Title

Essays on Housing Policy and Economic Disruptions

ISBN

9798247998143

Author

Pan, Yingru

Publication Date

2026-06-10

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

Essays on Housing Policy and Economic Disruptions

A dissertation submitted in partial satisfaction

of the requirements for the degree

Doctor of Philosophy in Management

by

Yingru Pan

2026

ABSTRACT OF THE DISSERTATION

Essays on Housing Policy and Economic Disruptions

by

Yingru Pan

Doctor of Philosophy in Management

University of California, Los Angeles, 2026

Professor Mark J. Garmaise, Chair

In Chapter 1, I examine how Los Angeles citywide discontinuous increases in real estate transfer taxes, aimed at preventing homelessness and funding affordable housing, actually affected housing supply. Using the city boundary as a natural experiment, I find that the tax reduced both home sales and new construction. I assemble a novel dataset of building permits, using AI-powered textual analysis of permit descriptions to extract granular information on construction activity. Sellers clustered transactions just below the tax threshold, while developers broadly reduced construction activity rather than pivoting to affordable or multi-family housing. Luxury homeowners increasingly chose remodeling over selling, further constraining supply. These findings suggest that taxing high-end properties may undermine rather than improve housing affordability. This paper offers novel insights into the effects of high-value property transaction taxes.

In Chapter 2, I find that when COVID-19 containment policies were imposed sequentially across U.S. states, the cumulative layering of restrictions produced compounding effects on house prices. Comparing adjacent county-pairs that straddle state borders, I find a two-phase dynamic response. In the first two months following an additional workplace-closure policy, house price growth in the treated county declines relative to its cross-border counterpart. This initial decline reverses over subsequent quarters: by nine to eleven months out, the treated county exhibits significantly higher price growth, and cumulative gains turn positive within

a year. The short-run contraction is consistent with demand suppression from uncertainty and reduced market activity. The medium-run appreciation is consistent with the view that workplace closures durably reshaped residential demand by accelerating the shift to remote work. The pattern is robust to one-to-one matching of county pairs.

In Chapter 3, I show that disaster risk propagates through supply chain networks to downstream customers, producing an asymmetric timing pattern: affected customers increase capital expenditure immediately but experience significant sales declines only several quarters later. The investment spike is consistent with emergency procurement and the search for alternative supply sources. The delayed revenue loss is consistent with the gradual depletion of inventory buffers. Both effects survive the exclusion of customers ever located in a disaster zone, confirming that the channel runs through supply chain linkages rather than direct exposure. Customers with more diversified supplier bases are partially insulated from the upstream shock.

The dissertation of Yingru Pan is approved.

Andrea Lynn Eisfeldt

Stuart A. Gabriel

Bernard Herskovic

Gregor Schubert

Mark J. Garmaise, Committee Chair

University of California, Los Angeles

2026

To my parents, friends, Johnny, and Yutou.

Contents

List of Figures

List of Tables

Acknowledgements

There are many people to thank.

My advisor, whose patience held steady even when mine did not. The timing was always gentle, the words always kind, arriving before the path was lost. My committee members, Andrea Eisfeldt, Bernard Herskovic, Gregor Schubert, and Stuart Gabriel, who offered their counsel with such generosity. Every suggestion has shaped this work, and shaped me.

The faculty at Anderson, who have been quietly tending to my thinking since the first seminar in 2019. What was passed along was not only the craft of research, but a kind of taste, a way of sensing what matters and what does not. That schooling is not yet finished, and perhaps such schooling never truly is.

My parents, who first suggested this doctorate almost in passing. My mentor, Prof. Gao, who encouraged me throughout and guided me in every way a person can be guided. The journey has come to its close, and the person I had hoped to become is not quite here yet. But the road toward her is clearer now than before.

My friends, the four of us in a group chat that changes its name so often no one keeps track anymore, and the wider circle that holds me beyond it, Mengxue Hou, Lijing Wang, Jitao Wang, Zitong Cheng, Kangying Zhou, Jiahong Shi, and Yan Yan. Through the pandemic, through the fires, through the long, quiet stretches that had no name but were difficult all the same, their voices kept me company.

The new friends found in Los Angeles, among the brightest and most tender-hearted people I have ever known. There is a pride in having found such company, and a hope in knowing the work, in time, will continue alongside theirs.

Los Angeles itself, restless and luminous and strange. In some way, everything that happened in this city became the soil in which the thinking took root.

And last, Yutou, the little dog. The day he came home, there was already a strange feeling that something was being given that would, in time, be taken away. The lesson has

proven harder than expected.

