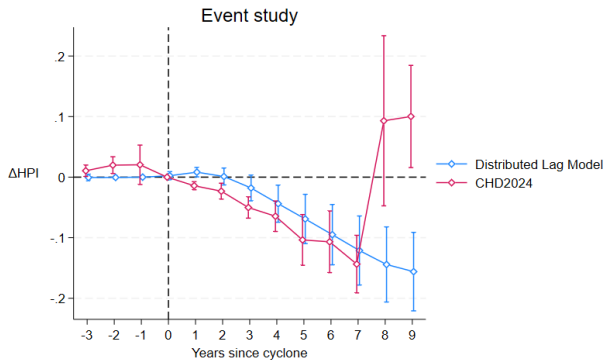


Note:

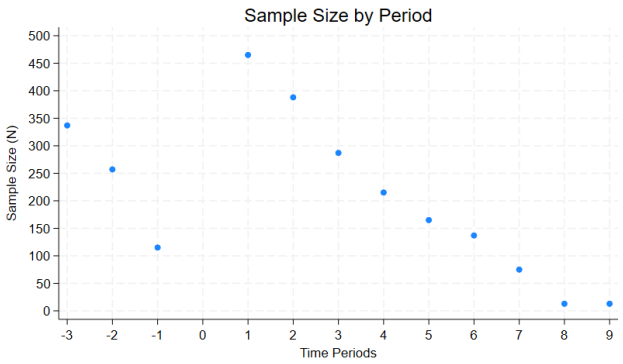
This result show negative estimates following cyclones using the Distributed Lag Model (DLM), which highlights the extent of the bias in the traditional TWFE estimator. These results imply that the previously observed positive effects of cyclones under TWFE were likely driven by specification bias rather than any actual "supply" effects in the housing market.

To further refine these findings, I applied the CHD2024 estimator. Figure 19a presents a comparison between DLM and CHD2024. Meanwhile, Figure 19b shows the number of observations available for each lag estimate, which highlights an important distinction: the CHD2024 estimator relies on significantly fewer observations beyond lag 7, and overall it uses fewer observations than DLM due to its stricter assumptions and cleaner control requirements.

Figure 19: Cross validation (MSE) by number of lags



(a) Comparison of DLM and CHD2024



(b) Observations by each lag estimate

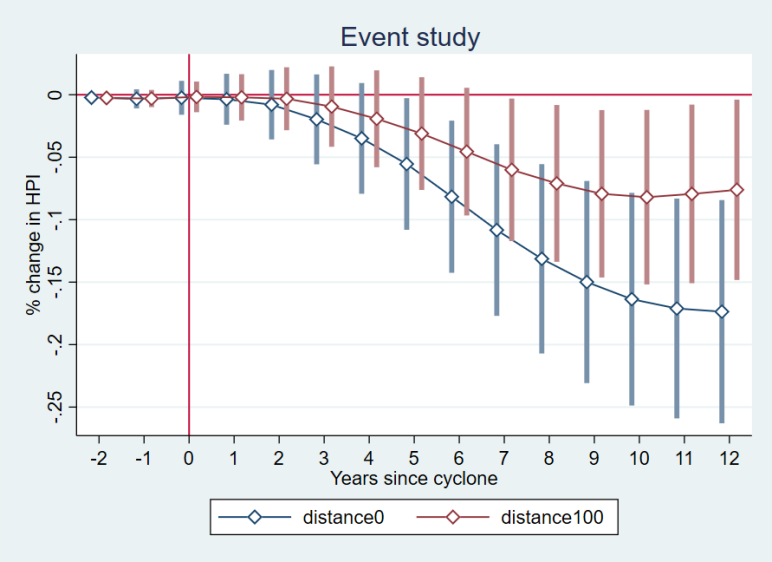
Similar to the LP-DID comparison, the CHD2024 estimator also confirms negative impacts of cyclones on the housing market. However, due to its reliance on fewer treatment units per period, it

imposes stricter assumptions and limits the number of observations per lag. Despite these limitations, the consistent finding of negative effects across models further supports the conclusion that cyclones negatively impact housing market prices, discounting any potential positive effects observed in traditional TWFE models as a result of bias.

4.3 Changes in the neighboring "missed counties"

Here, I offer empirical proof that even in counties narrowly spared from a direct hit, cyclones still exert lasting negative effects.

