Two Centuries of U.S. Innovation: Firms' Internal Networks and Resilience to Disasters*

Mathias S. Kruttli, Noah Stoffman, Sumudu W. Watugala†

December 2025

Abstract

Using advanced machine learning methods, we construct a comprehensive database of the universe of approximately 12 million U.S. patents from 1836 to 2023. We analyze the resilience of innovation to disaster shocks using hurricane landfall data spanning two centuries. Major hurricanes destroy local innovative capacity for up to a decade and lead to permanent losses relative to the counterfactual. Multi-location firms reallocate resources from establishments in the landfall region and increase innovation in establishments elsewhere in the aftermath of a hurricane. These positive spillovers along firms' internal networks increase aggregate innovation in counties distant from the hurricane.

JEL classifications: G30, N1, N5, O30, R11

Keywords: Innovation, disasters, extreme weather, firm internal networks, patents, LLM, big data

*We thank Juliane Begenau, Stefano Giglio, Lars Peter Hansen, Monika Piazzesi, Nora Pankratz, Brigitte Roth Tran, Derek Wenning, Dan Wilson, Mike Woeppel, Kaspar Zimmermann, and seminar participants at Indiana University, University of Melbourne, University of Sydney, University of Technology Sydney, Monash University, University of New South Wales, Deakin University, Stanford Institute for Theoretical Economics (SITE) Conference on Climate Finance and Banking, Toronto Financial Economics Conference, UNC Junior Finance Roundtable, and the Luxembourg Conference in Sustainable Finance for helpful comments. Naman Vipul Shah provided excellent research assistance.

†Kruttli: Kelley School of Business, Indiana University. Email: mkruttli@iu.edu. Stoffman: Kelley School of Business, Indiana University. Email: nstoffma@iu.edu. Watugala: Kelley School of Business, Indiana University. Email: sumudu@iu.edu.

1 Introduction

Innovation is a key driver of economic growth (Solow, 1957; Romer, 1990, 1994). The geography of innovation—its spatial distribution and clustering—plays a crucial role in how knowledge diffuses and evolves.[1] However, there is little understanding on the impact of destructive, local exogenous shocks like extreme weather events on aggregate innovation and its spatial distribution. The local vulnerability to such shocks and potential propagation across regions are important to assess the resilience of innovation and the economic costs of extreme weather events.

In this paper, we analyze the local impact of extreme weather shocks and these shocks propagate through multi-location firms. We present a conceptual framework highlighting two channels of shock propagation within firms' internal networks: (i) the capital channel and (ii) the innovation productivity channel. The first, in the style of Giroud and Mueller (2019), posits that an extreme weather shock negatively impacts the local (and hence aggregate) capital of a firm and predicts that innovation will decrease in establishments both inside and outside the disaster zone, generating negative spillovers.[2] The second channel predicts that firms will reallocate innovation activity to establishments outside the disaster zone because of higher relative productivity, generating positive spillovers. Which effect dominates is ultimately an empirical question.

To analyze these mechanisms, we first overcome two key empirical challenges. First, by harnessing recent advances in optical character recognition (OCR) and large language model (LLM) technologies, we obtain a reliable database of patent activity that captures, for the first time, a comprehensive view into inventor and assignee names, locations, and types for the universe of around 12 million U.S. patents from 1836. We harness these data to match establishments of the same firm to identify firms' internal "innovation" networks over this

[1] The externalities to the local distribution of people and firms have been long debated (Marshall, 1890; Jacobs, 1969) with implications for both regional and national place-based policy (Acemoglu, Alp, Bloom, and Kerr, 2018; Fajgelbaum and Gaubert, 2020).

[2] These spillovers can occur as corporate headquarters optimally allocate resources across different establishments (Stein, 1997; Gertner, Scharfstein, and Stein, 1994).

period. Second, we obtain a set of exogenous local shocks that impact different locations at different points in time by exploiting the quasi-exogenous spatial and temporal variation of natural disasters over a period spanning two centuries. Adapting the research design of Kruttli, Roth Tran, and Watugala (2025), we use hurricane landfalls as exogenous shocks that destroy local capital and generates reallocation. Using a long time series gives us a sample of 64 major hurricanes, which helps reduce selection bias and allows for subsample analysis.

We first analyze the direct effects of hurricanes on county-level innovation. We assign patents generated each year to counties by inventor location to obtain a county-year time series and estimate a panel version of the local projections method developed by Jordà (2005). The coefficient estimates imply large and long-lasting negative effects on patent output in counties that are within the landfall region of a hurricane. For hit counties, the annual patent output relative to the pre-period is up to ten percentage points lower than for control counties (i.e., counties outside of the landfall region) for close to a decade after landfall. The reduction in the growth rate of patents is most pronounced for counties located closest to the eye of the hurricane at landfall, which likely experience the most intense weather. Importantly, this reduction in the growth rate of patents for hit counties is a shortfall that does not appear to be recouped by a higher patent growth rate once the hit county recovers.

Next, we analyze how these local hurricane shocks to innovation can spill over to other regions in the U.S. through firms' internal firm networks, testing the predictions of our conceptual framework. Multi-location firms can have plants and establishments that span different regions in the U.S. with corporate headquarters allocating capital across establishments and adjusting to shocks. We find strong evidence of positive spillovers. Firms have a decrease in patent output in their establishments within a landfall region, but increase patent output in establishments elsewhere. The estimated effect is not only statistically but also economically significant. If half of a firm's establishments that produce patents are hit by a hurricane, the firm increases patent output in the other half of the establishments by up to

8.5%. These results are robust to including county interacted with year fixed effects, which controls for all time-varying and time-invariant county-level characteristics—like population and GDP—and allows us to compare establishments from different firms in the same county.

Interestingly, the magnitude of the effects are larger when we analyze longer horizons around the hurricane landfall. This finding is in line with agglomeration benefits and efficiency gains stemming from the within-firm capital reallocation. When a firm is forced to reallocate resources and capital across regions, they will choose the region with the largest agglomeration benefits and such benefits could be more likely to materialize at longer horizons.

These spillover effects hold for both the period from 1851 to 1975—the period for which we used OCR and LLM technologies to produce the novel patents data—and for the period from 1976 to 2023. Further, the effects are stronger when a hurricane makes landfall during a year in which the U.S. is in a recession. This result suggests that reallocating resources out of the hurricane landfall region is even more important during recessionary periods.

These spillovers via firms' internal networks are not only important at the firm-level but also at the county-level. We find that counties outside the hurricane landfall region but with links to the hurricane landfall region through local firms' internal networks experience increases in patent growth post-landfall. These results hold when including state interacted with year fixed effects, which allows for the comparison of two counties within the same state and controls for all time-varying and time-invariant state-level characteristics.

There are two major aspects to our contributions to the literature in this paper. First, by analyzing hurricane landfalls from 1851 to 2023, we show that extreme weather shocks have a large negative effect on innovation output lasting up to a decade for counties hit by the event. Further, firms' internal networks are important for the propagation of such exogenous local shocks, impacting innovation across regions. Our conceptual framework formalizes two potential channels of within-firm spillovers, the capital and innovation productivity channels, which we test empirically. Our findings highlight the importance of the

innovation productivity channel, which shows that firms can compensate for some of the lost innovation by reallocating resources out of the landfall region and increasing innovation in other counties. Second, by developing a new methodology and constructing a new, highly accurate, comprehensive database of patent information (including inventors and their locations, technological characteristics, and assignee information), we enable future research that answers important questions regarding two centuries of U.S. innovation.

[3] Firms are more likely to have greater access to capital that is also regionally diversified, with the ability to reallocate capital across internal firm networks (Giroud and Mueller, 2019). On the other hand, large firms can be less efficient in allocating resources (Scharfstein and Stein, 2000) and exhibit dampened innovation (Seru, 2014).

Our empirical findings contribute to several strands of economic research. Our analysis adds to previous studies on how local shocks can propagate across a firm's establishments. Giroud and Mueller (2019) show that for firms in non-tradable industries (e.g., retail and gastronomy), local shocks originating from collapse in house prices during the great recession lead to lower employment at a firm's establishments close to the local shock but also in other regions in which the firm operates. In contrast, we exploit multiple exogenous shocks yielding temporal and spatial variation that allow us to analyze the impact of local shocks on regional and firm-level innovation. We formalize two potential channels of generating negative and positive spillovers within firms. Our setting is unique given the benefits to agglomeration and local bias in knowledge spillovers and the constraints to reallocating the physical and human capital necessary for innovation across regions stemming from financial frictions.[4,5]

[3]Households can face significant financial constraints after experiencing major natural disasters like hurricanes. (Gallagher and Hartley, 2017; Billings, Gallagher, and Ricketts, 2022; Collier, Hartley, Keys, and Ng, 2024).

[4]Financial constraints affect access to the necessary capital and the development of innovative capacity (see, for example, Arrow, 1972; Hall and Lerner, 2010; Brown, Martinsson, and Petersen, 2012; Hsu, Tian, and Xu, 2014; Bernstein, 2015; Kerr and Nanda, 2015).

[5]Innovation and attendant knowledge spillovers can be highlight localized and sticky to place (see, for example, Jaffe, Trajtenberg, and Henderson, 1993; Audretsch and Feldman, 1996; Moretti, 2021; Atkin, Chen, and Popov, 2022).

Giroud, Liu, and Mueller (2025) find that large tech clusters have both high productivity in local inventors and distant inventors in plants linked via parent firms.[6]

We add to the literature on the economic risks associated with climatic shocks. The risks are broadly categorized into two types: physical and transition risks (Carney, 2015; Giglio, Kelly, and Stroebel, 2021). Physical risks encompass the destruction from more intense and frequent extreme weather events like hurricanes and heat waves (e.g., Hong, Li, and Xu, 2019; Kruttli, Roth Tran, and Watugala, 2025) or sea-level rise due to climate change. Transition risks refer to the risks associated with the transition to a low-carbon economy due to policies imposing a cost on firms and households (e.g., a cap-and-trade program). There exists an inherent trade-off between the two risks: imposing a price or a tax on greenhouse gas emissions increases transition risks but reduces greenhouse gas emissions and consequently future physical risks (e.g., Ivanov, Kruttli, and Watugala, 2024). Therefore, a precise estimate of the cost of physical risks is key to determining how high the price or tax on greenhouse gas emissions should be set. Beginning with the seminal work of Nordhaus (1977), there is an extensive literature on the economic cost—in the form of reduced economic growth—of greenhouse gas emissions. The empirical studies on physical climate risks and economic growth have focused on temperature shocks (e.g., Dell, Jones, and Olken, 2012 and Burke, Hsieng, and Miguel, 2015) and other natural disasters (e.g., Deryugina, 2017; Boustan, Kahn, Rhode, and Yanguas, 2020; Roth Tran and Wilson, 2024) on countries and regions. Prior work has shown that extreme weather events can propagate through (across-firm) supply-chain networks (e.g., Barrot and Sauvagnat, 2016; Pankratz and Schiller, 2024). In contrast, we focus on transmission via firms' internal networks. Importantly, a key gap in this literature is an in-depth empirical assessment of the impact of physical risks on innovation. Examining this innovation channel is particularly important given the integral role for innovation and technological change in economic growth. Noy and

[6]Other work in this area is not focused on innovation and include, for example, Cravino and Levchenko (2017); Bena, Dinc, and Erel (2022); Biermann and Huber (2024) on international shock propagtion via multinational firms.

Strobl (2023) analyze the impact of hurricanes on subsequent disaster-mitigating innovation using the patent data constructed by Petralia, Balland, and Rigby (2016). They find small, temporary increases in patents with the terms "hurricane" or "storm" in the text, but long-term reduction in general innovation. In contrast, we focus on analyzing spillovers across regions through firms' internal networks.[7]

By analyzing an extended history of the universe of U.S. patents, we complement other papers that study innovation for years not covered by the NBER patent database (Hall, Jaffe, and Trajtenberg, 2001), which starts in 1976. For example, Kogan, Papanikolaou, Seru, and Stoffman (2017) process patent information for the subset of patents issued since 1926 for which the assignee is a publicly traded firm. Kelly, Papanikolaou, Seru, and Taddy (2021) measure the novelty of historical patents using textual analysis methods but do not extract information on the inventor and assignee names and location, and rely on the text provided by Google Patents. Petralia, Balland, and Rigby (2016) also provide a database of historical patents with inventor and assignee names and location. By exploiting recent improvements in OCR and LLM technology to extract information, we are able to construct a database that appears to contain fewer errors and missing data than earlier databases.

Exploiting the full history U.S. patents together with natural disaster shocks, we shed light on factors driving the geographic distribution of U.S. innovation. We reveal the impact of *forces majeures* in determining where and why innovation thrives—with attendant spillovers and agglomeration benefits—or declines.

The remainder of the paper is organized as follows. In the next section, we present the empirical design. Section 3 describes the data, including the methodology for constructing the novel patent database covering two centuries. Section 4 presents the results. Section 5 concludes.

