Three Essays at the Intersection of Environmental Economics, Climate Finance, and Urban Economics

Davide Castellani

Ph.D. in Economics "Vilfredo Pareto"

XXXVII Cycle

Supervisors:
Giovanni Mastrobuoni
Elena Esposito

University of Turin & Collegio Carlo Alberto





October 2025

Abstract

I study how environmental and climate-related risks shape economic outcomes across households, firms, and local economies.

In the first chapter, I analyze the long-term effects of tropical cyclones on U.S. housing markets and show that housing values decline persistently after cyclone exposure, especially in areas with little prior experience of such events. I identify rising insurance costs as a key mechanism amplifying post-disaster economic burdens.

In the second chapter, I examine how floods affect firm credit ratings globally. I find that flood exposure leads to a deterioration in creditworthiness, with moderate exposure fostering some adaptive capacity and repeated extreme events revealing limits to resilience.

In the third chapter (co-authored), we study how the visibility of wind turbines affects tourism in Italy. The negative impacts are concentrated in rural, tourism-dependent areas, particularly for small private accommodations, though overall environmental benefits from renewable energy outweigh these local costs.

Together, these essays highlight the economic implications of both physical and perceived environmental risks and emphasize the importance of preparedness, adaptation, and equitable policy design in the face of climate change.

Statement of Joint Work: I declare that Chapter 3 is co-authored with Gabriele Pinto and that I contributed to 50% of this work.

Since high school, I have been fascinated by two things: philosophy and nature. Philosophy seeks to understand what matters to human beings and how we interact with the world around us. Nature, in turn, reminds us of the limits and interconnectedness of human action. As I explored these interests, I realized that economics has become a central force shaping philosophy since the Industrial Revolution. Over the past two centuries, societies have pursued production and innovation to improve well-being, often overlooking the negative externalities of rapid growth. Today, we face the consequences: climate change and environmental risks that challenge our societies. This awareness motivated my research, which sits at the intersection of environmental economics, climate finance, and urban development, aiming to understand how humans respond to and manage the economic consequences of environmental hazards.

This work would not have been possible without the support of many people. I am deeply grateful to my advisors, Giovanni Mastrobuoni and Elena Esposito, for their guidance and encouragement throughout this journey. I thank Paolo D'Odorico and Solomon Hsiang for hosting me at the University of California, Berkeley during the third year of my PhD. I am also grateful to Federico Galizia for establishing a new internship program at the International Finance Corporation with Collegio Carlo Alberto PhD students, which made the second chapter possible, and to Aurèle M. Houngbédji for his continuous guidance throughout its development. I extend my thanks to my co-author, Gabriele Pinto, for his invaluable support on the third chapter. Finally, I am profoundly thankful to my PhD colleagues and faculty at Collegio Carlo Alberto, who made my years in Turin both intellectually stimulating and personally enriching.

I dedicate this work to the memory of my grandmother Elisa, who showed me how to approach life with reason sometimes, without letting emotions entirely dictate my choices.

Contents

Introduction **4**

Chapter 1 **5**

Chapter 2 **51**

Chapter 3 **80**

Concluding Remarks **118**

Introduction

This dissertation studies the economic consequences of environmental and climate-related risks across households, firms, and local economies. It is composed of three empirical chapters, each addressing a different dimension of how physical and perceived environmental hazards shape economic behavior.

In the first chapter, *"The Long-Run Effects of Tropical Cyclones on the U.S. Housing Market"*, I examine how tropical cyclones affect housing values in the contiguous United States from 1988 to 2021. I find that while housing values may rise temporarily following a cyclone, they decline by 15% over an eight-year horizon. Areas with historically low cyclone activity experience the most persistent declines, reflecting both household sorting into seemingly safer locations and increased insurance uptake that raises post-disaster living costs.

The second chapter, *"The Impact of Flooding on Credit Risk: Evidence from a Global Financial Institution"*, investigates how floods influence firm creditworthiness globally, with a focus on private firms in emerging markets. I find that exposure to a single flood event leads to a deterioration in credit ratings of roughly 20% over three years. The results reveal a non-linear relationship: firms experiencing moderate exposure to floods show evidence of adaptive capacity, while those affected by more than five floods suffer sharply increased credit deterioration, indicating limits to resilience.

The third chapter, *"The Impact of Wind Turbine Visibility on the Local Economy"*, co-authored with Gabriele Pinto, examines the effects of wind turbine visibility on tourism across Italian municipalities. We show that turbine visibility can reduce tourist arrivals by up to 20%, with rural, tourism-dependent areas and small private accommodations bearing the largest costs. Despite these local economic costs, the broader environmental benefits of renewable energy outweigh the losses, highlighting the need for policies that balance local externalities with societal gains.

Together, these chapters illustrate how environmental and climate risks influence economic outcomes through multiple channels, property markets, firm credit, and local economic activity, and highlight the importance of preparedness, adaptation, and policy design in mitigating their impacts.

Chapter 1

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

Davide Castellani†

October 30, 2025

Abstract

Housing is a major component of wealth in the United States. Despite tropical cyclones causing approximately $1.3 trillion in damages since 1980, there is limited empirical evidence on their long-term effects on the housing market. This paper examines the impact of tropical cyclones on housing values in the contiguous U.S. from 1988 to 2021. Using data from the National Hurricane Center and a radial wind-speed model, I construct a new cyclone exposure measure and provide the first plausibly causal estimates of the aggregate impact of all recorded cyclones at the county-year level. Despite a short-term rise, housing values decline by 15% over an eight-year horizon following cyclone exposure. Within coastal counties, areas with lower historical cyclone activity tend to have more housing development and higher property values, suggesting households sort into seemingly safer locations. However, these less-experienced areas suffer the most persistent declines. One key mechanism of decreasing housing prices appears to be a sharp increase in insurance uptake, contributing to higher post-disaster living costs. These findings highlight that cyclone risk is driven not just by exposure but also by local preparedness.

JEL-Codes: Q54, R31, C21

*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, Òscar Jordà, Guglielmo Zappalà, Juan Vargas, Yagan Hazard, Joel Ferguson, 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. 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

Although 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 lead to elevated sea surface temperatures, which enhance the thermodynamic energy available to tropical cyclones, thereby increasing their intensity (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). The rise in sea level is projected to further intensify storm surges, increasing the potential for coastal damage (Collins et al., 2019).

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] According to FEMA, the federal government spent more than $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. In particular, predictions indicate that the most substantial economic losses due to climate change in coastal areas are expected to arise due to the intensification of cyclones and sea level rise (Hsiang et al., 2017). A significant portion of these damages can be attributed to the destruction of residential properties, which represent a substantial share of household wealth in the U.S. The housing market not only plays a critical role in personal financial stability, but also contributes significantly to the GDP of the country. For example, in 2023, the real estate sector accounted for 12.6% of the U.S. GDP. [2]

According to the U.S. Census, in 2017, approximately 60.2 million people—about 18.5% of the U.S. population—lived in counties along the Atlantic Ocean and Gulf of Mexico, the regions most vulnerable to tropical cyclones. Living in coastal areas offers numerous amenities, but also involves significant risks, particularly in light of climate change. Yet, empirical evidence on the influence of rising sea level risk on the real estate market is mixed (Hino and Burke, 2021; Bakkensen and Barrage, 2021; Murfin and Spiegel, 2020; Bernstein et al., 2019; Giglio et al., 2021b), while the overall effects of cyclones on

[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`

housing markets remain 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. However, empirically estimating the impact of cyclones is challenging for several reasons. First, measuring cyclone exposure is complex; cyclones vary in intensity, size, and duration, and their effects can differ widely depending on the specific characteristics of the affected areas. Second, there are identification issues: counties across the United States have varying probabilities of experiencing cyclones, which may be correlated with other regional characteristics, making it difficult to isolate the causal impact of cyclones on the housing market.

Combining a novel dataset source, HURDAT2, from the National Hurricane Center with a model for the radial profile of hurricane winds, I conduct a comprehensive analysis of the impact of cyclones on the continental United States housing market spanning the period from 1992 to 2021. The newly available information in HURDAT2 includes a raw measure of the maximum spatial extent of cyclones, defined as the radius at which wind speeds of at least 34, 50, and 64 knots occur. To refine these approximate wind speed estimates at various points within the cyclone, I use an interpolation method based on the dynamics of the radial wind profile described by Holland et al. (2010). This refinement process is based on key characteristics, such as pressure gradients, cyclone intensity metrics, and environmental factors, to produce an improved spatial distribution of wind speeds at each point in time. I then aggregate the wind speed into county-by-year units to align it with county-level housing value data.

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 have comparable probabilities of exposure to a cyclone. I restrict the sample to Atlantic coastal counties, those directly adjacent to the east coast, and apply propensity score weighting. Another issue is the presence of "unclean" comparisons, which occur when previously cyclone-affected areas are used as control groups for regions impacted by new cyclones (Sun and Abraham, 2021). To address this, I employ a recently developed local projection estimator for Difference-in-Differences event studies (Jordà et al., 2023).

The main finding reveals that experiencing a new cyclone results in a significant 15% decline in housing market values over an 8-year period, relative to unaffected counties. This result is robust to the

traditional event study approach using the Two-Way Fixed Effect model as well as an alternative Difference-in-Differences estimator designed for intertemporal treatment effects (de Chaisemartin and D'Haultfœuille, 2024). This decline reveals a far greater persistence compared to other natural disasters, like wildfires (McCoy and Walsh, 2018b). This finding builds on previous research on hurricanes, which highlighted positive short-run effects on housing values (Graff Zivin et al., 2023; Murphy and Strobl, 2010). By examining both short- and long-run effects, this study expands our understanding of the dynamic impacts on housing values.

Regarding heterogeneous effects, the impact is more pronounced and lasts longer in counties with a lower probability of cyclone occurrences, suggesting that counties with a higher likelihood of being affected may have implemented resilience strategies or infrastructure improvements to mitigate cyclone impacts. When I consider cyclone intensity rather than frequency, I find that hurricanes (cyclones with winds of at least 74 mph) have a more significant impact compared to weaker cyclones. However, excluding weaker cyclones from the sample overlooks some negative effects on the housing market. This analysis is crucial because it expands on previous studies that have focused solely on hurricanes, offering a broader perspective on the impact of all cyclones of varying intensities.

Next, I employ a spatial econometric model to demonstrate that "nearly missed counties", those in the control group that are adjacent to affected counties, are not indirectly impacted by cyclones. This is crucial for two reasons. First, it serves as a test for the Stable Unit Treatment Value Assumption (SUTVA), ensuring that there are no "spillover effects" where the treatment in one county influences housing values in neighboring control counties. Such spillovers could arise, for example, if residents of cyclone-affected areas relocate to nearby counties, thereby influencing housing prices in the control group. Second, I can determine whether there is evidence of increased perceived risk in nearby but unaffected counties. A close call cyclone event could act as a signal, prompting residents to anticipate future impacts and adopt adaptive strategies that ultimately influence the housing market. The findings indicate that cyclones have no significant impact on nearby municipalities, validating the SUTVA assumption.

To investigate the mechanisms behind the observed decline in the housing market over the eight years following cyclones, I examine the impact of cyclones on the labor market, population, and insurance market using the same empirical strategy applied to the housing market. Income, employment, and population data are sourced from the U.S. Bureau of Economic Analysis (BEA), while flood insurance policy data are obtained from the Federal Emergency Management Agency (FEMA) through the National Flood Insurance Program (NFIP). While I find relatively small effects on income, employment, and population, there is a substantial increase in the number of flood insurance policies signed in the years following cyclones. This finding suggests that higher living costs associated with cyclone-prone

areas, rather than displacement effects, may play a more significant role in driving the decline in housing prices. This warrants further investigation into the economic burden of insurance and other costs (e.g. maintenance costs) in these regions.

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 used 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 am 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. This reduction in measurement error may also help explain why this study finds more pronounced impacts compared to previous work, as it provides a more accurate representation of affected areas and their exposure to cyclone forces. Moreover, compared to Deryugina (2017) and Graff Zivin et al. (2023), this study includes all cyclones that occurred 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 characteristics, although it caused significant damage (Ortega and Taṣpınar, 2018).

The second noteworthy 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 and by Ortega and Taṣpınar (2018) after Hurricane Sandy. On the other hand, Deryugina et al. (2018) 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. Here, I find results that contrast with previous studies on climate risk capitalization in the housing market (Hino and Burke, 2021; Bakkensen and Barrage, 2021; Murfin and Spiegel, 2020), suggesting that cyclones have more pronounced effects than sea level rise or wildfires.

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; Strömberg, 2007; Kellenberg and Mobarak, 2008; Noy, 2009; Hsiang and Narita, 2012), 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 1988 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

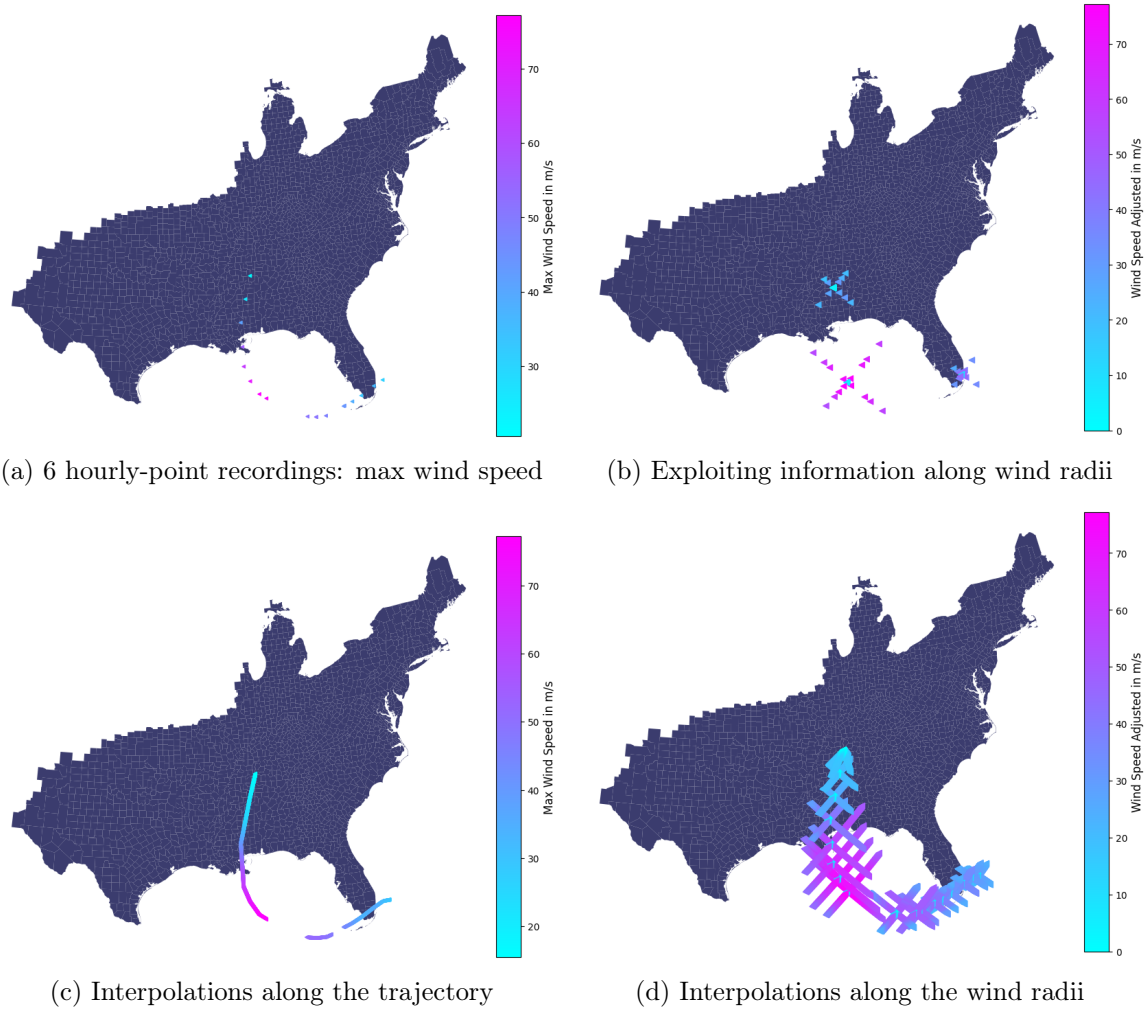
[3] This measure was developed as part of a short-term collaboration between myself and Professor Deryugina, with the mutual understanding that each of us could subsequently use the results independently in our respective research projects

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

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.

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).

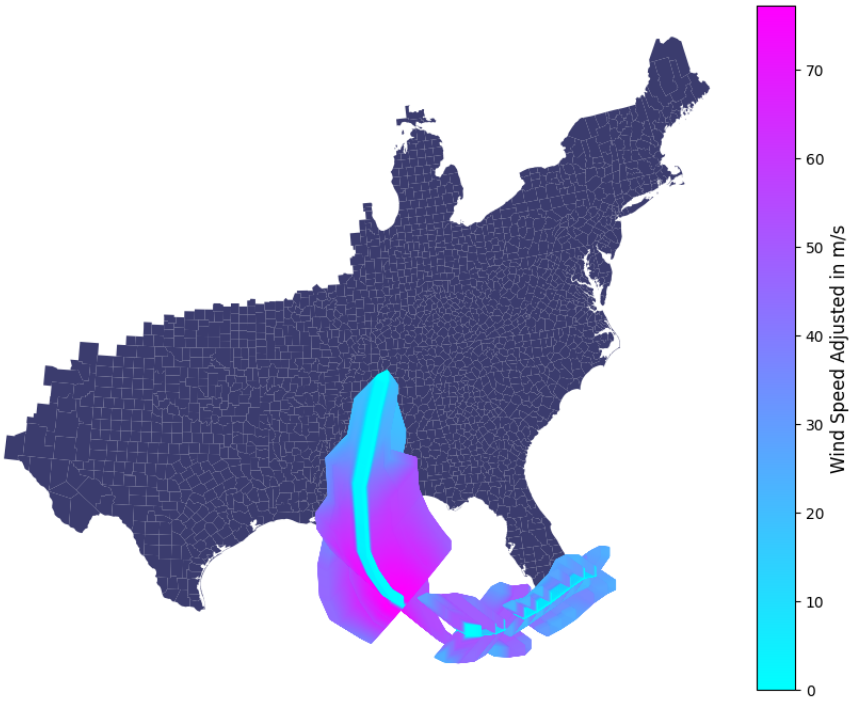
Panel 1a shows the 6-hourly recording points for the maximum wind speed, v_{ms} in Equation 1, excluding all the observation that are far away from the coast or with a wind speed below 17 m/s. These recording points represent the center of the cyclone and they come also with other data, like the minimum

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.

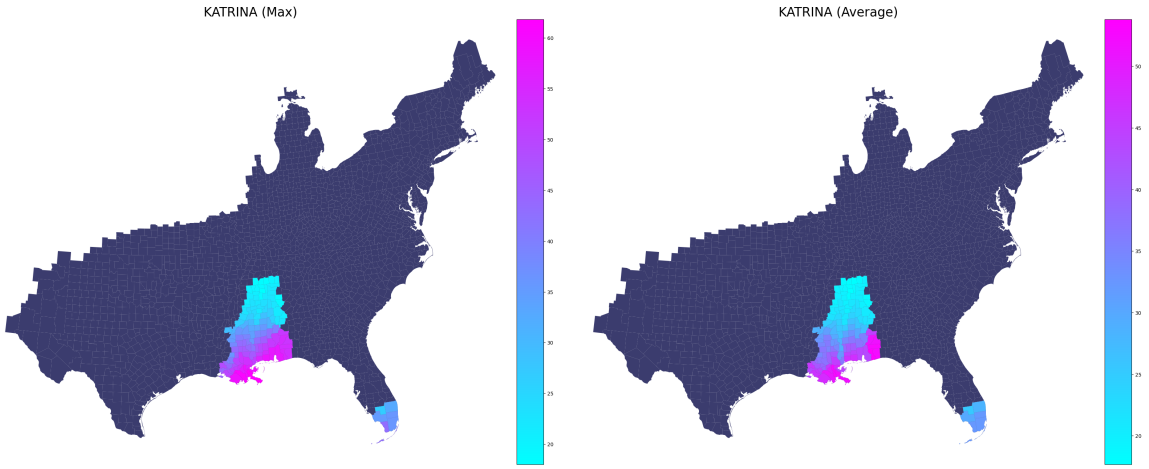
To fill the gaps between the 6-hourly recording points and between the different recordings of the maximum wind radii, I use linear interpolation respectively along the main trajectory (Panel 1c) and along the different wind extents (Panel 1d). The number of new points depends on the distance between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

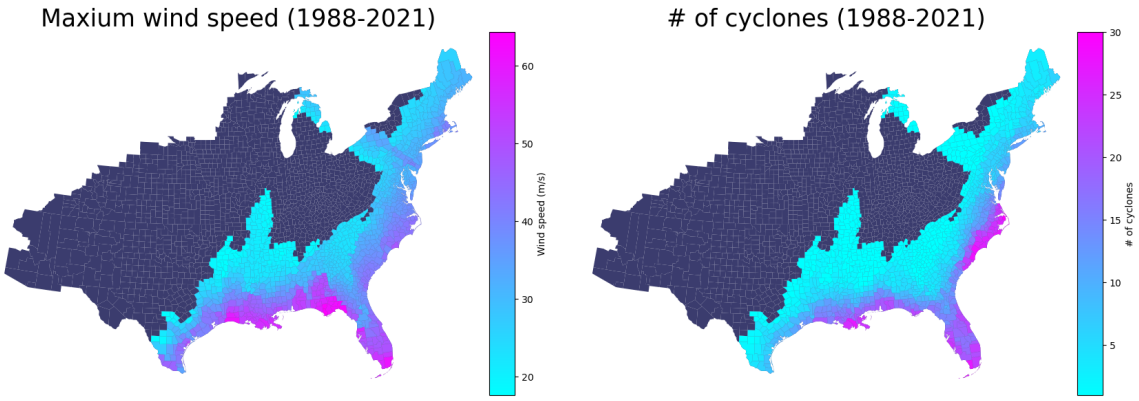


(a) Maximum wind speed of Katrina by county (b) Average wind speed of Katrina by county

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 1988. 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 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 B.1.

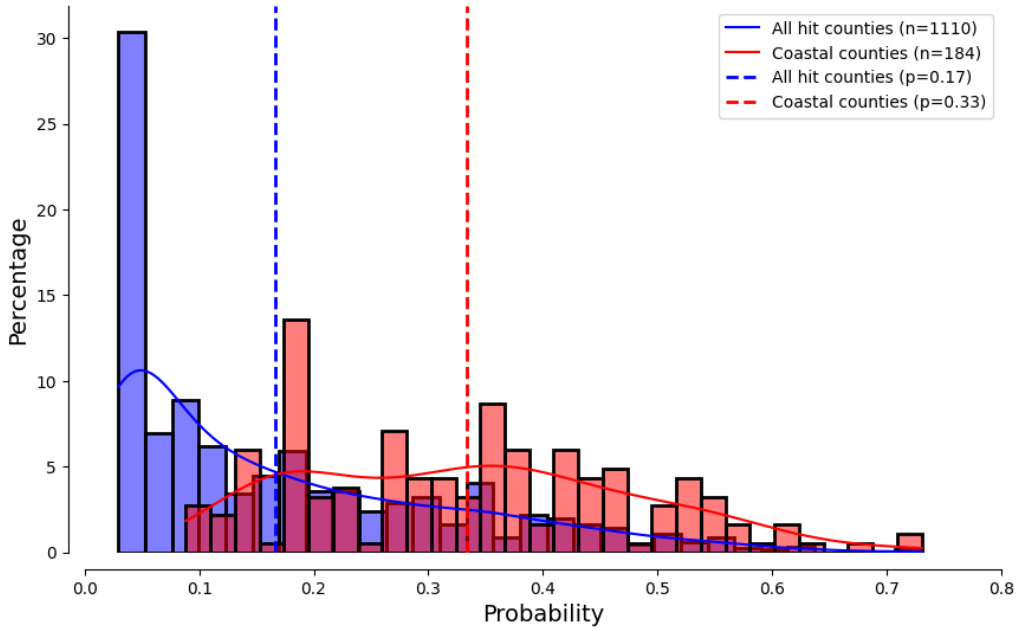
Figure 4: Aggregate the wind speed at county level



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

Figure 5 shows the distribution of the probability of receiving a cyclone from 1988 to 2021 for counties that experienced at least one cyclone (defined as those with wind speeds greater than 17 m/s, as shown in Figure 4), and only considering coastal counties. In particular, the sample size for coastal counties is much smaller. However, the probability distribution for coastal counties is more concentrated around the mean (0.33), indicating a more homogeneous pattern. In contrast, the distribution for all hit counties is skewed to the left, with a lower mean of 0.17, suggesting a less consistent pattern of cyclone probabilities in these counties.

Figure 5: Probability of Receiving Cyclones from 1988 to 2021



I focus the main analysis on coastal counties not only due to their higher probability and intensity of

cyclone events, but also because damage in these areas is often more pronounced. Coastal counties are particularly vulnerable to storm surge, a phenomenon in which rising sea levels, driven by high winds, low pressure, and the forward motion of the storm, flood coastal areas. Storm surge can significantly amplify the destructive impacts of cyclones, as the combination of strong winds and rising water levels leads to severe flooding and extensive damage to infrastructure and ecosystems. Given this increased vulnerability, it is crucial to emphasize coastal regions when analyzing cyclone impacts (Booth et al., 2021).

2.3 Housing value

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.[5] 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, 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) [6] 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

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

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.[7] This measure includes all separate living quarters intended 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).[8] This dataset, which spans from 1980 to the most recent reporting year, captures official authorizations for new residential construction.

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 1988 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

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

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

Table 1: Housing market averages 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
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	180	91	89

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 1988 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.

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 areas perceived to be less vulnerable to cyclone hazards.

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). This variable represents the number of flood insurance policies issued, serving as a proxy for a county's vulnerability

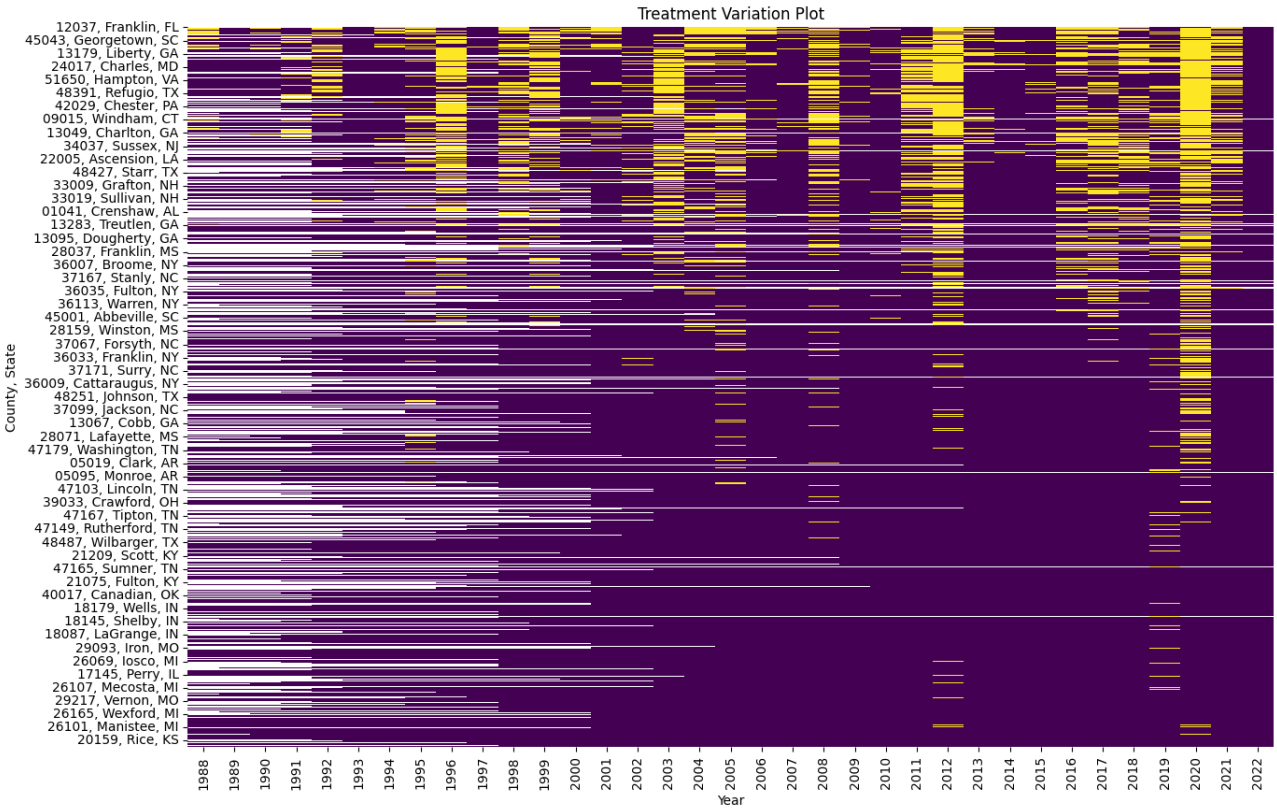
to flood-related risks and the level of preparedness among its residents and businesses. The National Flood Insurance Program (NFIP) is a federal initiative established by the U.S. Congress in 1968 to improve the availability and affordability of flood insurance in the United States.

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 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.

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.

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' units among coastal counties, as all experienced at least one cyclone between 1988 and 2021. Only a few coastal counties are untreated for more than ten years, as depicted in Figure 7.

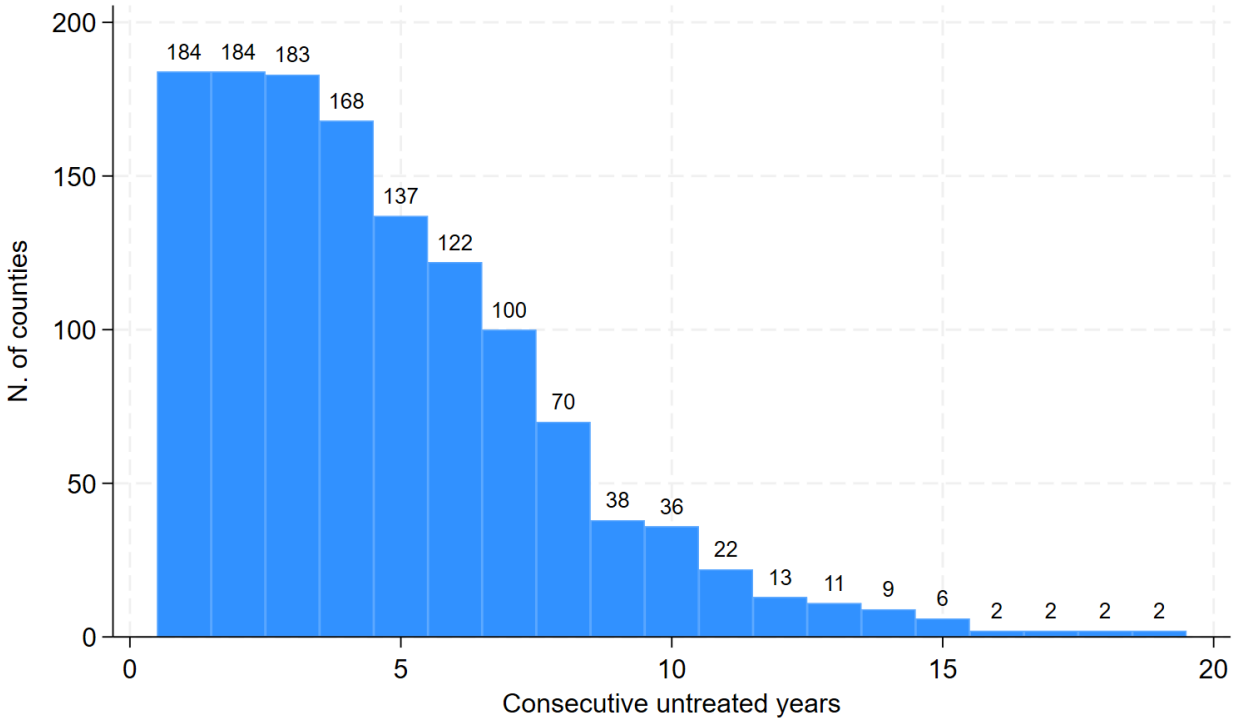


Figure 7: Consecutive untreated years for coastal counties

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 Jordà et al. (2023). This approach is well-suited for settings like mine, where treatment is non-absorbing, effects are dynamic, and all units eventually 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. The coefficient β_h^{LP-DID} captures the dynamic effect of cyclone exposure h years after the event.