Figure 20: Distributed Lag Model Result



Note: Red line represent the impact of being closed within

Even counties situated within a 100-kilometer radius of cyclone-affected areas exhibit a decline in housing values, despite not being directly impacted by the cyclones themselves.

4.4 Heterogeneous analyses for different intensities and frequency of being hit

Here, I aim to delve into the nuanced impacts of various cyclone intensities and frequencies on counties. Specifically, I intend to analyze how unexpected events, characterized by a low probability of cyclone occurrence, compare to anticipated events in counties that frequently experience cyclones. Furthermore, I seek to discern the contrasting effects of strong cyclones, such as hurricanes with wind speeds of at least 33 meters per second, versus weaker cyclones. The upcoming graph will illustrate the influence of hurricanes compared to cyclones with wind speeds below 33 meters per second on housing dynamics.

Figure 21: Impact of strong and weak cyclones



(a) Impact of hurricanes on housing

(b) Impact of weaker cyclones on housing

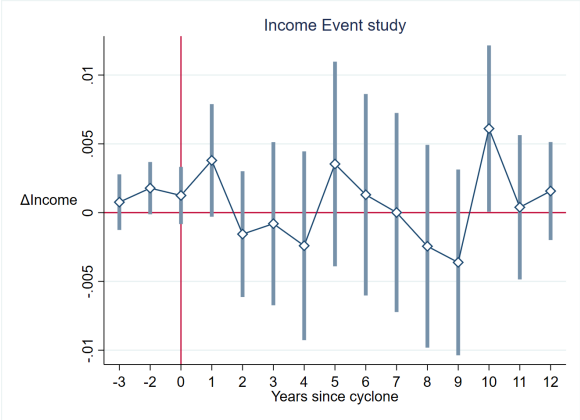
Here, it becomes apparent that both the immediate positive impact and the enduring negative repercussions are heightened. However, what's particularly notable is the swifter pace of recovery compared to the overarching results that encompass all cyclones in the main specification. Strikingly, even when excluding hurricanes and focusing solely on cyclones with wind speeds below 33 m/s, a detrimental effect on housing persists.

Probability of receiving a cyclone (frequency)

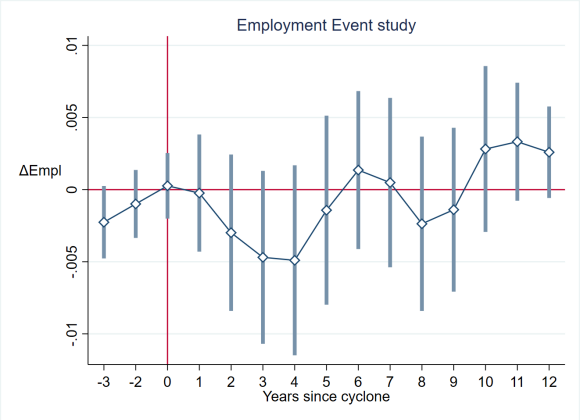
4.5 Changes on other outcomes

To understand the drivers of the decrease in housing market following cyclones, I estimated the impact of cyclones on other economic variables such as income, employment and population. With the same identification strategy of equation 7, I have obtained the following results.

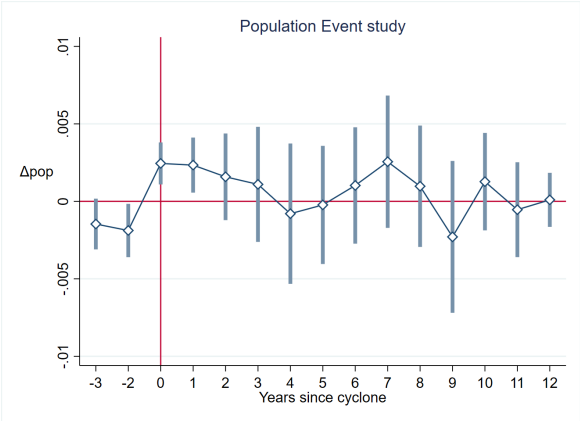
Figure 22: Impact of cyclones on other variables