[7]Separately, there is a literature on climate adaptation, green technology, and innovation. See, for example, Nanda, Younge, and Fleming, 2015; Cohen, Gurun, and Nguyen, 2020; Bena, Bian, and Tang, 2023; Howell, 2024. Our analysis of exogenous natural disaster shocks over two centuries also differs from papers like Nanda and Nicholas (2014); Babina, Bernstein, and Mezzanotti (2023); Mao and Wang (2023) that analyze the impact of historical bank access or financial crises on innovation.

2 Empirical strategy

Our identification strategy relies on hurricanes as exogenous shocks to regions and is adapted from Kruttli, Roth Tran, and Watugala (2025). Hurricanes make landfall on the Atlantic and Gulf Coasts of the U.S. and the landfall region typically spans several counties in one or more states. Hurricanes have made landfall over major population and economic centers in a range of states in this region. Figure 4 plots the landfall regions of four hurricanes in our sample, which we construct by processing data on hurricane paths that are available from 1851 onwards from NOAA.

Figure 5A presents a stylized example illustrating inventor locations within counties that are exposed and unexposed to a particular hurricane landfall. We consider counties located within a hurricane landfall region as treated and counties outside of it as controls. This spatial variation in the empirical design gives us cross-sectional variation. Because counties are hit by hurricanes infrequently, the time series variation allows us to analyze within county effects.

Using the hurricane landfall regions and the location of the inventors for a given patent, we first estimate how a hurricane affects innovation activity in a county. We adapt the local projection estimator of Jordà (2005) for our panel of county-year observations.[8]

The local projection regression specification is:

$$\log\left(\frac{NPatents_{c,t+h}}{NPatents_{c,t-1}}\right) = \beta_1 Hit_{c,t} + \sum_{r=-5, r\neq 0}^{h} \rho_r Hit_{c,t+r} + \sum_{r=1}^{5} \kappa_r \log NPatents_{c,t-r} + \mu_c + \theta_t + \epsilon_{c,t}. \quad (1)$$

The dependent variable is the change in the log of the number of patents, $NPatents$, generated in a county c from the year before the hurricane hit, $t-1$, to h years after the hurricane hit, $t+h$. A patent is assigned to county c if all the inventors of the patent are located in

[8] A similar regression specification is used by Roth Tran and Wilson (2024).

that county. A patent is assigned to the year in which the patent is issued.[9] The exceptions are patents issued between January and May. These patents count for the previous year to align patents to the correct event because Atlantic Coast hurricanes only make landfall during the hurricane season from June to November.[10] The variable $Hit_{c,t}$ is an indicator variable that takes the value one when a county c is in the landfall region of a hurricane in year t and zero otherwise. In addition to county and year fixed effects, we include several control variables. First, we control for a county being hit by hurricanes that make landfall during the preceding 5-year period from $t-5$ to $t-1$. We also control for a county being hit by hurricanes that make landfall during the years from $t+1$ to $t+h$. These controls are important to account for the possibility of staggered and multiple treatment of counties (De Chaisemartin and d'Haultfoeuille, 2020; Athey and Imbens, 2022; Baker, Larcker, and Wang, 2022; Dube, Girardi, Jorda, and Taylor, 2023). Second, we capture potential pre-trends in a county's innovation activity by including the log of $NPatents$ lagged up to five years.[11] The standard errors are double clustered at the county and year levels.

To estimate how a firm's internal innovation network reacts when some of the firm's establishments are in a hurricane landfall region, we consider changes in innovation at the firm-county-year level. For many firms innovation is geographically dispersed. Therefore, the question arises that if one innovation location experiences a hurricane landfall, will the firm move innovation to a different location or reduce innovation across its establishments.

We employ an event study regression specification where each hurricane enters as a separate event:

$$\log\left(\frac{\overline{\text{NPatents}}_{i,c,t:t+h-1}}{\overline{\text{NPatents}}_{i,c,t-h:t-1}}\right) = \beta_1 \text{Hit}_{c,t} + \beta_2 \sum_{k \neq c} w_{i,k,t-h:t-1} \text{Hit}_{k,t} + \mu_c + \theta_t + \varepsilon_{i,c,t+h} \qquad (2)$$

[9] The results are qualitatively similar when using the application instead of the issue date. For patents before 1900, no application date is recorded, and thus, we set the application date to three months before the issue date for these patents.

[10] Omitting this adjustment leads to qualitatively similar results.

[11] The estimates are qualitatively similar when including additional lags. Including additional lags does not materially improve the fit of the regression model.

The dependent variable is the change in the average number of patents of firm i in county c over h years after a hurricane landfall relative to the average number of patents over h years before landfall. The variable h is set to one, three, five, and ten years, respectively. Having windows of multiple years around the landfall accounts for the fact that firms do not generate a patent in a location every year. The variable $w_{i,k,t-h:t-1}$ is the share of firm innovation in county c over h years before the hurricane hit. The variable $\sum_{k \neq c} w_{i,k,t-h:t-1}\text{Hit}_{k,t}$ measures the share of a firm's innovation locations other than county c that are in the hurricane landfall region. For simplicity, we will denote this variable as $HitOther_{i,t}$. The weight is given by

$$w_{i,k,t-h:t-1} = \frac{\text{NPatents}_{i,k,t-h:t-1}}{\text{NPatents}_{i,t-h:t-1}}. \tag{3}$$

Therefore, the variable $HitOther_{i,t}$ ranges from zero to close to but not including one.

To control for unobservable variables we include county and year fixed effects. Further, we interact county and year fixed effects, which allows us to compare changes in innovation of one firm's establishment to changes in innovation of another firm's establishment in the same county and year. These fixed effects control, for example, for differences in economic growth across regions. The standard errors are double clustered by firm and year. To ensure that a firm hit by one hurricane does not enter as a control for a different hurricane, we exclude firms that had more then 50% of their innovation establishments hit from the controls for ten years before and after the hurricane.

3 Data

3.1 Novel patents database

> *[The Congress shall have Power . . .] To promote the Progress of Science and useful Arts, by securing for limited Times to Authors and Inventors the exclusive Right to their respective Writings and Discoveries.*

—U.S. Constitution, art. I, §8, cl. 8.

To facilitate studying the longest possible time series of hurricanes, we construct a new patent database. The data include information on grant date and inventor location (city and state) spanning 1836 through 2023. For data beginning in 1976, we obtain data from PatentsView, which is maintained by the Office of the Chief Economist at the USPTO. For data from 1836 through 1975, we extract data directly from the PDFs of scanned patent documents on the USPTO website by taking advantage of modern OCR algorithms and ChatGPT, as we explain below. Additional detail is provided in the Appendix.

Modern patent documents include the title and abstract of the invention; application and grant dates; the name(s) and locations of residence for all inventors and any assignees; a detailed description of the invention; drawings of the invention; and the claims, which define the scope of legal protection provided by the patent. An inventor must be a person, but patents can be assigned to a corporation or another person. Generally, when an invention is developed in a company, the inventor will assign the patent to their employer. Beginning in February 1947, patent documents began consistently including citations to existing patents as references to "prior art" (Nicholas, 2010). Older patents typically do not contain all of this information, although the grant date and inventor location have always been included.

While the layout of patent documents has changed over time, the required information has mostly remained constant. Since 1976, the USPTO has digitally recorded newly-granted patents and made machine-readable files publicly available. For the period prior to that, information must be extracted from scanned documents, an example of which is shown in Figure 1.

The launch of Google Patents in 2006 made it easier to search the text of the pre-1976 historical patent documents. These data serve as the source for several influential papers, including Kogan, Papanikolaou, Seru, and Stoffman (2017), who created a widely used database, and Kelly, Papanikolaou, Seru, and Taddy (2021). However, the quality of the text extracted from scanned documents using Optical Character Recognition (OCR)

available in Google Patents is quite variable, especially for patents granted before 1950. The importance of the errors in these data depends greatly on how the text is being used. For our analysis, correctly identifying the inventor's location is of paramount importance, so correctly capturing this information is critical.

As an example of the OCR errors that can be found in the Google data, Figure 2 shows how the text for the patent in Figure 1 is rendered by Google (left column). Note that this text does not include anything that resembles the name of the inventor, his city of residence, or the title of the invention. The entire first paragraph of the text is missing, as is the beginning of the next paragraph. Although not all of the OCR in Google Patents is of this poor quality, many hundreds of thousands of documents have significant degradation in the extracted text. Even the use of exceptionally flexible textual analysis techniques cannot overcome the "garbage-in-garbage-out" problems associated with such poor OCR quality.

OCR algorithms have vastly improved in recent years (Correia and Luck, 2023) but the text available from Google Patents has not been updated. Advances in computer vision have given OCR algorithms the ability to identify parts of documents, patterns in text layout, and words in ways that far surpass what is available in Google Patents. Moreover, ChatGPT is a game changer when it comes to extracting information from text. Even in the presence of OCR errors, spelling mistakes, typos, or missing information, it can draw on its vast training data to answer carefully-crafted questions about the (intended) meaning of text.

As a point of comparison with our novel data, we compared the locations of inventors in our data with those in the "HistPat" data created by Petralia, Balland, and Rigby (2016). Of the 3.9 million unique patents between 1836 and 1975, HistPat is missing over 600,000 patents, which appears to be mainly due to the exclusion of foreign inventors. Of the remaining 3.3 million patents, just over 400,000 have a discrepancy between inventor locations; this appears to be because the assignee's location is often misreported as the inventor's location in the HistPat data. Despite the requirement that all patents have an inventor, another 115,000 patents in HistPat have an assignee but no inventor. We also observe significant

time series variation in differences between the two datasets. For example, the rate of missing patents in HistPat relative to our dataset is below 10% through 1880 before jumping to the teens for most of the period between 1880 and 1965. This missing rate increases to 19%, 24%, and 50% in 1965, 1970, and 1975, respectively.

Using our new inventor location data, aggregated to the county level, Figure 3 illustrates the expected westward shift in innovation over the last two hundred years.

3.2 Firms' internal networks

The raw patent data provide no identifier that links assignees across patents. An important part of our analysis relies on our ability to identify which patents are assigned to which firms over time—and to know that the *same* firm has been assigned different patents. We accomplish this disambiguation by applying several machine learning techniques to the assignee names.

We begin by standardizing company names and extracting unique company names along with corresponding patent grant years. Our basic approach is to apply clustering techniques to these names, but we first extensively preprocess the text, by standardizing special characters, merging of single-letter sequences, and normalizing legal entity abbreviations such as Inc. or Corp. We then identify the 1,000 most common words across names to create a list of low-information tokens. We apply fuzzy matching techniques to correct misspelled legal entities, and correct inconsistent spacing, which can originate either in the original PDF document or from OCR transcription errors.

After producing this relativley clean list of unique company names, we divide the data into overlapping 40-year time windows for pairwise similarity scoring. We again apply a fuzzy matching algorithm to identify name similarities above a threshold that we determine from the data. (The pairwise nature of this procedure makes the computation extremely computationally intensive, but we are able to exploit parallel processing to complete the work in a reasonable time.) We use the *UnionFind* algorithm to merge related company

names into clusters, ensuring continuity across overlapping windows, so firms like IBM that are continually granted patents over long spans of time are identified as one firm. To refine these clusters and address false positives, we use hybrid TF–IDF cosine similarity metrics to evaluate coherence, with low-quality clusters re-split using an agglomerative clustering algorithm. Finally, we used ChatGPT to review any large clusters (groups where at least 15 variations of a name are identified as belonging to one firm). This allows us to ensure semantic consistency and correct any remaining misgroupings, resulting in a thoroughly cleaned set of company name clusters.

3.3 Two centuries of hurricane landfalls

We obtain the hurricane path data for all Atlantic and Gulf Coast hurricanes making landfall from 1851 to 2023 from NOAA's Atlantic HURDAT2 database, which covers all known tropical cyclones and subtropical cyclones and is a part of the Re-analysis Project (Landsea and Franklin, 2013). These data record the latitude and longitude of the eye of a tropical cyclone at least every 6 hours before it dissipates. This allows us to calculate hurricane paths and landfall dates and times. We use these data to calculate the list of counties that are within 100 and 200 miles of the landfall location of the eye of a hurricane. These radii line up with NOAA's measurements on the average distance from the hurricane eye within which wind speeds cause damage to infrastructure. Figure 4 shows the landfall regions for four hurricanes in our sample.

For our main analysis, we focus on the set of deadliest tropical storms with more than 25 fatalities during our sample period. We construct this set starting with NOAA's report on "The deadliest, costliest, and most intense United States tropical cyclones from 1851 to 2010 (and other frequently requested hurricane facts)" that covers the 1851-2010 period (Blake, Landsea, and Gibney, 2011). We manually augment these data for the more recent years to cover the period from 1851 to 2023. The list of 64 deadliest hurricanes in our sample is shown in Table B.3. We focus on this set of storms for two key reasons. First, these are

likely the hurricanes with the most reliable data in the early period of our sample. Second, these are likely the hurricanes that significantly impacted human and physical capital in the landfall regions. The number of fatalities is one reliable measure of the destructive impact of a hurricane, which does not rely on as many assumptions as, potentially, the monetary value of damages.