The LP-DiD estimator avoids "contamination effects" caused by overlapping treatments or prior exposures, which is critical in a setting where no county remains permanently untreated and cyclone exposure varies irregularly over time. To ensure valid estimation of β_h^{LP-DID}, I restrict the sample following the criteria in Jordà et al. (2023):

- Treated units must not have experienced a cyclone in the K years before treatment and must remain treated for h subsequent years.
- Control units must be untreated from $t-K$ to $t+h$.

This 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 remain untreated for at least eight years prior to being classified as newly treated. I also verify the robustness of the results by extending this window—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^*.

Given the limitations of Two-Way Fixed Effect under treatment effect heterogeneity and dynamic responses, I rely on the LP-DiD estimates as the main specification. For transparency, I report TWFE and Distributed Lag Model estimates in the Appendix C, which provide useful but potentially biased benchmarks.
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.

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. Here, I analyze spillover effects by using a Spatial Lag Model, taking into account both one's own treatment status and the treatment status of neighboring counties. The Spatial Lag Model can be easily implemented using the Ordinary Least Squares estimation (Ruttenauer, 2024). 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 implement the spatial lag 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).

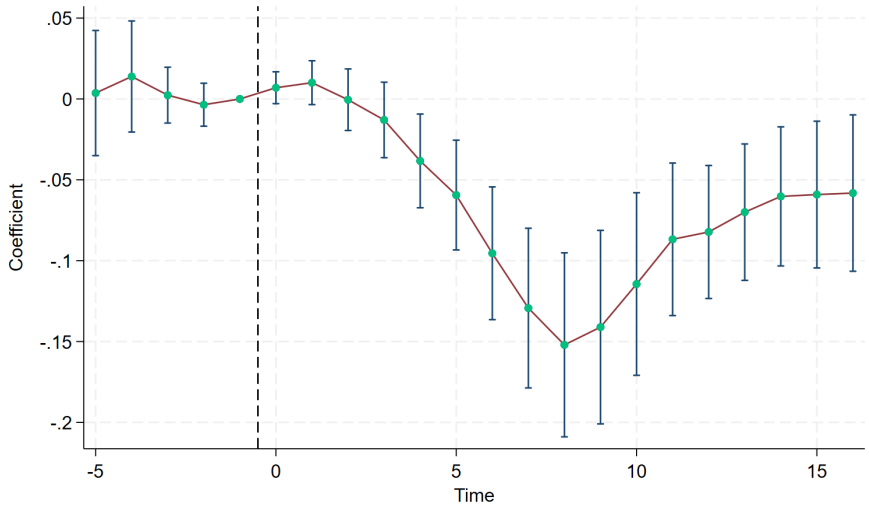
4 Results

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. Furthermore, I show that the effects dissipate for counties farther inland and test the optimal number of lags to include in Equation 7. Lastly, 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 Main result: changes in housing prices following cyclones

Figure 8 presents my main result: the effects of tropical cyclones on the value of housing relative to a pre-disaster baseline trend. 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, 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 the cyclone, indicating no anticipatory behavior or systematic pre-treatment divergence.

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



In this analysis, I assume that the effects of cyclones 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). Following a cyclone event, the housing market experiences a surge in value at first during the year of the cyclone and the subsequent year. However, from the second year onward, housing values begin to decline steadily over a period of approximately eight years compared to a counterfactual scenario in which the cyclone never occurred. The housing market shows signs of recovery by the ninth year, with values stabilizing around the 14th year. Given this dynamic pattern, setting $K^* = 8$ allows me to exclude units still in the declining phase of prior cyclone effects from the control group. This choice balances two goals: capturing the full medium-run impact of a new cyclone, and retaining enough untreated observations to maintain statistical power. I further verify that results remain stable for alternative windows $8 < K^* < 15$, indicating robustness to the exact cutoff. A cyclone 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, the counties experiencing cyclones witness an average decrease of around \$38,000 in property value after 15 years. Additionally, we can analyze the same effect using wind speed as a continuous variable. In Figure D.2b, an increase of 1 meter per second in wind speed corresponds to a housing inflation reduction of 0.008. 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.

This finding represents the key contribution of this paper, as previous estimates of cyclone impacts on the housing market identified only positive effects within the two years following a cyclone, without detecting any long-term effects (Graff Zivin et al., 2023). The observed decline indicates a much greater

persistence compared to other natural disasters, such as wildfires (McCoy and Walsh, 2018a).

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 specifications—such as a comparison between the Local Projection Difference-in-Differences (LP-DiD) estimator and the Distributed Lag Model (DLM)—as well as estimators that restrict the control group to never-treated units. I also replicate the analysis using an alternative housing market dataset (ZHVI) and explore sample selection criteria by varying geographic proximity to the Atlantic coast. 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 in the main specification. To guard against model misspecification or spurious correlations, I implement multiple forms of randomization tests. Furthermore, I explicitly address concerns about spatial spillovers by estimating a Spatial Lag Model and examining the outcomes in neighboring but untreated counties. Collectively, these robustness checks not only confirm the statistically significant and economically meaningful negative effects of cyclones on local housing markets but also inform key modeling choices—such as lag structure and sample definition—used in the main empirical specification.

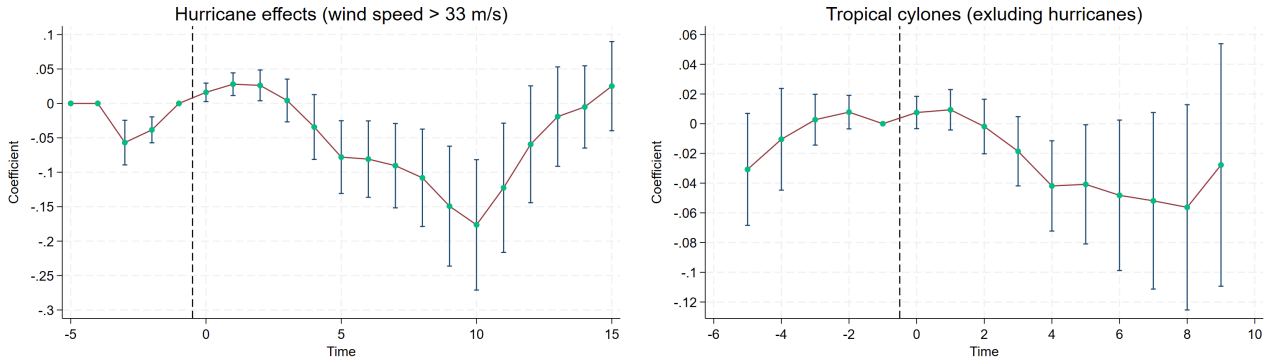
To further understand the short-run dynamics of housing markets following cyclone events, I examine the response of rental prices, which tend to adjust more quickly than home values in the face of temporary demand or supply shocks. Due to data limitations—rent data are only available from 2015—the analysis focuses on short-term effects. The results, presented in Appendix F, show no significant changes in rents in the months leading up to a cyclone, lending support to the assumption of no pre-trends. Following a cyclone, rents decline by up to 2% over the subsequent 15 months, suggesting localized market disruptions or temporary reductions in housing demand. From month 17 onward, rents begin to recover, indicating a gradual return to pre-disaster conditions. While the limited time span and smaller sample size restrict broader inference, the findings are consistent with the observed effects on housing values and highlight the transitory nature of post-disaster rental adjustments.

I also assess the response of local housing supply to cyclone exposure using two complementary measures: new residential construction permits and estimated total housing units. The event study results, shown in Appendix F, reveal no significant pre-trends for either measure, supporting the identification strategy. Following a cyclone, permit issuance declines sharply—by approximately 50% relative to the sample average—and remains subdued for about seven years. A sharp rebound in permits is observed from year 14 onward, which may signal delayed reconstruction or renewed development interest. Consistent with these patterns, the total number of housing units also falls in the aftermath of cyclones, indicating that the contraction in new construction is not fully offset by rebuilding in the short-to-medium term. Together, these results underscore a persistent negative effect on housing supply, with long-lasting implications for post-disaster recovery and local housing availability.

4.2 Heterogeneous impacts of cyclone by intensity and frequency

Here, I examine how the effects of cyclones on coastal counties vary by both intensity and frequency. Leveraging the continuous nature of my wind speed exposure variable, I am able to investigate heterogeneous impacts based on cyclone intensity. Specifically, I contrast the effects of strong cyclones—defined as hurricanes with maximum wind speeds of at least 33 meters per second—with those of weaker storms.

Figure 9: Impact of strong and weak cyclones on housing



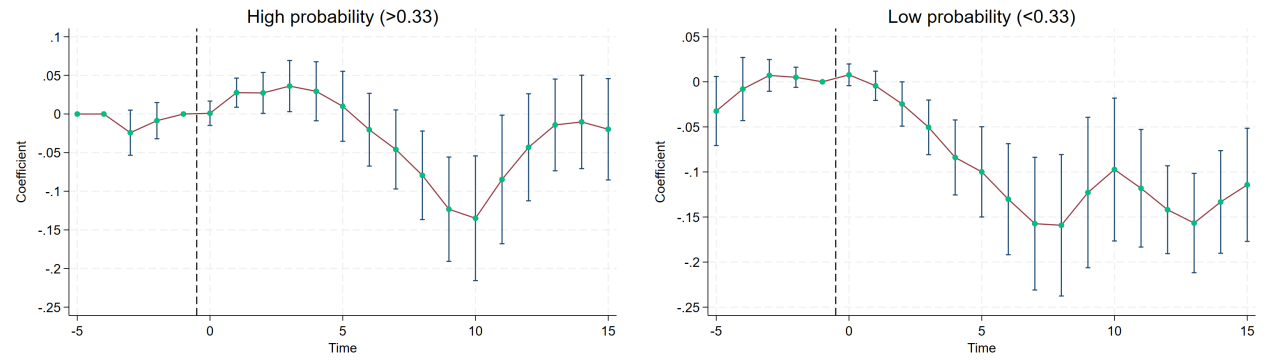
(a) Impact of hurricanes on housing

(b) Impact of weaker cyclones on housing

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. 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—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 the mean probability of cyclone occurrence for coastal counties (0.33), shown in Figure 5, I divide the sample into two groups: counties with a higher probability of experiencing cyclones and those with a lower probability.

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



(a) Impact of cyclones in high-probability zones

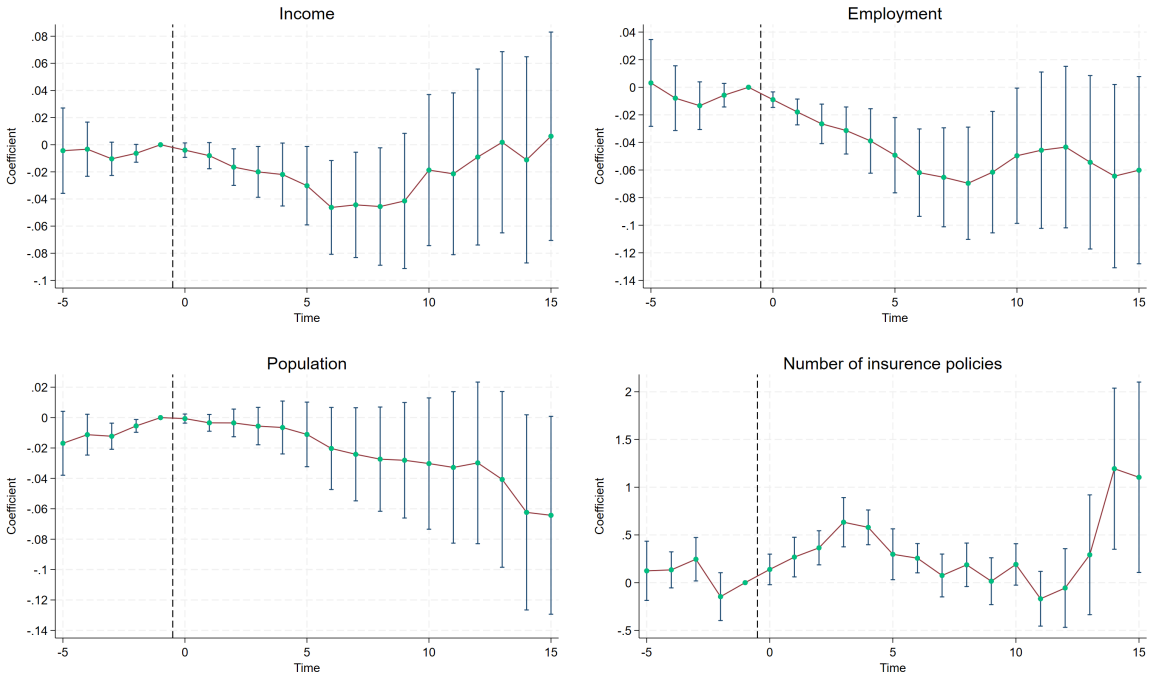
(b) Impact of cyclones in low-probability zones

Counties with a higher probability of cyclone exposure exhibit a more gradual response: housing values persistently increase for up to three years post-cyclone. Starting in the fourth year, prices drop, culminating in a reduction of 15% in the tenth year, followed by a recovery phase. In contrast, counties where cyclones are less frequent experience an immediate housing market decline beginning the year after the event, with no signs of recovery in the years following. This means that in low-probability areas, the negative impacts of cyclones on housing markets are more permanent. This suggests that highly exposed areas, which experience cyclones more frequently, may have stronger adaptation measures in place, such as resilient infrastructure, stricter building codes, and insurance mechanisms, that help mitigate long-term economic losses. This finding aligns with Hsiang and Jina (2014), who observed that the economic impacts of cyclones tend to be more severe in countries with limited historical exposure to such events.

4.3 Potential mechanism

To understand the drivers of the decrease in housing market after cyclones, I estimate the impact of cyclones on other economic variables such as income, employment, and population. With the same identification strategy described in Equation 2, I obtain the following results.

Figure 11: Impact of cyclones on other variables



The findings indicate that the long-term decline in housing values can be attributed more to an elevated perception of living in a high-risk area than direct economic damage. As is evident from the graphs, cyclones have only modest effects on employment, income, or population. On the other hand, the notable increase in flood insurance adoption suggests an added financial burden on affected counties, which contributes to the decline in the housing market. This outcome is in line with previous findings

on flood events, which showed an increase in flood insurance take-up (Gallagher, 2014).

5 Conclusion

This study provides the first plausibly causal, long-term estimates of the impact of tropical cyclones on the U.S. housing market, revealing a substantial and persistent 15% decline in housing values over an eight-year period. These effects extend well beyond the short-term disruptions commonly associated with natural disasters, and are notably more enduring than those linked to events such as wildfires (McCoy and Walsh, 2018b).

The magnitude and persistence of this decline are especially pronounced in counties with historically low exposure to cyclones, suggesting a lack of preparedness and adaptive capacity. In contrast, more frequently exposed areas appear to have developed resilience mechanisms that buffer housing markets from long-term devaluation. In addition, I find a marked increase in insurance uptake and premiums in cyclone-prone areas, highlighting a key channel through which market expectations may adjust to climate risk. Labor market effects, by contrast, are modest, indicating that the housing response is not primarily driven by income or employment shocks.

This work advances the literature by moving beyond studies of individual hurricanes to incorporate all cyclones over more than three decades, using a newly constructed wind speed exposure measure. The results underscore that climate events — particularly tropical cyclones — can induce deep and durable economic disruptions that are often underestimated in both research and policy (Carleton and Hsiang, 2016; Hsiang and Jina, 2014; Young and Hsiang, 2024).

Several limitations remain. First, while the evidence points to insurance dynamics as a potential mechanism, the precise channels linking cyclone exposure to housing market decline merit further investigation. Second, my analysis focuses on wind speed and excludes other relevant dimensions of cyclone damage, such as storm surge or flooding. Finally, though the 30-year cyclone probability weighting helps contextualize long-run risk, these probabilities are not static; climate change is likely to shift the distribution and frequency of future exposure.

In addition to documenting persistent effects of cyclones on local housing markets, the findings highlight important spatial disparities in the cost of living associated with climate risk. While housing units and values are generally higher in coastal areas compared to inland counties, a more nuanced pattern emerges within coastal regions themselves. Specifically, areas with historically lower cyclone exposure tend to have both more housing development and significantly higher average home values and rents. This suggests that, all else equal, people tend to concentrate in coastal zones perceived to be safer from cyclone-related disruptions. However, the event study reveals that these same low-exposure areas

suffer the most severe and persistent impacts when cyclones do occur, likely due to limited preparedness and weaker institutional or infrastructural resilience. One mechanism examined in this study involves post-disaster shifts in insurance behavior: areas with less frequent historical exposure may have lower pre-event insurance coverage, leading residents to adopt new policies after a shock—thereby increasing their long-run insurance costs.

Yet, the magnitude of the negative effect on housing prices suggests that insurance is only part of the story. As shown in this paper, cyclone exposure also depresses local labor markets, and likely entails a broader set of economic and logistical costs—particularly in places with limited prior experience managing such events. These may include higher maintenance and retrofitting costs, delays in rebuilding, or inefficiencies stemming from less established emergency response systems. A more granular, household-level analysis would help uncover these additional burdens.

Policymakers and insurers should recognize that counties with infrequent cyclone histories are not necessarily low-risk in economic terms. Even when the probability of landfall is relatively low, the lack of preparedness can translate into disproportionately high costs when disasters strike. Improving communication, insurance design, and resilience investments in these areas could help mitigate long-term losses and foster more equitable climate adaptation.

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:

$$v_s = v_{ms} \left(\left(\frac{r_{v_{ms}}}{r} \right)^{b_s} \exp \left[1 - \left(\frac{r_{v_{ms}}}{r} \right)^{b_s} \right] \right)^{x} \tag{4}$$

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):

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

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[9], and $x_1 = 0.6 \times (1 - (\Delta_{ps}/215))$. It

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

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:

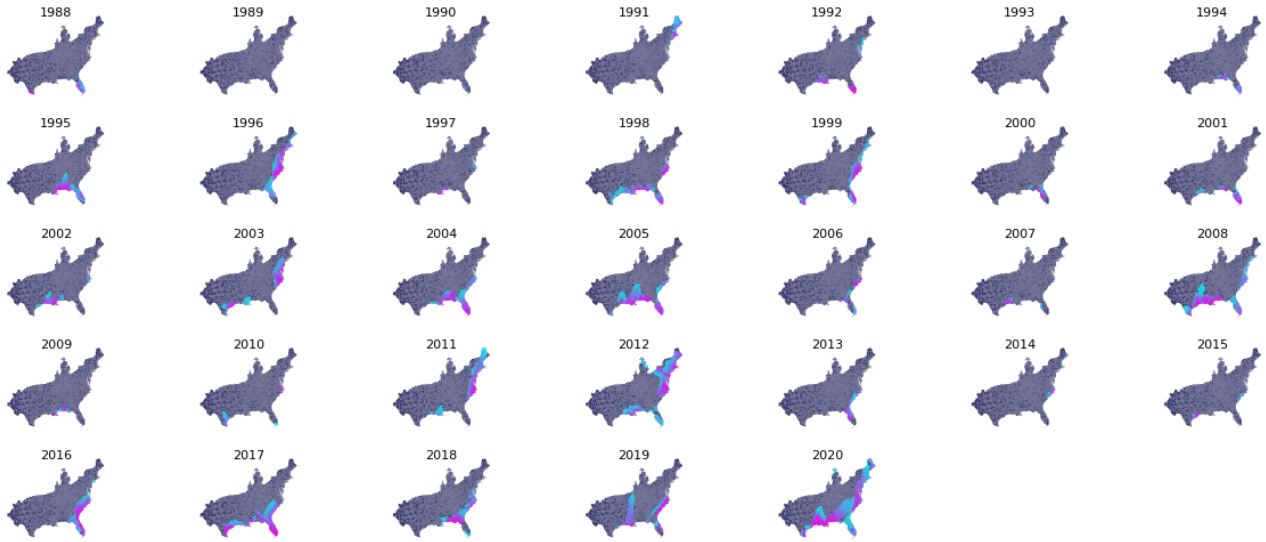
$$x_d = \frac{\ln\left(\frac{v_s}{v_d}\right)}{\ln\left(\left(\frac{r_{v_d}}{r_d}\right)^{b_s} e^{\left[1-\left(\frac{r_{v_d}}{r_d}\right)^{b_s}\right]}\right)} \tag{6}$$

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.

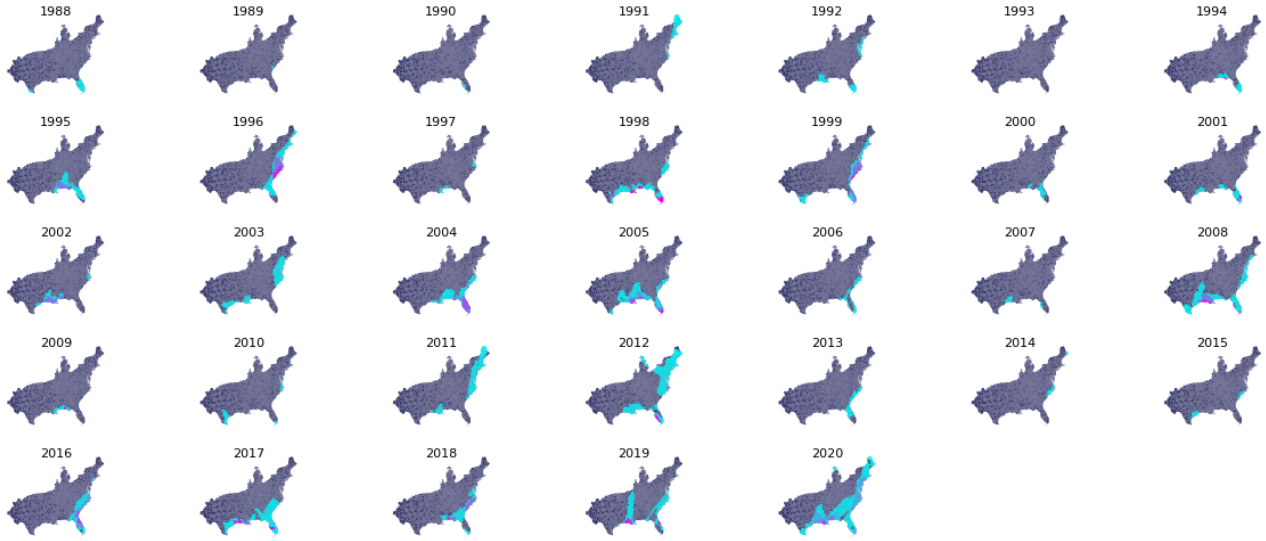
B Supplementary information on data variability

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:

$$\ln y_{i,t} = \alpha_i + \delta_t + \sum_{l=-Q}^{L} \gamma_l^{TWFE} D_{i,t}^l + e_{i,t} \tag{7}$$

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 7), 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:

$$\beta_k = \sum_{L=0}^{k} \gamma_L \tag{8}$$

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).[10] 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 7, 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:

- 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.

[10]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

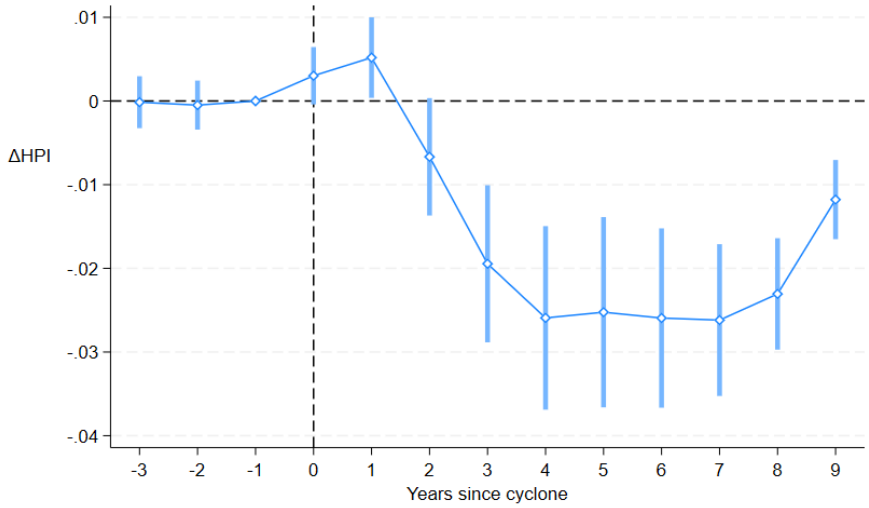
D Robustness of the main result

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.

Results from traditional event study approach

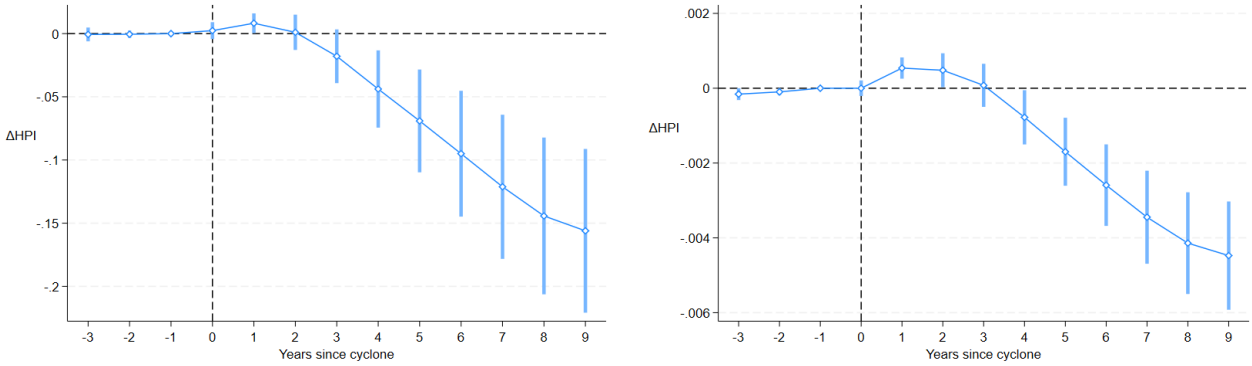
Figure D.1 presents the result of a traditional event study within Dynamic Two-Way First Difference Model: The plot depicts $\gamma_l \in [-3, 9]$ after Equation 7 is estimated for the coastal counties. It's important to note that the leads serve as placebo tests, demonstrating that there are no significant 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.1: 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 8). Figure D.2a 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.2: Cumulative effects of cyclone using Distributed Lag Model



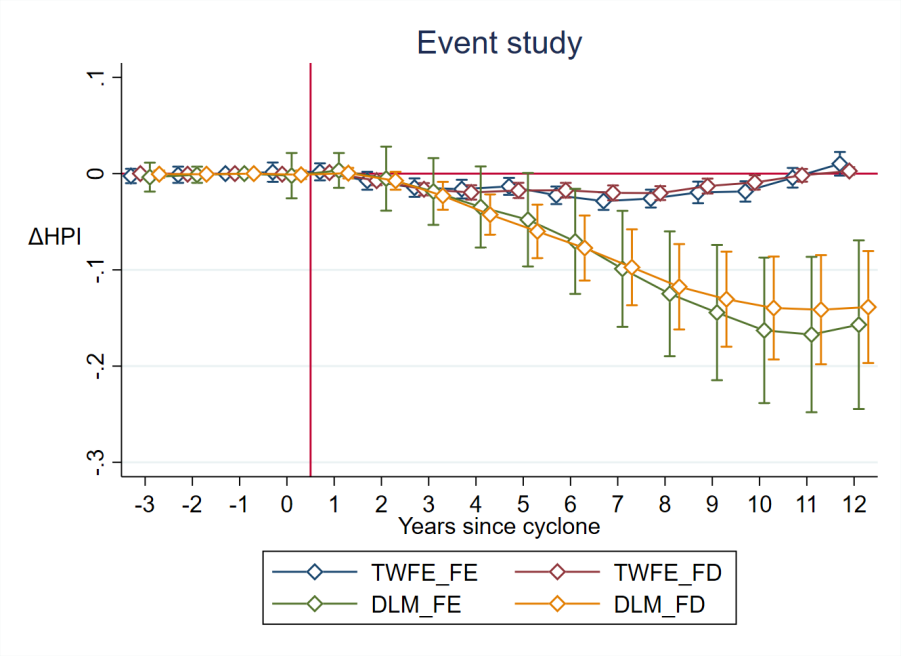
(a) Long-run effects of a new cyclone on hosing market

(b) Wind speed as a continuous variable

In Figure D.2a, 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. In the figure, an increase of 1 meter per second in wind speed corresponds to a housing inflation reduction of 0.008. 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.

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 7) and distributed lag model (DLM) (illustrated in equation 8), employing both Fixed Effect (FE) and first difference (FD) models.

Figure D.3: Comparison between fixed effect model and first difference model



The event study graph D.3 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.

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 8 .