Vita

Education

2016 - 2017 University of Chicago

Master of Science, Financial Mathematics

2012 - 2016 Hong Kong University of Science and Technology

Bachelor of Science, Quantitative Finance

Chapter 1

Taxing the Top, Building the Bottom? The Impact of Los Angeles' Mansion Tax on Housing Supply and Affordability

Real estate transfer taxes imposed on the value of real estate transactions play an important role in the housing market. In the City of Los Angeles, the new tax was approved by a ballot initiative, Measure ULA (United to House L.A.), in the election of 8 November 2022, and is commonly known as the 'mansion tax'. The Ordinance imposes a 4% tax on all real property sales priced or valued from $5 million up to $10 million and a 5.5% tax on real property sales priced or valued at $10 million or greater, effective April 1, 2023. The LA mansion tax was expected to generate an estimated $1.1 billion annually, with the majority of the proceeds going to increase the supply of affordable housing and reduce housing costs for low-income residents.

The tax applies only to properties located within the City of Los Angeles and does not directly affect the other 87 incorporated cities in Los Angeles County, such as Beverley Hills, Long Beach, Pasadena and Santa Monica. This study leverages tax-rate and geographic discontinuities in transfer taxes in Los Angeles County before and after the "mansion tax" to evaluate the effectiveness of this tax reform and assess how the construction industry's new housing supply responded to these policy-induced changes. The ULA Tax is expected to discourage some transactions, and developers may seek opportunities outside the City of Los Angeles (Green et al., 2025). However, as demand for affordable housing increases, certain lower-cost housing developments may become more profitable.

Some may argue that luxury housing represents a relatively small segment of the housing market, casting doubt on whether a mansion tax could meaningfully affect broader housing supply or economic activity. This skepticism underlines a key motivation for the study. Even niche markets can exert a disproportionate influence in broader environments. High-value developments often set price benchmarks, shape developer strategies, and drive investment in adjacent sectors such as architecture, materials, and skilled labor. A decline in luxury construction, even from a modest baseline, could disrupt these interconnected networks, amplifying indirect effects on housing availability and employment. The introduction of a mansion tax can create major challenges for developers in attracting investors and securing bank loans, leading to a decline in construction projects. Investors prioritize high returns, often favoring luxury real estate where rents help them meet their financial targets, whereas affordable housing with lower rents may offer diminished profitability (Flemming, 2024). A mansion tax may reduce the appeal of luxury developments, prompting investors to seek opportunities elsewhere and leaving developers without funding. Meanwhile, banks, which typically finance 50–65% of construction costs, rely on stable rental income and investor backing to mitigate risk. If rents drop or investors withdraw due to the tax, banks may view projects as higher risk and withhold loans. As a result, developers could face a dual financing crisis: dwindling investor interest and restricted access to bank capital. This would stifle new construction, worsening housing shortages, particularly in high-demand markets like Los Angeles, where the need for affordable and high-end housing remains acute.

Furthermore, luxury projects frequently anchor large-scale urban developments, subsidizing infrastructure upgrades or mixed-income housing components. A tax-induced reduction in such projects may unintentionally stall wider community investments. Although known as the "mansion tax," except for rare exceptions, the tax applies to all properties sold for more than $ 5 million, including commercial properties such as gas stations, strip malls, apartment buildings, or actual mansions. Apartment developers and real estate brokers claimed that the additional costs imposed by ULA make it even harder to earn a reasonable profit in a risky business. That's because when building apartments, developers often sell their finished product, which would likely trigger the ULA tax for any building with more than 15 units. [1]Even developers who hold onto their properties typically need to take out a mortgage on the finished building, and lenders may be less willing to finance real estate because they, too, would need to pay the tax if they foreclose and sell the property.

By examining these cascading risks, this paper challenges the assumption that small market shares equate to negligible consequences. Luxury developments sometimes fund community improvements and amenities, such as roads or affordable housing, through profits from high-end projects. The loss of such developments may stall the realization of these benefits. The analysis demonstrates that even a small segment of the market can generate substantial ripple effects.

Figure 1.2 presents a time series of the 95th percentile of home sales prices from 2018 to 2023. The graph shows separate trends for the City of Los Angeles (the City) and the County. Each line represents the 95th percentile sale price in a given month. The two series move closely together. Both show similar trends before the mansion tax takes effect. Housing prices in the City are consistently higher than in the County, with prices fluctuating between $3 million and $4 million. Just before the mansion tax takes effect, I observe a sharp increase in the City's 95th percentile sale prices up to 6 million dollars, followed by a steep decrease to below 3 million dollars. This jump is not observed in the County. The figure highlights a strong short-run response among high-end properties inside the City.

As shown in the graph, luxury homes sometimes account for more than 5% of total

[1]https://www.latimes.com/california/story/2024-04-02/california-is-building-fewer-homes-the-state-could-get-even-more-expensive

housing transactions, signaling their growing influence on the market. A mansion tax targeting high-value properties could change construction trends, as developers often prioritize luxury projects for their profitability. Introducing tax penalties might reduce incentives for such projects, potentially slowing construction activity. This could tighten overall housing supply, worsening existing affordability challenges, particularly for middle-income buyers in increasingly competitive markets.

The ripple effects might extend beyond housing. Construction delays or scaled-back projects could lead to job losses in the sector, straining local economies reliant on development. In contrast, the tax could generate significant public revenue to fund affordable housing programs or infrastructure.