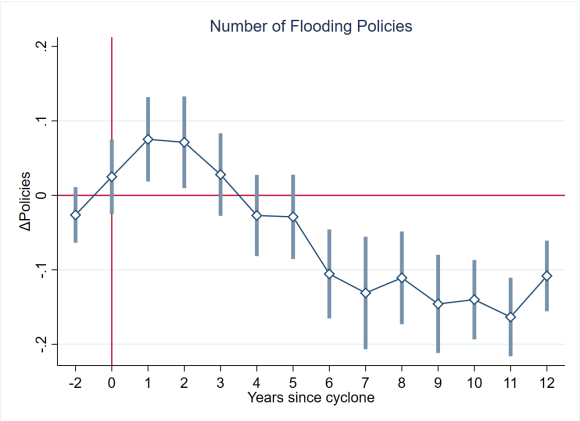
(a) Changes in income



(b) Changes in employment



(c) Changes in population



(d) Changes in the number of flooding insurance policies

The findings indicate that the long-term decline in housing values may be attributed more to an elevated perception of living in a high-risk area rather than direct economic damages. As evident in the graphs, cyclones show no substantial impact on employment, income, or population. Conversely, the observed rise in the adoption of flood insurance policies implies an additional financial burden for affected counties, contributing to the housing market downturn (result similar to Gallagher, 2014). It is noteworthy that a parallel analysis, focusing solely on non-mandatory insurances (see Appendix), yielded comparable results, suggesting that the surge in flood insurance uptake is rooted in an augmented perceived risk rather than a response to new mandatory homeowner requirements.

5 Conclusion

This study provides critical insights into the long-term impacts of tropical cyclones on the U.S. housing market, highlighting several key findings. In the short run, there are initial positive effects attributed to supply-side dynamics; however, employing advanced difference-in-difference estimators reveals that these positive effects dissipate over time. In contrast, the long-run effects are characterized by a significant decline in housing market values, specifically a 15% decrease over a nine-year period in

counties directly affected by cyclones. Notably, these negative impacts also extend to neighboring counties that did not directly experience cyclone landfalls, suggesting a broader economic influence.

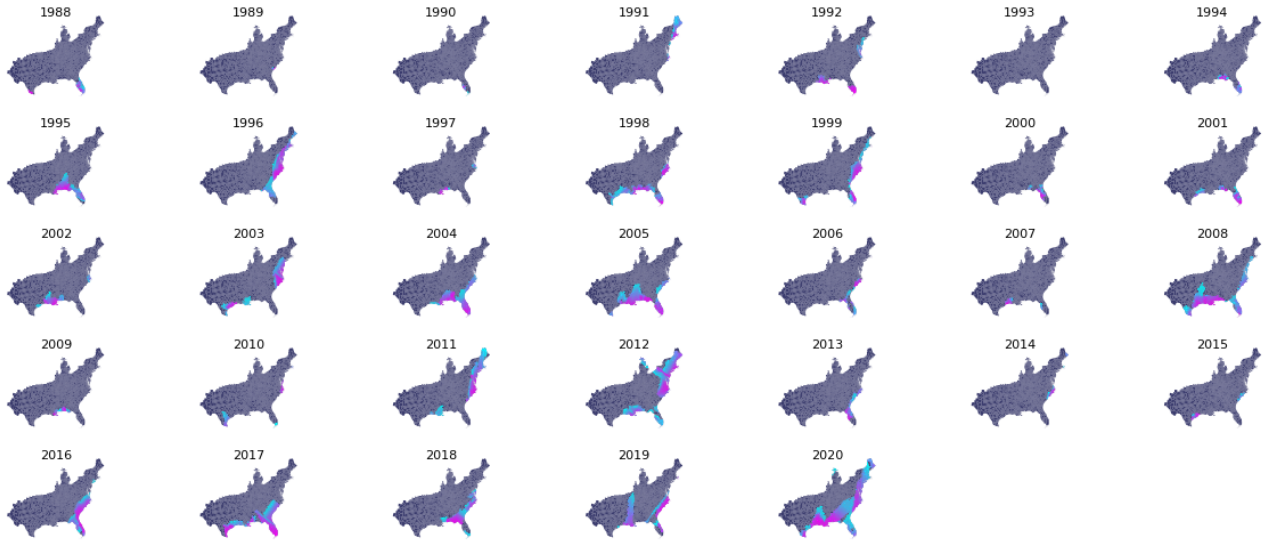
The evidence indicates that the persistence of negative effects is primarily driven by heightened perceived risks associated with living in cyclone-prone areas. Notably, following cyclones, there is a marked increase in both the number and cost of flood insurance policies, reflecting the heightened concern among homeowners. Despite these significant shifts in perception and behavior, the study finds no substantial impacts on local economic indicators such as employment, income, or population growth, indicating that the broader economic ramifications of cyclones may be limited.