Figure D.4: Comparison between DLM and LP-DiD

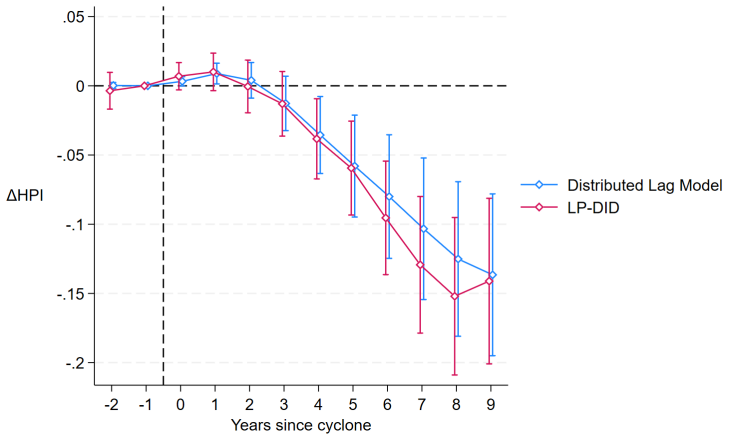
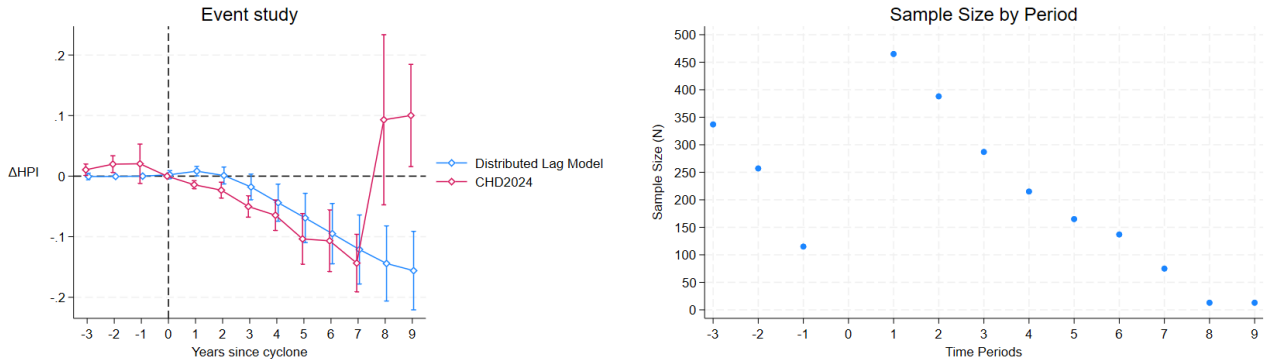


Figure D.4 demonstrates that the estimates obtained from the Distributed Lag Model (DLM) and

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.5a compares the results from the Dynamic Linear Model (DLM) with those from the de Chaisemartin and D'Haultfœuille (2024) estimator, while Figure D.5b 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.5: Comparison between DLM and CHD2024



(a) Comparison of DLM and CHD2024

(b) Observations by each lag estimate

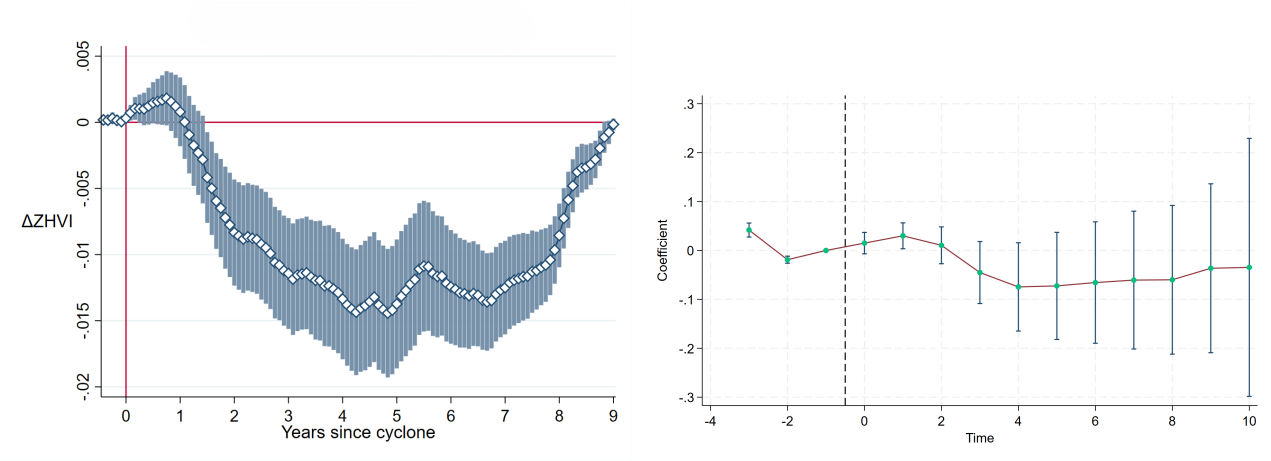
Similar to the LP-DiD approach, the de Chaisemartin and D'Haultfœuille (2024) estimator also identifies negative impacts of cyclones on the housing market. However, due to its reliance on fewer treatment units per period, it imposes stricter assumptions and limits the number of observations available 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.6a). I then aggregate the ZHVI data to the yearly level (Figure D.6b) and I use the LP-DiD estimator to facilitate direct comparisons with the HPI.

Figure D.6: Zillow Event Study



(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 1988 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.6a, 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.

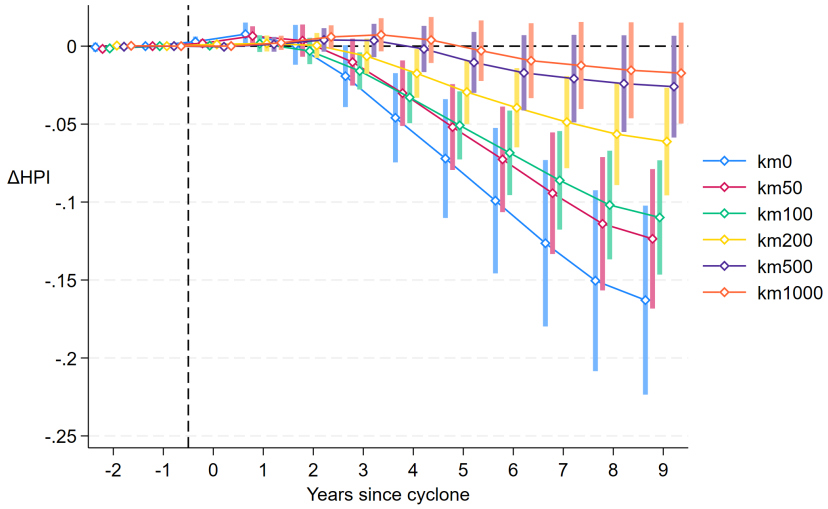
Apart from these differences and limitations, the ZHVI results broadly align with the narrative of the housing market response to cyclone events, reinforcing the findings derived from the HPI analysis.

Sample selection

Although Murphy and Strobl (2010) limit the analysis exclusively to coastal counties, other studies, such as Deryugina (2017), include all counties affected by hurricanes. To examine how the impact of cyclones varies with distance from the coast, I analyze samples that progressively incorporate inland

counties.

Figure D.7: Cyclone Impact Based on Proximity to the Atlantic Coast

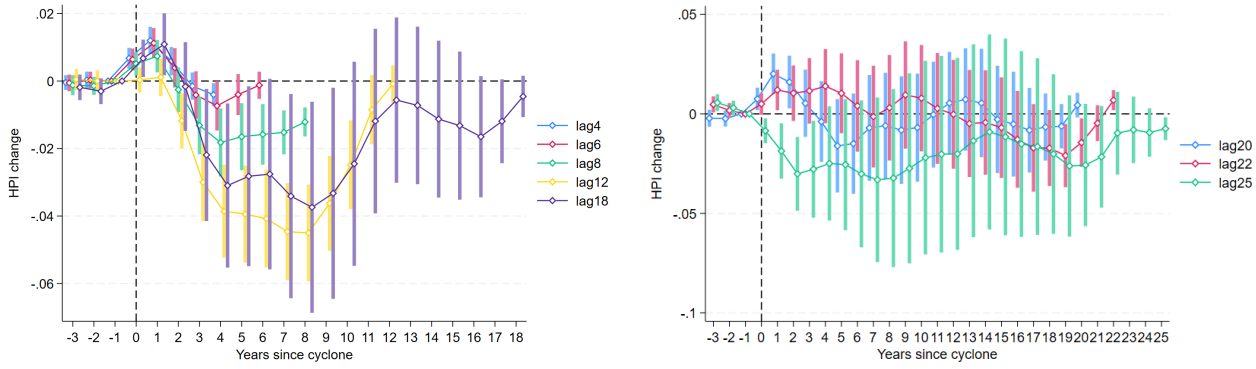


As shown in Figure D.7, the effects of cyclones diminish as the sample expands to include counties farther from the Atlantic coast. Coastal counties experience the strongest impacts, as they endure not only storm winds but also the added devastation of storm surges. Given that the most severe and economically significant effects are concentrated in coastal regions, I focus on these areas.

Lag length selection

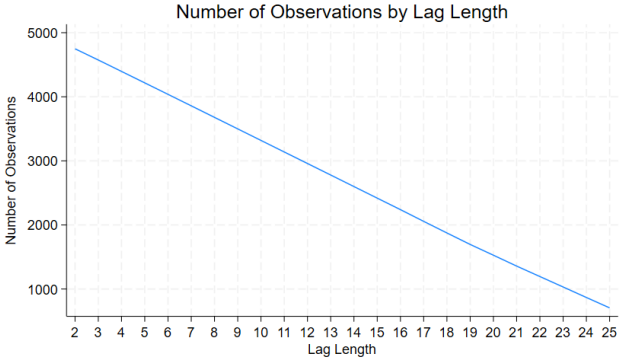
To determine the optimal number of lags to include in Equation 7, I plot several models with varying lag lengths. For example, the model with "lag12" includes all lags from 0 to 12 in Equation 7.

Figure D.8: 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.8b, from lag22, the placebo leads become statistically significant, leading to a lack of comparable estimates between the treated and control units.

In panel D.8a, 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 approach involves predicting housing prices for each county by training on the time series data from all other counties. The complete methodology is detailed in Appendix E. Based on cross-validation analysis, I conclude that the long-run effects of cyclones are best captured with a lag of 8.

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 7. 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.9: Randomization of cyclones

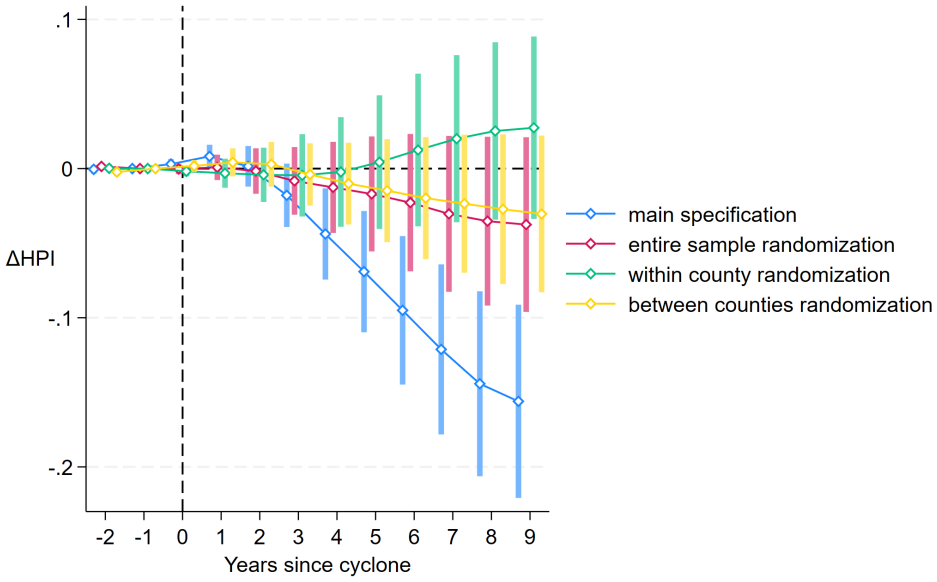


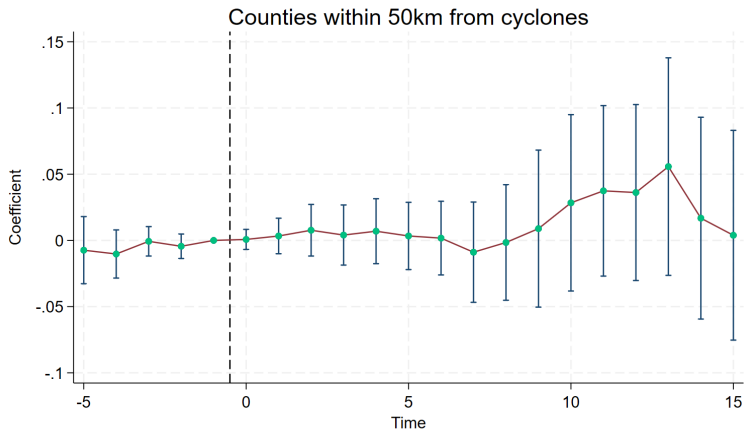
Figure D.9 shows that the three types of randomization yield insignificant results, thus reinforcing the robustness of my findings.

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.10: 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 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.

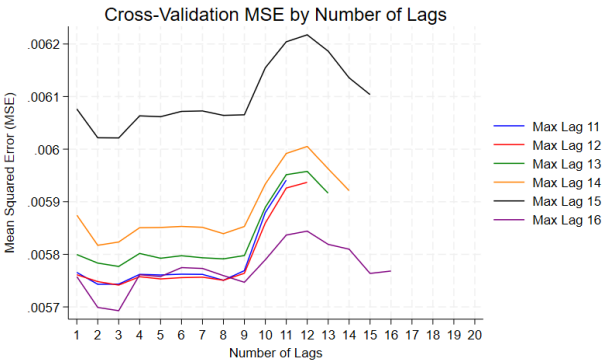
2. **Leave-One-Out Cross-Validation**: For each geographic unit i (e.g., a county), I treat it as the validation set while using all other units as the training set. The model is fit on the training set, then predictions are made for the validation set, computing the squared prediction error for that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

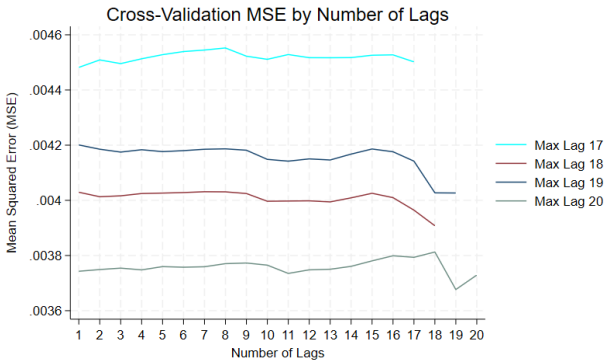
$$\text{MSE}(L) = \frac{1}{N} \sum_{i=1}^{N} \left(Y_i^t - \hat{Y}_i^t(L) \right)^2$$

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.

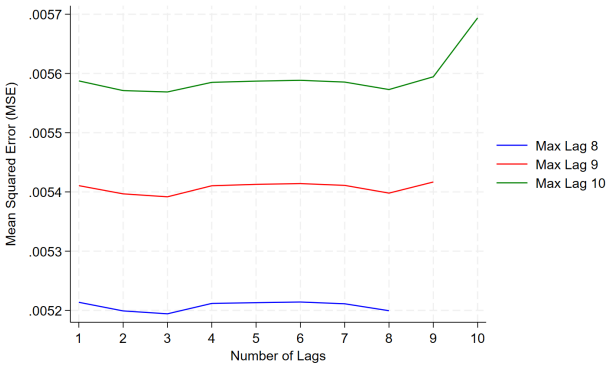
Figure E.1: Cross validation (MSE) by number of lags



(a) MSE from lag11 model to lag16



(b) MSE from lag17 to lag20



(c) MSE from lag8 to lag10

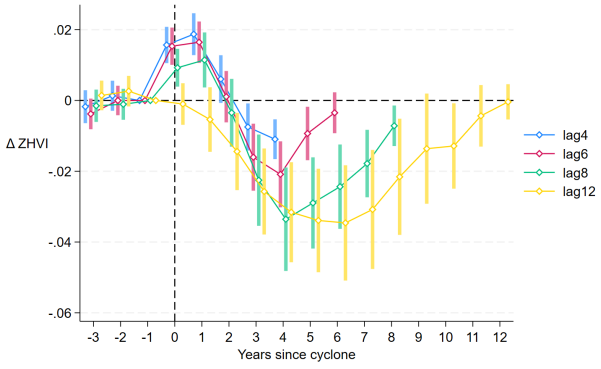
To improve clarity, I focus on results beginning from lag 11, as models with shorter lags exhibit relatively stable Mean Squared Error (MSE). As shown in panel E.1a, MSE values begin to increase after lag 8, indicating a decline in predictive accuracy across models as lag length extends up to lag 16. 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 E.2: Lag Selection with Zillow Data



F 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 F.1: Impact of cyclones on rent

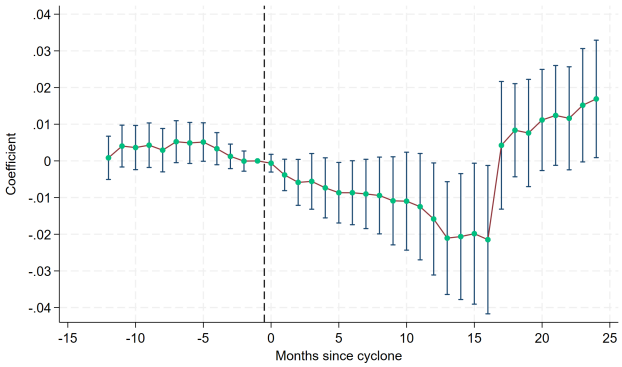
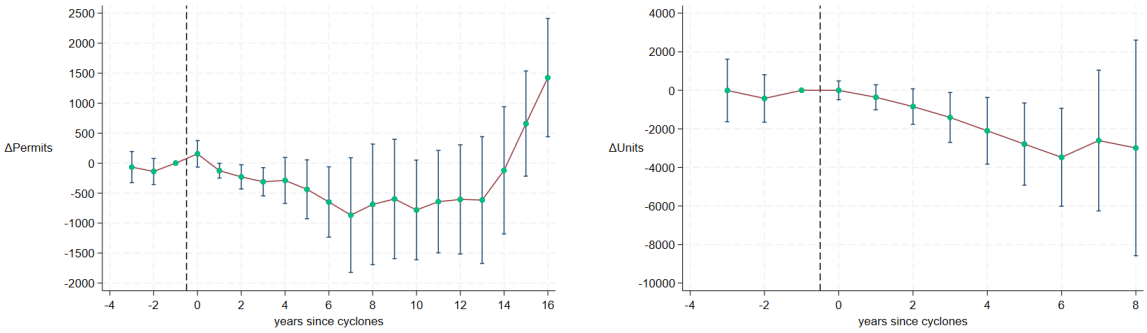


Figure F.1 presents an event study of the effect of cyclones on rent prices. The months leading up to the cyclone landfall exhibit no statistically significant changes in rent, supporting the assumption of no anticipation effects. Following a cyclone, however, rents decline noticeably, with a cumulative 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 F.2: Impact of cyclones on new housing permits and total housing units



(a) New construction permits

(b) Total housing units

Figure F.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 F.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.

References

Laura A Bakkensen and Lint Barrage. Going underwater? flood risk belief heterogeneity and coastal home price dynamics. *The Review of Financial Studies*, 35(8):3666–3709, 11 2021. doi: 10.1093/rfs/hhab122.

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/`.

Asaf Bernstein, Matthew T. Gustafson, and Ryan Lewis. Disaster on the horizon: The price effect of sea level rise. *Journal of Financial Economics*, 134(2):253–272, 2019. ISSN 0304-405X. doi: https://doi.org/10.1016/j.jfineco.2019.03.013.

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

James F. Booth, Veeshan Narinesingh, Katherine L. Towey, and Jeyavinoth Jeyaratnam. Storm surge, blocking, and cyclones: A compound hazards analysis for the northeast united states. *Journal of Applied Meteorology and Climatology*, 60(11), 2021. doi: 10.1175/JAMC-D-21-0062.1.

Brantly Callaway and Pedro H.C. Sant'Anna. Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230, 2021. ISSN 0304-4076. doi: https://doi.org/10.1016/j.jeconom.2020.12.001. Themed Issue: Treatment Effect 1.

Tamma A. Carleton and Solomon M. Hsiang. Social and economic impacts of climate. *Science*, 353 (6304):aad9837, 2016. doi: 10.1126/science.aad9837.

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.

Clément de Chaisemartin, Xavier d'Haultfoeuille, and Gonzalo Vazquez-Bare. Difference-in-differences estimators with continuous treatments and no stayers. *working paper*, February 8 2024. doi: 10.2139/ssrn.4720535. URL `https://ssrn.com/abstract=4720535`.

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.

Susana Ferreira. Extreme weather events and climate change: Economic impacts and adaptation policies. *Annual Review of Resource Economics*, 16(Volume 16, 2024):207–231, 2024.

Justin Gallagher. Learning about an infrequent event: Evidence from flood insurance take-up in the united states. *American Economic Journal: Applied Economics*, 6(3):206–33, July 2014. doi: 10.1257/app.6.3.206. URL https://www.aeaweb.org/articles?id=10.1257/app.6.3.206.

Stefano Giglio, Bryan Kelly, and Johannes Stroebel. Climate finance. *Annual Review of Financial Economics*, 13:15–36, 2021a. doi: https://doi.org/10.1146/annurev-financial-102620-103311.

Stefano Giglio, Matteo Maggiori, Krishna Rao, Johannes Stroebel, and Andreas Weber. Climate change and long-run discount rates: Evidence from real estate. *The Review of Financial Studies*, 34 (8):3527–3571, 03 2021b.

Andrew Goodman-Bacon. Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277, 2021. ISSN 0304-4076. doi: https://doi.org/10.1016/j.jeconom.2021.03.014. Themed Issue: Treatment Effect 1.

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.

Miyuki Hino and Marshall Burke. The effect of information about climate risk on property values.

Proceedings of the National Academy of Sciences, 118(17):e2003374118, 2021. doi: 10.1073/pnas.2003374118.

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.

Marynia Kolak and Luc Anselin. A spatial perspective on the econometrics of program evaluation. *International Regional Science Review*, 43(1-2):128–153, 2019. doi: 10.1177/0160017619869781. Original work published 2020.

Ricardo Carvalho de Andrade Lima and Antonio Vinícius Barros Barbosa. Natural disasters, economic growth and spatial spillovers: Evidence from a flash flood in brazil. *Papers in Regional Science*, 98 (2):905–924, 2019.

Shawn J. McCoy and Randall P. Walsh. Wildfire risk, salience & housing demand. *Journal of Environmental Economics and Management*, 91:203–228, 2018a. ISSN 0095-0696. doi: https://doi.org/10.1016/j.jeem.2018.07.005. URL https://www.sciencedirect.com/science/article/pii/S0095069617305181.

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

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.

Donald B. Rubin. Formal mode of statistical inference for causal effects. *Journal of Statistical Planning and Inference*, 25(3):279–292, 1990. ISSN 0378-3758. doi: https://doi.org/10.1016/0378-3758(90)90077-8.

Tobias Ruttenauer. Spatial data analysis. 2024. doi: https://doi.org/10.31235/osf.io/mq7te.

Kurt Schmidheiny and Sebastian Siegloch. On event studies and distributed-lags in two-way fixed effects models: Identification, equivalence, and generalization. *Journal of Applied Econometrics*, 38 (5):695–713, 2023. doi: https://doi.org/10.1002/jae.2971.

Eric Strobl. The economic growth impact of hurricanes: Evidence from u.s. coastal counties. *The Review of Economics and Statistics*, 93(2):575–589, 05 2011. ISSN 0034-6535. doi: 10.1162/REST_a_00082.

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`.

Liyang Sun and Sarah Abraham. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199, 2021. ISSN 0304-4076. doi: https://doi.org/10.1016/j.jeconom.2020.09.006.

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`.

Rachel Young and Solomon Hsiang. Mortality caused by tropical cyclones in the united states. *Nature*, 635(8037):121–128, Nov 2024. doi: 10.1038/s41586-024-07945-5.

Chapter 2

The Impact of Flooding on Credit Risk: Evidence from a Global Financial Institution*

Davide Castellani†

October 22, 2025

Abstract

Climate change is reshaping financial markets. Quantifying the impact of physical climate risks, such as floods, on firm credit ratings is essential for understanding firms' access to capital, particularly in emerging markets where financial systems are less developed and institutional lending plays a crucial role. Here, I quantify how flood events affect firm credit risk globally, using credit rating data for 3,243 firms from 1994 to 2023. The analysis cover private firms in emerging markets, regions that are both highly exposed to climate hazards and underrepresented in financial dataset. The findings show that exposure to a single flood event leads to a deterioration in credit ratings of approximately 20% over a three-year period. This effect is driven by flood events that displace more than 8,000 people. Moreover, I find a non-linear relationship between the frequency of flood exposure and credit deterioration. Firms experiencing two to five flood events over the 30-year period exhibit a smaller decline in credit ratings compared to those affected by a single flood, suggesting evidence of adaptive capacity. However, this pattern reverses for firms exposed to more than five floods, where the deterioration in creditworthiness becomes significantly larger, indicating potential limits to adaptation and resilience.

JEL-Codes: Q54, G32

Keywords: physical climate risk, credit risk, adaptation

*The views expressed in this paper are solely those of the author and do not represent the views of the International Finance Corporation (IFC) or its affiliates. All credit rating and default statistics are based on publicly available information (International Finance Corporation, 2024a). I thank Aurèle M. Houngbédji for his continuous support throughout the development of this work. I also thank my colleagues of the Climate Risk Team at IFC for their invaluable support in the data collection process and for the collaborative environment that made this research possible. I also want to acknowledge the insightful feedback and engaging discussions with Giovanni Mastrobuoni, Elena Esposito, Solomon Hsiang, Paul Smeets, Guglielmo Zappalà

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

1 Introduction

In recent years, there has been a global increase in the frequency and intensity of extreme precipitation events, a trend that has become particularly pronounced over the last decade and is closely linked to human-induced climate change (IPCC, 2021). One of the most severe consequences of these changes is flooding, which remains the most frequent natural disaster worldwide. Approximately 1.81 billion people, or 23% of the global population, are directly exposed to 1-in-100-year flood events (Rentschler et al., 2022). These physical climate shocks not only reduce firm productivity (Carleton and Hsiang, 2016; Dell et al., 2014; Burke et al., 2015), but also threaten firms' financial stability (Giglio et al., 2021a), especially in emerging markets (Aguilar-Gomez et al., 2024). Empirical studies show that climate risk affects credit and firm valuation: corporate bond spreads widen (He et al., 2025), municipal and sovereign borrowing costs rise (Painter, 2020; Goldsmith-Pinkham et al., 2023), bank loan spreads increase (Huang et al., 2022), and equity returns fall (Venturini, 2022; Engle et al., 2020; Faccini et al., 2023). However, these measures are often limited to publicly traded firms in advanced economies, reflect investor expectations and can capture rapid changes in perceived risk driven by short-term sentiment, making it difficult to isolate firm-level credit risk. In contrast, credit ratings offer a more direct and forward-looking measure of credit risk, based on systematic assessments of a firm's financial health, governance, and ability to meet debt obligations. Credit ratings play a central role in determining firms' access to capital, particularly in emerging markets, where financial systems are less developed and firms depend more heavily on institutional lending. Despite increasing attention to climate physical risk, there is still a significant empirical gap in understanding how specific hazard, such as flooding, affect credit ratings over time and whether firms adapt to such risks.

This gap stems largely from data limitations. Firm location linked to credit ratings are rarely available, particularly in emerging markets where disclosure practices are limited. Using proprietary data from the International Finance Corporation (IFC),[1] this paper provides the first evidence linking flood events to firm-level credit ratings and default in emerging markets. Moreover, the analysis uncovers a non-linear relationship between the frequency of flood exposure and credit deterioration: while moderate exposure appears to induce partial adaptation, repeated or severe flooding leads to renewed declines in creditworthiness. Studying repeated flood events is crucial for understanding credit risk and financing dynamics because firms' financial resilience depends not only on their immediate response to a single disaster but also on their capacity to recover and adapt to recurrent shocks.

The IFC is the private-sector development arm of the World Bank Group, supporting investment and growth in emerging markets. The dataset contains credit ratings for 3,243 IFC clients between 1994 and 2023. Ratings are recorded on a 15-point ordinal scale, where a higher value indicates greater

[1] To the best of my knowledge, this is the first academic study using IFC credit ratings. Access to these data was made possible through a fellowship at the IFC, as acknowledged in section 7.

credit risk, and a rating of 15 corresponds to default (International Finance Corporation, 2024a). Thus, increases in credit rating reflect a deterioration in a firm's creditworthiness. The IFC dataset provides a uniquely valuable foundation for this analysis, as it includes headquarter coordinates for each client, enabling precise mapping of flood exposure to individual firms over time. Unlike major cradit rating agency (CRAs) such as S&P Global, Moody's, and Fitch, which primarily cover large, publicly traded firms in advanced economies, IFC credit ratings capture the financial health of private firms operating across emerging and developing markets, regions that are both highly exposed to climate hazards and underrepresented in global financial datasets. The IFC ratings are derived from a rigorous, standardized internal process involving annual reviews, on-site due diligence, and approval by credit committees, similar in rigor to CRAs but applied to private clients. Combined with nearly three decades of longitudinal coverage, these features make IFC data exceptionally well-suited to identify how physical climate shocks, such as flooding, affect firm-level creditworthiness.

To identify clients exposed to flooding, I link the IFC credit rating data with two complementary sources of flood information: the Dartmouth Flood Observatory, which compiles global flood events from news and official reports dating back to 1985, and the Global Flood Database, which provides satellite-based measurements for 192 major flood events worldwide between 2000 and 2017. While the Dartmouth dataset offers broader temporal coverage, the Global Flood Database provides much finer spatial resolution—up to 250 meters—allowing for more precise identification of localized flood exposure.

Despite the originality of the data, several key empirical challenges arise. First, asset locations are often not co-located with firm headquarters. Here I can focus only on headquarters locations under the assumption that they remain a critical component of corporate operations, decision-making, and strategic exposure. Headquarters often concentrate key functions—such as finance, governance, and risk management—that are central to firm performance and resilience. Moreover, they typically anchor firms' reputation and logistical networks, meaning that disruptions at headquarters can have firm-wide consequences even if other assets are geographically dispersed. Under this interpretation, restricting attention to headquarters locations is likely to bias the results downward, since additional vulnerabilities at the asset level are not captured. Thus, my findings can be viewed as a conservative estimate of the relationship between flood exposure and firm outcomes. Second, the repeated occurrence of floods in the same geographic areas makes it difficult to clearly distinguish between treated and control firms over time. To overcome this challenge, I use the local projection estimator for Difference-in-Differences event studies recently developed by Jordà et al. (2023), which allows for flexible dynamic treatment effects while ensuring that control observations remain unaffected by treatment across a specified event window. I begin the analysis by comparing firms that experience a single flooding event with those that have never been exposed to flooding. This allows for an initial estimation of the average effect of a

one-time events. I then extend the analysis by incorporating multiple flooding exposures, considering alternative assumptions about the duration of impacts and the degree to which subsequent events can be treated as isolated shocks. This stepwise approach ensures both a baseline comparison and a richer understanding of how repeated climate shocks affect firms over time.

The analysis reveals that both the intensity and frequency of flood exposure have significant and lasting impacts on firm creditworthiness. Firms exposed to a single severe flood, defined as an event displacing more than 8,000 people, experience a deterioration in credit ratings of approximately 20% within three years, while no significant effect is observed for less intense events. The negative effects are largely driven by clients in emerging markets, where the deterioration in creditworthiness is significantly greater than in high-income countries. Repeated exposure compounds these risks: clients affected by more than five flood events during the 30-year period exhibit outcomes comparable to those with only one, suggesting that while some degree of adaptation may occur with moderate exposure (two to four events), this resilience breaks down under sustained climate stress. Sectoral analysis shows that firms in commerce, such as retail and banking, suffer a more persistent credit deterioration than those in industry. Although physical damage is more intuitively associated with industrial activity, the stronger and more persistent deterioration observed in commerce likely reflects the concentration in urban, flood-prone areas, where even moderate flood events can disrupt operations, cash flows, and client solvency, whereas large industrial facilities may often more spatially dispersed or equipped with protective infrastructure.

This paper contributes to the literature in three ways. First, it provides novel empirical evidence on the financial implications of physical climate risk at a global scale, including emerging markets. Previous research has examined how such risks are capitalized in asset prices, from housing markets (Giglio et al., 2021b; Bernstein et al., 2019; Murfin and Spiegel, 2020) to corporate outcomes (Giglio et al., 2021a). For instance, Bernstein et al. (2019) show that homes exposed to sea level rise sell at a discount compared to unexposed properties in the US. Giglio et al. (2021a) highlight that climate risk can increase firms' default probabilities and financing costs, emphasizing its relevance for financial stability and firm valuation. Other studies document sector-specific impacts, including reduced earnings in certain industries (Addoum et al., 2023), lower industrial output (Chen and Yang, 2019), and decreased liquidity alongside higher debt levels for manufacturing plants (Elliott et al., 2019). Most of this evidence comes from the US or China, and focuses on particular sectors. By examining how firm credit ratings respond to flooding events across different countries and industries, this paper provides a more systematic, globally relevant assessment of how climate physical risks affect financial stability, enhancing the external validity of this literature.