Table 1 presents the summary statistics for the main variables used in our empirical analysis.

4 Results

This section presents the results of our analyses. First, we discuss the baseline effects of hurricanes on local innovation and spillovers of these local innovation shock through firm's internal networks. Then we discuss robustness tests, effects of financing constraints, and county-level spillovers.

4.1 Baseline effects

4.1.1 Local effects

We first estimate the regression in equation (1) to test if a hurricane hit adversely affects innovation in that county, how large the effects are, and whether the effects are transitory or permanent. When identifying hit counties, that is, the counties for which the variable $Hit_{c,t}$ is one, we measure landfall regions based on two radii (100 and 200 miles) around the eye of the hurricane. We estimate the regression separately for each radius. For estimations that rely on the 100 mile radius, we exclude from the control group any counties located within the 200-mile but not the 100-mile landfall region.

Figure 6 plots the coefficient estimate of the variable $Hit_{c,t}$ for different horizons h. The negative effects on patents do not appear immediately after landfall. The drop in the number of patents generated in a hit county, relative to the year before landfall, begins three years

after landfall. However, the effect intensifies as time passes and the decrease in innovation is most pronounced between five and eight years after the hurricane landfall. Ten years after landfall, the coefficient estimates become statistically insignificant.

The effects are larger for counties closer to the eye of the hurricane at landfall. The magnitude of the coefficients for the 100-mile radius reaches as low as -0.069 in seven years after landfall. This magnitude implies that the number of patents in a hit county is 6.9% lower in the seventh year after the hurricane hit than in the year before the hurricane hit. For the 200-mile radius, the coefficient magnitude decreases but still reaches -0.044. The lower coefficient magnitude is consistent with the idea that hurricanes are more destructive closer to the eye of the storm.

The negative and significant coefficient estimates for several years after landfall show that a county in a hurricane landfall region experiences a substantial shortfall in patents.

Table B.1 in the appendix reports the estimates of the coefficients in equation (1). The first row of the table shows the coefficient on the independent variable of interest, $Hit_{c,t}$, which is plotted in Figure 6. The coefficient estimates on the lagged patent variables are highly significant. The coefficient on the first lag is negative with a magnitude between -0.67 and -0.78. This estimate implies that the time series of county-level patents is no longer explosive after differencing the dependent variable. However, the positive and significant coefficient estimates on the additional patent lags suggest that the undifferenced time series is a unit root process. We confirm the unit root in county-level patents with an augmented Dickey-Fuller test (unreported) for all U.S. patents from 1936 to 2023. Therefore, differencing the dependent variable is necessary when estimating regression models with county-level time series of patents. We also estimate an event study regression as an alternative to the local projection model. The results are qualitatively similar and presented in Table B.2.

4.1.2 Spillovers via firms' internal networks

Based on the results that hurricanes constitute large negative effects on local innovation, we investigate how such shocks affect innovation output in firms' internal networks. On the one hand negative shock to one location might lead to firms reducing R&D investment in other regions to better absorb the financial shock. This mechanism would be in line with Giroud and Mueller (2019), who document that local shocks to some of a firm's establishments lead to a decrease in employment in other establishments of the firm. On the other hand, firms might be able to move some of their innovation capacity to counties outside the hurricane landfall region and prevent the loss of innovation at least to some extent.

The estimates of the regression model in equation (2) are presented in Table 2. The coefficient estimates of the variable $Hit_{c,t}$ are negative and significant. These estimates imply that when a county is in a hurricane landfall region, a firms' patent output is up to 4% lower than for counties not in the landfall region. These estimates are in line with the local innovation impact discussed in Section 4.1.1.

Interestingly, the coefficient estimate of the variable $HitOther_{i,c,t}$ is positive and strongly significant across all specifications. These results show that firms increase innovation in other counties when part of their operations are hit by a hurricane. Across the different specifications, the coefficient estimates range from 0.05 to 0.17. These magnitudes imply that a firm with nearly all of its establishments in a landfall region will increase innovation in the remaining establishments by 5% to 17%.

The magnitude of the coefficient estimates is larger for longer horizons around the hurricane landfall. For both landfall region radii, the specification with the largest estimate is the one where we compare ten years after to ten years before landfall. The larger increase at longer horizons could be explained by agglomeration benefits and efficiency gains that come with the relocation. When a firm is forced to allocate resources and capital to other regions, they will choose the region where the agglomeration benefits are largest and such agglomeration benefits could be more likely to materialize at longer horizons.

Panel B shows the results when including county times year fixed effects. Including the interacted fixed effects strengthens the identification. We can now compare the change in innovation of two establishments of two different firms in the same county. The fact that the coefficient estimates remain strongly significant and largely unchanged in magnitude confirms that our results hold after controlling for time-varying unobservable county characteristics as, for example, economic growth.[12]

Based on these regression specifications, it is unclear if firms simply reallocate resources within counties in the hurricane landfall region. Table 3 presents the results for the specification with the county times year fixed effects, but we exclude counties in the hurricane landfall region from the sample. The magnitude and significance of the coefficient estimates remains qualitatively the same. These results confirm that hurricane shocks create spillovers through firms' internal networks, and this leads to innovation growth in firm establishments outside the landfall region.

4.2 Effects across time

The documented results are for our total sample of 173 years. A natural question is whether this baseline effect in firms' internal networks changed through our sample period. Particularly, are the results different in the patents data up to 1975 that we generated compared to the patents data from USPTO that start in 1976?

The first regression we estimate includes an indicator term that takes the value one for hurricanes that make landfall between 1851 and 1975 and zero otherwise. We interact this indicator variable with our variable of interest $HitOther_{i,c,t}$. The estimates are reported in Table 4.

The magnitude of the coefficient estimate on the unconditional variable $HitOther_{i,c,t}$ remains strongly significant and positive for all specifications. The magnitude decreases slightly

[12]County-level GDP growth is only available from the 1950s.

compared to the estimates in Table 2. However, the coefficient estimates for $HitOther_{i,c,t} \times I_{t \in 1851-1975}$ are also positive, strongly significant, and larger in magnitude.

These estimates show that the mechanism we document is present across both patent datasets. However, the magnitude of the effects are larger in the sample up to 1975. A potential reason could be that infrastructure became more resilient and rebuilding efforts quicker, and therefore, the need to move innovation activity into other regions became somewhat mitigated. Improved and faster rebuilding efforts could have been partly driven by the Federal government. The Federal Disaster Relief Act of 1950 started efforts for centrally organized and improved disaster relief.

The second analysis on variation of the effects along the time-series dimension focuses on U.S. recessions. Recessions increase the financial constraints of firms. Therefore, firms might be more likely to reallocate innovation activity when financial constraints are tight, and they do not have the capital to rebuild in the hurricane landfall region.

To test this hypothesis, we estimate the same regression as presented in Table 4, but we replace the time indicator variable with an indicator variable that takes the value of one if the year of the hurricane landfall is in a recession year. To identify recession years and use the NBER recession data to identify years for which at least six months were in a recession.

The results are reported in Table 5. The unconditional regression estimate remains positive and strongly significant, but the spillovers within firms' internal networks are more pronounced when a hurricane makes landfall in a recession year. The coefficient estimates for $HitOther_{i,c,t} \times I_{t \in Recession}$ are always positive and mostly strongly significant.

4.3 Extensive margin

In the regression in equation (2), the dependent variable is the log change in the average annual number of patents of a firm in a county from h years pre- to h years post-hurricane landfall. Using the percent difference as the dependent variable ensures that the variable is well behaved despite large differences in the number of patents across firms, and we control

for unobservable firm-county level characteristics. However, the estimates do not capture if a firm completely stops to innovate in a county in the post-hurricane landfall period or starts innovating in a county in which it had no innovation in the pre-hurricane landfall period.

To address the first case, we estimate a logit regression where the dependent variable takes the value of one if the firm innovates in a county in the post-hurricane landfall period and zero otherwise. The regression specification is given by

$$\begin{aligned} logit(P(\overline{\text{NPatents}}_{i,c,t:t+h-1} > 0 | \overline{\text{NPatents}}_{i,c,t-h:t-1} > 0)) = & \beta_1 \sum_{k \neq c} w_{i,k,t-h:t-1} \text{Hit}_{k,t} \\ & + \psi_{c,t} + \varepsilon_{i,c,t+h}. \end{aligned} \tag{4}$$

Table 6 presents the estimates. A firm for which other counties are in a hurricane landfall region is much more likely to keep an existing establishment open. The coefficient estimate for $HitOther_{i,c,t}$ is strongly significant and positive for all specifications. The coefficient magnitude is between 0.4 and 0.5, which implies a marginal effect between 8% and 12%. Therefore, the probability of a firm to keep innovation activity in a county is 8% to 12% higher if the $HitOther_{i,c,t}$ variable increases by one unit. Compared to the baseline estimates in Table 2, the magnitude of the coefficient estimates do not increase for longer horizons. This finding is intuitive, as the dependent variable in the logit regression is capped at one and does not measure agglomeration benefits.

To address the second case—firms with exposure to a hurricane are more likely to start generating patents in counties in which they did not innovate before—we estimate a firm-year-level regression given by

$$\frac{NrNewCounties_{i,t:t+h-1}}{NrTotalCounties_{i,t-h:t-1}} = \beta_1 SharePatentsHit_{i,t} + \theta_t + \varepsilon_{i,t+h}. \tag{5}$$

The dependent variable measures the number of counties in which the firm generated patents in the post-landfall but not the pre-landfall period relative to the total number of coun-

ties in which the firm generated patents in the pre-landfall period. We include a year fixed effects and the standard errors are double clusterd by year and firm. The variable $SharePatentsHit_{i,t}$ measures the share of a firms' innovation location in the landfall region of the hurricane:

$$\text{Hit}_{i,t} = \frac{\sum_c \left(\text{NPatents}_{i,c,t-h:t-1} \times \text{Hit}_{c,t} \right)}{\sum_c \text{NPatents}_{i,c,t-h:t-1}}. \tag{6}$$

Table 7 reports the results. The coefficient estimate of $SharePatentsHit_{i,t}$ is positive for all eight specifications and strongly significant for six of the eight specifications. The magnitude of the coefficient estimate goes up to 0.04, which suggests that a firm that had all of its establishments in the landfall region adds 4% new counties to its innovation locations. This number is relatively small and suggests that while firms do shift resources and innovation activity into new counties post-hurricane landfall, our baseline regression specification in equation (2) captures the major effects of the internal network spillovers.

4.4 Firm-level substitution

Our results raise the question of whether firms can perfectly substitute any impact from a hurricane on their innovation output by moving resources to other locations. To test this hypothesis of perfect substitution, we estimate a firm-level regression that is similar to equation (5):

$$\log \left(\frac{\overline{\text{NPatents}}_{i,t:t+h-1}}{\overline{\text{NPatents}}_{i,t-h:t-1}} \right) = \beta_1 SharePatentsHit_{i,t} + \theta_t + \varepsilon_{i,t+h}. \tag{7}$$

The dependent variable is the change in the average annual number of patents pre- and post-hurricane. We again include a year fixed effects and double cluster the standard errors by year and firm.

In Table 8, the results show that firms can almost perfectly substitute the effects from a hurricane. The coefficient estimates are negative and significant for some specification but the magnitude is small. The magnitude ranges from 0 to -0.05, which implies that a firm that has all of its establishments hit by a hurricane experiences a decline in patents of only up to 5%.

4.5 County-level spillovers

When firms reallocate resources in their internal networks in response to a hurricane shock, we could observe changes in innovation output at the aggregate county-level. In other words, if a county is located outside of the hurricane landfall region but firms with establishments in this county also have have establishments in the hurricane landfall region, does the county see an increase in innovation?

To test such county-level spillovers we estimate the regression model given by:

$$\log\left(\frac{\overline{\text{NPatents}}_{c,t:t+h}}{\text{NPatents}_{c,t-1}}\right) = \beta_1 CountyHitOther_{c,t} + \mu_c + \phi_{s,t} + \epsilon_{c,t+h}. \tag{8}$$

The dependent variable is the change in the average annual number of patents in county c over h years after the hurricane landfall relative to the number of patents in the year before landfall. We include county and state interacted with year fixed effects. The standard errors are clustered by county and year. Counties in the hurricane landfall region are excluded from the regression. The independent variable $CountyHitOther_{c,t}$ measures the share of patents for firms with establishments in county c that are hit by a hurricane in counties other than county c:

$$CountyHitOther_{c,t} = \frac{\sum_{i \in c} \sum_{k \neq c} \text{NPatentsFirm}_{i,k,t-5:t-1} \times \text{Hit}_{k,t}}{\sum_{i \in c} \sum_{k} \text{NPatentsFirm}_{i,k,t-5:t-1}} \tag{9}$$

The results are presented in Table 9. Panel A shows the estimates with county and year fixed effects. The estimates are positive for all specifications. The magnitude of the estimates are increasing for longer horizons h and are strongly significant for all specifications where h is larger than one year. The magnitude increases from 0.05 at the one year horizon to 0.18 at the ten year horizon. This increase is in line with the increasing estimates at the firm-county level (e.g., Table 2) and suggest that agglomeration benefits also occur at the county-level.