Overall, these findings underscore the need for policymakers to consider the lasting economic burdens posed by tropical cyclones on housing markets. Understanding the dynamics of demand-side effects and their relation to perceived risk can guide more effective strategies for disaster preparedness, housing policy, and community resilience in the face of climate change. As the frequency and intensity of tropical cyclones continue to rise, addressing these challenges will be crucial for safeguarding coastal communities and ensuring sustainable housing markets.

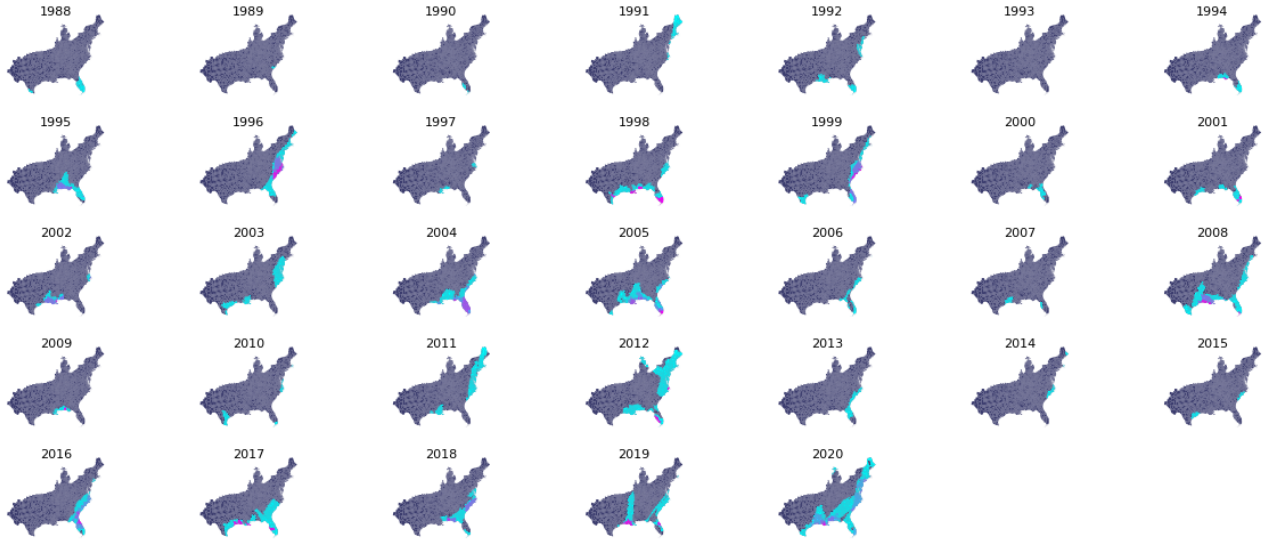
A Data variability

Variability of wind speed intensity and frequency by year:

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



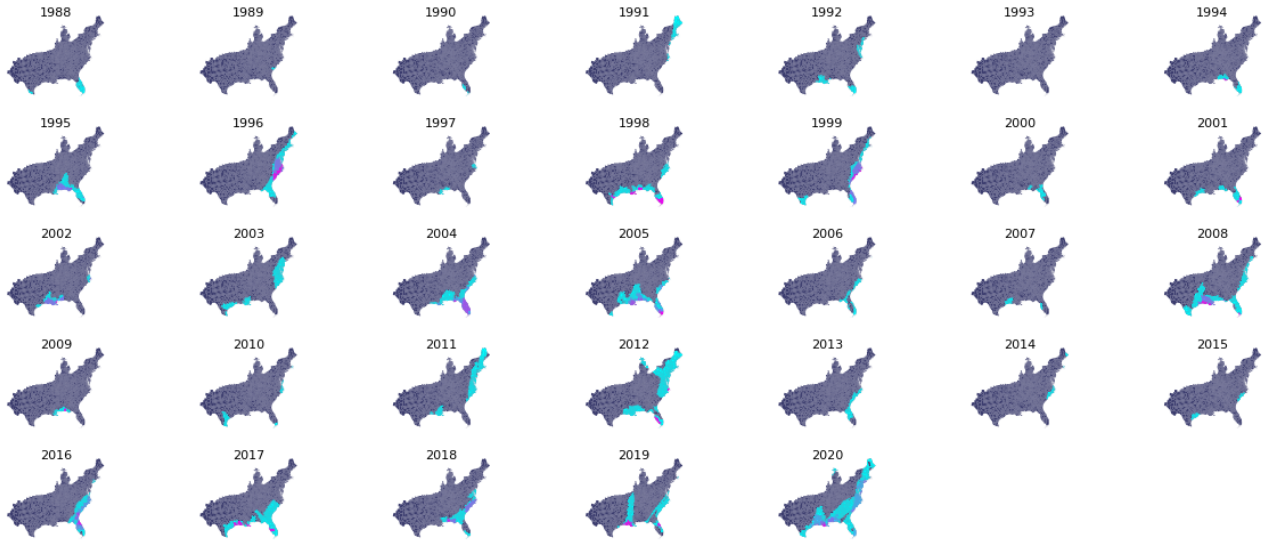
(a) Maximum wind speed of cyclones by year and location