Second, my findings offer novel empirical evidence on how extreme weather events affect credit risk, measured directly through credit ratings and realized defaults. Previous research has focused primarily

on how climate risk affects corporate bonds (He et al., 2025; Huynh and Xia, 2021), municipal bonds yields (Painter, 2020; Goldsmith-Pinkham et al., 2023), bank loan spreads (Huang et al., 2022) or trade credit (Lai et al., 2022). For example, He et al. (2025) find that physical climate risk increases bond credit spreads in China by reducing internal cash flows and increasing external financing constraints. Other studies model climate risk via stock market responses or risk premia, highlighting how extreme weather events and long-term climate hazards can influence firm valuation, cost of capital, and investment decisions (Venturini, 2022; Engle et al., 2020; Faccini et al., 2023). For example, Engle et al. (2020) and Faccini et al. (2023) show that firms exposed to higher physical climate risk experience lower expected returns, while investors demand a risk premium for climate vulnerability. Kilincarslan et al. (2025) demonstrate that firms with better climate-related disclosures have improved credit scores, highlighting the importance of accounting for climate risk in credit assessments. However, no study has causally examined how direct physical exposure to extreme weather events affects corporate creditworthiness as measured by credit ratings.[2] This paper addresses this gap by linking firm-level credit ratings to observed flood events, providing the first causal evidence on how direct physical climate shocks affect corporate creditworthiness.

Third, this paper contributes to the emerging literature on firm adaptation to climate risk by documenting a non-linear relationship between flood exposure frequency and credit deterioration. While existing studies in macroeconomics and environmental economics find non-linear responses of productivity and output to climate shocks (Burke et al., 2015; Dell et al., 2014; Carleton and Hsiang, 2016), the financial literature has largely treated physical risk as a linear determinant of asset prices and financing costs (He et al., 2025; Huang et al., 2022; Painter, 2020). Only a few recent papers explore firms' adaptive responses to climate risk through changes in operations or investment behavior (Addoum et al., 2023; Giglio et al., 2021a), yet none have examined how adaptation manifests in credit risk outcomes. This paper fills that gap by providing the first evidence that firm-level creditworthiness responds non-linearly to repeated climate shocks, highlighting both adaptive capacity and its limits under persistent exposure.

The rest of the paper proceed as follow. Section 2 describes the background and data. Section 3 introduces the research design to identify the impact of flooding on credit risk. Section 4 presents the results. Section 5 provides the conclusion, and Section 6 acknowledgments.

[2]Only Nguyen et al. (2025) show that firms in countries with higher climate vulnerability tend to receive lower ratings, but their analysis relies on country-level aggregates and cannot identify causal effects at the firm level.

2 Background and Data

2.1 Physical risk

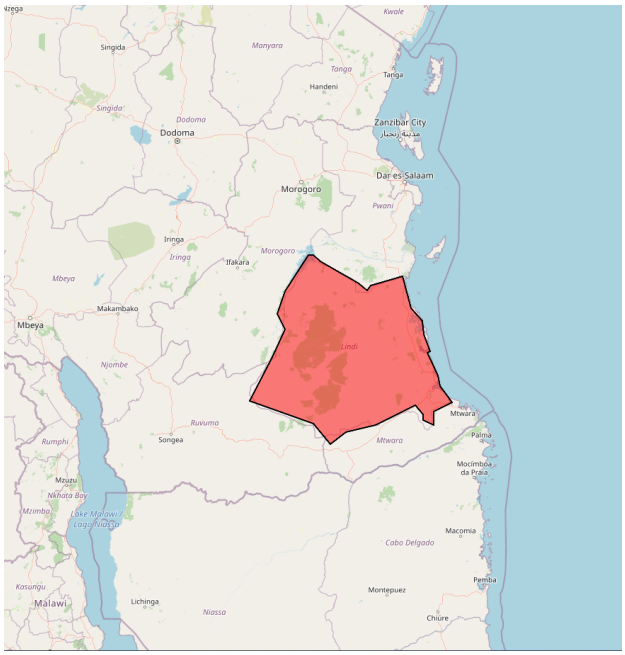
Physical risks—defined as the direct consequences of climate-related hazards such as floods, hurricanes, wildfires, and droughts—are increasingly recognized as material threats to financial stability. These risks can damage physical assets, disrupt business operations, and impair a firm's ability to meet financial obligations, ultimately influencing credit risk and valuation. Recent work has highlighted how such risks can affect asset prices, lending practices, and firm default probabilities, yet empirical evidence remains limited, particularly in emerging markets. In this study, I focus on one of the most common and damaging physical hazards—flooding—as a lens through which to examine how environmental shocks shape credit outcomes.

2.1.1 Flooding exposure

To operationalize exposure to physical risk, I focus on flooding, the most frequent and economically costly natural disaster worldwide. This analysis draws on two complementary data sources: the Dartmouth Flood Observatory (DFO) and the Global Flood Database (GFD). The DFO provides a comprehensive global record of flood events since 1985, compiled through media reports, government sources, and satellite imagery. It offers broad temporal and geographic coverage, making it particularly useful for long-run event analysis. In contrast, the GFD, developed using satellite-based remote sensing, includes 192 major flood events between 2000 and 2017 and offers high-resolution spatial data (250-meter resolution). While the GFD captures fewer events, its precision enables more accurate identification of firm-level exposure to flooding. By combining these datasets, I am able to capture both the long-term frequency and the fine-grained geography of flood events, allowing for a robust examination of their effects on corporate credit risk.

Figure 1 illustrates examples of flood polygons constructed by the Dartmouth Flood Observatory (DFO). These polygons represent the estimated geographic extent of each flood event, based on a combination of satellite imagery and information collected from news reports and official sources.

Figure 1: Examples of flooding extent from DFO



(a) Example 1

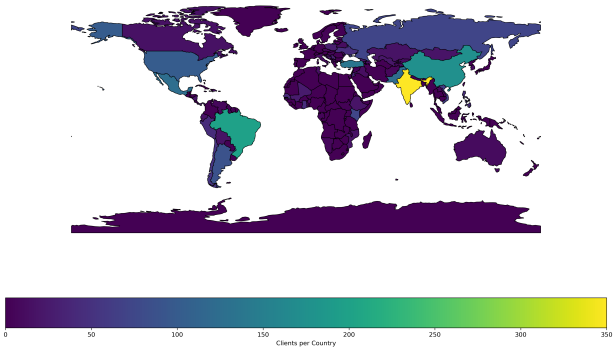


(b) Example 2

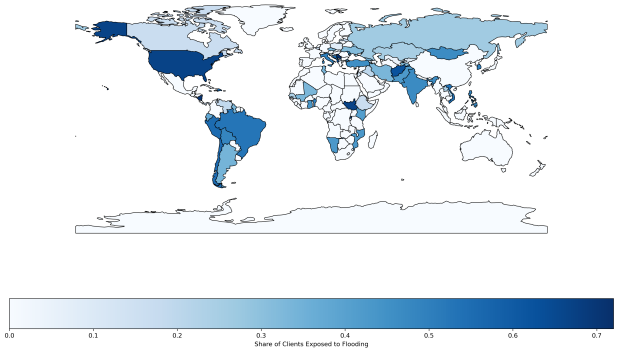
To assess client exposure, I spatially matched the DFO flood polygons with the geographic coordinates of the IFC client headquarters. For each year, I identified whether a client was located within a flood-affected area. I find that 2,165 out of 3,243 IFC clients, over 60%, experienced at least one flood event during the sample period.

The global distribution of IFC clients is illustrated in the maps below, which showcase the number of client headquarters per country and the proportion of clients that have experienced at least one flooding event.

Figure 2: Comparison of Clients' Headquarters and the Share of Clients Exposed to Flooding



(a) Number of Clients' Headquarters per Country



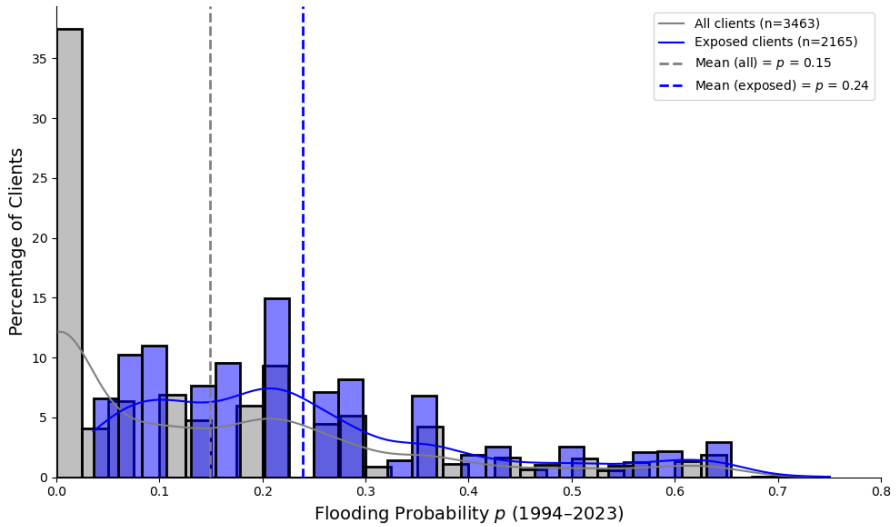
(b) Share of Clients Exposed to Flooding

As shown in Figure 2a, India ranks first in terms of the number of client headquarters, with approximately 350 clients, followed by countries like Brazil, China, Argentina, the United States, and Mexico. Interestingly, Figure 2b reveals that countries such as Brazil, the United States, India, and

Turkey—home to a significant concentration of client headquarters—also have a large proportion of clients who have experienced at least one flooding event.

To analyze how the likelihood of experiencing a flood varies across clients, I calculate, for each client, the share of years since 1994 during which their headquarters were located within a flooded area. This frequency serves as a proxy for the probability of flood exposure, denoted by p. Figure 3 presents the distribution of p across all clients, as well as the subset of clients who experienced at least one flood during the sample period.

Figure 3: Probability of Receiving Flooding from 1994 to 2024



Note: This figure plots two distributions of the probability of experiencing a flood event over 1994–2024, computed as years with flooding divided by total years in the panel. The grey distribution covers all clients; the blue distribution is restricted to clients that experienced at least one flood event. Dotted vertical lines indicate the mean of each distribution.

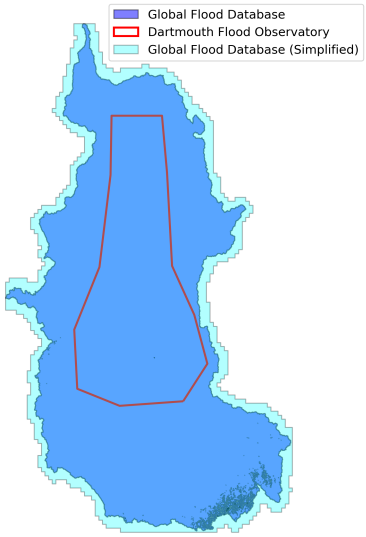
The average probability of experiencing at least one flood event over the 30-year period is 0.24, indicating that affected areas face a flood approximately once every four years. In Table 1, I use this probability as a benchmark to compare how credit risk varies depending on whether a region is above or below this average. Additionally, I conduct subsample analyses in the empirical section to examine how the impact of flooding differs across varying probabilities of occurrence, providing insights into potential adaptation to these events.

In addition to geographic information, the DFO dataset includes qualitative details for each flood event, such as the number of displaced people, fatalities, and infrastructure damage, which I use to report descriptive statistics in Table 1.

While the DFO provides broad coverage over time, the delineated flood polygons tend to underestimate the true spatial extent of flooding. This becomes evident when comparing DFO polygons to those from the Global Flood Database (GFD), which is based on satellite-detected surface water changes at a much

finer resolution (250 meters). Figure 4 illustrates this comparison for a representative event: the red polygon (DFO) lies entirely within the broader blue-shaded area derived from the GFD. To manage the high granularity of the GFD—comprising millions of pixels—I generate simplified flood boundaries by extracting only the outermost pixels, which I refer to as "GFD simplified" in the figure.

Figure 4: Comparison of flood extent: DFO vs. GFD



This pattern holds across multiple events (additional examples are provided in the Appendix A.1), and the discrepancy between datasets has quantitative implications: despite covering only 192 events between 2000 and 2017, the GFD identifies flood exposure in approximately 18% of firm-year observations—higher than the DFO's estimate. This comparison highlights the importance of spatial resolution in measuring physical climate risk and motivates the use of both datasets in the analysis. While the DFO provides longer historical coverage, with data dating back to 1994, the GFD offers more spatially granular information but is only available from 2000 onward—illustrating a trade-off between historical coverage and spatial precision.

2.2 Credit risk

Credit risk refers to the possibility that a borrower will fail to meet its financial obligations, potentially resulting in financial losses for the lender. In this study, I use firm-level internal credit ratings and default events as proxies for credit risk, consistent with established practices in the finance literature (Altman, 1968; Merton, 1974; Duffie and Lando, 2001). Credit ratings offer an assessment of a firm's creditworthiness, incorporating both quantitative financial indicators and qualitative assessments of business risk. Default, in turn, represents the most severe manifestation of credit risk. Together, these measures provide complementary views on a firm's evolving financial health and exposure to adverse shocks—such as those induced by natural disasters.

2.2.1 International Finance Corporation

The International Finance Corporation (IFC) is a member of the World Bank Group and the largest global development institution dedicated to the private sector in emerging and developing markets. Founded in 1956, IFC plays a critical role in promoting inclusive and sustainable economic growth by providing investment, advisory, and asset management services. Its mission is to encourage private sector development as a means to reduce poverty and improve lives. IFC works with businesses, financial institutions, and governments to unlock private investment, foster innovation, and build markets that are resilient and environmentally sustainable. With a strong focus on sectors such as infrastructure, manufacturing, agribusiness, financial services, and climate solutions, IFC helps mobilize capital, create jobs, and expand access to services in regions that need it most.

The IFC's disbursed investment portfolio consists of loans (57.0%), equity investments (29.0%), and debt securities (14.0%), with broad diversification across countries, regions, industries, and project types (International Finance Corporation, 2024b). For the purposes of this analysis, I restrict the sample to IFC clients that received loan-based financing. This focus reflects both empirical and conceptual considerations: loans provide a more direct and interpretable link between climate shocks and credit risk. In contrast, equity positions introduce valuation complexities and strategic interactions that obscure causal identification. Moreover, credit ratings and repayment performance are more consistently defined and observed in the loan portfolio, facilitating a cleaner analysis of climate-related credit dynamics.

2.2.2 Credit ratings

To assess credit risk, this study begins by examining credit ratings. I use firm-level credit rating data from the International Finance Corporation (IFC). While IFC investments are targeted toward projects in developing countries, credit ratings are based on financial performance of the client and the headquarters can be in a different location with respect to the projects. [3]

Each observation in the dataset corresponds to a unique client-year pair, enabling a panel-level analysis of credit ratings over time. The dataset includes annual credit ratings from 1994 to 2023 for 3,463 institutions, totaling 21,595 observations, and is further complemented with monthly data from 2017 onward.

According to International Finance Corporation (2024a), ratings are reviewed at least once annually and reflect a broad spectrum of credit events. The IFC's internal credit rating system assesses clients on a 15-point ordinal scale, where 1 denotes the best credit risk and 15 indicates a default. A counterparty

[3]Notably, approximately 60% of IFC clients are headquartered in developing markets.

is considered to have defaulted if either of the following occurred: (i) it was assigned non-accrual status or (ii) a portion or all the loan balance was written-off.

When considering the number of unique clients that have ever defaulted, I found that 570 out of 3,243 clients (approximately 17%) experienced at least one default event from 1994. This is comparable to IFC's published statistics, which report 784 defaults out of 3,680 clients since 1986 (International Finance Corporation, 2024a). This discrepancy reflects my shorter sample period (1994 onward) and incomplete location data for some clients.

The methodology used by IFC to compute defaults is broadly aligned with the approaches adopted by major credit rating agencies, such as Moody's and Standard & Poor's. This alignment enables meaningful benchmarking of empirical default rates and rating transitions. For details on the mapping between IFC's internal credit ratings and S&P classifications, see International Finance Corporation (2024a).

3 Descriptive statistics

Table 1 presents descriptive statistics for clients grouped by their historical exposure to flooding. The grouping is based on the share of years between 1994 and the end of the sample period in which a client's headquarters fell within a flooded area. I classify clients into three categories: those with no exposure ($p = 0$), those with occasional exposure ($p < 0.24$), and those with repeated exposure ($p > 0.24$), where 0.24 represents the average flood exposure probability across exposed clients (Figure 3.

Table 1: Means of Key Variables by Flood Exposure

	No Exposure ($p = 0$)	**Occasional Exposure** ($p < 0.24$)	**Repeated Exposure** ($p > 0.24$)
# of Clients	1,823	782	702
Observations	8,551	7,225	4,308
Default (0/1)	0.078	0.136	0.103
Credit Rating (1–15)	8.830	9.047	8.788
Displacements	0	200,04	456,883

Interestingly, the default rate is highest among clients with occasional exposure to flooding, while it is slightly lower for those with repeated exposure, suggesting a potential adaptation or resilience effect. Credit ratings are relatively stable across groups, with a slight dip for repeatedly exposed clients. The average number of displacements—used as a proxy for flood severity—increases markedly with exposure probability. This likely reflects the greater chance of encountering a severe event due to more frequent exposure.

While Table 1 provides a long-run view of clients' historical exposure to flooding, it is also important to examine the nature of the flood events themselves in terms of both their intensity and their frequency. This complementary perspective links the probability-based exposure groups to the actual characteristics of the shocks they face. To this end, I draw on event-level data from the Dartmouth Flood Observatory. Among the available indicators, the most consistent and viable measure of *intensity* is the number of people displaced due to flooding events.

The right panel of Table 2 presents summary statistics for this variable. Given the highly skewed nature of displacement (with extreme right-tail values), I use the median displacement value of 8,000 people to split the sample into low- and high-intensity events. This threshold allows for a meaningful comparison in downstream analysis while mitigating the influence of extreme outliers.

In addition to intensity, I examine the frequency of flood exposure per client over the observation window. The left panel of Table 2 reports the number of observations by event frequency and the corresponding number of unique clients exposed at each level. Approximately 41% of the observations pertain to clients who were never affected by any flood, while another 9% were affected only once. A significant share of the dataset—over 50% of the observations—involves clients exposed to multiple flood events. This distribution is central to the research design. The variation in treatment frequency supports a two-pronged analytical strategy. First, I evaluate the causal effect of a single flooding event on outcomes using a difference-in-differences design comparing clients with and without exposure before and after the event. Second, I extend the analysis to account for multiple treatments—a more complex but realistic setting given the recurrence of flooding in many regions.

Frequency of Flooding Events

Events	**Freq.**	**Percent**	**Cum.**	**Clients**
Never treated	12,925	40.83	40.83	1,274
Single event	2,858	9.03	49.86	125
2	2,518	7.95	57.82	199
3	2,170	6.86	64.67	212
4	1,676	5.29	69.97	156
5	2,099	6.63	76.60	194
6	837	2.64	79.24	292
7	1,365	4.31	83.55	144
8	674	2.13	85.68	165
9	1,277	4.03	89.72	27
10	209	0.66	90.38	136
11	53	0.17	90.54	37
12	139	0.44	90.98	54
13	585	1.85	92.83	20
14	319	1.01	93.84	46
15	421	1.33	95.17	20
16	1,488	4.70	99.87	40
17	41	0.13	100.00	57
Total	31,654	100.00		3,248

Intensity of Flooding Events
(Displacements from Flooding)

Statistic	**Value**
Observations	2,810
Mean	397,842
Std. Deviation	2,384,811
Percentiles	
1st	0
5th	0
10th	0
25th	180
50th (Median)	**8,000**
75th	81,600
90th	452,000
95th	1,166,000
99th	9,500,000
Maximum	40,000,000
Minimum	0

Table 2: Summary Statistics: Frequency and Intensity of Flood Exposure

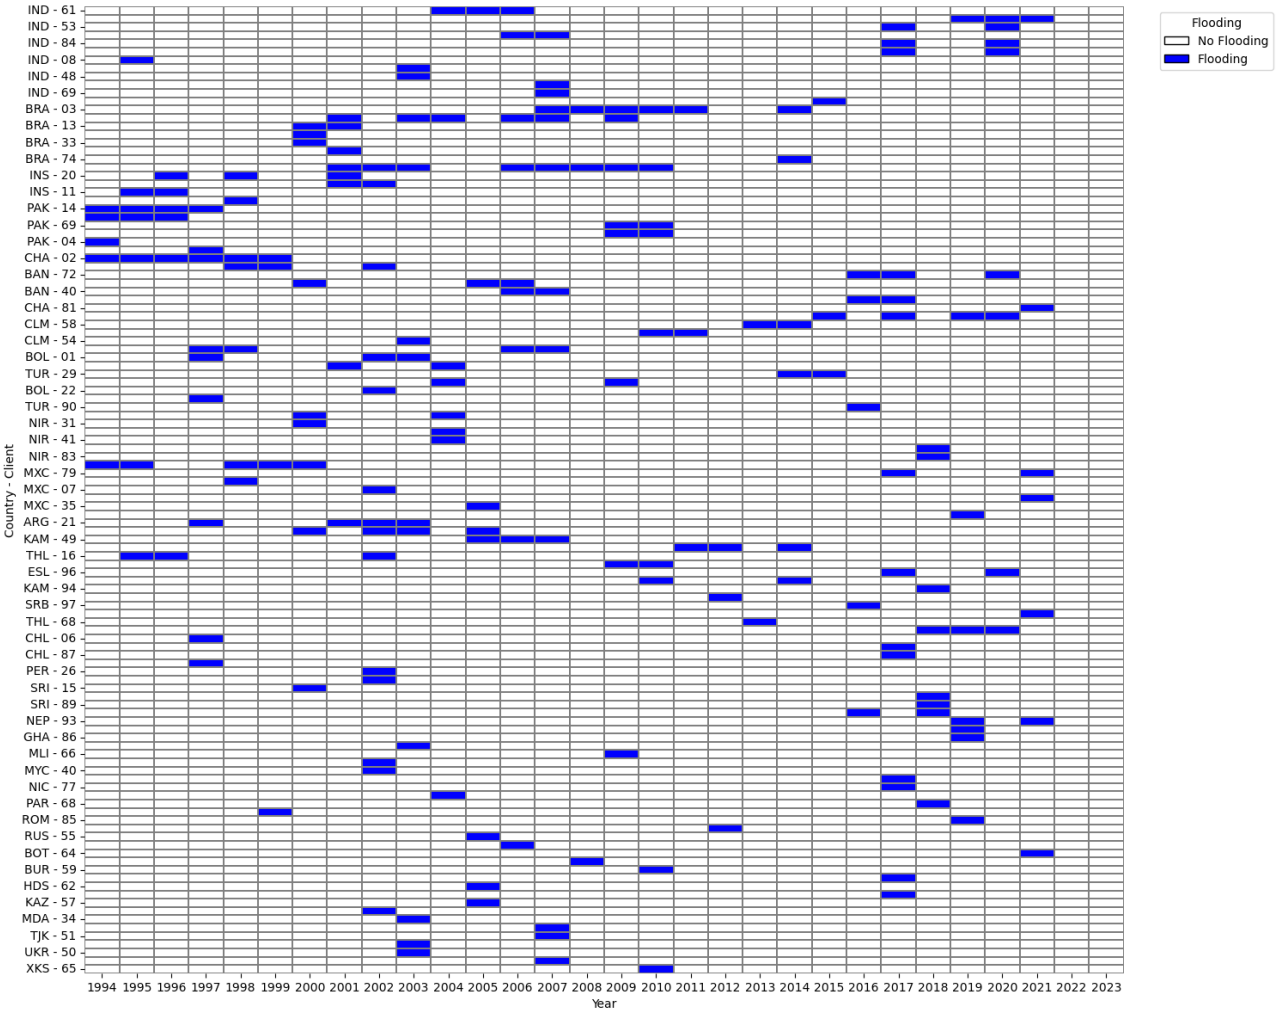
4 Research design

To estimate the impact of flooding on the creditworthiness of private-sector borrowers, I exploit the quasi-random timing and location of flood events across the globe. The dataset comprises client-level data from the International Finance Corporation (IFC), spanning multiple countries and years. This diverse, multinational setting offers a unique opportunity to study climate risk in emerging and developing markets, where credit infrastructure and resilience may vary substantially.

Flood events are matched to clients based on geolocation and year. Conditional on client and year fixed effects, the incidence of flooding in the client's location can be considered as-good-as random, enabling a natural experiment framework for causal inference.

To better understand the impact of flooding on clients, I construct a treatment variation plot. I randomly select 100 clients who experienced at least one flooding event and track their exposure over time. For each client-year, the plot indicates whether the client was affected by flooding. The y-axis is ordered by country, with countries experiencing more frequent flooding events appearing at the top.

Figure 5: Flood exposure across 100 random clients



Note: The heatmap displays flood exposure for 100 random clients over time with at least one flooding event. Countries of clients are ordered by total flood exposure (top = more frequent flooding). Blue indicates exposure in year t; white indicates no exposure.

As shown in Figure 5, some clients remain unaffected by flooding for multiple consecutive years, while others experience repeated exposure over time. This heterogeneity in treatment patterns complicates identification, as it introduces challenges in clearly distinguishing between treated and control groups.

Impact of flooding on Credit Ratings

To estimate the average treatment effect on the treated (ATT) of flooding exposure on municipal credit ratings, I begin with a staggered difference-in-differences (DiD) design, comparing municipalities that have experienced a flooding event with those that are either not yet treated or never treated. I employ a staggered DiD framework following Callaway and Sant'Anna (2021) and de Chaisemartin and D'Haultfœuille (2024).

For a more complex scenario involving multiple treatments, I adopt the Local Projection Difference-in-Differences (LP-DiD) estimator proposed by Jordà et al. (2023):

$$\Delta^h y_{i,t} = \delta_t^h + \beta_h^{LP-DID} \Delta D_{i,t} + e_{i,t}^h \tag{1}$$

In Equation 1, $\Delta^h y_{i,t}$ denotes the change in credit ratings from year $t-1$ to $t+h$, while $\Delta D_{i,t}$ captures the change in flooding exposure. The coefficient β_h^{LP-DID} measures the dynamic impact of a flood occurring at time t on credit ratings h years later.

The LP-DiD framework is particularly well-suited for this context, where treatment exposure is intermittent, and units are treated at varying times—often more than once. It avoids the "contamination effects" that bias traditional fixed-effects models in the presence of overlapping treatments and dynamic responses. To ensure unbiased estimation, I follow the conditions proposed by Jordà et al. (2023):

1. 'clean control condition': Control units must be untreated throughout the interval from $t-K$ to $t+h$.

2. Treated units must not have experienced flooding in the K years prior to treatment and must remain treated for the h years following it.

I set the stabilization window to $K^* = 3$ in the baseline specification, meaning units must be flood-free for at least 3 consecutive years to entering the control group. I confirm robustness by extending this window and find that the results remain stable, suggesting the estimates are not sensitive to the choice of K^* and the control group is sufficiently clean.

One challenge with the second condition is the possibility of overlapping flood events within the h years following treatment. To address this, I isolate a subset of treated units that experienced only a single flooding event (model equivalent toto Callaway and Sant'Anna (2021) and de Chaisemartin and D'Haultfœuille (2024) . I then conduct a sensitivity analysis by varying the restriction on K for treated units as well—requiring a minimum number of flood-free years before a unit can be treated again ('clean treatment condition'). This approach allows me to estimate the average treatment effect of a flooding event on credit ratings more accurately, eliminating potential confounding from repeated or overlapping events.

Because the likelihood of flooding is not uniform across countries, in all the analyzes, I apply Inverse Probability Weighting (IPW) to correct for differential treatment propensities. For each client, I compute the historical flood probability p based on the share of flood years. Weights are then assigned as follows:

- $1/p$ for treated clients (those exposed to floods in year t),

- $1/(1-p)$ for untreated clients (in years without flood exposure).

Lastly, placebo tests examining pre-flood trends support the no-anticipation assumption, confirming that treated and control groups were on similar trajectories before exposure.

Impact of flooding on probability of default

To identify the causal effect of flooding exposure on the likelihood that an IFC client defaults, I estimate logistic regression models with and without fixed effects. Since the outcome variable is binary—equal to 1 if an IFC client is in default in a given year and 0 otherwise—a logistic model is appropriate for this setting.

The baseline specification is a pooled logit model of the form:

$$\Pr(\text{Default}_{i,t} = 1) = \Lambda\left(\beta_1 \cdot \text{Flood}_{i,t} + \mathbf{X}'_{i,t}\boldsymbol{\gamma}\right) \tag{2}$$

where $\Lambda(z) = \frac{e^z}{1+e^z}$ is the logistic function, $\text{Default}_{i,t}$ is a binary indicator for default of client i in year t, $\text{Flood}_{i,t}$ is a binary indicator equal to 1 if the client was exposed to a flood event in year t, and $\mathbf{X}_{i,t}$ is a vector of time-varying controls.

To account for unobserved heterogeneity across IFC clients or institutions—such as differences in risk profiles, sectors, or geographic exposures—I estimate a fixed-effects logistic regression model:

$$\Pr(\text{Default}i, t = 1) = \Lambda\left(\alpha_i + \alpha_t + \beta_1 \cdot \text{Flood}i, t + \mathbf{X}'_{i,t}\boldsymbol{\gamma}\right) \tag{3}$$

where α_i represents time-invariant fixed effects at the client level and α_t represents year fixed effects. The identification assumption in both specifications is that, conditional on observables (and fixed effects, where included), flood exposure is exogenous to other unobserved shocks that may influence default. The fixed-effects specification strengthens this assumption by controlling for all time-invariant unobserved heterogeneity. However, it comes with the trade-off that clients with no variation in the outcome (i.e., always in default or never in default) are excluded from the estimation sample.

A central limitation of this default analysis is that it does not account for lagged treatment effects, which play an important role in capturing the persistence of flood impacts over time, as shown by the credit rating variation. To date, there is no established event-study framework that employs logistic estimators, which constrains the ability to fully capture dynamic adjustment paths following flood events for the default analysis.

5 Results

Impact of Flooding on Credit Rating (CR)

Here, I assess the dynamic impact of flooding events on firms' creditworthiness. Following Local Projections Difference-in-Differences (LP-DiD) presented in equation 1, I estimate how credit ratings evolve relative to the year prior to a flood event. .

To estimate the average treatment effect (ATE) of a single severe flooding event, I restrict the analysis to clients who experienced exactly one flood over the entire 30-year panel. This allows for a clean comparison between treated and untreated firms, free from confounding effects related to repeated or overlapping events. While this is the most restrictive sample, it offers the clearest identification strategy by comparing firms with one isolated event to firms that were never treated. To explore heterogeneity in treatment intensity, I further divide the treatment group based on the median value of displacement (8,000 people), as shown in Table 2. This results in two subgroups: clients experiencing a single flood that displaced more than 8,000 people, and those exposed to a single flood with fewer than 8,000 displacements.