In Figure 1.3, the time series of the 90th percentile of listed home prices in the City and County between 2018 and 2023 falls between 2 million and 6 million, reflecting a similar trend, though less extreme, compared to that of the 95th percentile of home sale prices. Although these list prices may be higher than actual sale prices, they still highlight a significant segment of high-value housing. Even at this level, properties could be swept into a mansion tax bracket, depending on how thresholds are set.

To study the effect of the mansion tax on housing supply, proxied by the number of filed permits for new construction, a difference-in-differences model is employed in this quasi-experimental setting. The data contain two sources of variation: (1) time-series variation within the same neighborhood before or after the implementation of the "mansion tax"; and (2) cross-sectional variation across heterogeneous neighborhoods in the treatment (City of Los Angeles) and control (County of Los Angeles outside the City of Los Angeles) groups. The difference-in-differences model is well-suited for exploiting exogenous time-series variation and identifying the causal effect of treatment on housing supply outcomes. Difference-in-differences regression models are estimated to compare properties on either side of the City of Los Angeles border before and after the policy change. The near-border properties are organized into 22 relatively homogeneous regions. The standing hypothesis is that, after the imposition of the mansion tax on the City of Los Angeles, treated sellers were incentivized to sell their luxury properties at a premium to compensate for the burden of paying the tax, while buyers may have preferred to continue searching in order to benefit from locally depressed prices that were bunched below the notch. Together, there is a strong incentive for buyers and sellers near the notch to avoid transacting. With the local housing market discouraged, developers and contractors can leave affected neighborhoods in search of opportunities elsewhere.

Did developers and contractors opt out of affected neighborhoods? The evidence suggests that the answer is yes. The baseline results reveal that the total number of new construction permits in the City of Los Angeles decreased by 21% following the tax announcement in November 2022, compared to pre-tax trends. This decline contrasts starkly with stable permit volumes in adjacent regions outside the city, which serve as a control group unaffected by the policy. This exodus reflects rational economic behavior. Facing higher costs and uncertainty from the tax, developers likely redirected capital and labor to jurisdictions unaffected by the policy.

Subsequent sections explore mechanisms (e.g., permit-type heterogeneity, price-tier shifts) and assess whether the policy achieved its goals.

After opting out, contractors may have shifted their focus to constructing ADUs (Accessory Dwelling Units, which are secondary homes on a property). The effect of the mansion tax on ADU housing supply is examined using a difference-in-differences analysis. This approach compares neighborhoods that implemented the mansion tax with those that did not, before and after the tax was introduced, while controlling for property, neighborhood, and local housing market characteristics. The identification assumption is that, in the absence of the tax, neighborhoods would exhibit similar trends in ADU development, conditional on these local characteristics. No evidence of violations of the parallel trend assumption is found. Following the implementation of the mansion tax, ADU housing supply decreased significantly in the City of Los Angeles's border areas.

Next, I examine how implementing a mansion tax affects the supply of more affordable multifamily units by conducting a similar difference-in-differences analysis. I find a marginally negative impact of the tax on the supply of multi-family constructions. Contractors may exercise greater caution in light of the tax, as it is difficult to estimate property values three, five, or ten years out, which are typical hold periods for multifamily developers and owners. And for construction loans specifically, lenders may not want to convert the loan to a traditional commercial loan upon completion, given the uncertainty about how much the asset will sell for.

The uneven impact of the mansion tax becomes clear when dissecting the effects of housing permits by price levels. The grouping of properties into low-, mid-, and high-priced segments reveals divergent trends:

Single-family ground-up construction (with project square feet larger than 1,200) rose in lower-priced tiers post-tax (below 4 million dollars). This suggests that developers pivoted toward cheaper projects to avoid tax exposure. Remodeling permits for high-priced homes increased as owners opted to upgrade rather than sell, potentially serving as a tax-avoidance strategy that could exacerbate supply constraints.

Although the tax curbed luxury ground-up construction as intended, it failed to channel resources toward affordability. Luxury remodels increased, potentially promoting inequality by allowing high-value properties to accrue additional premiums. The outcomes reinforce a paradox: policies targeting luxury markets risk collateral damage to the very segments they aim to uplift. The mansion tax may deepen disparities by discouraging moderate-density solutions.

1.1 Related Literature

In recent years, policymakers across the globe have responded to affordable housing supply with a range of affordable housing policies, such as rent control, tax credits, housing subsidies, or rezoning policies. A comprehensive welfare analysis of affordable housing is done by

(Favilukis et al., 2023). In the United States, rent control measures are a strongly debated topic in local ballots, particularly in larger agglomerations such as the San Francisco Bay Area, New York York, New Jersey, or Maryland, and have been in place in some areas for multiple decades (Diamond & McQuade, 2019). The findings include a negative impact on the mobility of tenants (Diamond & McQuade, 2019), misallocation in housing markets (Glaeser & Luttmer, 2003), price appreciation in decontrolled segments of the housing market (Autor et al., 2014; Whamond, 2024), as well as a negative impact on housing investment and residential construction in the controlled segments of the market. However, recent papers such as (Mense et al., 2019) also highlight an increase in construction activity and investment in the unregulated segments of the market.