(b) Frequency of cyclones by year and location

Variability of wind speed intensity and frequency by year:

Figure A.2: Variability of intensity and frequency by year



B Robustness checks for identification

References

Laura A. Bakkensen, Xiangying Shi, and Brianna D. Zurita. The impact of disaster data on estimating damage determinants and climate costs. *Economics of Disasters and Climate Change*, 2(Volume 8, 2018):49–71, 2018. doi: https://doi.org/10.1007/s41885-017-0018-x.

Eli Beracha and Robert S. Prati. "how major hurricanes impact housing prices and transaction volume". *Real Estate Issues*, 33(1), 2008. URL `https://cre.org/real-estate-issues/major-hurricanes-impact-housing-prices-transaction-volume/`.

Okmyung Biny and Stephen Polasky. Effects of Flood Hazards on Property Values: Evidence Before and After Hurricane Floyd. *Land Economics*, 80(4), 2004.

M. Collins, M. Sutherland, L. Bouwer, S.-M. Cheong, et al. Chapter 6: Extremes, abrupt changes and managing risks. *IPCC Special Report on the Ocean and Cryosphere in a Changing Climate*, page 603, 2019.

Clément de Chaisemartin and Xavier D'Haultfœuille. Difference-in-Differences Estimators of Intertemporal Treatment Effects. *The Review of Economics and Statistics*, pages 1–45, 2024.

Tatyana Deryugina. The fiscal cost of hurricanes: Disaster aid versus social insurance. *American Economic Journal: Economic Policy*, 9(3):168–98, August 2017. doi: 10.1257/pol.20140296. URL `https://www.aeaweb.org/articles?id=10.1257/pol.20140296`.

Tatyana Deryugina, Laura Kawano, and Steven Levitt. The economic impact of hurricane katrina on its victims: Evidence from individual tax returns. *American Economic Journal: Applied Economics*, 10(2):202–33, April 2018. doi: 10.1257/app.20160307.

Olivier Deschênes and Michael Greenstone. The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather. *American Economic Review*, 97(1):354–385, March 2007. doi: 10.1257/aer.97.1.354. URL `https://www.aeaweb.org/articles?id=10.1257/aer.97.1.354`.

Joshua Graff Zivin, Yanjun Liao, and Yann Panassié. How hurricanes sweep up housing markets: Evidence from florida. *Journal of Environmental Economics and Management*, 118:102770, 2023. ISSN 0095-0696. doi: https://doi.org/10.1016/j.jeem.2022.102770. URL `https://www.sciencedirect.com/science/article/pii/S0095069622001231`.

J. Edward Graham, William W. Hall, and Peter W. Schuhmann. Hurricanes, catastrophic risk, and real estate market recovery. *The Journal of Real Estate Portfolio Management*, 13(3):179–190, 2007. ISSN 10835547. URL `http://www.jstor.org/stable/24883107`.

Greg J. Holland, James I. Belanger, and Angela Fritz. A revised model for radial profiles of hurricane winds. *Monthly Weather Review*, 138(12):4393 – 4401, 2010. doi: 10.1175/2010MWR3317.1. URL https://journals.ametsoc.org/view/journals/mwre/138/12/2010mwr3317.1.xml.