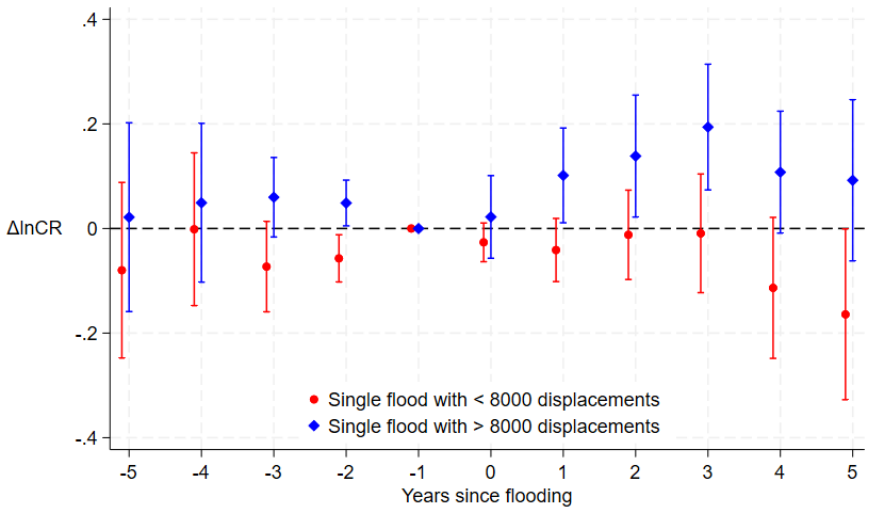


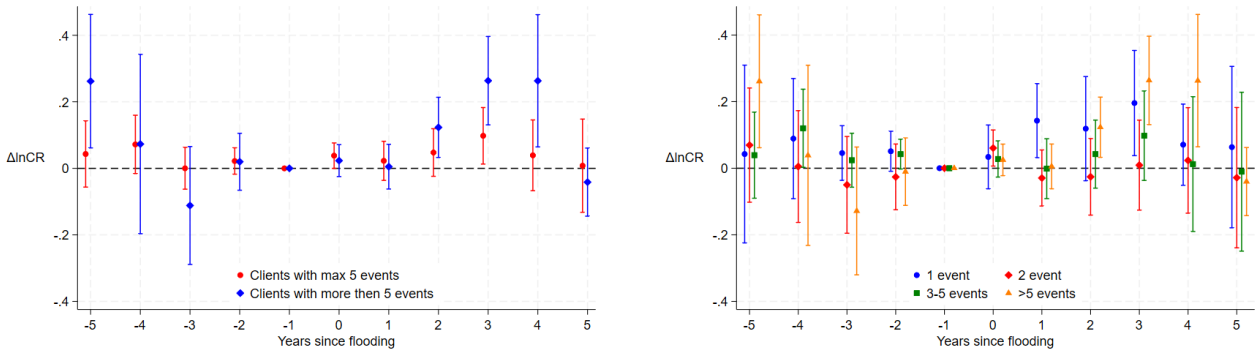
Figure 6: ATE of flooding on credit rating (single-event sample)

Figure 6 shows that credit ratings deteriorate significantly following a severe flood (i.e., more than 8,000 displacements), with an average increase of 20% in credit rating values observed three years after the event. In contrast, no significant effect is found for floods below the median displacement threshold. These findings are robust to alternative difference-in-differences estimators such as (de Chaisemartin and D'Haultfœuille, 2024; Callaway and Sant'Anna, 2021)(Appendix B).

While this restrictive sample improves causal identification, it limits statistical power due to the small number of treated units, as reflected in the wider confidence intervals. To address the trade-off between bias and efficiency, I extend the analysis to include clients with multiple flood exposures. Appendix B

presents a sensitivity analysis using varying exclusion windows to define "clean" treatment episodes. Specifically, I consider values of K ranging from 1 to 5, where K indicates the number of years before and after the focal event during which no other flood is allowed. For example, $K = 3$ requires that a client not experience another flood three years before or after the event. Results under this relaxed definition closely resemble the baseline estimates from the single-event sample, supporting the robustness of the main findings. This broader sample allows me to incorporate clients who were exposed to multiple floods and examine how the number of events affects outcomes.

Figure 7: Heterogeneous effects by number of flood events



(a) Split by median number of events, 5

(b) Further breakdown by number of events

Figure 7 presents heterogeneous effects based on the number of flood events experienced. In panel (a), I split the sample at the median number of flood events (five). Clients exposed to more than five events over the 30-year panel suffer larger negative impacts on credit ratings compared to those with fewer than five events.

Panel (b) further decomposes the sample by number of events. Interestingly, clients exposed to two or three floods experience smaller impacts than those exposed to only a single event, suggesting some degree of adaptation. However, the protective effect appears to diminish with more repeated exposure. Clients with more than five events exhibit similar outcomes compared to those with only a single exposure. This suggests that while some resilience may emerge with moderate exposure, it eventually breaks down under repeated stress.

In addition to the overall treatment effects, I conduct a sector-level analysis to examine whether certain industries are more vulnerable to flood-related credit deterioration.

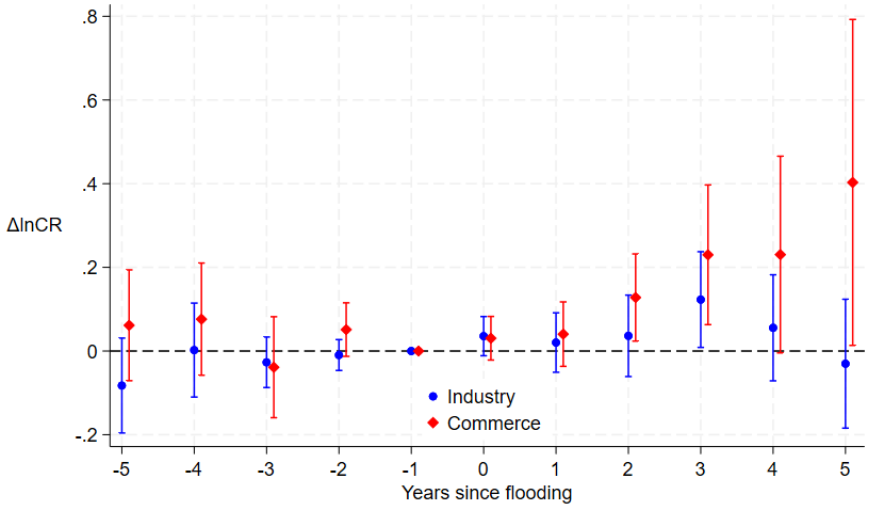


Figure 8: Sector analysis

Figure 8 illustrates that the commerce sector (e.g., retail and banking services) experiences significantly greater declines in credit ratings following flooding events compared to the industrial sector. Moreover, the negative effects in commerce persist for a longer duration—up to five years after the event—whereas impacts on industrial clients tend to be more transitory. This differential vulnerability may be attributed to several factors. Commercial firms are more likely to rely on localized infrastructure (e.g., storefronts, customer foot traffic, and physical banking branches), which can be severely disrupted by floods. In contrast, industrial firms may be better equipped to absorb or mitigate flood shocks due to stronger asset bases, diversified supply chains, or the ability to temporarily suspend operations without losing customer relationships. Additionally, firms in commerce may have tighter cash flows and greater dependency on short-term revenue, making them more financially fragile when physical disruptions occur. These findings suggest that credit risk assessments in the face of climate shocks should not only consider the intensity and frequency of exposure, but also sectoral heterogeneity in vulnerability and recovery capacity.

Impact of Flooding on the Probability of Default (PD)

This section examines whether exposure to flooding significantly affects a firm's probability of default. Assessing this outcome provides a more direct measure of financial distress than credit ratings alone. I estimate binary response models using both standard logit and fixed effects (FE) logit specifications, where the dependent variable is a binary indicator equal to one if the firm is in default in a given year. The primary independent variable is an indicator for whether the firm's location was affected by a flood event in that year.

To ensure robustness, I rely on two distinct flood datasets: the Dartmouth Flood Observatory, which documents major riverine flood events globally based on media and governmental sources, and the Global Flood Database (GFD), which uses satellite imagery to detect flood extents at high spatial

resolution. Table 3 reports the results using both sources.

Table 3: Impact of Flood Exposure on Probability of Default (PD)

	Dartmouth Flood Observatory		**Global Flood Database**	
	Logit	FE Logit	Logit	FE Logit
Log-odds Coefficient	0.121	0.158	0.226	0.395
Standard Error	0.060	0.095	0.055	0.096
p-value	0.045	0.099	0.000	0.000
Marginal Effect (dy/dx)	1.14%	3.94%	2.12%	9.79%
Observations	21,606	4,940	21,606	4,940
Number of Groups	–	570	–	570

The findings indicate that flooding significantly increases the likelihood of default in the years of events. According to the logit estimates using the Dartmouth dataset, firms in flood-affected areas are 1.14 percentage points more likely to default, with the effect rising to 3.94 percentage points in the fixed-effects specification. However, the stronger effects are observed when using the Global Flood Database: the probability of default increases by 2.12 percentage points in the standard logit model and nearly 9.8 percentage points when accounting for institution fixed effects.

Notably, these results are obtained in a context where approximately 10% of observations in the full sample are in default, underscoring the economic relevance of the estimated effects. In the fixed-effects models, roughly 2,893 clients (16,655 observations) are omitted due to a lack of variation in the dependent variable.

Taken together, these findings reinforce the view that flooding significantly compromises firms' financial health. The stronger effects observed with the Global Flood Database also underscore the importance of high-resolution, remotely sensed data in disaster impact assessments. These results align with and complement the credit rating analysis, providing consistent evidence that flood exposure leads to a meaningful deterioration in firms' creditworthiness and repayment capacity.

6 Conclusion

In response to increasing climate risks, investors and regulators are relying on Environmental, Social, and Governance (ESG) indicators to inform financial decision-making (Krueger et al., 2020). ESG metrics are now used not only to assess sustainability practices of a company but also to assess creditworthiness in the face of climate-related risk. Major credit rating agencies (CRAs) such as S&P Global, Moody's, and Fitch have begun to systematically integrate climate-related risks into their credit rating models. Companies that lack robust climate risk management or transparent disclosure practices are increasingly exposed to the risk of downgrades, higher financing costs, and declining market confidence. Despite increasing attention to climate physical risk, there is still a significant empirical gap in understanding how specific hazard, such as flooding, translate into firm-level credit risk.

This paper provides new empirical evidence on how extreme weather events—specifically floods—affect firm-level credit risk, with a focus on emerging markets. By linking high-resolution geospatial data on flood exposure from the Dartmouth Flood Observatory to proprietary credit rating data from the International Finance Corporation (IFC), the analysis offers one of the first direct assessments of how physical climate shocks are transmitted into credit risk as measured by firm-level ratings.

The findings highlight two central dimensions of physical climate risk. First, regarding intensity, only severe floods—defined as events displacing more than 8,000 people—lead to a significant deterioration in credit ratings, equivalent to roughly a 20

These results are robust across alternative difference-in-differences estimators (including Callaway–Sant'Anna and de Chaisemartin–D'Haultfœuille) and multiple exposure definitions, confirming that the link between flood exposure and credit deterioration is not sensitive to model specification. The persistence of results even without detailed asset-level data suggests that physical climate risks are already embedded—albeit implicitly—into the credit assessment process. Future research integrating asset-level exposure and damage estimates could further refine the accuracy of climate-adjusted credit risk models.

The sectoral heterogeneity uncovered in the analysis adds a critical nuance. Firms in commerce and banking show stronger and more persistent deterioration in credit ratings than industrial firms, likely reflecting their higher leverage, limited collateral, and greater concentration in urban, flood-prone areas. These findings highlight that credit risk linked to physical climate exposure is uneven across sectors and that incorporating industry-specific sensitivities is vital for effective risk management.

From a broader perspective, the results emphasize that physical climate shocks can materially affect financial stability, even in the absence of explicit climate disclosures. For financial institutions and investors—especially multilateral development banks like the IFC—this evidence supports the integration of physical risk metrics into credit evaluation, pricing, and portfolio management frameworks. As investors increasingly recognize and price climate risks into financial assets (Huynh and Xia, 2021; Engle et al., 2020; Faccini et al., 2023), understanding how these risks translate into creditworthiness becomes essential for maintaining resilient capital flows in vulnerable economies.

Overall, this study contributes to the growing literature on the financial implications of climate change by providing causal evidence that floods—one of the most frequent and costly climate hazards—deteriorate firms' credit quality. The results underline the importance of proactive physical risk assessment for both public and private credit markets, particularly in emerging economies where adaptation capacity remains limited. Integrating granular climate risk information into credit modeling is thus not only a matter of financial prudence but a necessary step toward sustainable and climate-resilient development.

7 Acknowledgements

I am grateful to Federico Galizia and Rapti A. Goonesekere for establishing a new fellowship program at the International Finance Corporation (IFC), which gave me the opportunity to join the organization and develop this project. I am also especially thankful to Aurèle M. Houngbédji, Sylvia Leonor Garcia De Villafuerte and Couro Kane Janus for their support throughout the development of this work. I also thank my colleagues at the Corporate Risk Management department for their invaluable support in the data collection process and for the collaborative environment that made this research possible.

References

Jawad M. Addoum, David T. Ng, and Ariel Ortiz-Bobea. Temperature shocks and industry earnings news. *Journal of Financial Economics*, 150(1):1–45, 2023. ISSN 0304-405X. doi: https://doi.org/10.1016/j.jfineco.2023.07.002.

Sandra Aguilar-Gomez, Emilio Gutierrez, David Heres, David Jaume, and Martin Tobal. Thermal stress and financial distress: Extreme temperatures and firms' loan defaults in mexico. *Journal of Development Economics*, 168:103246, 2024. ISSN 0304-3878. doi: https://doi.org/10.1016/j.jdeveco.2023.103246.

Edward I. Altman. Financial ratios, discriminant analysis and the prediction of corporate bankruptcy. *The Journal of Finance*, 23(4):589–609, 1968. doi: 10.2307/2978933.

Asaf Bernstein, Matthew T. Gustafson, and Ryan Lewis. Disaster on the horizon: The price effect of sea level rise. *Journal of Financial Economics*, 134(2):253–272, 2019. ISSN 0304-405X. doi: https://doi.org/10.1016/j.jfineco.2019.03.013.

Marshall Burke, Solomon M. Hsiang, and Edward Miguel. Global non-linear effect of temperature on economic production. *Nature*, 527:235–239, 2015. doi: 10.1038/nature15725.

Brantly Callaway and Pedro H.C. Sant'Anna. Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230, 2021. ISSN 0304-4076. doi: https://doi.org/10.1016/j.jeconom.2020.12.001. Themed Issue: Treatment Effect 1.

Tamma A. Carleton and Solomon M. Hsiang. Social and economic impacts of climate. *Science*, 353 (6304):aad9837, 2016. doi: 10.1126/science.aad9837.

Xiaoguang Chen and Lu Yang. Temperature and industrial output: Firm-level evidence from china. *Journal of Environmental Economics and Management*, 95:257–274, 2019. ISSN 0095-0696. doi: https://doi.org/10.1016/j.jeem.2017.07.009. URL `https://www.sciencedirect.com/science/article/pii/S009506961730520X`.

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.

Melissa Dell, Benjamin F. Jones, and Benjamin A. Olken. What do we learn from the weather? the new climate-economy literature. *Journal of Economic Literature*, 52(3):740–98, September 2014. doi: 10.1257/jel.52.3.740. URL `https://www.aeaweb.org/articles?id=10.1257/jel.52.3.740`.

Darrell Duffie and David Lando. Term structures of credit spreads with incomplete accounting information. *Econometrica*, 69(3):633–664, 2001. doi: 10.1111/1468-0262.00214.

Robert J.R. Elliott, Yi Liu, Eric Strobl, and Meng Tong. Estimating the direct and indirect impact of typhoons on plant performance: Evidence from chinese manufacturers. *Journal of Environmental Economics and Management*, 98:102252, 2019. ISSN 0095-0696. doi: https://doi.org/10.1016/j.jeem.2019.102252. URL https://www.sciencedirect.com/science/article/pii/S009506961830651X.

Robert F Engle, Stefano Giglio, Bryan Kelly, Heebum Lee, and Johannes Stroebel. Hedging climate change news. *The Review of Financial Studies*, 33(3):1184–1216, 02 2020.

Renato Faccini, Rastin Matin, and George Skiadopoulos. Dissecting climate risks: Are they reflected in stock prices? *Journal of Banking Finance*, 155:106948, 2023. ISSN 0378-4266. doi: https://doi.org/10.1016/j.jbankfin.2023.106948. URL https://www.sciencedirect.com/science/article/pii/S037842662300153X.

Stefano Giglio, Bryan Kelly, and Johannes Stroebel. Climate finance. *Annual Review of Financial Economics*, 13:15–36, 2021a. doi: https://doi.org/10.1146/annurev-financial-102620-103311.

Stefano Giglio, Matteo Maggiori, Krishna Rao, Johannes Stroebel, and Andreas Weber. Climate change and long-run discount rates: Evidence from real estate. *The Review of Financial Studies*, 34 (8):3527–3571, 03 2021b.

Paul Goldsmith-Pinkham, Matthew T Gustafson, Ryan C Lewis, and Michael Schwert. Sea-level rise exposure and municipal bond yields. *The Review of Financial Studies*, 36(11):4588–4635, 05 2023.

Feng He, Xingzi Ren, Yueren Wang, and Xingfan Lei. Climate risk and corporate bond credit spreads. *Journal of International Money and Finance*, 154:103297, 2025. ISSN 0261-5606. doi: https://doi.org/10.1016/j.jimonfin.2025.103297. URL https://www.sciencedirect.com/science/article/pii/S0261560625000324.

Henry He Huang, Joseph Kerstein, Chong Wang, and Feng (Harry) Wu. Firm climate risk, risk management, and bank loan financing. *Strategic Management Journal*, 43(13):2849–2880, 2022. doi: https://doi.org/10.1002/smj.3437.

Thanh D. Huynh and Ying Xia. Climate change news risk and corporate bond returns. *Journal of Financial and Quantitative Analysis*, 56(6):1985–2009, 2021. doi: 10.1017/S0022109020000757.

International Finance Corporation. Ifc portfolio default rate analysis. Technical report, International Finance Corporation, 2024a.

International Finance Corporation. Ifc annual information statement. Technical report, International Finance Corporation, 2024b. URL https://www.ifc.org/en/about/investor-relations/financial-statements.

2021 IPCC. Chapter 11: Weather and climate extreme events in a changing climate. In Valérie Masson-Delmotte, Panmao Zhai, Anna Pirani, Sarah L. Connors, Clotilde Péan, Sophie Berger, Nathalie Caud, Yele Chen, Leah Goldfarb, Melissa I. Gomis, Maike Huang, Katherine Leitzell, Elise Lonnoy, John B. R. Matthews, Thomas K. Maycock, Tim Waterfield, Ozge Yelekçi, Rong Yu, and Biao Zhou, editors, *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 1513–1766. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, 2021. doi: 10.1017/9781009157896.013. URL https://www.ipcc.ch/report/ar6/wg1/chapter/chapter-11/.

Òscar Jordà, Arindrajit Dube, Daniele Girardi, and Alan M. Taylor. A local projections approach to difference-in-differences event studies. CEPR Discussion Paper DP18141, Centre for Economic Policy Research, May 2023. URL https://ideas.repec.org/p/cpr/ceprdp/18141.html.

Erhan Kilincarslan, Zezeng Li, and Jiafan Li. Firm-level climate change exposure and credit ratings. *Economics Letters*, 255:112512, 2025. ISSN 0165-1765. doi: https://doi.org/10.1016/j.econlet.2025.112512. URL https://www.sciencedirect.com/science/article/pii/S0165176525003490.

Philipp Krueger, Zacharias Sautner, and Laura T Starks. The importance of climate risks for institutional investors. *The Review of Financial Studies*, 33(3):1067–1111, 02 2020.

Shaojie Lai, Lihan Chen, Qing Sophie Wang, and Hamish Anderson. Natural disasters, trade credit, and firm performance. *Economic Modelling*, 116:106029, 2022. ISSN 0264-9993. doi: https://doi.org/10.1016/j.econmod.2022.106029. URL https://www.sciencedirect.com/science/article/pii/S026499932200267X.

Robert C. Merton. On the pricing of corporate debt: The risk structure of interest rates. *The Journal of Finance*, 29(2):449–470, 1974. doi: 10.2307/2978814.

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.

Harvey Nguyen, Anh Viet Pham, Man Duy (Marty) Pham, and Mia Hang Pham. Climate change and corporate credit worthiness: International evidence. *Global Finance Journal*, 64:101073, 2025. ISSN 1044-0283. doi: https://doi.org/10.1016/j.gfj.2024.101073. URL https://www.sciencedirect.com/science/article/pii/S1044028324001455.

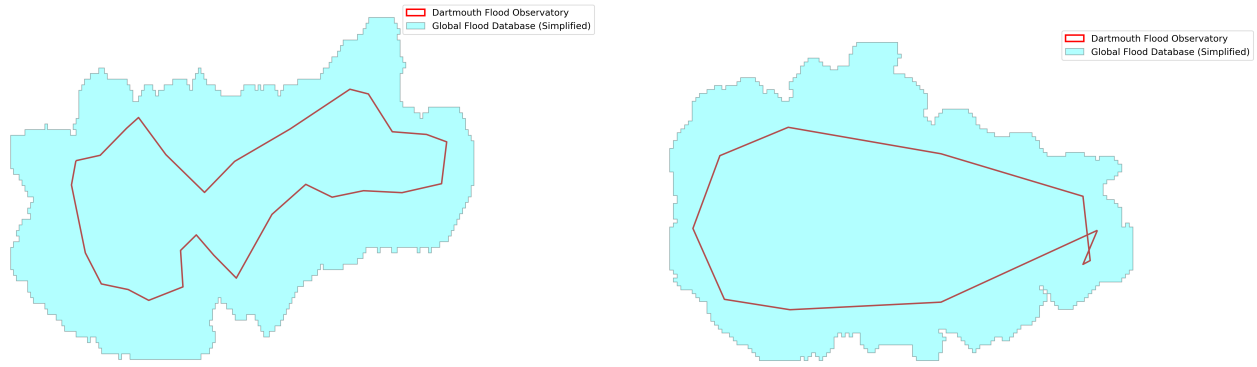
Marcus Painter. An inconvenient cost: The effects of climate change on municipal bonds. *Journal of Financial Economics*, 135(2):468–482, 2020. ISSN 0304-405X. doi: https://doi.org/10.1016/j.jfineco.2019.06.006. URL https://www.sciencedirect.com/science/article/pii/S0304405X19301631.

Jun Rentschler, May Salhab, and Bramka Arga Jafino. Flood exposure and poverty in 188 countries. *Nature Communications*, 13:3527, 2022. doi: 10.1038/s41467-022-30727-4. URL https://doi.org/10.1038/s41467-022-30727-4.

Alessio Venturini. Climate change, risk factors and stock returns: A review of the literature. *International Review of Financial Analysis*, 79:101934, 2022. ISSN 1057-5219. doi: https://doi.org/10.1016/j.irfa.2021.101934. URL https://www.sciencedirect.com/science/article/pii/S1057521921002568.

A Flooding maps

Figure A.1: Examples of comparisons between GFD and DFO

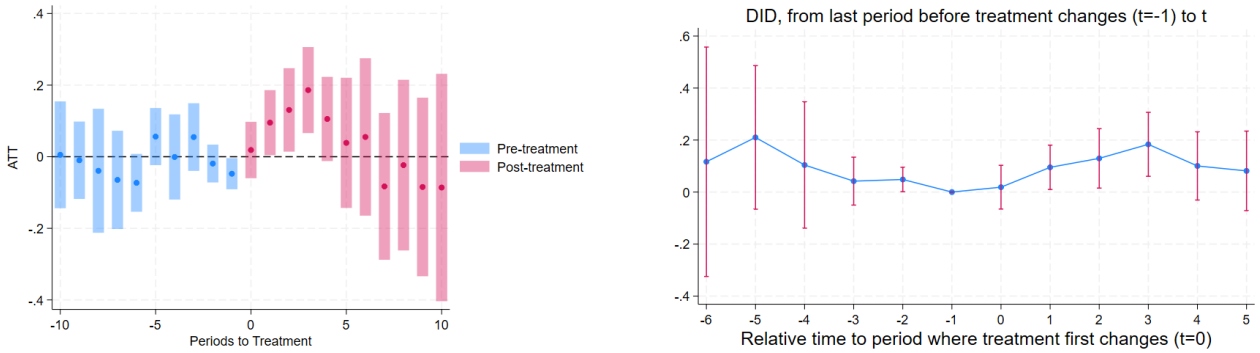


(a) Example 1 (b) Example 2

B Robustness of the main result

Results from other difference in difference literature

Figure B.1: Results from csdid and multipleg



(a) Example 1 (b) Example 2

Trade-off between bias and efficiency

To estimate the effect of flooding on credit ratings, I rely on different assumptions about the "cleanliness" of the treatment group—specifically, how isolated a flooding event must be from others. I define a parameter K to indicate how many years before and after the treatment window a client must remain flood-free to be included. For example, $K = 1$ excludes treated units that experience another flood in the year immediately following the initial event. At the other extreme, I consider only clients hit by a single flood over the entire 30-year panel (as shown in Figure 6).

Figure B.2: Comparing different clean treatment conditions

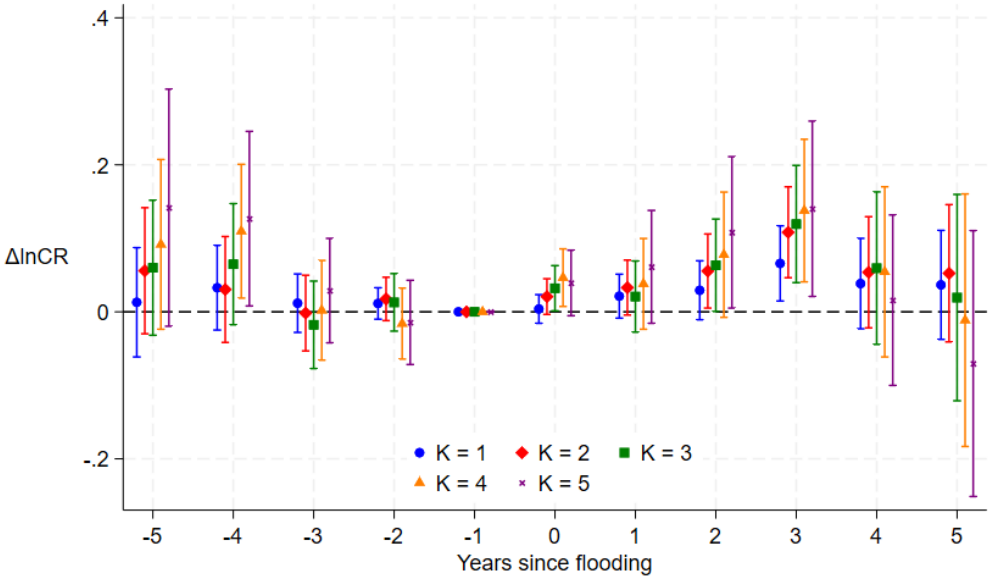


Figure B.2 illustrates how treatment effect estimates vary across different values of K (from 1 to 5). While the point estimates are broadly consistent, increasing K reduces the sample size and increases the variance of the estimates, reflecting a classic trade-off between bias and efficiency.

When $K = 1$, the treatment group is larger, resulting in tighter confidence intervals and better balance in pre-trends. However, this specification includes clients affected by overlapping flood events, which may attenuate the estimated effect of a single flood. In contrast, higher values of K (e.g., $K = 5$) produce cleaner comparisons by removing clients exposed to repeated events, thereby isolating the impact of a standalone flood. These cleaner estimates show a more immediate increase in credit risk following the event—an effect that is less visible in more inclusive samples. Because all the specification point out a significant effect after 3 year,and then insignificant reuslts, I have used K=3 in the main specification (Figure 6)

Chapter 3

The Impact of Wind Turbines Visibility on the Local Economy

Abstract

Tourism in Italy generated an estimated €215 billion in 2023, accounting for 10.5% of national GDP. This study assesses how wind turbine visibility affects tourism across Italian municipalities using a novel, high-resolution measure that combines turbine data with satellite-based elevation models. We find that visibility of wind turbines reduces tourist arrivals by up to 20%, with the effect peaking seven years after installation and gradually fading, suggesting local adaptation. The effect is driven by municipalities where more than 80% of the territory is exposed, particularly in rural, tourism-dependent areas. We uncover a non-linear relationship between the number of visible turbines and the tourism drop, indicating that large wind farms (>65 wind turbines) may be less disruptive than smaller, scattered installations. Tourist spending mirrors these patterns, declining by about 15% (€1.8 million annually), or 1.6% of local GDP. Yet these local costs are outweighed by environmental benefits: 50 turbines generate roughly €5 million annually in CO_2 savings, over three times the tourism-related losses. Our findings call for policies that mitigate local externalities while preserving broader environmental gains, including the strategic siting of large wind farms and attention to environmental justice, as rural communities and small touristic operators bear disproportionate costs.

JEL-CODES: Q42, Q51, Q56, Z32

Keywords: Environmental externalities, environmental justice, wind energy, tourism

Davide Castellani
Amsterdam Business School
University of Amsterdam
d.castellani@uva.nl

Gabriele Pinto
Department of Social and Economic Sciences
Sapienza University of Rome
gabriele.pinto@uniroma1.it

1. Introduction

In late summer 2024, a series of arson attacks targeted wind turbine construction sites across Sardinia, a renowned Mediterranean tourist destination (Ansa, 2024). In the village of Orgosolo, located in the island's interior and historically known for its defiance during the "Pratobello Uprising"[1], residents are once again mobilizing in protest, this time against a new wind power project (Locci, 2024). Orgosolo, celebrated for its striking murals and rugged landscapes, has become a symbol of local resistance, with opponents arguing that wind turbines would undermine the area's scenic beauty and reduce its appeal to tourists, a vital source of community income. These demonstrations form part of a broader backlash against the central government's plans to install over 6.2 GW of new wind capacity on the island. Such "Not In My Backyard" (NIMBY) protests against wind energy projects are not unique to Italy; similar movements have emerged worldwide (Boyle et al., 2019; Daniele et al., 2024; Devlin, 2005; Stokes et al., 2023; Töller et al., 2024; Wolsink, 2000). Even high-profile figures have voiced opposition: U.S. President Donald Trump, for instance, famously criticized the visibility of wind turbines from his Turnberry golf course (The Guardian, July 2025).

However, the construction of new wind turbines remains a crucial step toward achieving ambitious climate objectives, including the European Union's commitment to carbon neutrality by 2050 (Wind Europe, 2022). While some local environmental drawbacks are acknowledged, the overall global impact of wind energy is overwhelmingly positive (Meng et al., 2024). Proponents further argue that wind power can stimulate job creation and foster economic growth, although existing evidence suggests that such local economic effects are generally short-lived (Fabra et al., 2024; Costa & Veiga, 2021). Another common argument in favor of wind energy is that the spatial footprint of turbines is relatively limited[2]. Yet Italy's experience offers a compelling counterexample. The country's roughly 7,000 wind turbines occupy just 0.002% of the national territory, concentrated in only 311 municipalities out of nearly 8,000. Nevertheless, as we show in Section 2.1, these turbines are visible

[1] In June 1969, the residents of Orgosolo, a village in Sardinia, organized a peaceful protest against the Italian military's plan to establish a firing range in Pratobello, a common grazing area for local shepherds. Approximately 3,500 people occupied the land, leading to the military's withdrawal and highlighting the community's commitment to nonviolent resistance and territorial autonomy.

[2] For example, with reference to the Sardinian disputes against wind turbines, the local representative of *Legambiente* (the most important environmental association in Italy) claimed that: "*even if we did 10GW, just 1% of the Sardinian territory would be enough. Then of course, the wind turbines will be seen, but this is inevitable for any infrastructure"* (Tecleme, 2024).

from more than 1,950 municipalities, affecting around 25% of Italy's landscape[3]. These figures highlight that the visual footprint of existing and planned wind farms is far from negligible.

In this work we aim to study how the visibility of onshore wind turbines[4] can impact local economies, with a specific focus on tourism, a sector that highly depend on the aesthetic and attractiveness of a pristine landscape. In fact, the visual and acoustic presence of wind turbines may influence tourists' choices (Krekel & Zerrahn, 2017; Schmidt & Klokker, 2014). Our analysis centers on Italy, the fourth most popular touristic destination in the world, where Travel & Tourism sector contributed to €215B in 2023, representing 10.5% of GDP (WTTC, May 2024). The relevance of tourism is not confined to Italy; in the United States for example, the Travel & Tourism market contribute a record-breaking $2.36 TN to the nation's economy in 2023 (WTTC, September 2024).