This paper is also closely related to the broad literature on behavioral responses to discontinuous tax schemes, such as (Mense et al., 2019) and (de Silva, 2023), using clumping at discontinuities in tax rates to identify income elasticities. (Chen et al., 2021) analyze the effects of tax incentives for investment in R&D over a threshold or a 'notch'. (Londoño-Vélez & Avila-Mahecha, 2024) study the behavioral responses to the wealth tax in Colombia. (Chen et al., 2025) shows that property taxes will lower housing prices at the state borders. (Horton et al., 2024) examines property tax reduction as a tool to increase housing affordability. This paper contributes by using a quasi-experimental setting in LA County and uniquely explores housing supply's response to the notched taxation scheme. My paper has the advantage of exploiting both cross-sectional and temporal variance in housing construction. This study examines whether the "mansion tax" intended to improve affordable housing improves local housing affordability and evaluates the effectiveness of this taxation scheme.

Specifically, this paper is closely related to the literature on transfer tax and construction. A strand of literature focuses on the effect of transfer taxes on the functioning of the real estate market ((Benjamin et al., 1993); (Van Ommeren & Van Leuvensteijn, 2005); (Dachis et al., 2012)). Tax notches not only distort prices near thresholds but also suppress transactions broadly, highlighting the unintended costs of such policies ((Kopczuk & Munroe, 2015);(Best

& Kleven, 2018);(Han et al., 2022)).

An empirical analysis of Measure ULA's impact in Los Angeles is underway. Initial projections from the UCLA Lewis Center ((Phillips & Ofek, 2022)) estimated that approximately 4% of property transactions would qualify for taxation. Post-implementation evaluations uncover two significant market adjustments: a 50% reduction in sales exceeding $5 million and evident price clustering below the tax threshold ((Ward et al., 2025); (Manville & Smith, 2025)). These findings collectively document measurable market responses to the policy intervention. (Green et al., 2025) provides an empirical quantification of the fiscal externality of transaction taxes on acquisition value property tax systems.

This paper is the first to empirically link transfer taxes (e.g., mansion taxes) to housing construction outcomes, or more specifically, to intended housing constructions, proxied by permit applications. First, previous work focuses on the price and volume effects of transaction taxes (e.g., (Kopczuk & Munroe, 2015)) but neglects their impacts on the supply side. By analyzing permit filings across price tiers and housing types, I show how such taxes distort developer incentives, reducing both luxury and missing-middle housing. Second, this study delivers the first comprehensive analysis of intended construction activity across Los Angeles County using complete permit application data from ATTOM. By integrating highly accessible municipal records from the City of LA with previously underutilized city-level data from unincorporated areas, I extract permit information with greater precision using AI-powered textual analysis.

1.2 Background

This tax scheme applied to real estate transactions greater than $5 million. It imposed a 4% tax on transactions priced between 5 million and $10 million, and a 5.5% tax on transactions over $10 million. Before Measure ULA, Los Angeles, like many cities, had a real estate transfer tax, but it was very low: its rate was 0. 45%, on transactions greater than $100,000.

For high-end sales, Measure ULA represented an almost ninefold increase in the tax rate. The goal of ULA is to help alleviate housing burdens, and the revenue it raises is dedicated to projects such as housing for the homeless affordable housing and tenant assistance.

In April 2024, after the measure had been in place for a year, the proponents declared it a success. ULA, they observed, was already raising far more revenue than any other housing program in Los Angeles and had withstood a number of legal challenges from the real estate industry. 'ULA has transformed our city', a city council member said. "We want a rising tide to raise all boats. If people are going to have multimillion-dollar homes, they should also contribute to making sure that everyone has a home." Similarly, a group of academics concluded a report on ULA by saying, "Los Angeles voters made the right decision when they approved Measure ULA in November 2022. Despite efforts by the real estate industry to undermine its implementation, Measure ULA is already a success." [2] When proponents first placed the "Mansion Tax" on the ballot, they did so estimating that it would raise between $600 million and $1.1 billion annually. [3] However, it has not come close to hitting that mark. According to the city's ULA Dashboard, from April 2023 to December 2024, the measure actually raised $480 million, or roughly $288 million per year, which was well below what was anticipated. [4]

The mansion tax appears to be reducing sales of higher-value real estate in Los Angeles. A reduction in such sales is troubling because these sales directly and indirectly drive the construction of much-needed new housing and help create new commercial and manufacturing opportunities for workers. A new report from Hilgard Analytics shows a 23 % drop in citywide residential permitting during 2024 compared to 2023. "The persistent shortage of affordable housing restricted by deeds, the lingering effects of redlining, exclusionary zoning policies, and