S. M. Hsiang and D. Narita. Adaptation to cyclone risk: Evidence from the global cross-section. *Climate Change Economics*, 3(2):1–28, 2012. ISSN 20100078, 20100086. URL http://www.jstor.org/stable/climchanecon.3.2.08.

Solomon Hsiang. Climate econometrics. *Annual Review of Resource Economics*, 8(Volume 8, 2016): 43–75, 2016. doi: https://doi.org/10.1146/annurev-resource-100815-095343. URL https://www.annualreviews.org/content/journals/10.1146/annurev-resource-100815-095343.

Solomon Hsiang, Robert Kopp, Amir Jina, James Rising, Michael Delgado, Shashank Mohan, D. J. Rasmussen, Robert Muir-Wood, Paul Wilson, Michael Oppenheimer, Kate Larsen, and Trevor Houser. Estimating economic damage from climate change in the united states. *Science*, 356 (6345):1362–1369, 2017. doi: 10.1126/science.aal4369. URL https://www.science.org/doi/abs/10.1126/science.aal4369.

Solomon M. Hsiang and Amir S. Jina. 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, July 2014.

Intergovernmental Panel on Climate Change (IPCC). Climate change 2021: The physical science basis. contribution of working group i to the sixth assessment report of the intergovernmental panel on climate change. pages 8–9, 15–16, 2021.

Òscar Jordà, Arindrajit Dube, Daniele Girardi, and Alan Taylor. A Local Projections Approach to Difference-in-Differences Event Studies. CEPR Discussion Papers 18141, May 2023. URL https://ideas.repec.org/p/cpr/ceprdp/18141.html.

Matthew E. Kahn. The Death Toll from Natural Disasters: The Role of Income, Geography, and Institutions. *The Review of Economics and Statistics*, 87(2):271–284, May 2005. URL https://ideas.repec.org/a/tpr/restat/v87y2005i2p271-284.html.

Derek K. Kellenberg and Ahmed Mushfiq Mobarak. Does rising income increase or decrease damage risk from natural disasters? *Journal of Urban Economics*, 63(3):788–802, 2008. ISSN 0094-1190. doi: https://doi.org/10.1016/j.jue.2007.05.003. URL https://www.sciencedirect.com/science/article/pii/S0094119007000770.

Shawn J. McCoy and Randall P. Walsh. Wildfire risk, salience housing demand. *Journal of Environmental Economics and Management*, 91:203–228, 2018.

Justin Murfin and Matthew Spiegel. Is the Risk of Sea Level Rise Capitalized in Residential Real Estate? *The Review of Financial Studies*, 33(3):1217–1255, 02 2020. URL https://doi.org/10.1093/rfs/hhz134.

Anthony Murphy and Eric Strobl. The impact of hurricanes on housing prices: Evidence from us coastal cities. Technical report, Working Paper, 2010.

William D. Nordhaus. Economic aspects of global warming in a post-copenhagen environment. *Proceedings of the National Academy of Sciences*, 107(26):11721–11726, 2010. doi: 10.1073/pnas.1005985107. URL https://www.pnas.org/doi/abs/10.1073/pnas.1005985107.

Ilan Noy. The macroeconomic consequences of disasters. *Journal of Development Economics*, 88(2): 221–231, 2009. ISSN 0304-3878. doi: https://doi.org/10.1016/j.jdeveco.2008.02.005. URL https://www.sciencedirect.com/science/article/pii/S030438780800031X.

Francesc Ortega and Süleyman Taşpınar. Rising sea levels and sinking property values: Hurricane sandy and new york's housing market. *Journal of Urban Economics*, 106(C):81–100, 2018. URL https://EconPapers.repec.org/RePEc:eee:juecon:v:106:y:2018:i:c:p:81-100.

David Strömberg. Natural disasters, economic development, and humanitarian aid. *Journal of Economic Perspectives*, 21(3):199–222, September 2007. doi: 10.1257/jep.21.3.199. URL https://www.aeaweb.org/articles?id=10.1257/jep.21.3.199.

Jill C. Trepanier, John Nielsen-Gammon, Vincent M. Brown, Derek T. Thompson, and Barry D. Keim. Stalling north atlantic tropical cyclones. *Journal of Applied Meteorology and Climatology*, 2024. doi: 10.1175/JAMC-D-23-0229.1. URL https://journals.ametsoc.org/view/journals/apme/aop/JAMC-D-23-0229.1/JAMC-D-23-0229.1.xml.