We assess the impact of wind turbine visibility on tourism using a staggered difference-in-differences approach, applied to a 20-year dataset (1998–2019) on municipal-level tourism. The data are sourced from the Bank of Italy and the Italian National Institute of Statistics (ISTAT). We demonstrate that our estimates are robust to heterogeneous treatment effects by comparing the results from the classic difference-in-differences event study with those obtained using more recent estimators (Callaway and Sant'Anna, 2021; de Chaisemartin & D'Haultfœuille, 2024). To measure exposure, we combine satellite-based digital elevation data with the exact location of turbines installed allowing us to identify areas from which turbines are visible and to construct municipality-level indicators of visual exposure.[5] Wind turbines are not randomly distributed, which can introduce selection biases and challenge the identification of robust and externally valid causal effects. To mitigate this, we restrict our sample to municipalities meeting climate and regulatory criteria, ensuring treatment and control units have a comparable likelihood of hosting wind turbines.

Our results show that, over a ten-year horizon, wind turbine visibility reduces tourist arrivals by 20%, with the effect peaking seven years after turbines first become visible and gradually diminishing thereafter, suggesting local adaptation. The effect is driven primarily by municipalities where between 9 and 65 turbines are visible. We observe a non-linear relationship with the number of visible wind turbines, likely because larger wind farms are typically located in areas with limited tourist

[3] These figures represent municipalities with at least a small part (i.e., >0 per cent) of their territory from which a wind turbine is visible. If we set a minimum threshold of 10% of visibility, the total number of municipalities decreases to 1,530. For more details, refer to Section 2.1.5.

[4] We focus exclusively on onshore wind turbines, as Italy's first offshore wind farm was only commissioned in Taranto in 2022

[5] The approach of using the digital elevation model and the viewshed is similar to what have been done in Guo et al. (2024) and Stegmaier & Krause (2023) to measure respectively the effect of wind turbine visibility on housing prices in US and electoral support for green parties in Germany.

presence. The impact is strongest in municipalities where more than 80% of the territory is exposed, particularly in rural, tourism-dependent areas. Tourist expenditure mirrors this pattern: affected municipalities experience an average annual decline of 15%, approximately €1.8 million per year, or about 1% of local GDP. This contraction extends to broader economic activity, with local income falling by around 1%. On the supply side, B&Bs and private accommodations experience negative effects from wind turbine visibility, while hotels show no decline, reflecting their greater capacity to absorb fluctuations in demand.

On the other hand, wind turbines have been a cornerstone of the green energy transition in Italy. In 2023 alone, they generated a record 23.6 TWh of electricity, accounting for 20% of renewable energy production and 8.8% of total electricity generation (GSE, January 2025). Substituting coal-based electricity with wind energy results in an average reduction of roughly 600 gCO_2 per kWh (Pryor et al., 2020). Given the market price of CO_2 under the EU Emissions Trading System, these reductions can be translated into substantial monetary benefits. We find that 50 wind turbines generate approximately €5 million per year in value from CO_2 savings. In our sample, municipalities are, on average, exposed to 50 turbines, meaning that the environmental benefits are more than three times greater than the estimated annual loss of €1.8 million in tourism expenditure for a municipality facing similar exposure. Given these substantial gains, we discuss policy options aimed at mitigating local negative externalities while still supporting the expansion of renewable energy, and crucially, the strategic selection of locations and quantity of wind turbines. For instance, we find a non-linear relationship between the number of visible turbines, suggesting that constructing large wind farms at once (more than 65 turbines) may be preferable to dispersing smaller numbers of turbines across multiple sites.

This paper contributes to the literature in several ways. First, it provides causal empirical estimates of the impact of wind turbines visibility on tourism activity. Our findings contribute to a relatively extensive body of research examining the relationship between wind turbines and tourists' perceptions (Fooks et al., 2017; Frantál & Kunc, 2011; Landry et al., 2012; Maehr et al., 2015; Parsons & Firestone, 2018; Trandafir et al., 2020; Tverijonaite & Sæþórsdóttir, 2023). This literature generally suggests that tourists tend to prefer unspoiled landscapes and often favor destinations where wind turbines are not visible[6]. Evidence from regions such as Iceland where tourism relies heavily on natural, untouched environments, indicates that tourism providers overwhelmingly perceive wind farms as detrimental, primarily because they "ruin the view" and transform natural areas into

[6] However, there is significant variation depending on individual characteristics, such as being or not a resident, or the individual attitudes toward green issues (Broekel & Alfken, 2015).

"industrial zones" (Sæþórsdóttir et al., 2021). However, most existing studies rely on surveys, field experiments, or qualitative assessments of visitor preferences, which may be influenced by social desirability biases and may not fully capture actual tourist behavior[7]. Complementing this literature, research on tourism in Italy underscores the economic importance of heritage sites and pristine rural landscapes in enhancing tourist appeal and supporting local development (Bertacchini et al., 2024; Nocito et al., 2023). In this context, our study provides robust empirical evidence that municipality characteristics matter, with tourism-oriented and rural municipalities being most affected by visibility of wind turbines.

Second, this paper contributes to the literature documenting a negative relationship between wind turbines visibility and housing values. Previous studies have consistently found that the presence of wind turbines visibility can reduce property values (Dröes & Koster, 2016; Sunak & Madlener, 2016; Guo et al., 2024). Several factors influence the magnitude of this effect. First, the size of the turbines is important, with larger turbines generally associated with greater reductions in property values (Dröes & Koster, 2021). Second, proximity matters: properties located closer to wind turbines typically experience larger declines (Guo et al., 2024; Dröes & Koster, 2021). Finally, the timing of the effect exhibits heterogeneity. For instance, Guo et al. (2024) show that the negative impact on property values may diminish over time or for more recent installations, suggesting that homebuyers may adjust their preferences as they become accustomed to wind turbines. Our results reveal a negative effect of wind turbine visibility on tourism, providing a benchmark for assessing long-term local impacts, while also suggesting adaptation, as the effect diminishes after approximately seven years.

Third, our findings contribute to the literature by highlighting the effects of wind turbines on local economic outcomes, particularly income, in ways that contrast with previous evidence on employment and income impacts from wind energy development. While prior studies generally report modest or short-lived employment gains and limited changes in income (e.g., Brunner & Schwegman, 2022; Costa & Veiga, 2021; Fabra et al., 2024), our results indicate a negative and significant effect on local income. This suggests that the economic spillovers from wind turbine installations may extend beyond temporary job creation, reflecting broader, longer-term impacts on local economies.

[7] The only notable exception is the work from Broekel & Alfken (2015) who study actual touristic flows in Germany in relation with distance from wind turbines, finding a negative effect. Although the authors themselves recognize that their approach cannot be defined as causal, their work can be considered the most robust among existing ones regarding wind turbines and tourism. Compared to them, at least as far as identification is concerned, we make considerable progress thanks to the use of an econometric model and approach that reflects the current state of the art, the details of which are described later in Section 2.2.

Crucially, these income effects align with our tourism findings, implying that wind turbines influence local economic activity not only through labor market channels but also via tourism-related pathways, such as reductions in visitor expenditure and diminished attractiveness of municipalities to tourists.

Finally, this study contributes to the literature on the cost-benefit analysis of wind power generation, aligning with the consensus that the overall global environmental impact of wind energy is decisively positive, despite the recognized presence of localized environmental externalities (Meng et al., 2024). Substantial prior literature has focused on quantifying these benefits, primarily through the estimation of CO2 emissions savings achieved by deploying wind turbines (Pryor et al., 2020; O'Mahoney & Denny, 2011; Beltrami et al., 2021). Utilizing these previous findings as a foundational methodology, this study reports the calculated CO2 savings specific to Italy over the considered time frame and compares them against the cost on the local economies. This analysis consolidates the fact that the global environmental benefits significantly outweigh the overall costs. We extend the discussion by proposing specific policy interventions designed to compensate for negative local externalities, ensuring the sustained realization of net environmental gains.

The remainder of the paper is structured as follows. Section 2 describes the data, and Section 3 outlines the empirical strategy. Section 4 presents the results, while Section 5 compares the costs associated with tourism to the environmental benefits arising from CO_2 savings. Section 6 concludes with a discussion of policy implications and potential directions for future research. Finally, Section 7 provides the acknowledgments.

2. Data

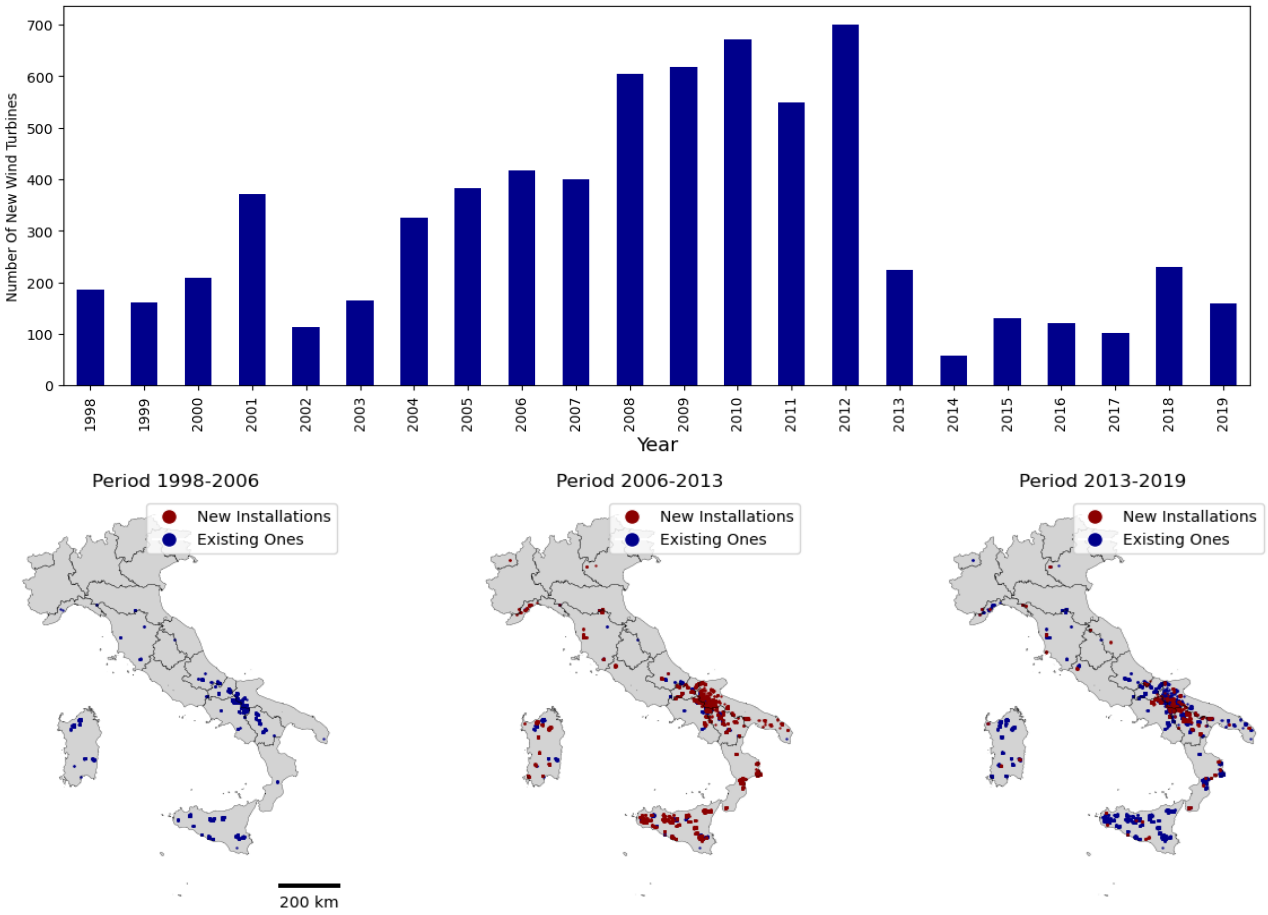
2.1 Wind Turbines Data

Data on wind turbines, producibility, and eligible areas are provided by the research unit of the Authority for the Management of Energy Services – GSE - (*Gestore dei Servizi Energetici*). Specifically, we use the latest (i.e. 2019) full database of all wind turbines in Italy, that reports for each turbine, the energy class, the geocoded location, and the year in which the turbine was linked to the national grid.[8] In Figure 1, we show the location and timing of construction of wind turbines over years. As we can see, the period between 2008 and 2012 was undoubtedly the booming years for new wind turbines installations, that suddenly dropped after 2013. In terms of territorial distribution, most

[8] All the dataset provided by the GSE are publicly and freely available, with the exception of the information regarding the year in which the wind turbine was linked to the grid that has been kindly provided by the GSE under a non-disclosure-agreement.

of new installations are concentrated in the same area of existing ones, tough we can observe a significant expansion in the southern part of the Puglia region (e.g. on the heel of Italy) as well in Sicily in the 2006-2013 period.

Figure 1: Wind Turbines Installations

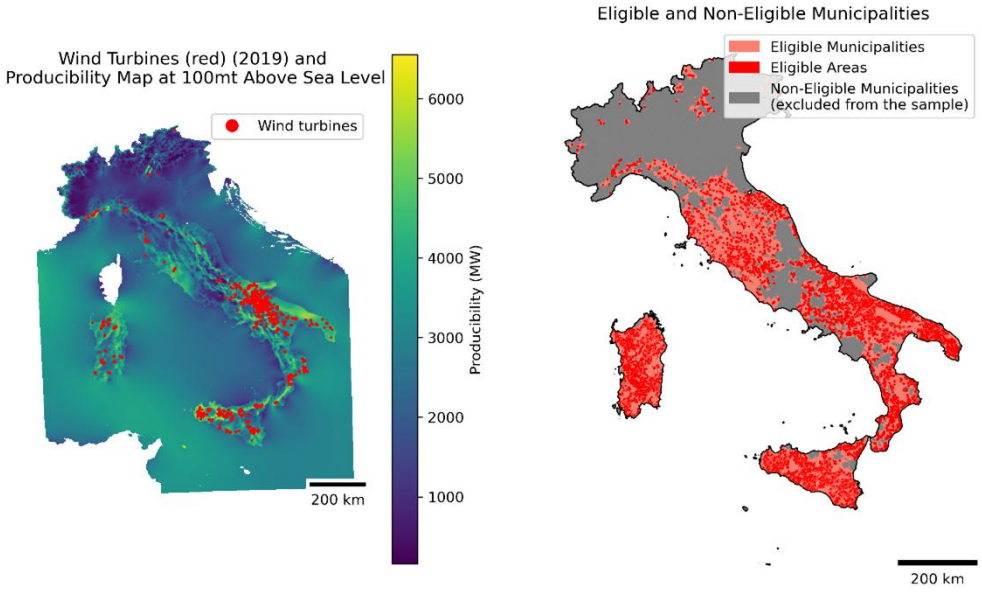


Note: The figure shows on the top-panel the distribution of the year of installations of wind turbines. Please note that year 1998 also include Pre-1998 installations. On the bottom panel we show the same data on the map, where the red markers highlight the turbines built in the period reported in the title.

From the GSE, we also collect the producibility grid that estimates how much wind energy can be produced by a wind turbine in a certain location (according to historical climate data). In addition, we also use another precious information provided by the GSE, that is a collection of geo-polygons that identify eligible areas (*Aree Idonee*) over the Italian territory, that are areas suitable for new onshore wind power installations, calculated starting from the areas with wind power production at a height of 100m above 2250 MWh/MW and excluding urbanized areas and areas subject to environmental and cultural restrictions.

The left-hand panel of Figure 2 shows the spatial distribution of wind turbines across Italy alongside the country's wind potential. In a natural experiment, the treatment should approximate random assignment, meaning that both the treatment and control groups should have a similar likelihood of receiving the treatment. To approach this ideal, we ensure that municipalities not exposed to wind turbines (the control group) and those within the turbines' viewshed (the treatment group) are comparable in terms of factors influencing turbine construction. To this end, we focus on eligible municipalities. The right-hand panel of Figure 2 highlights the municipalities included in our sample in pink, while the grey areas indicate municipalities excluded from the analysis.

Figure 2: Producibility and Eligible Areas



Note: On the left-panel of the figure we show the producibility grid at 100mt above the sea level expressed in MWh/MW and the existing wind turbines (red markers). On the right-panel we highlight eligible areas ("aree idonee") in red, and eligible municipalities in pink (e.g. municipalities which include a portion of eligible area).

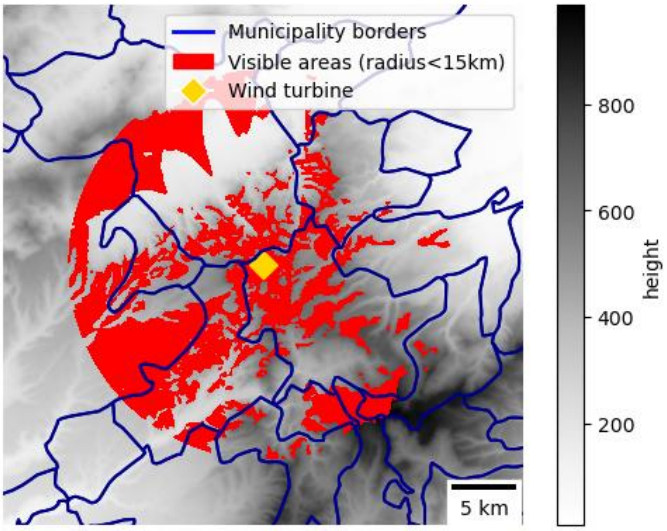
We also use Digital Elevation Data from the European Digital Elevation Model (EU-DEM) ("Copernicus DEM," 2022). These data provide the elevation of each 30-meter pixel across Italy, representing the first surface captured by the sensors (i.e., a digital surface model). Using this information, we calculate wind turbine visibility exposure for each location.

To assess municipal-level exposure, we generate a viewshed for each individual turbine, determining the areas from which it is visible. This analysis uses each turbine's location along with all 30-meter resolution pixels from the EU-DEM covering the surrounding municipalities as input.[9] This approach

[9] We take a square of 60X60 km centered around each wind turbine.

allows us to estimate, for each municipality, the proportion of its territory from which a wind turbine is visible. The viewshed analysis accounts for the terrain's morphology as well as artificial structures such as buildings, walls, and similar obstacles. For instance, if a wind turbine is obscured by a forest, mountain, or other impediment, this will be reflected in the analysis. We implement this procedure using the GDAL package, following the methodology proposed by Wang et al. (2000). Figure 3 illustrates the computed viewshed (in red) for the wind turbines (gold markers), highlighting the areas within each municipality where the turbines are visible. The viewshed is calculated with a maximum visibility distance of 15 km[10], an observer height (i.e., the wind turbine) of 90 meters and a target height of 10 meters. After computing the viewshed for each individual turbine, we aggregate the viewsheds within each municipality for each year (Figure 4). This allows us to calculate both the proportion of each municipality's area from which at least one wind turbine is visible and the expected number of visible turbines per municipality.

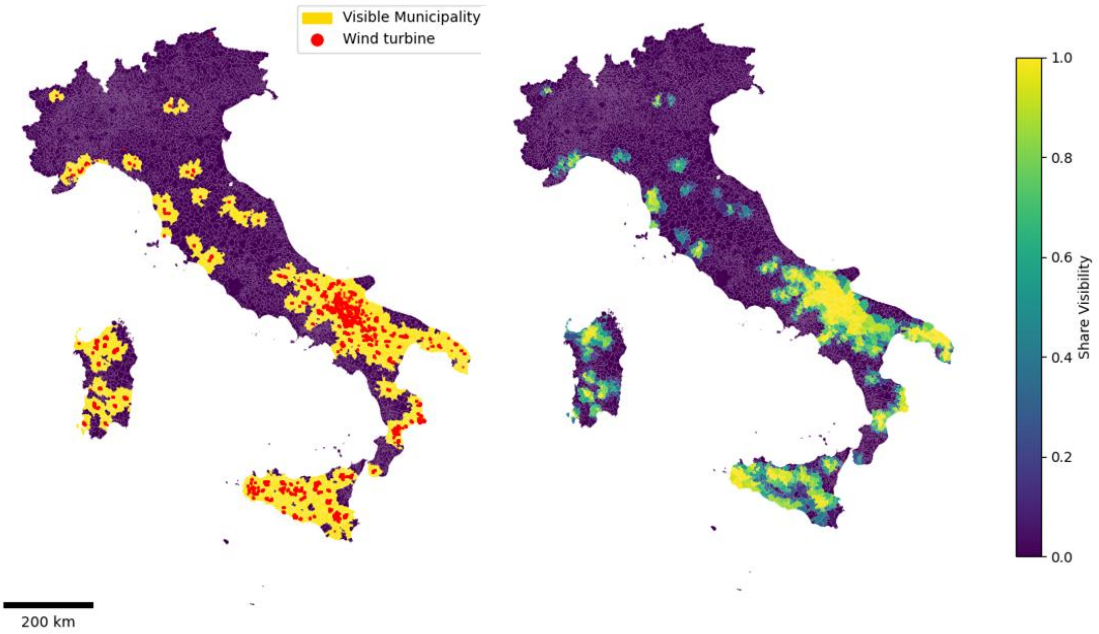
Figure 3: Viewshed Example



Note: The figure shows an example of the viewshed for one wind turbine. The red areas highlight the portion of territory from which the wind turbine (yellow polygon) is visible. The blue lines highlight municipal borders. The grey colormap reports the values of the digital elevation model (in meters).

[10] The maximum visible distance clearly depends on the weather and light conditions (e.g. time of day). These can obviously vary from period to period and within a day. The value we used (15 km), can be defined as a conservative value, in the sense that it is possible that visibility is lower in many cases with the consequence that the intensity of our treatment could be overestimated. We also made this choice taking into account any problems of treatment spillover to the closest areas. 15 km is also the default value used in the ARCGIS viewshed tool (ARCGIS, 2024).

Figure 4: Wind Turbines Visibility



Note: On the left panel we have wind turbines (in red), and municipality that have at least a portion of their territory where the wind turbines are visible. On the Right panel we show the share of visibility for each municipality (share of the territory from which at least one turbine is visible).

To conduct a heterogeneous analysis based on proportion of each municipality's area from which at least one wind turbine is visible and the expected number of visible turbines per municipality, we refer to the boxplots in Figure 5.

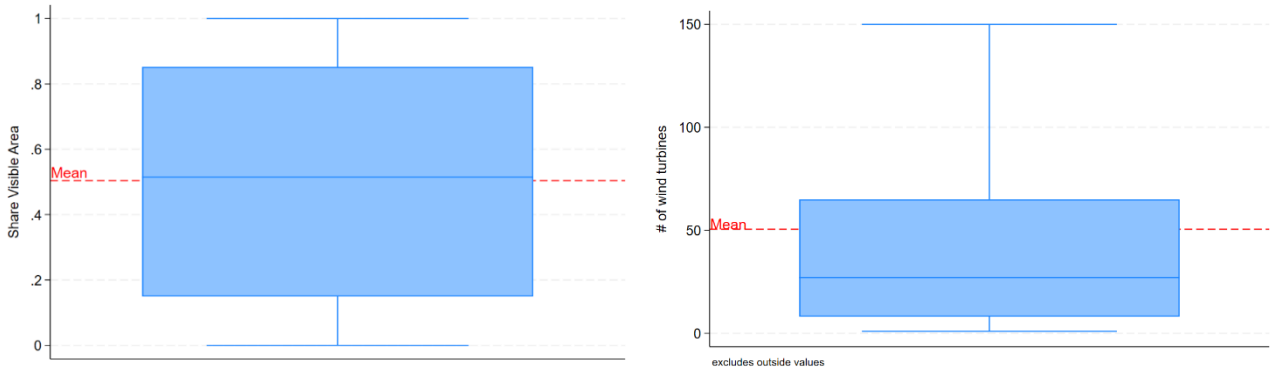


Figure 5: Boxplots of share of visible area and numbers of wind turbines

Note: In the right-hand side panel, the median and mean value are close to 0.5,while the first quartile is 0.15 and third quartile is 0.85. For the number of visible turbines per municipality the median=27 and mean=50, while the first quartile is 8 and third quartile is 65.

Observing a symmetric distribution of visible areas, we divide the sample into three bins for the analysis. This approach allows us to preserve statistical power while avoiding excessive subsampling.

Moreover, the Italian National Institute of Statistics (ISTAT) classifies municipalities according to their level of tourism and key characteristics, such touristic versus non-touristic[11] status. Touristic municipalities are further categorized based on tourism type, including maritime, mountain, historic, cultural, and lake destinations. Later, we leverage this classification to conduct heterogeneous analyses and identify which types of municipalities are most affected by wind turbine visibility. In particular we exploit as urban versus rural status. Table 1 presents descriptive statistics summarizing the characteristics of municipalities where wind turbines are visible.

Table 1. Number of municipalities by type and wind turbine visibility

Wind Turnine Visibility	Urban	Rural	Touristic	Non-Touristic	Total
With Wind Turbine	577	482	763	296	1,059
Without Wind Turbine	412	284	566	130	696
Total	989	766	1,329	426	1,755

The distribution shows that both rural and touristic municipalities represent a notable share of areas where wind turbines are visible, providing context for subsequent analyses of their potential local impacts.

2.2 Tourism Data

Tourism data in Italy come from two main sources. The Bank of Italy conducts an annual survey on international tourism, collecting information through interviews and traveler counts at border crossings—including roads, railways, ports, and airports—supplemented with administrative records and, since late 2020, mobile phone data. The second source is the Italian National Institute of Statistics (ISTAT), which carries out a monthly survey tracking tourist flows and hospitality infrastructure, including the number of beds and types of available accommodations. For tourist flows, we rely on data from the Bank of Italy, which offer a longer time coverage (available since 1998), whereas ISTAT data at the municipal level are only available from 2014 onward. For hospitality infrastructure, we use data provided by ISTAT. In addition, we draw on ISTAT's classification of municipalities by tourism characteristics (see Table 1 for further details).

[11] Municipalities where accommodation facilities are absent and/or which are characterized by negligible tourist flows. These comprise 1,704 municipalities (21.5% of the total), where 6% of the population resides.

The Bank of Italy provides the raw data of the interviews, together with the expansion factors that can be used to recover the total estimated number of tourists passing through each municipality. Clearly, the weakness of the Bank of Italy's data is that being a survey, although numerically important[12], it does not include all the municipalities every year. Following what was done also in Nocito et al. 2023, we impute to those municipalities for which not even one interview is present a value of zero tourists, but we exclude from the sample all those municipalities that never appear in the sample in any year of the survey (13 per cent of the total).

In Appendix 9.1, we present the distribution of average tourist flows according to the two data sources, along with the overall changes between the first and second periods of our study. This confirms substantial heterogeneity at both the national and regional levels. We also report the distribution of tourist accommodations (i.e., the supply side of tourism), which closely corresponds to the observed patterns of tourist demand.

2.3 Other variables

To compare our findings with other socioeconomic outcomes, we also consider house prices, income, and population. Data on house prices come from the Osservatorio del Mercato Immobiliare (OMI), managed by the Italian Revenue Agency (Agenzia delle Entrate), and are available from 2004 onward. The dataset reports minimum and maximum prices per square meter across different property categories (e.g., economic, luxury). We use these figures to compute the annual average housing price at the municipal level. Income data are obtained from the Italian Ministry of Economy and Finance (MEF), providing municipal-level estimates. Population figures are sourced from ISTAT, the Italian National Institute of Statistics.

Table 2 presents the mean values and sample sizes for the outcomes analyzed, restricted to the 'Eligible Areas' outlined in Section 2.1.

Table 2. Descriptive Statistics Table for Eligible Areas

	Observations	Mean
Tourist arrivals	67,938	6,987.535
Tourist expenditure	23,236	10,500,000
Average income	54,473	12,397.97

[12] Approximately 130,000 interviews are carried out annually, equal to approximately 1 per thousand of Italian and foreign travelers crossing the country's borders, and over 1.5 million qualified traveler counts.

Population	46,980	9,025.044
# of B&B	52,192	14.43
# of hotels	52,192	3.37
# of visible wind turbines	14,648	50
% territory with viewshed	14,648	0.5

Note: The number of visible wind turbines and the percentage of territory within the viewshed (i.e., exposed to wind turbine visibility) are reported only for municipalities in the treatment group'.

3. Empirical Strategy

The main challenge we face in terms of identification concerns the non-random timing and location of wind turbine construction. Turbines are rarely built arbitrarily; their placement often reflects a combination of economic, social, political, and environmental considerations. Importantly, many of these factors may be correlated with the very outcomes we aim to study, such as tourist flows or broader local economic indicators. For instance, in areas with a strong tourism industry, municipalities or tourist operators' association might actively oppose wind turbine projects to preserve the attractiveness of the landscape. Beyond socio-economic factors, regulatory constraints also play a role: certain territories, due to landscape, cultural, or heritage protection, are subject to restrictions that can prohibit or complicate turbine construction. These considerations illustrate that simply comparing municipalities with and without wind turbines can lead to significant biases, as the placement of turbines is influenced by factors that also affect economic outcomes. To mitigate these potential distortions, we take several steps to make the sample as homogeneous as possible and to construct a measure of treatment that is less dependent on local economic conditions. Once these steps are applied, we exploit the staggered timing of treatment to implement a difference-in-differences design, as wind turbines are built in different years.

A key point is that our measure of treatment, the visibility exposure to wind turbines, means that the group of treated municipalities may include some that do not actually host wind turbines within their territory. While these exposed municipalities may have some influence, formal or informal, over decisions in neighboring areas regarding turbine construction, it is reasonable to assume that their decision-making power is limited. Including municipalities that did not actively seek or have strong control over treatment helps attenuate selection-into-treatment bias.

To further enhance sample homogeneity, we first exclude municipalities that are ineligible for wind turbine construction. We define eligible territory using the criteria provided by GSE (as discussed in Section 2.1.3), ensuring that no municipality in the control group could have hosted a turbine even if willing. By removing municipalities with a very low probability of treatment, we avoid including "bad controls" and improve the comparability between treated and control groups.[13]

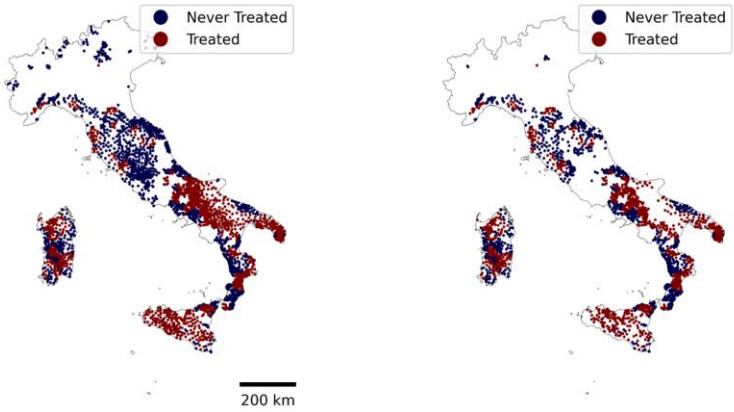
After having excluded non-eligible municipalities, we employ the following sub-sampling procedure: for each Treated Municipality (TM) that has ever been treated, we search all the Non-Treated municipalities (NTMs), that have never been treated and that match the following conditions:

1. Falls in the same region of TM
2. Falls in the same population bracket of TM [14]
3. Is closer than 30km from TM[15]

We keep all the NTMs that are matched at least once, and all TM that have at least one matched control. We drop all other municipalities.

Figure 8 shows the location of Treated and Non-Treated municipalities before and after the matching procedure. Clearly, we lose some observations especially in the North and Center of Italy, tough, our final clean sample still looks representative with respect to eligible territories and potential territories where wind turbines could be constructed (i.e. see Figure 5).

Figure 6: Sample Before and After the



Note: The two maps show the centroids of Never Treated (blue) and Ever Treated municipalities (red), before and after the matching procedure described above. Thus, the observations shown in the right-panel comprise the subsample that we will use in the analysis. Clearly, the visible consequence of the procedure is to have treated and control municipalities that are much closer.

[13] We also exclude a small number of municipalities (N=32) which are not eligible according to the GSE measures but do have wind turbines in their territory or are exposed to wind turbines of nearby municipalities.