[2] Peter Dreier, Joan Ling, Scott Cummings, Regina Freer, Manuel Pastor, Ananya Roy, Chris Tilly. "Measuring LA's Mansion Tax: An Evaluation of Measure ULA's First Year". April 2024. https://www.oxy.edu/about-oxy/community-engagement/uepi/publications/ula- report

[3] Peter Dreier, Joan Ling, Shane Phillips, Scott Cummings, Manuel Pastor, Seva Rodnyansky, and Jackson Loop. 2022. An Analysis of Measure ULA. September. https://www.lewis.ucla.edu/research/an-analysis-of-measure-ula-a-ballot-measure-to-reform-real-estate- transfer-taxes-in-the-city-of-los-angeles/

[4] https://housing.lacity.gov/ula-dashboard

the intensifying impacts of climate change all ensure that the city's housing and homelessness crises will remain critical challenges well beyond this year," writes the author Joshua Baum. [5]

What are the most affordable types of new housing that this tax scheme is targeting? There are three types of affordable housing construction most likely to be affected: single-family ADU construction, large government-supported affordable housing projects, and other "missing-middle" multi-family construction.

1.2.1 Accessory Dwelling Units (ADU)

An accessory dwelling unit is a self-contained residential unit on a lot that shares a lot line with a primary residential unit. The primary residential unit may be of any type: a single-family home, a duplex, a fourplex, or even an apartment building. In a single-family home or duplex with a large backyard, an ADU might be a detached structure. With less backyard space, building an ADU can involve adding a new room on top of one's garage. An ADU could be built in the basement garage of an apartment building that sees little need for extra parking spaces. An ADU can be detached or attached to the primary unit, may have its own pathway and porch, or may simply be an extra bedroom converted into a studio apartment. Crucially, the ADU must be an independent housing unit. The California Department of Housing and Community Development cites numerous benefits to building an ADU: that they are cost-effective, provide extra space for the extended family, and serve as housing for the elderly. However, the state's primary hope for ADUs is to increase the supply of affordable housing.

Proposing upzoning single-family residences for ADU construction is intended to address the housing supply shortage. Many desirable cities in the United States face a housing affordability crisis. To buy a home in Miami, the median household must part ways with 85.6% of their monthly income, while those in Los Angeles must use 83% and New York

[5] https://la.urbanize.city/post/report-la-residential-permitting-dropped-23-2024

78%. Land use regulations, particularly single-family zoning, protect the preferences of local homeowners but seriously constrain the supply of housing. However, reforming land use regulation is especially difficult, as local residents are often concerned about congestion, privacy, and, of course, property values.

In 2016, the state of California established a right to build an ADU on a single-family home, overriding local regulations against such construction ((Gilgoff, 2023)[6]. This approach to increased density has the goal of increasing housing supply without drastically changing the landscape of single-family neighborhoods ((Simpson et al., 2024)). Proponents hope that ADUs can provide affordable housing to alleviate the state's affordability crisis ((Chapple et al., 2023)). However, it is not clear whether individual homeowners can close the gap in housing supply.

1.2.2 Affordable Housing Projects

Executive Directive 1 (ED 1) issued by the City of Los Angeles expedites the processing of shelters and 100% affordable housing projects in the City of Los Angeles. Eligible projects receive expedited processing, clearances, and approvals through the ED1 Ministerial Approval Process. ED1 was most recently updated on July 1, 2024. ED 1 applies to all Shelter projects and 100 Percent Affordable Housing Projects with an active or valid City Planning application or referral form filed with City Planning, and any ED 1 eligible projects under review by LADBS or LAHD. [7]

[6]According to the California Department of Housing Community Development ("HCD"), in the last ten years, California built an average of 80,000 new homes per year though the housing need is roughly 180,000 homes per year. CALIFORNIA' S HOUSING FUTURE : CHALLENGES AND O PPORTUNITIES , CALIF . DEP' T OF HOUS . & CMTY DEV . 3 (Feb. 2018), https://www.hcd.ca.gov/policy- research/plans-reports/docs/sha_final_combined.pdf. As of 2018, the gap in affordable rentals for extremely low- and very low-income households stood at roughly 1.5 million. Id. at 29-30.

[7]For the purposes of implementing ED 1, "100 Percent Affordable Housing Project" is defined as:

A housing project with five or more units, and with all units affordable either at 80% of Area Median Income or lower (U.S. Department of Housing and Urban Development (HUD) rent levels), or at mixed income with up to 20% of units at 120% AMI (California Department of Housing and Community Development (HCD) rent levels) and the balance at 80% AMI or lower (HUD/TCAC rent levels), as technically described here: A Housing Development Project, as defined in California Government Code Section(§) 65589.5, that includes 100% restricted affordable units (excluding any manager's units) for which rental or mortgage amounts are limited so as to be affordable to and occupied by Lower Income households, as defined in California Health

The city has plans to expand its reliance on incentives to encourage developers of market-rate housing to include a handful of affordable units for low- or moderate-income tenants in their projects. Developers who do this get to build more units than the city would otherwise allow. Through the first half of 2024, 588 ED1-approved residential units were fully permitted citywide, which represents 11% of all units, and that number is significantly up from the 84 ED1 eligible residential units permitted through the same time period of last year. There are likely to be more ED1-approved units receiving final construction permits going forward, as Hilgard Analytics and Zenith Economics identified 25 projects that received ED1 approval and have also received permits for demolitions or grading, which are often intermediate steps taken before the final construction permit, alongside many more ED1 eligible projects that are still going through the approval and permitting processes.