Panel B includes state interacted with year fixed effects in the regression model. These fixed effects allow us to compare counties within the same state but different exposure to the hurricane landfall region through firms' internal networks. The results are largely unchanged in terms of magnitude and significance.

5 Conclusion

Innovation supports continued economic growth and is an essential driver of an economy maintaining a competitive edge. In this paper, we analyze the impacts on innovation from quasi-exogenous local disasters. To do so, we first construct a comprehensive, nearly-error-free database of all U.S. patents from 1836 to 2023 using advanced OCR and LLM technologies. We then analyze how hurricanes affect local innovation output and lead to spillovers across the U.S. through firms' internal networks.

We find that major hurricanes destroy local innovative capacity for up to a decade following landfall and lead to permanent counterfactual losses for the landfall region. Firms that produce patents in the region before landfall, reallocate resources out of the landfall region and increase their innovation output in counties outside of the landfall region. This finding contrasts with the existing literature, which finds that for firms in non-tradable industries regional cash flow shocks propagate in a firm's internal networks and lead to reduction in employment across a firm's establishments.

Our results are robust to different time-periods and become stronger during recessions. Further, the spillovers along the firms' internal networks lead to increases in aggregate innovation output for counties located outside of the hurricane landfall region.

Harnessing almost two centuries of data, we reveal a mechanism that can generate divergence in regional economic prosperity. This channel is important for local and national policymakers to factor in when choosing policies to promote economic growth and prosperity, especially those focused on protecting and promoting innovative activity. Our results on how the capital destruction and constraints induced by local natural disaster shocks have long-lasting effects on innovation are of particular importance to current debates regarding how to ensure the competitiveness of the U.S. economy while keeping it resilient to unexpected disasters and shocks. We quantify the impact of *forces majeures* in determining where and why innovation thrives—with attendant agglomeration benefits and local spillovers—or declines.

Further, if extreme weather events like hurricanes become more damaging in the future given trends in demographics, construction, infrastructure, or the environment, academics and policymakers alike will need to better understand the impact of extreme weather events on innovation to comprehensively assess the economic costs of disaster shocks. Our findings will be crucial for policymakers who seek to balance the costs and benefits of adaptation and resilience to natural disasters.

References

Acemoglu, Daron, Harun Alp, Nicholas Bloom, and William R. Kerr, 2018, Innovation, reallocation, and growth, *American Economic Review* 108, 3450–3491.

Arrow, Kenneth Joseph, 1972, *Economic welfare and the allocation of resources for invention* (Springer).

Athey, Susan, and Guido W Imbens, 2022, Design-based analysis in difference-in-differences settings with staggered adoption, *Journal of Econometrics* 226, 62–79.

Atkin, David, M Keith Chen, and Anton Popov, 2022, The returns to face-to-face interactions: Knowledge spillovers in silicon valley, *NBER Working Paper*.

Audretsch, David B, and Maryann P Feldman, 1996, R&d spillovers and the geography of innovation and production, *American Economic Review* 86, 630–640.

Babina, Tania, Asaf Bernstein, and Filippo Mezzanotti, 2023, Financial disruptions and the organization of innovation: Evidence from the great depression, *Review of Financial Studies* 36, 4271–4317.

Baker, Andrew C, David F Larcker, and Charles CY Wang, 2022, How much should we trust staggered difference-in-differences estimates?, *Journal of Financial Economics* 144, 370–395.

Barrot, Jean-Noël, and Julien Sauvagnat, 2016, Input specificity and the propagation of idiosyncratic shocks in production networks, *The Quarterly Journal of Economics* 131, 1543–1592.

Bena, Jan, Bo Bian, and Huan Tang, 2023, Financing the global shift to electric mobility, Working paper.

Bena, Jan, Serdar Dinc, and Isil Erel, 2022, The international propagation of economic downturns through multinational companies: The real economy channel, *Journal of Financial Economics* 146, 277–304.

Bernstein, Shai, 2015, Does going public affect innovation?, *Journal of Finance* 70, 1365–1403.

Biermann, Marcus, and Kilian Huber, 2024, Tracing the international transmission of a crisis through multinational firms, *The Journal of Finance* 79, 1789–1829.

Billings, Stephen B, Emily A Gallagher, and Lowell Ricketts, 2022, Let the rich be flooded: the distribution of financial aid and distress after hurricane harvey, *Journal of Financial Economics* 146, 797–819.

Blake, Eric S, Chris Landsea, and Ethan J Gibney, 2011, The deadliest, costliest, and most intense United States tropical cyclones from 1851 to 2010 (and other frequently requested hurricane facts), NOAA technical memorandum NWS NHC-6 National Weather Service, National Hurricane Center.

Boustan, Leah Platt, Matthew E. Kahn, Paul W. Rhode, and Maria Lucia Yanguas, 2020, The effect of natural disasters on economic activity in us counties: A century of data, *Journal of Urban Economics* 118, 103257.

Brown, James R, Gustav Martinsson, and Bruce C Petersen, 2012, Do financing constraints matter for R&D?, *European Economic Review* 56, 1512–1529.

Burke, Marshall, Solomon Hsieng, and Edward Miguel, 2015, Global non-linear effect of temperature on economic production, *Nature* 527, 235–239.

Carney, Mark, 2015, Breaking the tragedy of the horizon - climate change and financial stability, *Speech given at Lloyd's of London (29 September)*.

Cohen, Lauren, Umit G Gurun, and Quoc H Nguyen, 2020, The esg-innovation disconnect: Evidence from green patenting, NBER working paper.

Collier, Benjamin L, Daniel A Hartley, Benjamin J Keys, and Jing Xian Ng, 2024, Credit when you need it, NBER working paper.

Correia, Sergio, and Stephan Luck, 2023, Digitizing historical balance sheet data: A practitioner's guide, *Explorations in Economic History* 87, 101475.

Cravino, Javier, and Andrei A Levchenko, 2017, Multinational firms and international business cycle transmission, *The Quarterly Journal of Economics* 132, 921–962.

De Chaisemartin, Clément, and Xavier d'Haultfoeuille, 2020, Two-way fixed effects estimators with heterogeneous treatment effects, *American Economic Review* 110, 2964–2996.

Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken, 2012, Temperature shocks and economic growth: Evidence from the last half century, *American Economic Journal: Macroeconomics* 4, 66–95.

Deryugina, Tatyana, 2017, The fiscal cost of hurricanes: Disaster aid versus social insurance, *American Economic Journal: Economic Policy* 9, 168–198.

Dube, Arindrajit, Daniele Girardi, Oscar Jorda, and Alan M Taylor, 2023, A local projections approach to difference-in-differences event studies, NBER Working Paper.

Fajgelbaum, Pablo D., and Cecile Gaubert, 2020, Optimal spatial policies, geography, and sorting, *The Quarterly Journal of Economics* 135, 959–1036.

Gallagher, Justin, and Daniel Hartley, 2017, Household finance after a natural disaster: The case of hurricane katrina, *American Economic Journal: Economic Policy* 9, 199–228.

Gertner, Robert H., David S. Scharfstein, and Jeremy C. Stein, 1994, Internal versus external capital markets, *Quarterly Journal of Economics* 109, 1211–1230.

Giglio, Stefano, Bryan Kelly, and Johannes Stroebel, 2021, Climate finance, *Annual Review of Financial Economics* 13, 15–36.

Giroud, Xavier, Ernest Liu, and Holger Mueller, 2025, Innovation spillovers across u.s. tech clusters, .

Giroud, Xavier, and Holger M Mueller, 2019, Firms' internal networks and local economic shocks, *American Economic Review* 109, 3617–3649.

Hall, Bronwyn H, Adam B Jaffe, and Manuel Trajtenberg, 2001, The NBER patent citation data file: Lessons, insights and methodological tools, NBER working paper.

Hall, Bronwyn H, and Josh Lerner, 2010, The financing of R&D and innovation, in *Handbook of the Economics of Innovation*, vol. 1 . pp. 609–639 (Elsevier).

Hong, Harrison, Frank Weikai Li, and Jiangmin Xu, 2019, Climate risks and market efficiency, *Journal of Econometrics* 208, 265–281.

Howell, Sabrina T, 2024, Government intervention in innovation, *Annual Review of Financial Economics* 16.

Hsu, Po-Hsuan, Xuan Tian, and Yan Xu, 2014, Financial development and innovation: Cross-country evidence, *Journal of Financial Economics* 112, 116–135.

Ivanov, Ivan T., Mathias S. Kruttli, and Sumudu W. Watugala, 2024, Banking on carbon: Corporate lending and cap-and-trade policy, *Review of Financial Studies* 37, 1640–1684.

Jacobs, Jane, 1969, *The Economy of Cities* (Random House: New York).

Jaffe, Adam B, Manuel Trajtenberg, and Rebecca Henderson, 1993, Geographic localization of knowledge spillovers as evidenced by patent citations, *The Quarterly journal of Economics* 108, 577–598.

Jordà, Òscar, 2005, Estimation and inference of impulse responses by local projections, *American Economic Review* 95, 161–182.

Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy, 2021, Measuring technological innovation over the long run, *American Economic Review: Insights* 3, 303–320.

Kerr, William R, and Ramana Nanda, 2015, Financing innovation, *Annual Review of Financial Economics* 7, 445–462.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman, 2017, Technological innovation, resource allocation, and growth, *Quarterly Journal of Economics* 132, 665–712.

Kojima, Takeshi, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa, 2024, Large language models are zero-shot reasoners, in *Proceedings of the 36th International Conference on Neural Information Processing Systems* NIPS '22 Red Hook, NY, USA. Curran Associates Inc.

Kruttli, Mathias S., Brigitte Roth Tran, and Sumudu W. Watugala, 2025, Pricing poseidon: Extreme weather uncertainty and firm return dynamics, *Journal of Finance* 80, 783–832.

Landsea, Christopher W, and James L Franklin, 2013, Atlantic hurricane database uncertainty and presentation of a new database format, *Monthly Weather Review* 141, 3576–3592.

Mao, Yifei, and Jessie Jiaxu Wang, 2023, Access to finance and technological innovation: Evidence from pre-Civil War America, *Journal of Financial and Quantitative Analysis* 58, 1973=2023.

Marshall, Alfred, 1890, *Principles of Economics* (Macmillan: London).

Moretti, Enrico, 2021, The effect of high-tech clusters on the productivity of top inventors, *American Economic Review* 111, 3328–3375.

Nanda, Ramana, and Tom Nicholas, 2014, Did bank distress stifle innovation during the great depression?, *Journal of Financial Economics* 114, 273–292.

Nanda, Ramana, Ken Younge, and Lee Fleming, 2015, Innovation and entrepreneurship in renewable energy, *The changing frontier: Rethinking science and innovation policy* 199.

Nicholas, Tom, 2010, The role of independent invention in U.S. technological development, 1880-1930, *Journal of Economic History* 70, 57–82.

Nordhaus, William D., 1977, Economic growth and climate: The carbon dioxide problem, *American Economic Review* 67, 341–346.

Noy, Ilan, and Eric Strobl, 2023, Creatively destructive hurricanes: Do disasters spark innovation, *Environmental and Resource Economics* 84, 1–17.

Pankratz, Nora MC, and Christoph M Schiller, 2024, Climate change and adaptation in global supply-chain networks, *The Review of Financial Studies* 37, 1729–1777.

Petralia, Sergio, Pierre-Alexandre Balland, and David L. Rigby, 2016, Unveiling the geography of historical patents in the United States from 1836 to 1975, *Scientific Data* 3, 160074.

Romer, Paul M., 1990, Endogenous technological change, *Journal of Political Economy* 98, 71–102.

——— , 1994, The origins of endogenous growth, *Journal of Economic Perspectives* 8, 3–22.