[14] We use the following population brackets: [0,3000,5000,15000,30000,100000,250000,3000000]

[15] We compute distance using the aerial distance between the centroids of the two municipalities

After further restricting the sample, we begin by examining the dynamic Two-Way Fixed Effects (TWFE) model, considered the "traditional event study approach."

Equation 1

$$Y_{i,t} = \alpha_i + \alpha_t + \sum_{k=-K}^{-2} \gamma^{lead}{}_k D^k{}_{i.t} + \sum_{k=0}^{L} \gamma^{lags}{}_k D^k{}_{i.t} + X_{i,t} + \epsilon_{i,t} \quad (eq.1)$$

where $Y_{i,t}$ is the outcome variable (e.g. number of tourists, housing prices), α_i is the municipality fixed effect, α_t is the year fixed effect, $D^k{}_{i.t}$ is a dummy=1 if a municipality *i* is under the viewshed of wind turbines in year *t*; γ_k is the effect of being exposed to wind turbines *k* years before (leads) and after (lags) the wind turbine has been implemented and the contemporaneous effect (when *k*=0), $X_{i,t}$ is a set of covariates that are varying across municipalities, such as population, and finally $\epsilon_{i,t}$ is the error terms that we clustered at municipality level. We avoid normalizing tourist arrivals by population, as population itself may respond to wind turbine installation, potentially biasing per capita measures[16]. Instead, we use log tourism outcomes and include population as a control only in robustness checks to account for demographic differences across municipalities.

However, Sun & Abraham (2021) showed that coefficients from dynamic Two-Way Fixed Effects (TWFE) models can generate negative weights, potentially leading to misleading interpretations, as the estimated effects may combine treatment effects from different periods. To address this issue, we compare the classic event-study estimates with the Callaway & Sant'Anna, 2021 and de Chaisemartin & D'Haultfœuille, 2024 estimators, finding that in our case there are no statistically significant differences between the approaches.

For the number of tourists and tourist expenditure, we use Poisson regressions to account for the large number of zero observations and to obtain a consistent relative measure of the effects (Chen & Roth, 2024). Indeed, municipalities exhibit substantial variation in their baseline levels of tourist arrivals, implying that the interpretation of the event-study coefficients γ_k in absolute terms would not be scale-invariant. As a robustness check, we therefore re-estimate the models using the dependent variable in absolute levels and compare these estimates to the mean outcome of the treatment group, focusing only on changes along the extensive margin. Population is expressed in logarithmic form, while for other outcome (income and housing prices), we use the change in logarithm.

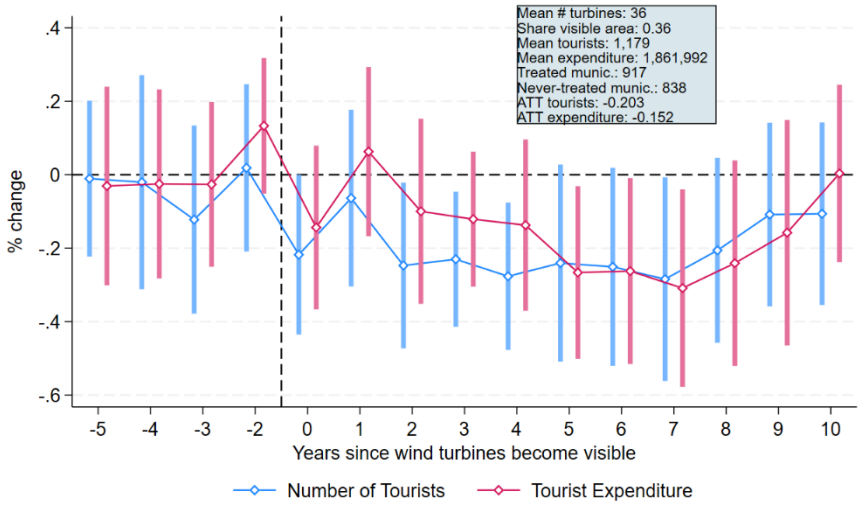
[16] Normalizing by a variable that is endogenous to the treatment can introduce post-treatment bias (see Angrist & Pischke, 2009; Montgomery et al., 2018), leading to misleading inferences about causal effects.

4. Results

4.1 Main result: effect of wind turbines visibility on tourism

We begin by examining the impact of wind turbine installations on the number of tourists and tourist expenditure. We compare municipalities that experience the installation with those that are not yet treated or never treated, using a classic staggered difference-in-differences event-study design. We restrict the analysis to the matched units, as outlined in Section 3 and relying on Poisson regressions to address potential biases introduced by zero values (Chen & Roth, 2024).

Figure 7: Impact of wind turbines visibility on tourist arrivals and expenditure



Note: Standard errors are clustered at municipality level; we keep only municipality that have balanced observations; The average number of tourists and tourist expenditure refer to the treatment group, since the ATT refers to the Average Treatment effect of Treated units,

The blue event-study shows that municipalities exposed to wind turbines visibility experience an average decline in tourist arrivals of approximately 20% relative to comparable matched controls. This effect grows over time, reaching a peak reduction of 28% around seven years after turbines become visible. From roughly the eighth year onward, the decline starts to moderate, suggesting that tourists may gradually adapt to the presence of visible wind turbines, rather than the effect being entirely permanent. The years before wind turbines become visible serve as placebo tests, confirming that there are no systematic differences in tourism trends prior to turbine installation. Appendix 9.2 presents a series of robustness checks to assess the sensitivity of our findings. First, re-estimating the model in absolute terms rather than using the Poisson specification yields consistent results, confirming the presence of negative effects after wind turbines become visible. Next, we show that the traditional TWFE event-study estimates are statistically different from those

obtained using more recent estimators (Callaway and Sant'Anna, 2020; de Chaisemartin and D'Haultfœuille, 2024), suggesting that our results are not driven by bias arising from the staggered timing of wind turbine implementation. We further extend the analysis to include a wider set of pre- and post-installation periods, again finding consistent negative effects. Expanding the sample to include all municipalities in Italy also confirms the robustness of the negative impact. Finally, a randomization test demonstrates that the estimated effects are not the result of spurious correlations.

When examining tourist expenditure (red line), we observe a pattern like that for tourist numbers after wind turbines become visible. On average, municipalities experience a 15% decline in tourist expenditure in the years after installation. Given that tourism accounts for approximately 10.5% of local GDP in Italy (WTTC, May 2024), this estimated 15% reduction in tourist expenditure corresponds to roughly 1.6% of municipal GDP. In Appendix 9.2, we further examine the local economic effects by analyzing income growth, measured as the change in log income to account for differences in baseline levels. We find a negative impact of around 1% in the six to seven years following wind turbine installation, confirming that turbine visibility has an adverse effect on the local economy. Both population and housing market indicators also decline after installation. Consistent with evidence from other countries, we observe pre-existing downward trends in income and population even before the turbines are built. This pattern suggests that residents may begin relocating in anticipation of construction, leading to a gradual reduction in local income. Alternatively, it may indicate that wind turbines are more likely to be sited in areas already experiencing demographic and economic decline. Interestingly, the housing market also shows a negative effect at the municipal level, with prices beginning to fall in the year prior to turbine visibility, consistent with the anticipation effects documented by Dröes and Koster (2016) for the Netherlands.

In Appendix 9.2, we show the impact of wind turbine installations on tourism supply metrics, specifically the number of hotels, bed and breakfasts (B&Bs), and private apartments. For analytical clarity, we group these into two categories: hotels and other accommodations. This distinction is crucial for capturing differences in their capacity to adapt to fluctuations in demand. Larger establishments, such as hotels, benefit from economies of scale, enabling them to absorb demand shocks by adjusting prices. In contrast, smaller establishments, such as B&Bs, may face greater challenges, often opting for closure when faced with sustained drops in demand. Private accommodations, such as bed and breakfasts and short-term rentals, decline following wind turbine installation, while the number of hotels remains stable or slightly increases, suggesting some degree of substitution effect. These patterns indicate that the negative impact on tourism is primarily demand-

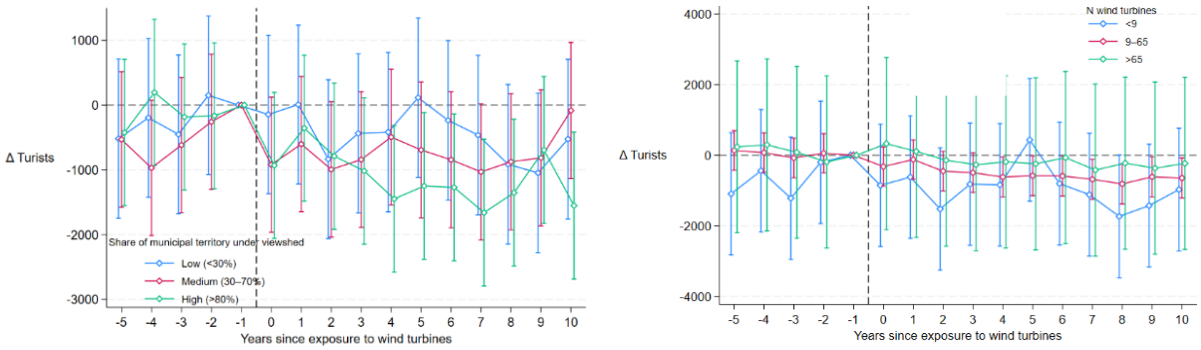
driven, with supply adjustments occurring gradually, six to seven years after installation, in line with declines in overall tourist expenditure (red line in Figure 7).

4.2 Heterogeneity analyses

While the main results in Figure 7 report average effects, municipalities differ in both the share of territory exposed to wind turbine visibility and the number of turbines visible. On average, treated municipalities have 0.36% of their territory exposed, with 36 turbines visible. In this section, we explore how the impact of wind turbines varies across municipalities along these two dimensions, examining whether larger exposure or a higher number of visible turbines amplifies the observed effects.

To examine heterogeneous effects, we first decompose the impact of wind turbine installation by the share of each municipality exposed to turbine visibility, and second, by the number of turbines visible from the municipality (see Section 2 for summary statistics). For the first categorization, we split municipalities into three groups: low exposure (<30% of territory), medium exposure (30–70%), and high exposure (>80%).

Figure 8: Impact of wind turbines visibility on tourist arrivals by intensity of treatment



Our findings in Figure 8 show that the negative effect of wind turbines on the number of tourists, as reported in Figure 7, is primarily driven by municipalities with high exposure. In this group (green line), the effect is the only one statistically different from zero between four and seven years after the turbines become visible. Moreover, municipalities with medium and high exposure levels experience a more immediate impact, with a significant decline in tourism (at the 10% level) already observable in the year of exposure. On the right-hand panel of Figure 8, we examine the heterogeneous effects associated with different numbers of wind turbines visible from each municipality. The results indicate that when more than 65 turbines are visible, the effect on tourist numbers is never statistically significant, likely because large wind farms are typically located in remote areas with limited tourist presence. In contrast, we find significant effects for municipalities

where between 9 and 65 turbines are visible, consistent with main result in Figure 1. For municipalities with fewer than 9 visible turbines, the estimated effects are even more negative, but are not statistically different from zero. This pattern suggests that the relationship between the number of visible wind turbines and tourist activity is likely non-linear.

In addition to the intensity of exposure, the visual characteristics of the territory may also play a crucial role. The visual impact of wind energy infrastructure varies considerably depending on the characteristics of the surrounding landscape and its perceived appropriateness for such developments (Johansson & Laike, 2007). This means that the type of landscape plays a central role in shaping both public perception and acceptance of wind turbines. For instance, Wolsink (2007a), in a study around the Wadden Sea, found that industrial and military zones were generally regarded as suitable locations for wind farms, whereas natural reserves, scenic landscapes, and recreational areas were considered far less appropriate. Here, to understand which kind of municipalities are most affected, we use the ISTAT classification (see section 2 for more detail). First, we find that the negative effect is largely driven by touristic municipalities[17]. Among these, rural municipalities, especially those located in mountain or maritime areas, experience the strongest negative effects on tourism from visibility of wind turbines (see Appendix 9.3 for further details). This finding contributes to the literature by showing that not only turbine height and proximity matter (Guo et al., 2024; Dröes & Koster, 2021), but also the type of surrounding landscape. In doing so, our analysis provides empirical evidence for a relationship that had previously been suggested only at a theoretical level.

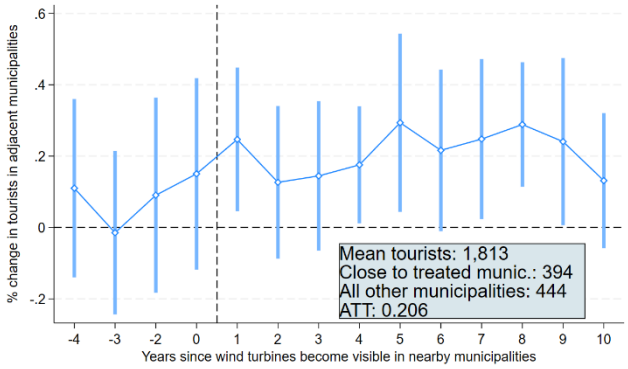
4.3 Potential spillover

To better understand the dynamics in bordering control municipalities, we conduct a separate analysis with two objectives. First, we assess whether bordering control municipalities serve as valid controls by ensuring they are not significantly influenced by wind turbine installations in adjacent treated municipalities ("no spillover effects"). Second, we examine the potential for displacement effects, hypothesizing that tourists might shift their preferences to nearby municipalities where wind turbines are not visible. If such displacement occurs, we would expect an increase in tourist numbers in bordering municipalities as they offer proximity to the treated areas without the perceived visual disamenities of wind turbines. So, we conduct an analysis in which the treatment group includes

[17] "Non-tourist municipalities" are municipalities where no accommodation facilities are present and/or tourist flows are negligible. These include 1,704 municipalities (21.5% of the total), which together account for 6.2% of the population.

municipalities adjacent to, but not directly affected by, wind turbine visibility, while the control group is restricted to municipalities without wind turbine visibility and located farther away.

Figure9: Impact of wind turbines visibility on tourist arrivals in adjacent municipalities



Our analysis shows that, in the years after a new wind turbine installation, tourist numbers increase in nearby municipalities that are not exposed to turbine visibility. This suggests that tourists may prefer nearby areas without visible wind turbines.

5 Cost-Benefit analysis and compensation

A comprehensive cost-benefit analysis of wind turbine implementation requires careful estimation of both the associated costs and the potential benefits. While data on the negative impact of wind turbines on tourism help quantify certain costs, they do not provide insights into the corresponding benefits. Although conducting a full cost-benefit analysis is beyond the scope of this paper, this section explores a key factor involved in estimating the benefits. Specifically, we estimate the amount of CO_2 emissions avoided through wind turbine installations at the municipal level, based on varying wind energy potentials. We then assign a monetary value to these emission reductions and compare this figure to the estimated negative impacts. Finally, we propose a theoretical compensation model for local externalities, incorporating the global environmental benefits into the analysis.

To calculate how much CO2 is saved by wind turbines, we start by understanding how much 'clean energy' (energy that is produced with 0 emissions) is produced by wind turbines in Italy. We are not considering the cost of producing wind turbines in terms of carbon emissions because are so low compared to other energy sources that they are irrelevant (cite some papers). To understand how much energy is produced by each wind turbine, we use the producibility at 100mt above the sea level (map on the left hand side of figure 5, section 1) expressed in MWh/MW. Producibility is the measure of how much electricity a wind turbine can produce on average in one year, given the wind resource at the site. We took the producibility grid is formed by 5km by 5km squares , and so for each wind

turbine we associate the producibility of the closest centroid. We then compute the annual production of each turbine by taking into consideration the power class (provided by Gestore dei Servizi Energetici S.p.A) in the following way:

Annual production (MWh) = producibility (MWh/MW)*Power class (MW)

There are four power classes in the dataset '200-1000 kW', '1000-2000 kW, '2000-3000 kW', >3000 kW' and we took the midpoint for each classes to calculate the annual production (for the last class we consider 3500) for each wind turbine.

The average annual production for a wind turbine in Italy in the considered time period (1998-2019) is around 7.000 MWh[18].

One challenge from engineering perspective is to calculate how much gCO_2 per kWh is saved using wind turbines. Researchers often estimate the climate benefits of renewable energy by comparing the life-cycle emissions intensity of different electricity sources. According to the IPCC Sixth Assessment Report (AR6 WGIII) (Skea et al., 2022), electricity generated from wind emits approximately 9.3 g CO_2-eq per kWh, whereas coal-based generation releases about 965 g CO_2-eq per kWh, and natural gas around 465 g CO_2-eq per kWh. These findings are consistent with prior analyses that examine emission reductions from replacing fossil fuel–based electricity with renewables (O'Mahoney & Denny, 2011; Beltrami et al., 2021). Building on these sources and following the approach of Pryor et al. (2020), we assume that substituting coal-fired power with wind energy results in an average reduction of roughly 600 g CO_2 per kWh.

We then consider the value in monetary terms of the CO_2 saved. We take the price of CO2 in SendeCO2, which publicly report the average yearly CO_2 price following EU ETS (emission trading system) market price. The average CO_2 price from 2008 to 2023 has been 23.90 €/t.

Converting everything in kg, we have $0.6\,\text{kg/kwh} * 0.0239\,€/\text{kg} = 0.01434\,€/\text{kwh}$ saved for every kwh. Considering that for our sample the average annual production is 7,000,000 kwh for a wind turbine, we found that on average a wind turbine saved 100,380€ in one year. For example, considering that wind turbines in Italy in 2019 were around 6000, this means savings of 602 million € in 2019.

According to Bank of Italy data, the average annual tourist expenditure in a municipality is around €12 million. As reported in Section 2.3, municipalities exposed to wind turbine visibility are, on average, exposed to about 50 wind turbines. Given the estimated 15% decrease in tourism spending

[18] We exploit the classes given by GSE for each wind turbine in Italy to calculate the mean value

presented in the results section, this corresponds to an annual loss of roughly €1.8 million. However, this loss remains lower than the estimated €5 million in revenue generated through CO_2 savings associated with wind energy production. While this is substantial, it remains well below the total environmental benefits generated by 50 turbines, which amount to roughly €1.8 million. To reconcile local losses with global gains, municipalities could be compensated through a dedicated "Renewable Energy Tourism Fund," financed by a small levy on wind energy revenues proportional to the CO_2 savings achieved. Additional funding could come from EU green funds, national climate subsidies, or direct contributions from wind farm developers via a "local impact fee," with government top-ups ensuring full compensation if revenues are insufficient. Rather than relying solely on compensation mechanisms, policymakers should prioritize the optimal siting of wind turbines to minimize local costs while maintaining overall energy and environmental benefits. Careful planning can take into account factors such as landscape characteristics, tourism intensity, and population density to reduce negative local externalities. By strategically locating turbines in areas where their visibility and impact on economic and social activities are limited, it is possible to reconcile the expansion of wind energy with the protection of local communities and their economic well-being.

6 Conclusions

Our results contribute to the often heated and polarized debate surrounding the "green–green" dilemma posed by the construction of wind turbines. Globally, wind energy clearly delivers environmental benefits, and at the national level, similar advantages arise in terms of energy independence and system resilience. In particular, wind energy represents an effective strategy to mitigate the volatility of energy prices and to strengthen resilience against disruptions in traditional fossil fuel supply chains, which are increasingly exposed to international tensions. However, while the benefits of wind energy tend to be widely shared, the environmental and visual costs are highly concentrated, falling primarily on the communities that host wind farms. This imbalance is partly due to technological and geographical constraints, which make wind energy development viable and profitable only in specific locations. Moreover, the direct local economic benefits of wind turbines are minimal (Fabra et al., 2024), as these projects require little ongoing labor and are largely limited to the construction phase, offering limited spillover effects for local economies. In this context, the emergence of NIMBY-type conflicts is almost inevitable.

Our findings provide compelling evidence that wind turbine visibility has substantial local economic effects. We show that tourist arrivals decline by up to 20%, peaking seven years after installation and

gradually diminishing thereafter, suggesting that communities may adapt over time. The effect is particularly pronounced in municipalities with 9 to 65 visible turbines and in areas where more than 80% of the territory is exposed, highlighting the importance of both the scale of exposure and the spatial distribution of turbines. Rural, tourism-dependent municipalities are disproportionately affected, underscoring the vulnerability of areas that rely heavily on natural and cultural amenities for their economic activity. This is also consistent with Daniele et al. (2024) who find that more touristic municipalities tend to have a greater political reaction to wind turbine projects.

Tourist expenditure reflects this pattern, with affected municipalities experiencing an average annual decline of 15%, equivalent to approximately €1.8 million or 1.6% of local GDP. The negative impacts extend to broader economic activity, with local income falling by around 1%. These results emphasize that, while wind energy generates substantial environmental benefits, careful policy design is essential to mitigate local economic costs, particularly in tourism-dependent and rural areas. In this context, our findings highlight the potential value of strategic turbine placement and compensatory measures to balance local economic outcomes with the broader environmental gains from renewable energy development. %.

However, these losses are outweighed by substantial environmental benefits: on average, 50 wind turbines generate €5 million per year in value from CO_2 emission savings, more than three times greater than the costs associated with wind turbine visibility for a municipality. Compensation mechanisms could take various forms, including dedicated funds, contributions from developers, or EU and national green subsidies. One promising policy direction is to develop complementary direct local benefits, such as enhanced electrification infrastructure, public charging stations, or improved grid access, to offset perceived landscape costs.

Overall, our results offer important guidance for the design and implementation of wind energy policies. They underscore the need for careful consideration of compensatory mechanisms and, crucially, the strategic selection of locations and quantity of wind turbines. For instance, we find a non-linear relationship between the number of visible turbines, suggesting that constructing large wind farms at once (more than 65 turbines) may be preferable to dispersing smaller numbers of turbines across multiple sites[19].

Beyond compensation and optimal allocation, our findings highlight the importance of minimizing the landscape impact of wind turbines and addressing the socio-economic heterogeneity of their

[19] At the same time, a careful assessment must also consider ecological impacts, since large wind farms may potentially degrade biodiversity and natural habitats over a broader area than smaller installations, highlighting the need to balance environmental benefits with potential local ecological costs.

effects. In Italy, potential wind farm locations frequently coincide with peripheral or ultra-peripheral municipalities, areas that are often economically disadvantaged, experiencing depopulation, and facing limited access to public services. This raises important environmental justice concerns: while the benefits of renewable energy are broadly shared, the costs are disproportionately borne by vulnerable communities. Our results reveal heterogeneity in impacts on the tourism sector. While we find no significant effect on traditional accommodations such as hotels, smaller operators, including B&Bs and private accommodation, experience notable declines, likely because they operate at smaller scales and target different tourist segments. Future research could further examine how renewable resources affects different demographic and socioeconomic groups, explore disparities in economic and social outcomes, and assess the effectiveness of compensation mechanisms in alleviating local burdens. Such work would provide valuable guidance for policies aiming to balance the environmental benefits of wind energy with equitable local outcomes.

Finally, our study underscores the need to consider how public perceptions evolve over time. While wind turbines are often initially seen as intrusive, awareness campaigns emphasizing their role in global environmental sustainability may gradually increase acceptance. Future work could also examine how technological and design innovations can minimize landscape degradation, providing actionable guidance to balance global environmental gains with localized social, economic, and aesthetic impacts.

7 Acknowledgements

The authors wish to acknowledge and sincerely thank Gestore dei Servizi Energetici – GSE S.p.A. for granting access to critical primary data. Their contribution of specific information on the location and characteristics of wind turbines was invaluable for the precision and granularity of this study.

8 References

Ansa. (2024, August 30). *Attentato a parco eolico in Sardegna, il secondo in pochi giorni*. Https://Www.Ansa.It/Sardegna/Notizie/2024/08/30/Attentato-a-Parco-Eolico-in-Sardegna-Il-Secondo-in-Pochi-Giorni_1a804ac9-5891-404c-9ca9-88ee4a96cc3f.Html.

Angrist, J. D., & Pischke, J.-S. (2009). *Mostly Harmless Econometrics: An Empiricist's Companion.* Princeton University Press.

ARCGIS. (2024, December 18). *Viewshed - ARCGIS*. Https://Pro.Arcgis.Com/En/pro-App/Latest/Tool-Reference/Ready-to-Use/Viewshed.Htm.

Beltrami, F., Fontini, F., & Grossi, L. (2021). The value of carbon emission reduction induced by renewable energy sources in the Italian power market. *Ecological Economics*.

Bertacchini, E., Revelli, F., & Zotti, R. (2024). The economic impact of UNESCO World Heritage: Evidence from Italy. *Regional Science and Urban Economics*, *105*, 103996. https://doi.org/10.1016/j.regsciurbeco.2024.103996

Boyle, K. J., Boatwright, J., Brahma, S., & Xu, W. (2019). NIMBY, not, in siting community wind farms. *Resource and Energy Economics*, *57*, 85–100. https://doi.org/10.1016/j.reseneeco.2019.04.004

Brunner, E. J., & Schwegman, D. J. (2022). Commercial wind energy installations and local economic development: Evidence from U.S. counties. *Energy Policy*, 165, 112993. https://doi.org/10.1016/j.enpol.2022.112993

Broekel, T., & Alfken, C. (2015). Gone with the wind? The impact of wind turbines on tourism demand. *Energy Policy*, *86*, 506–519. https://doi.org/10.1016/j.enpol.2015.08.005

Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-Differences with multiple time periods. *Journal of Econometrics*, *225*(2), 200–230. https://doi.org/10.1016/j.jeconom.2020.12.001

Canepa, C., & Taddei, M. (2024). Il turismo non genera il 18 per cento del PIL italiano. *Https://Pagellapolitica.It/Articoli/Santanche-Turismo-Italiano-Pil*.

Clément de Chaisemartin and Xavier D'Haultfœuille (2024). Difference-in-Differences Estimators of Intertemporal Treatment Effects. *The Review of Economics and Statistics*.

Chen, J., & Roth, J. (2024). Logs with Zeros? Some Problems and Solutions. *The Quarterly Journal of Economics*, *139*(2), 891–936. https://doi.org/10.1093/qje/qjad054

Copernicus DEM. (2022). In *Copernicus DEM*. https://doi.org/10.5270/ESA-c5d3d65

Costa, H., & Veiga, L. (2021). *Local labor impact of wind energy investment: An analysis of Portuguese municipalities*. *Energy Economics*, *94*, 105055. https://doi.org/10.1016/j.eneco.2020.105055

Daniele, F., de Blasio, G., & Pasquini, A. (2024). *Is local opposition taking the wind out of the energy transition?*

Devlin, E. (2005). Factors Affecting Public Acceptance of Wind Turbines in Sweden. *Wind Engineering*, *29*(6), 503–511. https://doi.org/10.1260/030952405776234580

Dröes, M. I., & Koster, H. R. A. (2016). Renewable energy and negative externalities: The effect of wind turbines on house prices. *Journal of Urban Economics*, *96*, 121–141. https://doi.org/10.1016/j.jue.2016.09.001

Fabra, N., Gutiérrez, E., Lacuesta, A., & Ramos, R. (2024). Do renewable energy investments create local jobs? *Journal of Public Economics*, *239*, 105212. https://doi.org/10.1016/j.jpubeco.2024.105212

Fooks JR, Messer KD, Duke JM, Johnson JB, Li T, Parsons GR. Tourist Viewshed Externalities and Wind Energy Production. *Agricultural and Resource Economics Review*. 2017;46(2):224-241. doi:10.1017/age.2017.18

Frantál, B., & Kunc, J. (2011). Wind turbines in tourism landscapes. *Annals of Tourism Research*, *38*(2), 499–519. https://doi.org/10.1016/j.annals.2010.10.007

Guo, W., Wenz, L., & Auffhammer, M. (2024). The visual effect of wind turbines on property values is small and diminishing in space and time. *Proceedings of the National Academy of Sciences*, *121*(13). https://doi.org/10.1073/pnas.2309372121

Gestore dei Servizi Energetici – GSE S.p.A. (2025). *Energia da fonti rinnovabili in Italia nel 2023: Rapporto Statistico*. GSE-Ufficio Statistiche e Monitoraggio Target.

Haan, P., & Simmler, M. (2018). Wind electricity subsidies — A windfall for landowners? Evidence from a feed-in tariff in Germany. *Journal of Public Economics*, *159*, 16–32. https://doi.org/10.1016/j.jpubeco.2018.01.011

ISTAT. (2022). *Classificazione Dei Comuni In Base Alla Densità Turistica*. Https://Www.Istat.It/Classificazione/Classificazione-Dei-Comuni-in-Base-Alla-Densita-Turistica/.

Johansson, M., & Laike, T. (2007). Intention to respond to local wind turbines: the role of attitudes and visual perception. *Wind Energy*, 10(5), 435-451.

Krekel, C., & Zerrahn, A. (2017). Does the presence of wind turbines have negative externalities for people in their surroundings? Evidence from well-being data. *Journal of Environmental Economics and Management*, *82*, 221–238. https://doi.org/10.1016/j.jeem.2016.11.009

Landry, C. E., Allen, T., Cherry, T., & Whitehead, J. C. (2012). Wind turbines and coastal recreation demand. *Resource and Energy Economics*, *34*(1), 93–111. https://doi.org/10.1016/j.reseneeco.2011.10.001

Locci, G. (2024, July 29). La rivolta di Orgosolo contro le pale eoliche: “Impediremo i progetti.” *La Stampa*.

Maehr, A. M., Watts, G. R., Hanratty, J., & Talmi, D. (2015). Emotional response to images of wind turbines: A psychophysiological study of their visual impact on the landscape. *Landscape and Urban Planning*, *142*, 71–79. https://doi.org/10.1016/j.landurbplan.2015.05.011

Meng, L., Liu, P., Zhou, Y., & Mei, Y. (2024). Blaming the wind? The impact of wind turbine on bird biodiversity. *Journal of Development Economics*, 103402. https://doi.org/10.1016/j.jdeveco.2024.103402

Montgomery, J. M., Nyhan, B., & Torres, M. (2018). How Conditioning on Posttreatment Variables Can Ruin Your Experiment and What to Do About It. *American Journal of Political Science*, 62(3), 760–775.

Nocito, S., Sartarelli, M., & Sobbrio, F. (2023). A beam of light: Media, tourism and economic development. *Journal of Urban Economics*, *137*, 103575. https://doi.org/10.1016/j.jue.2023.103575

OECD. (2022). *Housing Taxation in OECD Countries*. OECD. https://doi.org/10.1787/03dfe007-en

O'Mahoney, Amy & Denny, Eleanor, 2011. "The Merit Order Effect of Wind Generation on the Irish Electricity Market," MPRA Paper 56043, University Library of Munich, Germany.

Parsons, G., & Firestone, J. (2018). *Atlantic offshore wind energy development: values and implications for recreation and tourism*.

Pryor, S. C., Barthelmie, R. J., Bukovsky, M. S., & Leung, L. R. (2020). Climate change impacts on wind power generation. *Nature Reviews Earth & Environment*, *1*(12), 627–643. DOI: 10.1038/s43017-020-0101-7

Sæþórsdóttir, A. D., Wendt, M., & Tverijonaite, E. (2021). Wealth of Wind and Visitors: Tourist Industry Attitudes towards Wind Energy Development in Iceland. *Land*, *10*(7), 693. https://doi.org/10.3390/land10070693

Schmidt, J. H., & Klokker, M. (2014). Health Effects Related to Wind Turbine Noise Exposure: A Systematic Review. *PLoS ONE*, *9*(12), e114183. https://doi.org/10.1371/journal.pone.0114183

Serra, C. (2023). Blowing in the Wind: Revenue Windfalls and Local Responses from Wind Farm Development. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4638584

Stegmaier, V., & Krause, M. (2023). Headwind at the Ballot Box? - The Effect of Visible Wind Turbines on Green Party Support. *ZBW - Leibniz Information Centre for Economics, Kiel*.