1.2.3 Multi-family Constructions

State and local governments have enacted or are considering enacting policies, such as zoning revisions, that allow denser housing development. These policies are intended to incentivize builders to create "missing middle" housing, a range of house-scale buildings with multiple units, compatible in scale and form with detached single-family homes, located in walkable neighborhoods.

For example, in a traditional single-family zoned neighborhood, a builder can purchase a 0.25-acre lot for $220,000 and build one house. If zoning regulations allow for a duplex to be built on that same 0.25-acre lot, then theoretically, each house's land cost would be $110,000 rather than $220,000. This theoretical lower cost would be passed on to the buyer, and one side of a duplex would be priced lower than a single-family home.

However, a 0.25-acre lot in a single-family zoned neighborhood may be worth $220,000,

---

and Safety Code §50079.5, or that meets the definition of a 100% affordable housing development in CA Government Code §65915(b)(1)G), as determined by the Los Angeles Housing Department (LAHD).

This definition is limited to projects that are for rent, as opposed to affordable for-sale projects. Rent Schedules with the applicable rents by unit size and income category are maintained by LAHD and are updated annually.

but if two homes can be built on the lot (in the form of a duplex), then the lot will likely increase in value, and a builder will pay more for it. In other words, by allowing denser housing, the land itself becomes worth more per acre. This effect raises the cost of land, and thus the cost of building a house, rather than lowering it. Even if the cost per lot is lower, those savings will not be passed on to the buyer, since the builder will price the home based on market demand, as noted above. If there is additional profit from lower unit land costs, that profit will go to the builder. Additionally, attempts to increase housing density have proven to be politically divisive. Those who own homes in areas subject to increased zoning density fear the loss of neighborhood amenities due to overcrowding and increased traffic.

In summary, the mansion tax, often pitched as a revenue tool to fund affordable housing, operates alongside existing affordability programs like accessory dwelling units (ADUs), density incentives (e.g., ED1), and multi-family mandates. However, these programs face chronic challenges: ADUs struggle with cost barriers and zoning restrictions, multi-family projects confront financing gaps and community opposition, and density incentives often fail to scale.

This paper raises a critical question: Does the mansion tax complement these efforts by redirecting resources toward affordability, or does it undercut them by stifling broader housing supply?

1.3 Data

1.3.1 Building Permit Data

I obtained the permit data of 2018-2024 from ATTOM. Each observation contains a permit identification number, the date applied, the block and lot, coordinates, a street address, the status, and a description of the permit. Example descriptions from an October 2022 permit read: “New 2,000 sqft single-family dwelling with four bedrooms and 3.5 baths and 430 sqft attached garage,” “Remove & rebuild (e) front porch (8 x 15’) per engineer of record, ” and "detached garage: convert to adu. 1 story vertical addition with one bedroom and bathroom." It is through this description field that I determine whether a permit is for the construction of a new single-family residence, an affordable project, a new ADU, remodeling, or simply mechanical, etc. With the assistance of AI, I filter for new-construction-related language, including terms relating to new constructions or the relevant ordinances and state bills. I then use other variables, such as a permit category variable, to filter for whether the permit is for new construction or merely updating its wiring.

1.3.2 American Community Survey

I get census tract-level data from the American Community Survey (ACS), a yearly U.S. Census survey covering a broad range of social and economic information. I use ACS 5-year estimates of median income, population, median house prices, number of rental units, and household income. I merge the ACS data with the building permit data. I conduct this merge spatially, matching each permit to its corresponding ZIP code. From there, I can create a ZIP code-level measure on how many new constructions were permitted and constructed.

1.3.3 Los Angeles Property Value Data

For the empirical analysis, I use a comprehensive listing-level dataset of residential properties for sale in LA County, collected by ATTOM. The data comes from MLS (Multiple Listing Service) platforms. Each observation in the data represents a listing on the MLS platform, with variables describing detailed information on the property and the status of the listing. To estimate parcel-level market values, I combine MLS transaction prices (actual sale data), tax assessor records (historical assessed values), and the Zillow Home Price Index (regional price trends). This allows extrapolation of current prices by reconciling real-time sales with longer-term assessments and market-wide trajectories.