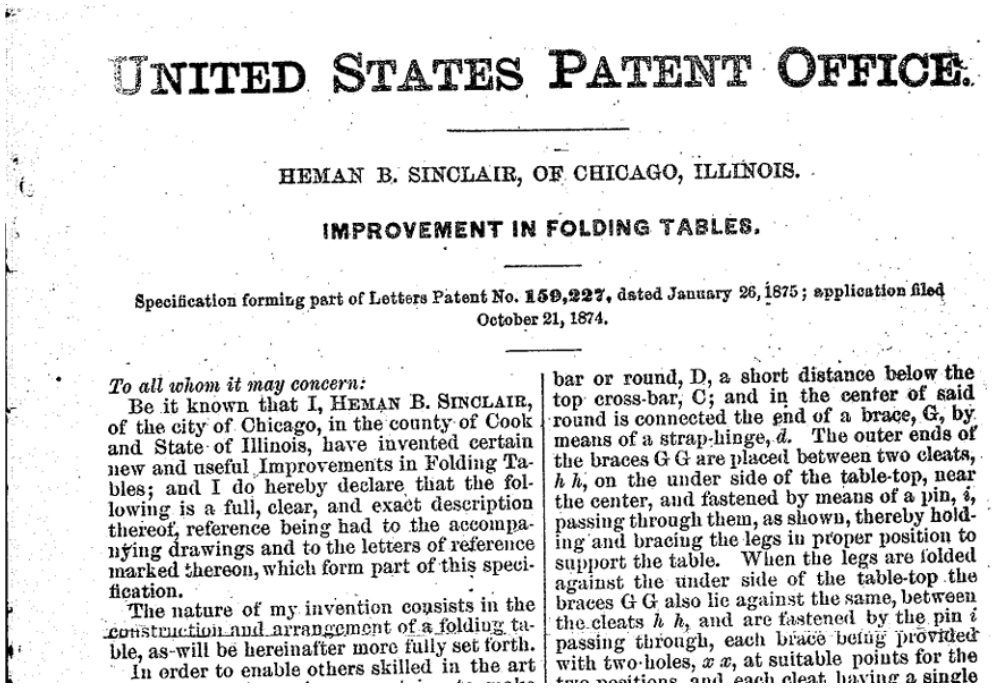
Roth Tran, Brigitte, and Daniel J Wilson, 2024, The local economic impact of natural disasters, Federal Reserve Bank of San Francisco working paper.

Scharfstein, David S, and Jeremy C Stein, 2000, The dark side of internal capital markets: Divisional rent-seeking and inefficient investment, *The Journal of Finance* 55, 2537–2564.

Seru, Amit, 2014, Firm boundaries matter: Evidence from conglomerates and R&D activity, *Journal of Financial Economics* 111, 381–405.

Solow, Robert M, 1957, Technical change and the aggregate production function, *Review of Economics and Statistics* 39, 312–320.

Stein, Jeremy C., 1997, Internal capital markets and the competition for corporate resources, *Journal of Finance* 52, 111–133.



UNITED STATES PATENT OFFICE.

HEMAN B. SINCLAIR, OF CHICAGO, ILLINOIS.

IMPROVEMENT IN FOLDING TABLES.

Specification forming part of Letters Patent No. **159,227**, dated January 26, 1875; application filed October 21, 1874.

To all whom it may concern:

Be it known that I, HEMAN B. SINCLAIR, of the city of Chicago, in the county of Cook and State of Illinois, have invented certain new and useful Improvements in Folding Tables; and I do hereby declare that the following is a full, clear, and exact description thereof, reference being had to the accompanying drawings and to the letters of reference marked thereon, which form part of this specification.

The nature of my invention consists in the construction and arrangement of a folding table, as will be hereinafter more fully set forth.

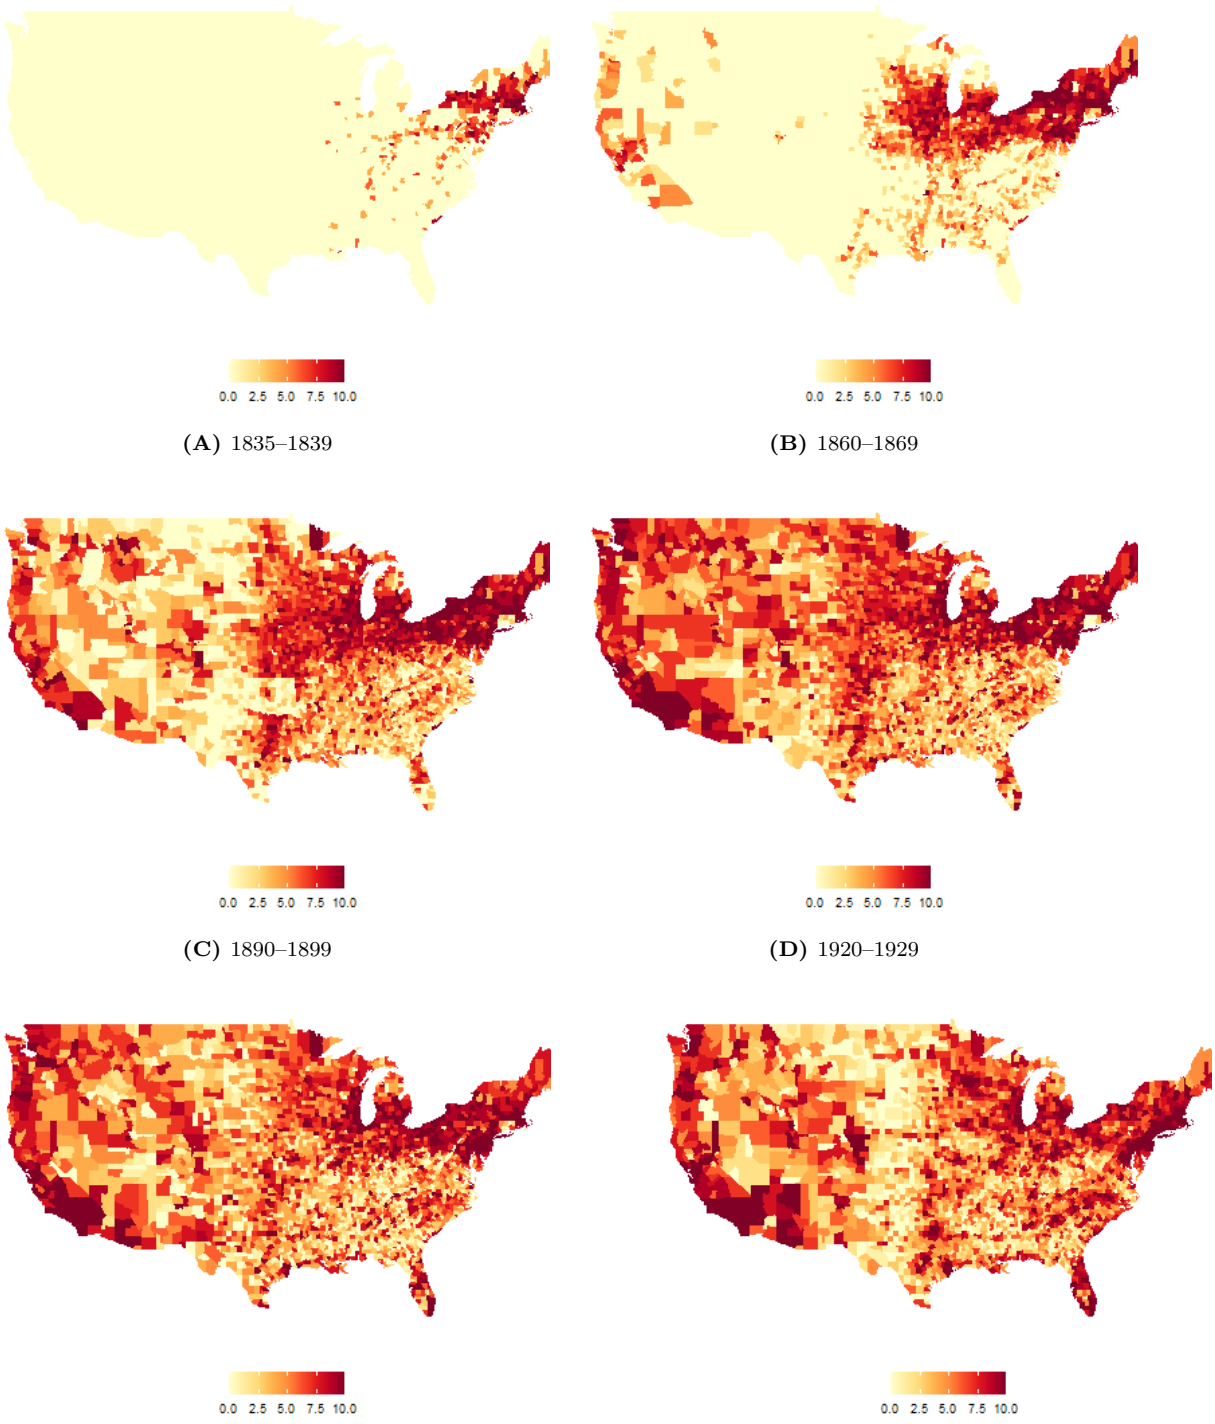
In order to enable others skilled in the art

bar or round, D, a short distance below the top cross-bar, C; and in the center of said round is connected the end of a brace, G, by means of a strap-hinge, *d*. The outer ends of the braces G G are placed between two cleats, *h h*, on the under side of the table-top, near the center, and fastened by means of a pin, *i*, passing through them, as shown, thereby holding and bracing the legs in proper position to support the table. When the legs are folded against the under side of the table-top the braces G G also lie against the same, between the cleats *h h*, and are fastened by the pin *i* passing through, each brace being provided with two holes, *x x*, at suitable points for the

Figure 1: Example USPTO scan of a patent (number 159,227)

Google Patents OCR	Our New OCR
iNITED,	UNITED STATES PATENT OFFICE.
' Bnl-uns n. srNoLAin, on, ourense, immers.	HEMAN B. SINCLAIR, OF CHICAGO, ILLINOIS.
inn-Psovsmsnr ns Forense tastes.	IMPROVEMENT IN FOLDING TABLES.
Specification formingpart of Letters Patent	Specification forming part of Letters Patent
No. 159,227. dated January 26, li875;	No. 159,227, dated January 26, 1875
application-filed October 21, 1874. f f	application filed October 21, 1874.
runstnnctionjindrrangement ofraQldiUE. tn-	To all whom it may concern:
	Be it known that I, HEMAN B. SINCLAIR,
	of the city of Chicago, in the county of Cook
	and State of Illinois, have invented certain
	new and useful Improvements in Folding Ta-
	bles; and I do hereby declare that the fol-
	lowing is a full, clear, and exact description
	thereof, reference being had to the accompa-
	nying drawings and to the letters of reference
	marked thereon, which form part of this speci-
	fication.
	The nature of my invention consists in the
	construction.and arrangement of a folding ta-
ble, as-will be hereinafter more fully set forth.	ble, as-will be hereinafter more fully set forth.

Figure 2: Comparison of OCR rendering of patent number 159,227



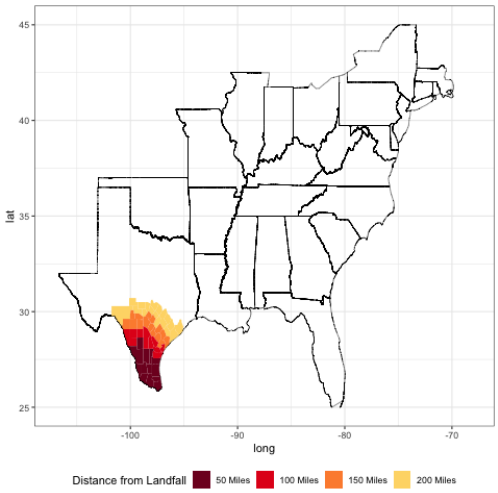
(A) 1835–1839 **(B)** 1860–1869

(C) 1890–1899 **(D)** 1920–1929

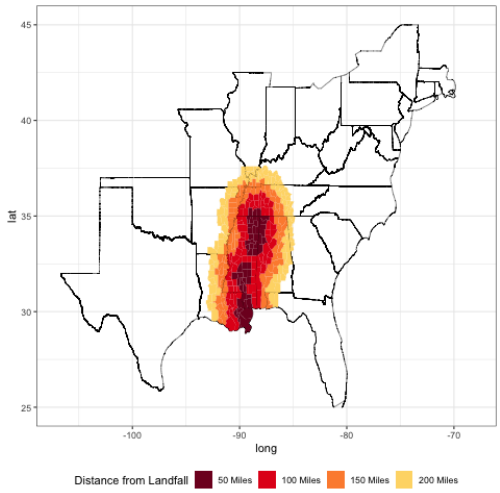
(E) 1950–1959 **(F)** 2010–2019

Figure 3: Locations of inventors

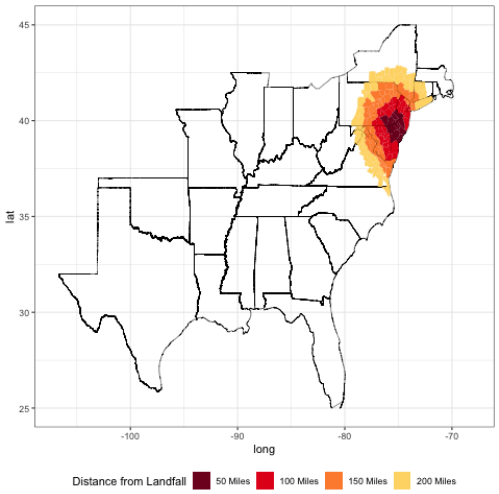
This figure shows the location of inventors who have patents granted during different periods over the last two centuries. Inventor locations are aggregated to counties that are then sorted into deciles in each period based on the number of inventors residing in each county. The darker the shade, the greater the number of active inventors in a county.