Stokes, L. C., Franzblau, E., Lovering, J. R., & Miljanich, C. (2023). Prevalence and predictors of wind energy opposition in North America. *Proceedings of the National Academy of Sciences*, *120*(40). https://doi.org/10.1073/pnas.2302313120

Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, *225*(2), 175–199. https://doi.org/10.1016/j.jeconom.2020.09.006

Sunak, Y., & Madlener, R. (2016). The impact of wind farm visibility on property values: A spatial difference-in-differences analysis. *Energy Economics*, *55*, 79–91. https://doi.org/10.1016/j.eneco.2015.12.025

Tecleme, L. (2024, August 12). Perché in Sardegna nessuno vuole l'energia eolica. *Fan Page*.

Töller, A. E., Garske, B., Rasch, D., Weigel, A., & Hahn, H. (2024). Failing successfully? Local referendums and ENGOs' lawsuits as challenges to wind energy expansion in Germany. *Zeitschrift Für Vergleichende Politikwissenschaft*. https://doi.org/10.1007/s12286-024-00610-1

The Guardian. (2025, July 28). *Tilting at windmills? Trump's claims about turbines fact-checked.* https://www.theguardian.com/environment/2025/jul/28/are-trump-claims-about-wind-power-correct

Trandafir, S., Gaur, V., Behanan, P., Uchida, E., Lang, C., & Miao, H. (2020). How Are Tourists Affected By Offshore Wind Turbines? A Case Study Of The First U.S. Offshore Wind Farm. *Journal of Ocean and Coastal Economics*, *7*(1). https://doi.org/10.15351/2373-8456.1127

Tverijonaite, E., & Sæþórsdóttir, A. D. (2023). Tourism and onshore wind turbines: Literature review. *Institute of Life and Environmental Sciences, University of Iceland*.

Wang, J., Robinson, G. J., & White, K. (2000). Generating viewsheds without using sightlines. *Photogrammetric Engineering and Remote Sensing*, *66*(1), 87–90.

Wind Europe. (2022). *European commision's Guidance on accelerating the deployment of renewables and facilitating PPAs uptake - WindEurope's response to the public consultation*.

Wolsink, M. (2000). Wind power and the NIMBY-myth: institutional capacity and the limited significance of public support. *Renewable Energy*, *21*(1), 49–64. https://doi.org/10.1016/S0960-1481(99)00130-5

Wolsink, M. (2007a). Planning of renewables schemes: Deliberative and fair decisionmaking on landscape issues instead of reproachful accusations of non-cooperation. *Energy Policy*, 35(5), 2692-2704. https://doi.org/10.1016/j.enpol.2006.12.002

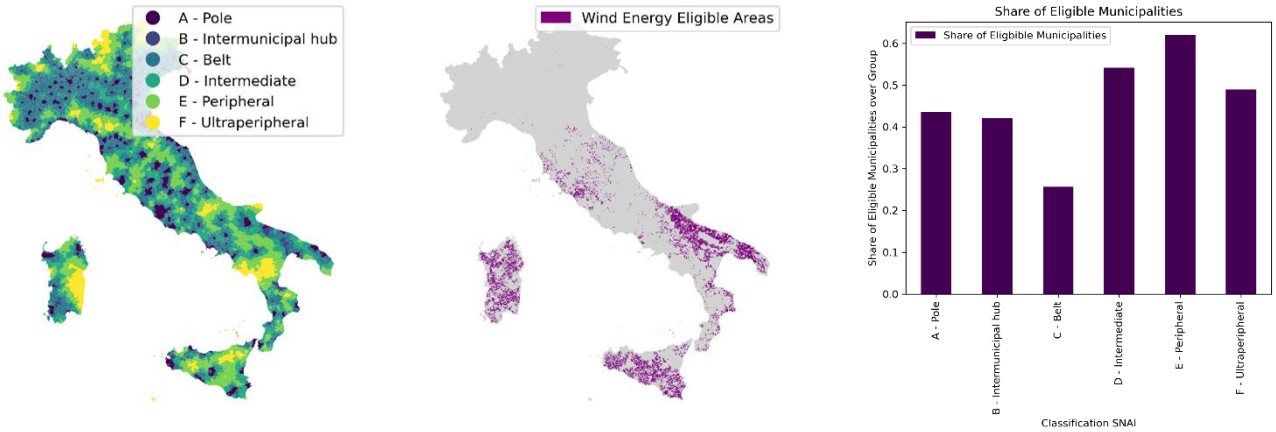
World Travel & Tourism Council. (2024, May 2). *Travel & tourism injected €215BN into Italy's economy*. https://wttc.org/news/travel-and-tourism-injected-215-euros-bn-into-italys-economy

World Travel & Tourism Council. (2024, September 4). *U.S. remains the world's most powerful travel & tourism market*. https://wttc.org/news/us-remains-the-worlds-most-powerful-travel-and-tourism-market

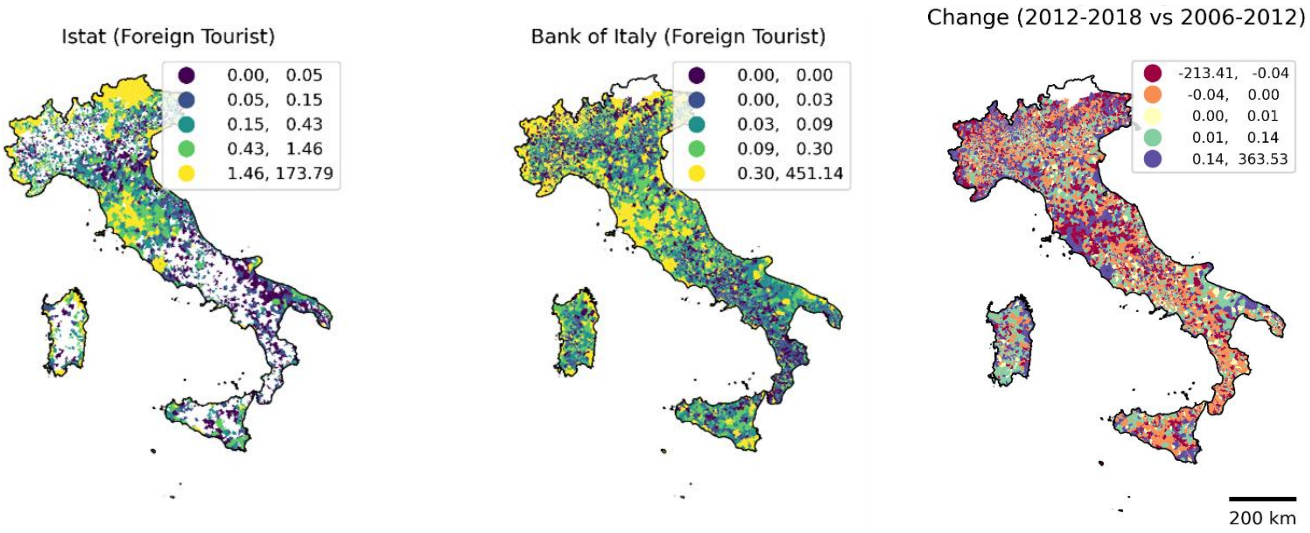
9 Appendix

9.2 Additional Descriptive Statistics

Figure A SEQ Figure_A * ARABIC 1: Classification of Aree Interne



Note: The figure shows on the left-panel the classification of Italian Municipalities according to the "Aree Interne Strategy", where municipalities are differentiated based on the level of accessibility of public services such as public transports and schools. Municipalities classified as Intermediated. Peripheral or Ultraperipheral are the "Internal Areas", territories that tend to be more economically disadvantaged and subject to depopulation, and for these reasons are designatory of specific public policies. For more details see https://www.agenziacoesione.gov.it/strategia-nazionale-aree-interne/. On the central panel we show eligible ("idonee") areas for wind turbines according to GSE (see more in Section 2.1.3 in the main text), on the right panel we show the share of municipalities which have eligible parcels in their territory according to each classification. Clearly, the "Internal Areas" appear to be more eligible compared to the other areas.

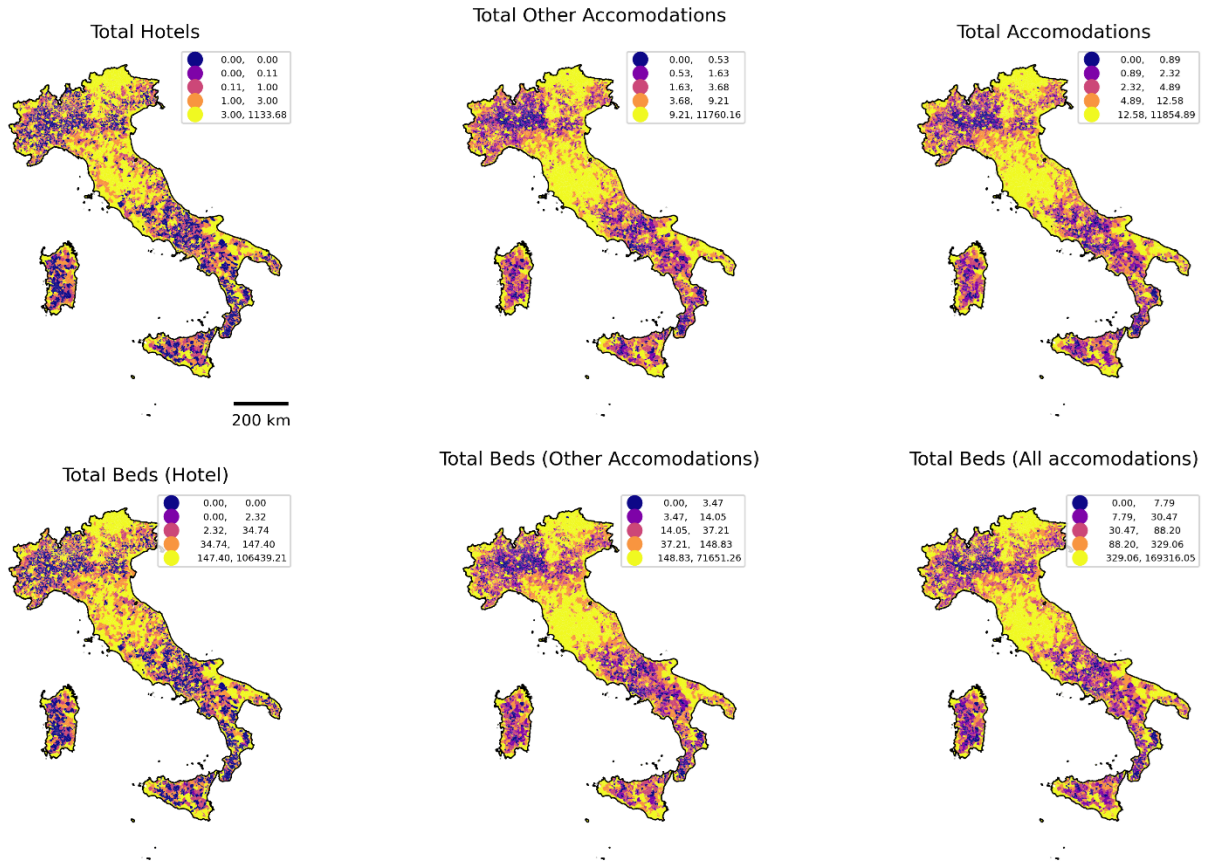


Note: The figure shows average (over all years) touristic flows per capita in all Italian Municipalities according to ISTAT and Bank of Italy data. Note that ISTAT data ranges between 2014 and 2019, while

Bank of Italy data between 1998 and 2019. White areas are missing data. Detailed descriptive statistics are provided in Table A 1 *in the appendix*

The figure shows average number of Touristic accommodations per capita in all Italian Municipalities according to ISTAT. Detailed descriptive statistics are provided in Table A 2 *in the appendix*

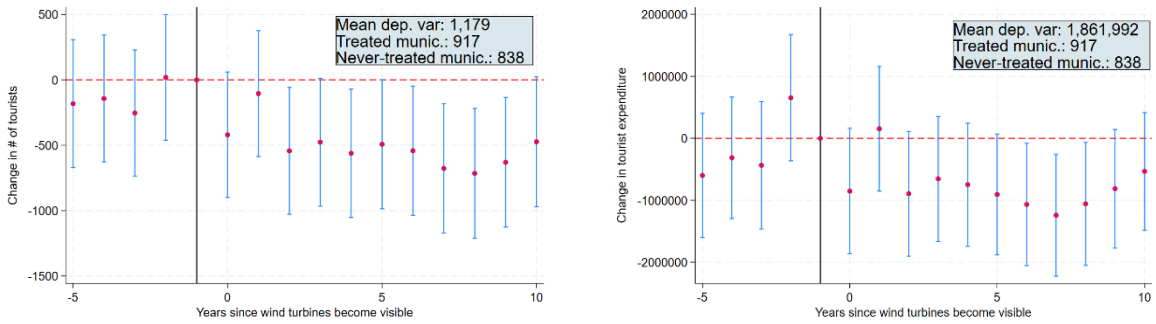
Figure SEQ Figure * ARABIC : Touristic Accomodations



9.3 Additional Results

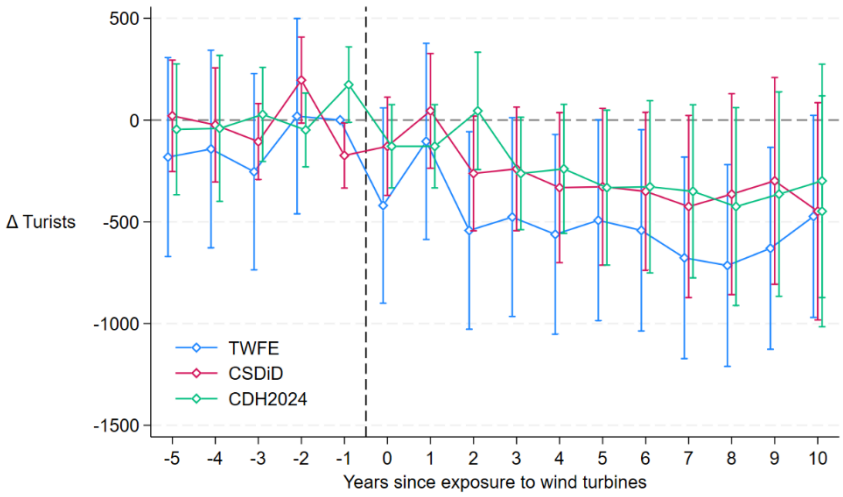
Rubustness on the main result

To assess the robustness of our main finding, we conduct a series of robustness checks. We begin by estimating the event study using the dependent variable in absolute terms, rather than relying on the Poisson specification.



Here, we observe a pattern very similar to the main results shown in Figure 7. In the left-hand panel, the visibility of wind turbines is associated with an average reduction of approximately 500 tourists per treated municipality relative to comparable matched controls. Given that the average number of tourists in our restricted sample is 1,400, this corresponds to a decline of roughly 35%. Nevertheless, the Poisson estimates provide a more reliable measure when considering relative effects (Chen & Roth, 2024). In right-hand side, toursits expenditure follow a similar patter.

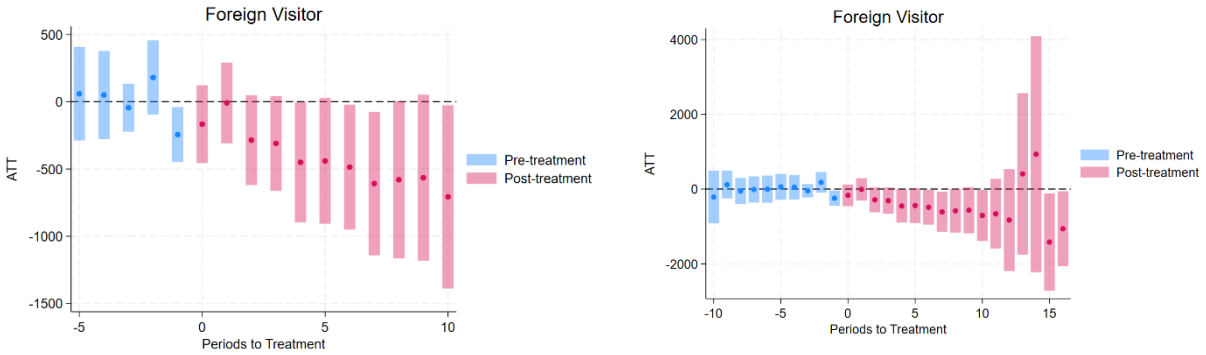
Then, we compare the traditional Two Way Fixed effect (TWFE) event study approach, with more recent estimator in literature (Callaway and Sant'Anna (2020); de Chaisemartin and D'Haultfœuille (2024), respectively called here CSDiD and CDH2024).



We show that the traditional TWFE event-study estimates are not statistically different from those obtained using the Callaway and Sant'Anna (2020, CSDiD) and de Chaisemartin and D'Haultfœuille (2024, CDH2024) estimators. This suggests that our results are not substantially affected by biases associated with the staggered timing of wind turbine implementation. Both the CSDiD and CDH2024 estimators produce slightly higher point estimates than the traditional event-

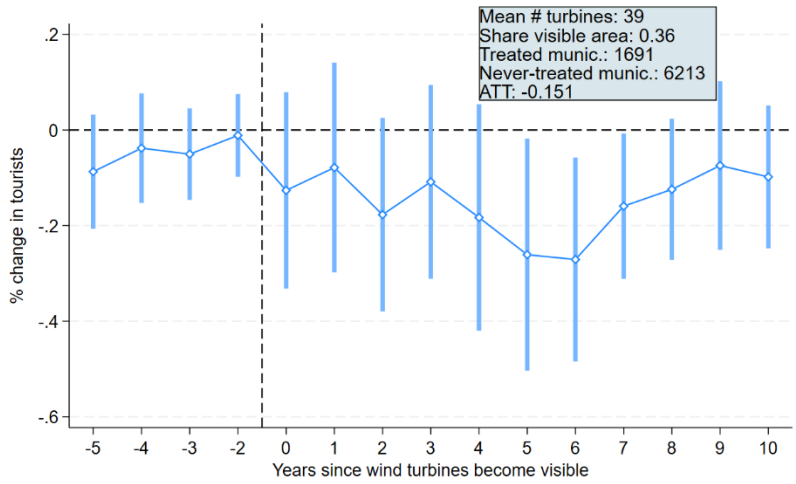
study approach, but they still capture the pronounced negative trend in tourist arrivals following turbine installation relative to the pre-treatment period. The relatively large standard errors reflect the limited sample size, so the point estimates are generally significant at the 10% level rather than at 5%.

Moreover, when controlling for log population and region fixed effects in the Callaway and Sant'Anna (2020) framework, the results closely align with the TWFE estimates:



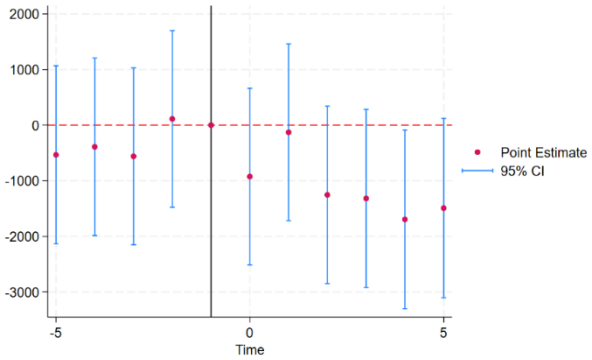
The estimated average treatment effect on the treated (ATT) is −419, statistically significant at the 5% level. Extending the analysis to include all possible pre- and post-treatment periods, right hand figure, confirms the robustness of these findings.

Here, we consider all municipalities in Italy to maximize statistical power, resulting in 1,691 treated municipalities and 6,213 never-treated municipalities.



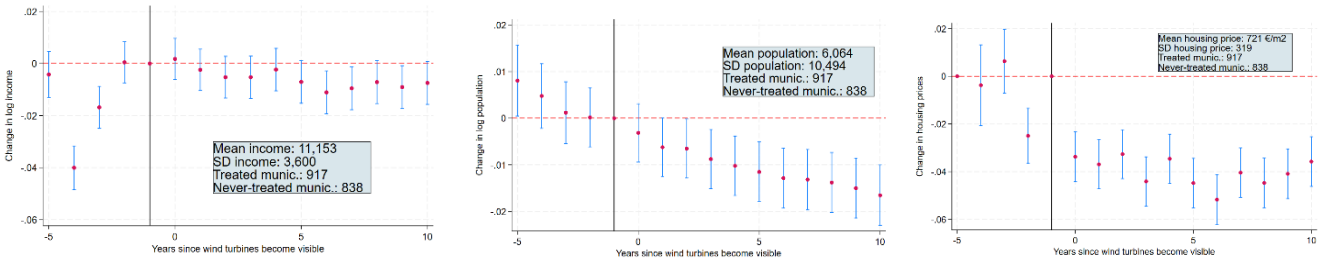
The negative effect of wind turbine visibility seems stronger in this broader sample. This likely reflects the fact that substantial differences remain between treated and never-treated municipalities, and the probability of receiving wind turbines varies across locations.

We also restrict the sample by considering only municipality with a stronger recoridng data about toursit, selecting only those municiplaity with at maxium 13 zeros (in the main analysis we select those with maximum 26 zeros).



As expected, our estimates are larger because we exclude many observations with zero values. However, in relative terms, the effect the same and the overall trajectory of the effect is unchanged. This indicates that our results are not driven by the presence of many zeros in certain municipalities.

Other outcomes: income, population, housing

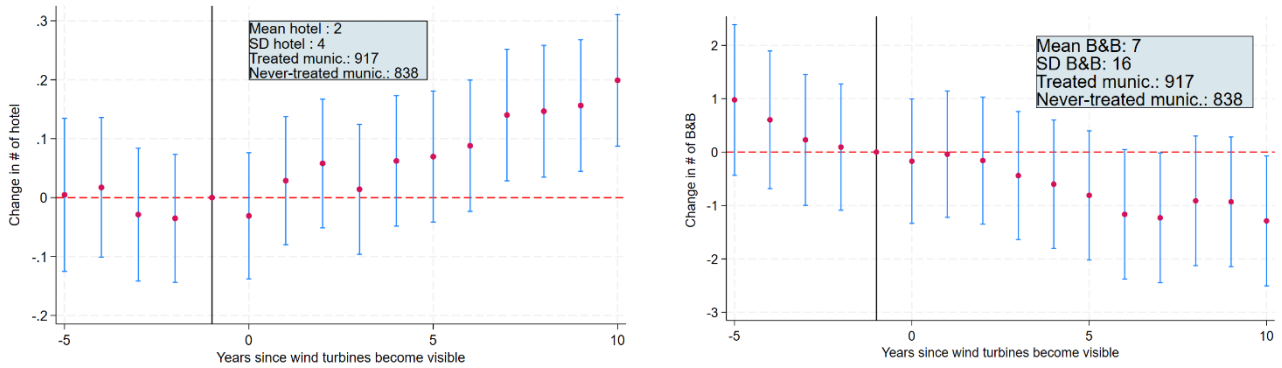


Figures illustrate the dynamic effects of wind turbine visibility on local economic outcomes. Panel (a) reports the event-study estimates for income growth, measured as the change in log income relative to 2010. The results indicate a decline of about 1% in the years following installation. Panel (b) shows similar dynamics for population, with a gradual decrease beginning even before turbine visibility, suggesting possible anticipatory migration or pre-existing demographic decline. Panel (c) presents the results for housing market outcomes, where a clear anticipation effect is observed: housing prices start to fall roughly one year prior to visibility, consistent with evidence from Dröes and Koster (2016). Overall, these figures support the main finding that wind turbine visibility is associated with localized economic and demographic contraction.

The effects of wind turbines on tourism supply

Here we show the impact of wind turbine installations on tourism supply metrics, specifically the number of hotels, bed and breakfasts (B&Bs), and private apartments. For analytical clarity, we group these into two categories: hotels and other accommodations. This distinction is crucial for capturing differences in their capacity to adapt to fluctuations in demand. Larger establishments, such as hotels,

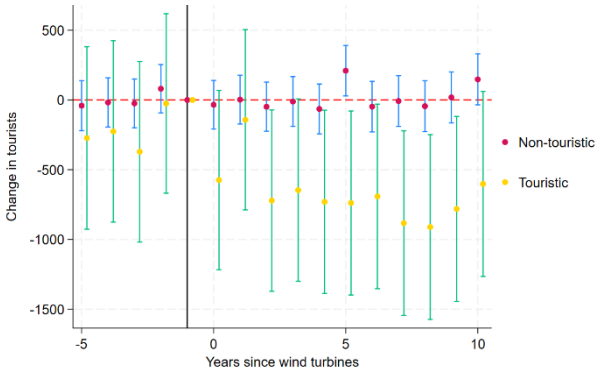
benefit from economies of scale, enabling them to absorb demand shocks by adjusting prices. In contrast, smaller establishments, such as B&Bs, may face greater challenges, often opting for closure when faced with sustained drops in demand.

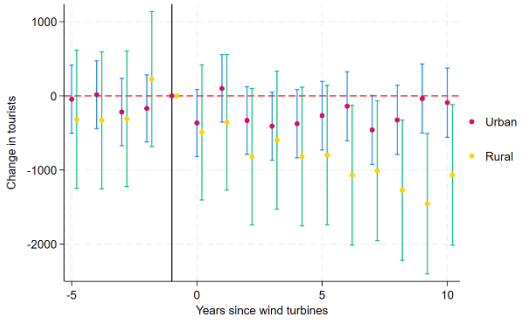


The findings reveal a decline in private accommodations, such as bed and breakfasts and short-term rentals. In contrast, we find no evidence of a negative effect on the number of hotels; if anything, their numbers slightly increase over time. This pattern may reflect a substitution effect, where the contraction in smaller, family-run accommodations is partially offset by the expansion of larger hotel establishments. Overall, these results suggest that the negative effects of wind turbines on tourism are primarily demand-driven, as the decline in tourist demand occurs more immediately after installation, while the supply side adjusts more slowly, with small decreases emerging only about 6 or 7 years after (when the overall tourist expenditure start to decrease in the main finding in figure 7).

Heterogeneous analysis by category

We begin by examining whether the impact of wind turbines differs between touristic and non-touristic municipalities (left-hand side graph). Our results show that the overall negative effect is largely driven by touristic areas, where the presence of wind turbines appears to have a stronger and more persistent impact on local economic outcomes. Next, we classify municipalities as either rural or urban (right-hand side graph). The rural category includes areas characterized by natural features such as mountains, coastlines, or lakes, while all others are classified as urban. Out of 1,755 municipalities in our sample, 766 fall into the rural category, which explains the relatively large standard errors observed in estimates for this subgroup (see Section 2 for additional descriptive statistics). Despite this, the evidence consistently points to rural municipalities as the primary drivers of the overall effect, suggesting that wind turbine visibility is particularly consequential in areas where natural landscapes are central to economic activity and local identity.





We find that main effect is driven by touristic municipality.

Next, we classify municipalities as either rural or urban. The rural category includes municipalities characterized by natural features such as mountain, maritime, or lake environments, while all others are classified as urban. Among the 1,755 municipalities in our sample, only 766 are categorized as rural, which contributes to relatively large standard errors in estimates for this subgroup (See section 2 for more descriptive statistics). Nevertheless, the point estimates indicate that the observed effects are primarily driven by rural municipalities.

Other potential mechanism

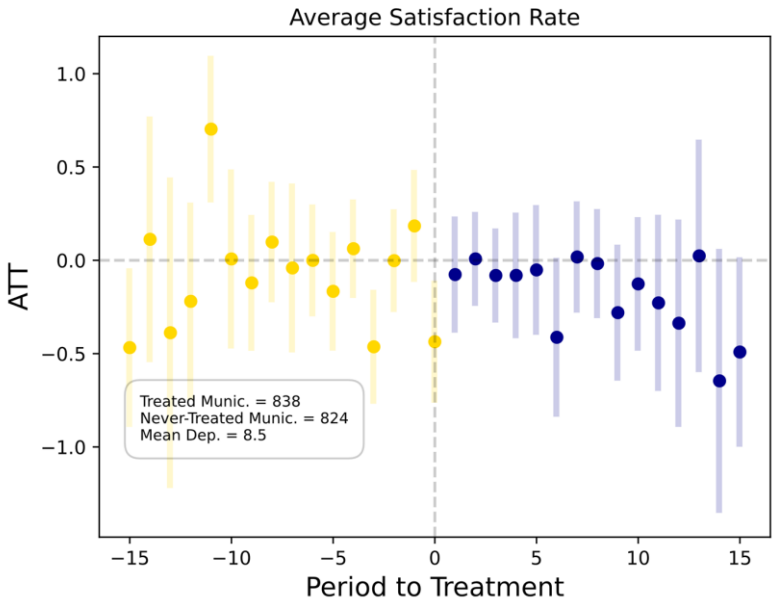
Satisfaction

To begin with, we tried to exploit the data relating to satisfaction contained in the survey of the Bank of Italy, i.e. an item that goes on a scale from 1 to 10, in which tourists are asked about their overall level of satisfaction with the holiday.[20] Results are reported in Figure 12, and appear clearly null and unsignificant, suggesting that satisfaction is not a mechanism driving the decline in the tourism demand. Tough, it is important to acknowledge that there can be some bias in this measure of reported satisfaction, in fact there is not even correlation between touristic flows and average touristic satisfactions. These results also raise important doubts regarding the interpretation and implications

[20] Specifically, the question asked is: On a scale from 1 (bad) to 10 (excellent), how much did you enjoy your visit to Italy?

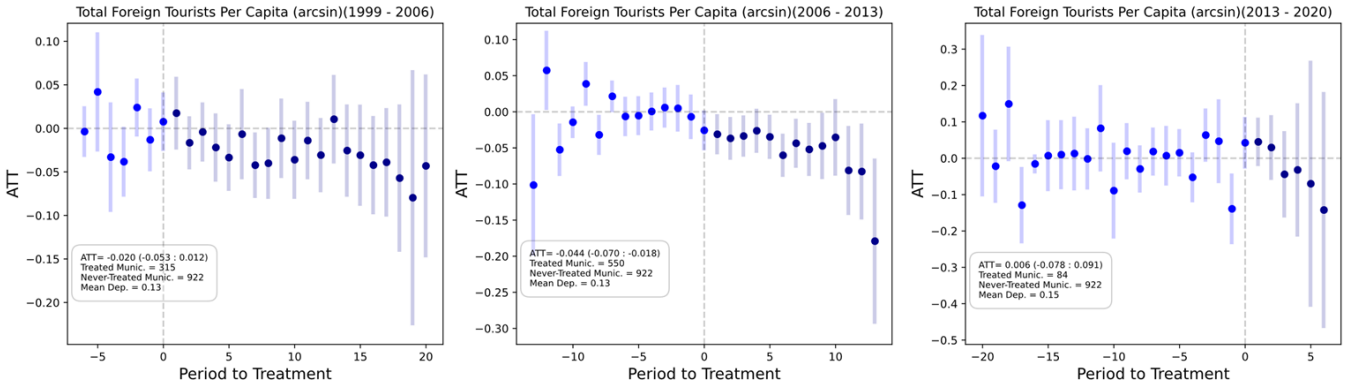
of qualitative or survey-based results suggesting that tourists decisions might not be affected by the presence of wind turbines.

Figure SEQ Figure * ARABIC 12: ATT of Wind Turbine Visibility on Tourists Satisfaction



Note: The figure shows the estimates of the ATT using the Callaway and Sant'Anna (2021) estimator. The average satisfaction rate is computed at the municipality level using the survey item contained in the Bank of Italy questionnaire.

Heterogeneous based on timing of implementation



Concluding Remarks

This dissertation has shown that environmental and climate risks have measurable economic impacts across households, firms, and local economies. Key findings include persistent housing value losses after cyclones, the nonlinear effects of floods on firm creditworthiness, and localized tourism losses offset by renewable energy benefits. Overall, the research underscores the importance of preparedness, adaptation, and policy interventions to manage both physical and perceived environmental risks.