1.3.4 Outcome Variables

There are two outcome variables I study in this paper. The first one is the log number of filed permits. The second one is the scaled number of permits. The scaled permit variable is constructed to normalize permit activity across ZIP codes, treatment or control areas, and time. For each ZIP code i, treatment or control area c, and quarter t, scaled permits are calculated as:

$$\text{Scaled Permit}_{i,c,t} = \frac{\text{Number of Permits}_{i,c,t}}{\text{Baseline Average Permits}_{i,c,2018-2019}} \tag{1.1}$$

Here, the baseline average permits for ZIP code i and area c is the mean number of permits filed in that ZIP code and area between 2018Q1 and 2019Q4. This period serves as a stable pre-tax, pre-pandemic, counterfactual reference.

The scaling acts as a counterfactual control: it implicitly answers, "How would permits in this ZIP code have trended without the mansion tax, based on its own history?" By scaling permits relative to their local historical averages, the metric adjusts for ZIP-level heterogeneity. For example, a ZIP code that typically issues 100 permits quarterly would see a value of 1.0 during the baseline, while a post-tax drop to 75 permits would register as 0.75. This approach standardizes comparisons across regions, ensuring that a 25% decline in permits reflects proportional changes, even in areas with vastly different construction volumes.

1.3.5 Summary Statistics

My final sample for building permits includes 2,019,858 observations. The sample covers 2016 to 2024. ADU construction began increasing in 2017. I use 2018 to 2019 as the base period, with 2022 to 2024 as the main observation window. Table 1.1 shows the summary statistics in my sample. Generally, over half of the permits come from the city. Typical construction projects span 523 square feet (median), while the average project spans 2322 square feet.

Project value and square footage both exhibit significant variation, as the sample covers all types of permits, from mechanical work to large apartment building projects. Lot sizes vary significantly, with a median of 7,705 square feet. Median household income is $6,807, and home values average $35,235. Areas typically contain 31 new buildings constructed after 2020. The price indices at the ZIP code level of Zillow reach $856, 659. The extrapolated prices are generally higher than the average house price index because the transaction prices are extrapolated using the 2024 house price indices. Demographic control variables are also included in the table.

Table 1.2 compares the areas inside and outside of Los Angeles City. Areas within Los Angeles City show a higher population density, with a median of 2,011 residents compared to 1,452 outside the city. Construction patterns diverge: the City of LA features smaller projects, with a median size of 25 square feet, versus 395 square feet elsewhere, but larger buildings, with a median footprint of 2,010 square feet, compared to 1,814 square feet outside the city. Income and housing values are considerably lower within the city. The median household income is $5, 710 versus $8, 402 and the median home values are $27,778 versus $37,652. The City of LA also has more recent construction activity, with a median of 40 new buildings post-2020 compared to 28 elsewhere. The price indices show the non-LA area with a median HP index of $1,090,753 versus $965,310 in the city. Both areas exhibit significant data dispersion, with standard deviations exceeding means across all metrics, indicating substantial variability within each region.

Figure 1.4 exhibits the different trends of permit applications between the City of LA and outside the city. Single-family residence construction permits reveal divergent trajectories between Los Angeles City and the surrounding areas from 2016 to 2024. The City of LA consistently maintained higher permit volumes, peaking at 5,315 in 2022, compared with outside the City, which peaked at 2,434 in 2022. The city demonstrated stronger early-period activity with permits exceeding 3,000 annually through 2021, while the suburbs remained below 2,500 until 2022. The data indicate the City of LA dominated early in the period but faced steeper declines, while non-LA areas exhibited stronger, sustained growth in later years despite starting from a lower baseline.

Multi-family construction permits show significantly lower volumes than single-family development in both regions, with neither area exceeding 1,500 annual permits compared to over 5,000 for single-family units. Los Angeles City demonstrated stronger mid-period momentum, peaking at 1,540 permits in 2022, while maintaining output above 1,000 permits in prior periods. Non-LA areas have lower permit volumes but exhibit lower volatility. This is comparable with single-family patterns.

In figure 1.5, ADU permit application and issuance began rising in 2017 and sustained growth through 2024 across both regions, though with divergent trajectories. Los Angeles City initiated the surge, jumping from 957 permits in 2017 to 1,871 by 2019, then sustained steady growth before peaking at 3,655 by 2022. Non-LA areas demonstrated a delayed but similar growth pattern, starting lower at 107 permits in 2017 before skyrocketing 15-fold to 1,555. The city now dominates ADU development, issuing 2.6 times as many permits as outside the City. Both regions have shown strong momentum since 2017. This contrasts with single-family and multi-family patterns, confirming that ADUs are the fastest-growing housing type across all landing zones.

The percentage change in ADU permits reveals a highly significant divergence between treated Los Angeles City and untreated non-LA areas across all nineteen groups. Non-LA areas demonstrate consistently positive growth trajectories in 16 of 19 groups. Both regions share significant post-2017 growth momentum, but non-LA's steadier expansion resulted in 280% higher cumulative permit growth by 2024 across comparable groups.