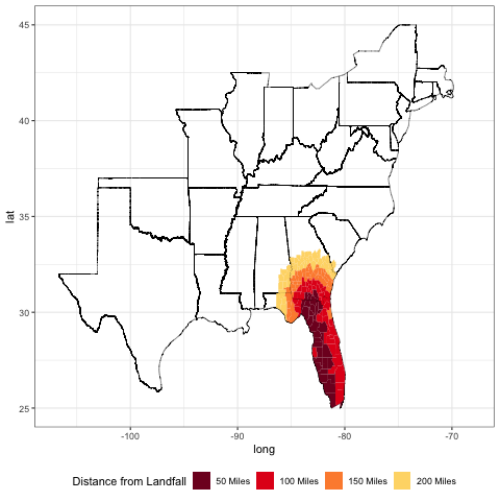
(A) 1999 Bret



(B) 2005 Katrina



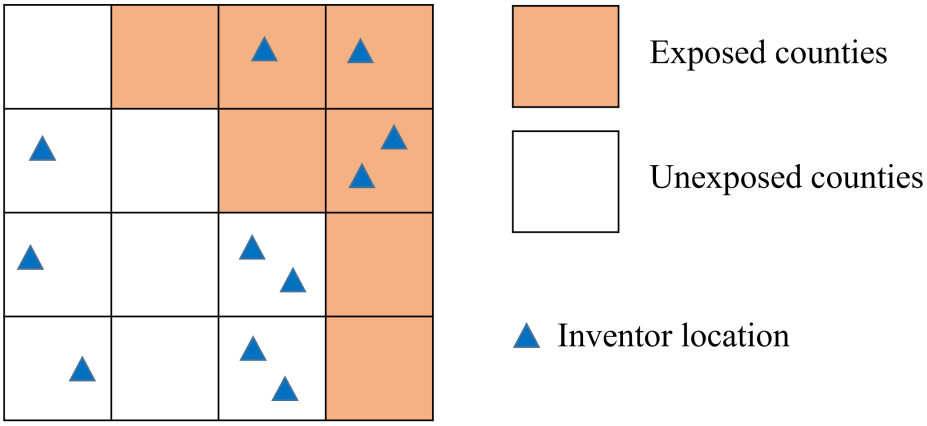
(C) 2012 Sandy



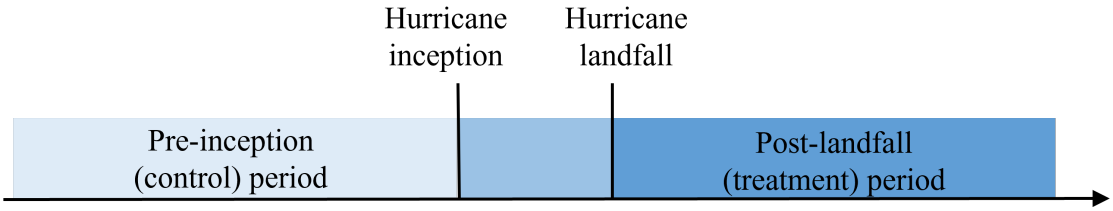
(D) 2017 Irma

Figure 4: Counties within a hurricane landfall region

This figure shows the counties within 50, 100, 150, and 200 miles of the hurricane eye for four hurricanes in our sample from 1851-2023.



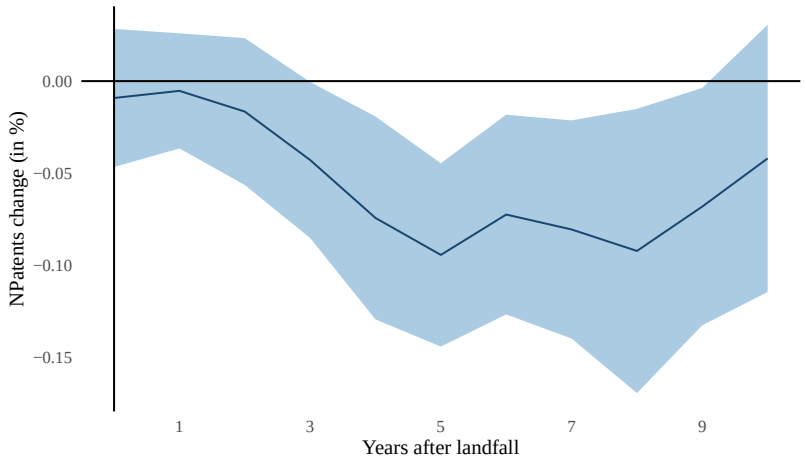
(A) Stylized example of spatial exposure (across variation)



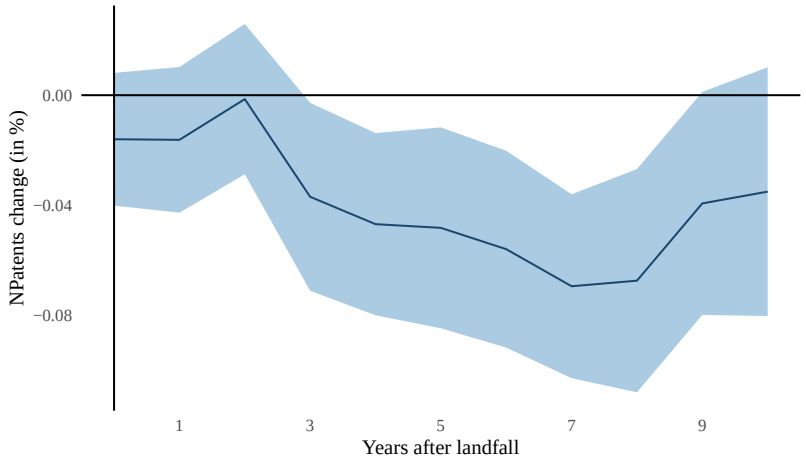
(B) The timeline of a hurricane (within variation)

Figure 5: Identification strategy

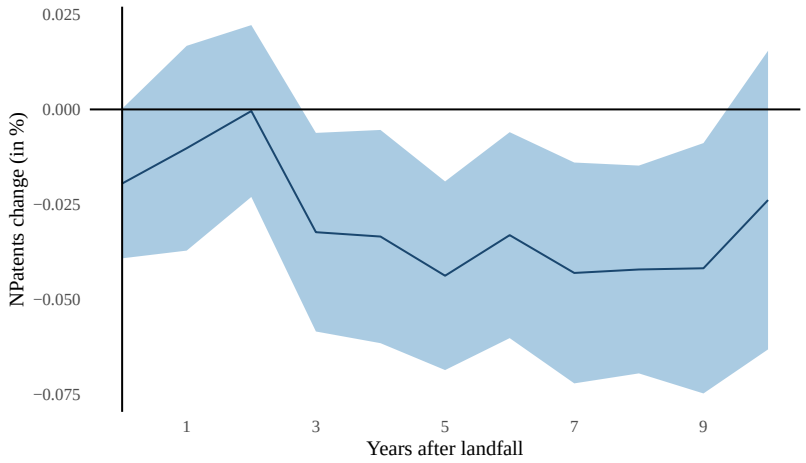
Panel A illustrates a stylized example of inventor locations and county exposure to a hurricane landfall region. Panel B illustrates the timeline of a hurricane.



(A) Counties within 50-mile radius



(B) Counties within 100-mile radius



(C) Counties within 200-mile radius

Figure 6: Change in local innovation following landfall

This figure presents the β_t estimate and its 95% confidence interval from estimating the local projection

Table 1: Patent Summary Statistics

This table shows the summary statistics for the main variables used in the paper. Panel A presents the data from 1851 to 2023 used in the county-year level local projection model given in equation (1). Panel B presents the summary statistics for the firm-county-year level event study regression given in equation (2). Panel C reports the summary statistics for the firm-level event study regression given in equation (7). Panel B and C show summary statistics for the landfall region based on a 100-mile radius around the eye of the hurricane and five years pre- and post-hurricane landfall. All variables are described in Appendix Table B.4.

Panel A: County-Year Level

	Observations	Avg.	St. dev.	10^{th}	25^{th}	50^{th}	75^{th}	90^{th}
$NPatents_{c,t}$	514,753	19.337	185.338	0.000	0.000	1.000	4.000	17.000
$\Delta log(NPatents_{c,t})$	226,493	0.014	0.633	-0.693	-0.336	0.000	0.375	0.693

Panel B: Firm-County-Year Level

$\overline{NPatents}_{i,c,t:t+4}$	394,053	2.035	12.915	0.200	0.200	0.400	1.200	3.200
$\overline{NPatents}_{i,c,t-5:t-1}$	394,053	1.721	10.520	0.200	0.200	0.400	1.000	2.600
$log(\overline{NPatents}_{i,c,t:t+4}/\overline{NPatents}_{i,c,t-5:t-1})$	394,053	0.079	0.961	-1.099	-0.560	0.000	0.693	1.322
$Hit_{c,t}$	394,053	0.093	0.290	0.000	0.000	0.000	0.000	0.000
$HitOther_{i,c,t}$	394,053	0.051	0.162	0.000	0.000	0.000	0.001	0.100

Panel C: Firm-Year Level

$\overline{NPatents}_{i,t:t+4}$	184,163	5.030	55.539	0.200	0.400	0.600	1.800	5.400
$\overline{NPatents}_{i,t-5:t-1}$	184,163	4.181	45.370	0.200	0.200	0.600	1.600	4.600
$log(\overline{NPatents}_{i,t:t+4}/\overline{NPatents}_{i,c,t-5:t-1})$	184,163	0.095	1.044	-1.099	-0.651	0.000	0.693	1.386
$SharePatentsHit_{i,t}$	184,163	0.082	0.248	0.000	0.000	0.000	0.000	0.227

Table 2: Spillovers via Firms' Internal Networks

This table presents results from estimating the regression specification given in equation (2). The dependent variable $\log\left(\frac{\overline{NPatents}_{i,c,t:t+h-1}}{\overline{NPatents}_{i,c,t-h:t-1}}\right)$ is at the firm-county-year level. The numerator (denominator) is the annual average number of patents over h years post (pre) hurricane landfall. The first independent variable is an indicator variable that takes a value of 1 if a county was in the landfall region of a hurricane in a given year ($Hit_{c,t}$). The dependent variable ($HitOther_{i,c,t}$) measures the share of a firm's patents in counties other than county c that were in the hurricane landfall region. The data span from 1851 to 2023. The specifications include year and county fixed effects (Panel A) and year interacted with county fixed effects (Panel B). Standard errors are double clustered by year and firm and shown in parentheses. The significance of each coefficient estimate is indicated by * for $p < 0.10$, ** for $p < 0.05$, and *** for $p < 0.01$.

Panel A: With County and Year Fixed Effects

Landfall radius	100				200			
Window length in years (h)	1	3	5	10	1	3	5	10
$Hit_{c,t}$	-0.024***	-0.028***	-0.031***	-0.027	-0.023***	-0.030***	-0.038***	-0.039*
	(0.006)	(0.008)	(0.011)	(0.019)	(0.005)	(0.008)	(0.012)	(0.022)
$HitOther_{i,c,t}$	0.054***	0.076***	0.101***	0.162***	0.058***	0.064***	0.100***	0.172***
	(0.019)	(0.021)	(0.025)	(0.040)	(0.015)	(0.019)	(0.022)	(0.036)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	232,955	406,105	451,265	384,621	210,375	407,701	409,129	346,958
R^2	0.012	0.015	0.017	0.023	0.013	0.015	0.017	0.024

Panel B: With County×Year Fixed Effects

$HitOther_{i,c,t}$	0.059***	0.072***	0.095***	0.154***	0.063***	0.069***	0.104***	0.168***
	(0.020)	(0.022)	(0.027)	(0.043)	(0.016)	(0.020)	(0.023)	(0.038)
County × Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	232,955	406,105	451,265	384,621	210,375	407,701	409,129	346,958
R^2	0.094	0.071	0.072	0.083	0.104	0.075	0.074	0.086

Table 3: Spillovers via Firms' Internal Networks (Excluding Hit Counties)

This table presents results from estimating the regression specification given in equation (2) but excluding hit counties from the sample. The dependent variable $\log\left(\frac{\overline{NPatents}_{i,c,t:t+h-1}}{\overline{NPatents}_{i,c,t-h:t-1}}\right)$ is at the firm-county-year level. The numerator (denominator) is the annual average number of patents over h years post (pre) hurricane landfall. The dependent variable ($HitOther_{i,c,t}$) measures the share of a firm's patents in counties other than county c that were in the hurricane landfall region. The data span from 1851 to 2023. The specifications include county times year fixed effects. Standard errors are double clustered by year and firm and shown in parentheses. The significance of each coefficient estimate is indicated by * for $p < 0.10$, ** for $p < 0.05$, and *** for $p < 0.01$.

Landfall radius	100				200			
Window length in years (h)	1	3	5	10	1	3	5	10
$HitOther_{i,c,t}$	0.054**	0.066**	0.087**	0.124**	0.052**	0.051*	0.087***	0.128***
	(0.023)	(0.032)	(0.035)	(0.047)	(0.022)	(0.028)	(0.028)	(0.041)
County × Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	192,735	369,023	357,472	295,421	185,009	355,942	346,952	284,358
R^2	0.100	0.072	0.073	0.086	0.102	0.074	0.075	0.089

Table 4: Spillovers via Firms' Internal Networks (Pre-1976 Period)

This table presents results from estimating the regression specification given in equation (2) but interacting the independent variable with an indicator that takes the value of one for the years from 1851 to 1975 and zero otherwise. The dependent variable $\log\left(\frac{\overline{NPatents}_{i,c,t:t+h-1}}{\overline{NPatents}_{i,c,t-h:t-1}}\right)$ is at the firm-county-year level. The numerator (denominator) is the annual average number of patents over h years post (pre) hurricane landfall. The dependent variable ($HitOther_{i,c,t}$) measures the share of a firm's patents in counties other than county c that were in the hurricane landfall region. The data span from 1851 to 2023. The specifications include county times year fixed effects. Standard errors are double clustered by year and firm and shown in parentheses. The significance of each coefficient estimate is indicated by * for $p < 0.10$, ** for $p < 0.05$, and *** for $p < 0.01$.

Landfall radius	100				200			
Window length in years (h)	1	3	5	10	1	3	5	10
$HitOther_{i,c,t}$	0.041**	0.059**	0.070**	0.111**	0.049***	0.053**	0.079***	0.115***
	(0.018)	(0.025)	(0.029)	(0.048)	(0.017)	(0.023)	(0.024)	(0.038)
$HitOther_{i,c,t} \times I_{t\in 1851-1975}$	0.128***	0.079**	0.134***	0.177***	0.095***	0.086**	0.127***	0.191***
	(0.025)	(0.036)	(0.036)	(0.058)	(0.033)	(0.035)	(0.039)	(0.054)
County × Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	207,623	398,787	394,053	333,573	210,375	407,701	409,129	346,958
R^2	0.010	0.071	0.072	0.082	0.104	0.074	0.074	0.085

Table 5: Spillovers via Firms' Internal Networks (during US Recessions)

This table presents results from estimating the regression specification given in equation (2) but interacting the independent variable with an indicator that takes the value of one if the year was a US recession year and zero otherwise. A year is classified as a recession year if at least six months were in a recession based on NBER data. The dependent variable $\log\left(\frac{\overline{NPatents}_{i,c,t:t+h-1}}{\overline{NPatents}_{i,c,t-h:t-1}}\right)$ is at the firm-county-year level. The numerator (denominator) is the annual average number of patents over h years post (pre) hurricane landfall. The dependent variable ($HitOther_{i,c,t}$) measures the share of a firm's patents in counties other than county c that were in the hurricane landfall region. The data span from 1851 to 2023. The specifications include county times year fixed effects. Standard errors are double clustered by year and firm and shown in parentheses. The significance of each coefficient estimate is indicated by * for $p < 0.10$, ** for $p < 0.05$, and *** for $p < 0.01$.

Landfall radius	100				200			
Window length in years (h)	1	3	5	10	1	3	5	10
$HitOther_{i,c,t}$	0.049**	0.067***	0.084***	0.135***	0.059***	0.063***	0.095***	0.146***
	(0.019)	(0.023)	(0.028)	(0.045)	(0.016)	(0.021)	(0.023)	(0.036)
$HitOther_{i,c,t} \times I_{t \in Recession}$	0.110***	0.064	0.102***	0.147***	0.069*	0.084***	0.121***	0.199***
	(0.032)	(0.047)	(0.033)	(0.042)	(0.035)	(0.029)	(0.023)	(0.033)
County × Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	207,623	398,697	394,029	333,595	210,375	407,645	409,171	346,958
R^2	0.099	0.072	0.072	0.083	0.104	0.075	0.075	0.086

Table 6: Probability of Firm Keeping Existing Innovation Location

This table presents results from estimating the logit regression specification given in equation (4). The dependent variable takes a value of one if the firm has patents in a county over h years post-hurricane landfall and zero otherwise. The sample includes only firm-county-year observations where the number of patents was positive over h years pre-hurricane. The dependent variable ($HitOther_{i,c,t}$) measures the share of a firm's patents in counties other than county c that were in the hurricane landfall region. The data span from 1851 to 2023. The specifications include county times year fixed effects. Standard errors are double clustered by year and firm and shown in parentheses. The significance of each coefficient estimate is indicated by * for $p < 0.10$, ** for $p < 0.05$, and *** for $p < 0.01$.

Landfall radius	100				200			
Window length in years (h)	1	3	5	10	1	3	5	10
$HitOther_{i,c,t}$	0.523***	0.426***	0.400***	0.403***	0.534***	0.435***	0.408***	0.403***
	(0.048)	(0.054)	(0.059)	(0.063)	(0.054)	(0.055)	(0.056)	(0.060)
Marginal effect	0.118	0.097	0.087	0.079	0.120	0.099	0.089	0.079
County × Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	541,060	1,037,433	1,157,467	1,185,779	553,164	1,060,368	1,192,796	1,211,923
Pseudo R^2	0.040	0.036	0.035	0.032	0.041	0.037	0.035	0.033

Table 7: Firms' Expansion Into New Counties

This table presents results from estimating the firm-level regression specification given in equation (5). The dependent variable is the number of new counties in which the firm generated patents over h years post-hurricane landfall divided by the total number of counties in which the firm generated patents over h years pre-hurricane landfall. The dependent variable ($SharePatentsHit_{i,t}$) measures the share of a firm's patents in counties that are in the hurricane landfall region. The data span from 1851 to 2023. The specifications include year fixed effects. Standard errors are double clustered by year and firm and shown in parentheses. The significance of each coefficient estimate is indicated by * for $p < 0.10$, ** for $p < 0.05$, and *** for $p < 0.01$.

Landfall radius	100				200			
Window length in years (h)	1	3	5	10	1	3	5	10
$SharePatentsHit_{i,t}$	0.024**	0.021**	0.029**	0.005	0.033***	0.031***	0.036***	0.013
	(0.010)	(0.009)	(0.012)	(0.010)	(0.009)	(0.009)	(0.010)	(0.008)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	232,955	406,105	451,265	384,621	210,375	407,701	409,129	346,958
R^2	0.094	0.071	0.072	0.083	0.104	0.075	0.074	0.086

Table 8: Within-Firm Substitution After Hurricane Shock

This table presents results from estimating the regression specification given in equation (7). The dependent variable $\log\left(\frac{\overline{NPatents}_{i,t:t+h-1}}{\overline{NPatents}_{i,t-h:t-1}}\right)$ is at the firm-year level. The numerator (denominator) is the annual average number of patents over h years post (pre) hurricane landfall. The dependent variable ($SharePatentsHit_{i,t}$) measures the share of a firm's patents in counties that are in the hurricane landfall region. The data span from 1851 to 2023. The specifications include year fixed effects. Standard errors are double clustered by year and firm and shown in parentheses. The significance of each coefficient estimate is indicated by * for $p < 0.10$, ** for $p < 0.05$, and *** for $p < 0.01$.

Landfall radius	100				200			
Window length in years (h)	1	3	5	10	1	3	5	10
$SharePatentsHit_{i,t}$	-0.020*	-0.029***	-0.029**	-0.050***	-0.002	-0.012	-0.005	-0.031***
	(0.012)	(0.010)	(0.011)	(0.010)	(0.008)	(0.011)	(0.011)	(0.010)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	93,582	176,664	184,163	178,620	98,079	183,568	192,548	185,843
R^2	0.0063	0.0101	0.0127	0.0177	0.0060	0.0103	0.0133	0.0173

Table 9: Aggregate County Patents Spillovers via Firms' Internal Networks

This table presents results from estimating the regression specification given in equation (8). The dependent variable $\log\left(\frac{\overline{NPatents}_{c,t:t+h-1}}{NPatents_{c,t-1}}\right)$ is at the firm-year level. The numerator is the annual average number of patents over h years post-hurricane landfall. The denominator is the number of patents in the year before landfall. The dependent variable ($CountyHitOther_{c,t}$) measures for firms located in county c the share of their patents that are outside county c and in the hurricane landfall region given in equation (8). The data span from 1851 to 2023. The specifications include county, year, and year interacted with state fixed effects. Standard errors are double clustered by year and county and shown in parentheses. The significance of each coefficient estimate is indicated by * for $p < 0.10$, ** for $p < 0.05$, and *** for $p < 0.01$.

Panel A: With County and Year Fixed Effects

Landfall radius	100				200			
Window length in years (h)	1	3	5	10	1	3	5	10
$CountyHitOther_{c,t}$	0.047	0.138***	0.139***	0.176***	0.045	0.123***	0.119***	0.150***
	(0.041)	(0.043)	(0.038)	(0.049)	(0.032)	(0.029)	(0.029)	(0.037)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	35,821	36,128	34,174	30,731	30,460	30,818	29,520	26,466
R^2	0.371	0.489	0.548	0.614	0.376	0.498	0.558	0.628

Panel B: With County and State×Year Fixed Effects

$CountyHitOther_{c,t}$	0.026	0.098**	0.119***	0.169***	0.032	0.105***	0.122***	0.147***
	(0.039)	(0.047)	(0.042)	(0.037)	(0.033)	(0.033)	(0.031)	(0.032)
Year×State FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	35,821	36,128	34,174	30,731	30,460	30,818	29,520	26,466
R^2	0.418	0.539	0.598	0.666	0.426	0.550	0.610	0.680

Appendix A Constructing the patents database

We extract information from pre-1976 patents in three steps. First, we identify the title page of each patent document. Second, we perform OCR (optimal character recognition) on the title page of each patent to obtain the text of each title page. Third, we use ChatGPT to extract inventor and assignee names and locations. We manually checked the accuracy of the extracted data of a random sample of 500 patents by comparing them to the patent documents. Only one patent had a minor discrepancy, indicating an error rate of 0.2%.

Combining these two technologies, we construct a new database of patents.[13] Our process is as follows:

1. Identify the title-page in each patent document. Beginning in the early 1970s, the first page of a patent document is the title page and contains relevant biographical and other information. Prior to that, the initial pages of a patent were often one or more pages of drawings, so the first page of text appears later.
2. Apply state-of-the-art OCR on the relevant PDF pages. We extract the layout of the document, making it possible to, for example, read columns of text in the correct order.
3. Use ChatGPT to extract information from the OCR text. To date, we have extracted inventor and assignee names and locations. We do this using a fine-tuned version of the GPT model and use OpenAI's Python API to submit millions of batch requests.

Consider again the patent shown in Figure 1; the OCR output we obtain is shown in the right column of Table 2. The quality improvement over the Google Patents version of the text is readily apparent; this rendering perfectly matches the original patent document.

Next, we use ChatGPT to extract the relevant information from the enhanced text. In particular, we use automated queries with questions about the text, as shown here along with ChatGPT's corresponding answers:

Input question	**ChatGPT response**
1) What is the title of this patent document?	IMPROVEMENT IN FOLDING TABLES
2) How many inventors are listed?	1
3) What is the name of each inventor?	Heman B. Sinclair
4) What city/state is each inventor from?	Chicago, IL
5) If the text mentions an "assignor" or "assignee," to what person or company was it assigned?	null
6) And if it was assigned, list any corresponding cities.	null

[13] To date, we have extracted information only for the title page of patent documents. We are in the process of securing funds to allow the creation of a complete and highly accurate historical database.

Given the improved quality of the OCR, one might wonder whether ChatGPT provides a significant benefit over simpler text analysis approaches. Determining the name and city/state of the inventor from this high-quality OCR text may appear to be quite straightforward using simple regular expression pattern matching. But ChatGPT can seamlessly handle far more complicated situations involving multiple inventors from different cities, multiple assignees, significant changes in formatting, and other edge cases that are hard to anticipate across millions of patents.

Our results were obtained using a recent ChatGPT model, `gpt-4o-mini`, which OpenAI describes as an "affordable and intelligent small model for fast, lightweight tasks." Out-of-the-box, this model performed quite well on our questions. We are able to elicit even higher-quality responses over a range of input types with two approaches. First, we adopt recent advances in prompt engineering. For example, we instruct the model to "take it step by step" before answering. Despite its apparent simplicity, this instruction has been shown to provide a significant improvement of an LLM's ability to "reason" through certain types of questions (Kojima, Gu, Reid, Matsuo, and Iwasawa, 2024). This approach is especially useful in patents with multiple inventors; in these patents we see an improvement in ChatGPT's answers once we require it *first* to count how many inventors there are, and *then* to name them and identify their locations.

Second, we further improve the output by "fine-tuning" the model to our particular needs. This is done by providing the model with additional training examples of questions along with our desired output; the model then learns to adjust its output to match the target. Fine-tuning alters the learning environment from "zero-shot" to "few-shot" by showing the LLM what an appropriate response looks like. After providing 100 fine-tuning examples, we verify that our model performs exceptionally well. We manually checked the accuracy of the extracted inventor and assignee names and locations of a random sample of 500 patents by comparing them to the patent documents. Only one patent had a minor discrepancy, indicating an error rate of 0.2%, which can likely be reduced further with additional fine-tuning of the model.

Appendix B Additional tables

Table B.1: Baseline Local Effects - Local Projection Estimation

This table presents results from estimating the regression specification given in equation (1) for horizons $h = 0, 1, 2, \ldots, 9, 10$. Panels A, B, C show the estimates with landfall regions determined at 50-, 100- and 200-miles, respectively, from the radius of the eye of a hurricane. The data span from 1851 to 2023. Standard errors are double clustered by year and county and shown in parentheses. The specifications include county and time fixed effects. The significance of each coefficient estimate is indicated by * for $p < 0.10$, ** for $p < 0.05$, and *** for $p < 0.01$.

Panel A: Country treatment within 50-mile radius

Years since hit	0	1	2	3	4	5	6	7	8	9	10
$Hit_{c,t}$	-0.009	-0.005	-0.017	-0.043**	-0.074***	-0.094***	-0.072***	-0.081***	-0.092**	-0.068**	-0.042
	(0.019)	(0.016)	(0.020)	(0.022)	(0.028)	(0.025)	(0.028)	(0.030)	(0.039)	(0.033)	(0.037)
$NPatents_{c,t-1}$	-0.677***	-0.700***	-0.721***	-0.730***	-0.737***	-0.750***	-0.749***	-0.765***	-0.760***	-0.771***	-0.774***
	(0.005)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.007)	(0.007)	(0.007)	(0.007)
$NPatents_{c,t-2}$	0.206***	0.190***	0.190***	0.185***	0.174***	0.178***	0.169***	0.172***	0.160***	0.159***	0.155***
	(0.004)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
$NPatents_{c,t-3}$	0.138***	0.143***	0.141***	0.134***	0.136***	0.130***	0.131***	0.126***	0.124***	0.122***	0.114***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
$NPatents_{c,t-4}$	0.112***	0.120***	0.113***	0.118***	0.111***	0.116***	0.106***	0.110***	0.105***	0.098***	0.103***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
$NPatents_{c,t-5}$	0.109***	0.109***	0.118***	0.113***	0.115***	0.109***	0.106***	0.104***	0.102***	0.104***	0.096***
	(0.004)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.006)	(0.006)	(0.006)
Other controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	80,813	79,263	77,797	76,288	74,993	73,589	72,388	71,197	69,936	68,531	67,264
R2	0.349	0.368	0.387	0.403	0.413	0.421	0.422	0.432	0.433	0.442	0.452

Panel A: Landfall Region Based on 100-Mile Radius Around Hurricane Eye

Years since hit	0	1	2	3	4	5	6	7	8	9	10
$Hit_{c,t}$	-0.016	-0.016	-0.001	-0.037**	-0.047***	-0.048**	-0.056***	-0.069***	-0.067***	-0.039*	-0.035
	(0.012)	(0.013)	(0.014)	(0.017)	(0.017)	(0.019)	(0.018)	(0.017)	(0.021)	(0.021)	(0.023)
$NPatents_{c,t-1}$	-0.678***	-0.705***	-0.724***	-0.734***	-0.741***	-0.752***	-0.755***	-0.767***	-0.764***	-0.774***	-0.778***
	(0.005)	(0.005)	(0.006)	(0.006)	(0.005)	(0.006)	(0.006)	(0.006)	(0.007)	(0.006)	(0.007)
$NPatents_{c,t-2}$	0.203***	0.190***	0.187***	0.183***	0.173***	0.174***	0.169***	0.170***	0.160***	0.158***	0.153***
	(0.004)	(0.005)	(0.004)	(0.004)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
$NPatents_{c,t-3}$	0.138***	0.142***	0.143***	0.136***	0.134***	0.131***	0.131***	0.128***	0.124***	0.120***	0.115***
	(0.005)	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
$NPatents_{c,t-4}$	0.112***	0.119***	0.114***	0.116***	0.113***	0.116***	0.107***	0.109***	0.105***	0.100***	0.103***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)	(0.004)	(0.005)	(0.005)	(0.005)	(0.004)
$NPatents_{c,t-5}$	0.109***	0.111***	0.117***	0.116***	0.117***	0.109***	0.108***	0.105***	0.101***	0.102***	0.096***
	(0.004)	(0.004)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Other controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	92,189	90,415	88,710	86,985	85,477	83,918	82,510	81,154	79,764	78,213	76,766
R2	0.348	0.369	0.387	0.403	0.413	0.419	0.422	0.430	0.432	0.440	0.450

Panel B: Landfall Region Based on 200-Mile Radius Around Hurricane Eye

Years since hit	0	1	2	3	4	5	6	7	8	9	10
$Hit_{c,t}$	-0.019*	-0.010	-0.000	-0.032**	-0.034**	-0.044***	-0.033**	-0.043***	-0.042***	-0.042**	-0.024
	(0.010)	(0.014)	(0.011)	(0.013)	(0.014)	(0.013)	(0.014)	(0.015)	(0.014)	(0.017)	(0.020)
$NPatents_{c,t-1}$	-0.679***	-0.705***	-0.724***	-0.735***	-0.743***	-0.751***	-0.756***	-0.766***	-0.763***	-0.774***	-0.777***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
$NPatents_{c,t-2}$	0.202***	0.192***	0.187***	0.183***	0.175***	0.173***	0.169***	0.170***	0.160***	0.158***	0.153***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.005)	(0.004)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
$NPatents_{c,t-3}$	0.139***	0.141***	0.142***	0.138***	0.134***	0.131***	0.133***	0.127***	0.123***	0.119***	0.116***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)	(0.004)	(0.004)	(0.004)	(0.005)
$NPatents_{c,t-4}$	0.113***	0.119***	0.115***	0.115***	0.113***	0.117***	0.108***	0.110***	0.105***	0.101***	0.103***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)	(0.005)	(0.004)
$NPatents_{c,t-5}$	0.109***	0.112***	0.117***	0.116***	0.118***	0.108***	0.106***	0.103***	0.101***	0.101***	0.093***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Other controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	100,938	98,997	97,075	95,156	93,443	91,772	90,203	88,688	87,185	85,533	83,818
R2	0.348	0.369	0.386	0.403	0.414	0.418	0.421	0.428	0.430	0.439	0.448

Table B.2: Baseline Local Effects - Event Study Specification

This table presents results from estimating event study regression specification. The dependent variable is $\log\left(\frac{\overline{NPatents}_{c,t+h}}{\text{NPatents}_{c,t-1}}\right)$, where the numerator is the average number of patents from 5 to 7 and 5 to 9 years, respectively, post hurricane hit. The main independent variable on interest is $Hit_{c,t}$. The data span from 1851 to 2023. Standard errors are clustered by year and county and shown in parentheses. The specifications include county and time fixed effects. The significance of each coefficient estimate is indicated by * for $p < 0.10$, ** for $p < 0.05$, and *** for $p < 0.01$.

Landfall radius	50		100		200	
Post-hit window length	3	5	3	5	3	5
$Hit_{c,t}$	-0.084**	-0.086**	-0.069**	-0.067**	-0.056**	-0.061**
	(0.034)	(0.035)	(0.029)	(0.029)	(0.023)	(0.025)
$logNPatents_{c,t-1}$	-0.735***	-0.752***	-0.739***	-0.756***	-0.740***	-0.747***
	(0.010)	(0.010)	(0.011)	(0.011)	(0.010)	(0.011)
$logNPatents_{c,t-2}$	0.182***	0.180***	0.183***	0.175***	0.175***	0.169***
	(0.009)	(0.009)	(0.009)	(0.009)	(0.010)	(0.009)
$logNPatents_{c,t-3}$	0.123***	0.128***	0.134***	0.132***	0.140***	0.137***
	(0.011)	(0.009)	(0.011)	(0.009)	(0.011)	(0.010)
$logNPatents_{c,t-4}$	0.132***	0.125***	0.115***	0.119***	0.133***	0.132***
	(0.011)	(0.010)	(0.011)	(0.011)	(0.011)	(0.011)
$logNPatents_{c,t-5}$	0.100***	0.103***	0.097***	0.101***	0.097***	0.098***
	(0.009)	(0.009)	(0.009)	(0.009)	(0.010)	(0.009)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	17,287	16,911	18,297	17,999	17,816	17,574
R2	0.523	0.568	0.526	0.570	0.531	0.570

Table B.3: List of Deadliest Hurricanes

This table presents the list of deadliest Atlantic and Gulf Coast hurricanes included in our sample based on data from NOAA. The first column indicates the main states of impact and/or, from 1954, the hurricane names.

Hurricane	Year	Fatalities
LA (Last Island)	1856	400
LA	1860	47
TX	1875	176
New England	1878	27
NC, VA	1879	46
GA/SC	1881	700
NC	1883	53
TX (Indianola)	1886	150
Texas	1886	27
Mid-Atlantic	1889	40
LA (Cheniere Caminanda)	1893	1100-1400
SC/GA (Sea Islands)	1893	1000-2000
SC, FL	1893	28
FL, GA, SC	1896	130
GA, SC, NC	1898	179
NC, SC	1899	50
TX (Galveston)	1900	8000
SE FL	1906	164
MS/AL/Pensacola	1906	134
LA (Grand Isle)	1909	350
TX (Velasco)	1909	41
SW FL	1910	30
LA (New Orleans)	1915	275
TX (Galveston)	1915	275
SW LA/Upper TX	1918	34
FL (Keys)/S TX	1919	287
FL (Miami)/MS/AL/Pensacola	1926	372
LA	1926	25
FL (SE/Lake Okeechobee)	1928	2500
TX (Freeport)	1932	40
S TX	1933	40
FL (Keys)	1935	408
New England	1938	256
GA/SC/NC	1940	50
Northeast U.S.	1944	64
SE FL/SE LA/MS	1947	51
Hazel (SC/NC)	1954	95
Carol (NE U.S.)	1954	60
Diane (NE U.S.)	1955	184
Connie (NC)	1955	25
Audrey (SW LA/N TX)	1957	416
Donna (FL/Eastern U.S.)	1960	50
Carla (N & Central TX)	1961	46
Hilda (LA)	1964	38
Betsy (SE FL/SE LA)	1965	75
Camille (MS/SE LA/VA)	1969	256
Agnes (FL/NE U.S.)	1972	122
Andrew (S FL, LA)	1992	26
Alberto (NW FL, GA, AL)	1994	30

Continued on the next page.

Table B.3: List of Deadliest Hurricanes (continued)

Hurricane	Year	Fatalities
Fran (NC)	1996	26
Floyd (Mid Atlantic & NE U.S.)	1999	56
Allison (SE TX)	2001	41
Ivan (NW FL, AL)	2004	25
Katrina (SE LA/MS)	2005	1200
Irene	2011	48
Sandy	2012	160
Matthew	2016	47
Harvey	2017	106
Irma	2017	96
Florence	2018	54
Michael	2018	59
Laura	2020	41
Ida	2021	92
Ian	2022	156

Table B.4: Variable Definitions

This table presents definitions of the main variables. The first column gives the variable name. The second column includes a short description.

Variable Name	Description	Source
$CountyHitOther_{c,t}$	This variable measures for firms located in county c the share of their patents over five years before hurricane landfall from counties in the hurricane landfall region and but not from county c.	NOAA, USTPO, Census
$Hit_{c,t}$	This is an indicator variable that takes a vlue of one if a county c is in the landfall region of a hurricane in year t and zero otherwise.	NOAA, Census
$HitOther_{i,c,t}$	This variable measures the share of firm i's patents over h years before hurricane landfall that are from counties in the landfall region but not from county c.	NOAA, USPTO, Census
$\text{NPatents}_{c,t}$	This variable counts the number of patents issued in year t for which all inventors reside in county c.	USPTO, Census
$\overline{\text{NPatents}}_{i,c,t:t+h-1}$	This variable is the average annual number of patents issued for firm i in county c over h years starting in year t.	USPTO, Census
$SharePatentsHit_{i,t}$	This variable measures the share of firm i's patents over h years before hurricane landfall from counties in the hurricane landfall region.	NOAA, USPTO, Census