1.4 Results

1.4.1 Methodology

To first present the stylized facts of the relationship between permit filings and the mansion tax, I compare permit-filing trends before and after the mansion tax. I estimate the following regression:

$$y_{i,c,t} = \alpha + \beta \times Post_t + X_{i,c,t} + \gamma_i + \varepsilon_{i,c,t}$$

where $y_{i,c,t}$ is the outcome variable in quarter t, treated or untreated city c, and ZIP code i. $Post_t$ is a dummy variable that equals 1 if quarter t is after the first quarter of 2023, which was after the mansion tax was announced. $X_{i,c,t}$ denotes a list of time- variant demographic and local economic characteristics, including income, population, and housing value, from the American Community Survey conducted by the U.S. Census Bureau. The specification also includes ZIP code fixed effects γ_i, which controls for ZIP-code-specific time-invariant features.

The results in Table 1.3 show a negative correlation between mansion tax and the likelihood- hood of builders filing a permit for new construction in the City. Column 1 shows that without any controls, filing ratios post-mansion-tax exhibit a 21% lower than those before the mansion tax. When I control for ZIP codes, fixed effects, and demographic characteristics, I find similar patterns before and after the mansion tax, with a decrease in the filing ratio of 29%. These findings are significant at the 1% level.

The mansion tax applies only within the City of Los Angeles. Properties outside the city, in Los Angeles County, are not subject to the tax. This difference creates a sharp policy boundary that I can use for causal inference. A potential confound is that Santa Monica and Culver City also enacted transfer taxes during this time. Culver City's Measure RE became effective on April 1, 2021. Santa Monica's Measure GS took effect on March 1, 2023. Measure RE and Measure GS have different thresholds, tax rates, and exemptions from Measure ULA.

In all my regression analyses, I drop the presence of Santa Monica. I keep Culver City in the main regressions because its tax policy uses marginal rates, which are markedly different from those of Measure ULA. I exclude the presence of Culver City in the robustness check in Table 1.15. To account for potential noise from permit duplications at the same address (repeated filings), I implement the following filter: remove duplicate permits issued to the same address within a 12-month window, thereby reducing the risk of overcounting permits for the same project.

I use a border discontinuity design to compare construction activity just inside the city with that just outside it. To structure the comparison, I group ZIP codes along the city boundary into 22 geographic clusters. Each group includes ZIP codes from both the City of Los Angeles and the County's neighboring parts. Figure 1.1 shows the geographic boundaries between the City of Los Angeles and the surrounding County areas. Each ZIP code is assigned to one of 22 groups. I use different colors to represent each group. This design helps ensure that treated and untreated areas are locally comparable. It also controls for broader housing market trends. I compare changes in construction permits before and after the mansion tax across these 22 groups. By focusing on properties near the city boundary, I minimize bias due to differences in demographics, land-use policies, and market dynamics. This strategy exploits the natural geographic discontinuity in tax exposure to estimate the policy's effect on new construction.

The diff-in-diff model is as follows:

$$y_{i,c,t} = \alpha + \beta * (City_c \times Post_t) + X_{i,c,t} + Group_m \times \delta_t + \gamma_i + \varepsilon_{i,c,t}$$

where $y_{i,c,t}$ is the outcome variable in quarter t, treated or untreated city c, and ZIP code i. $City_c$ is a dummy variable that equals 1 if a property is located within the City of Los Angeles, and $Post_t$ is a dummy variable that equals 1 if quarter t is after the first quarter of 2023, which was when the mansion tax was introduced. $X_{i,c,t}$ denotes a list of time-variant demographic and local economic characteristics, including income, population, and housing value, from the American Community Survey conducted by the U.S. Census Bureau. $Group_m$ represents the group m to which one ZIP code is assigned. By using group-by-quarter fixed effects, $Group_m \times \delta_t$, I focus on comparing permit filings within each group and time period. This reduces bias from unobserved factors that change differently across space. The specification also includes fixed effects γ_i, which control for ZIP-code-specific time-invariant features. Standard errors are clustered at the group level. The key identification assumption of my regression model is that, in the absence of the tax policy change, the treatment and control regions would have evolved in the same way. To test this assumption, I need to show how the two groups evolved prior to treatment.

Figure 1.6 provides a visual representation of the regression coefficients for $City_c \times Quarter_t$, revealing no signs of nonparallel trends during the pre-treatment period for the total number of new construction permits. The difference-in-differences (DiD) analysis compares trends in permit filings between Los Angeles City (treatment group, subject to the mansion tax) and areas outside the city (control group, unaffected by the tax), using 2022Q3 as the baseline (pre-tax period).

In the first two quarters post-baseline (2022Q4–2023Q1), permit filings in the City of LA declined modestly. Starting in 2023Q2, however, the gap widened sharply: the City of LA permits fell significantly relative to the control group, with the negative trend persisting through 2024Q3. Outside the city, permit levels remained stable or dipped only slightly, aligning with broader market conditions.

The coefficient, β, turns significantly negative starting in 2023Q1, aligning with the mansion tax approval in November 2022. This signals a lagged response, as developers took time to adjust their construction plans and likely delayed permit filings after the initial period of policy uncertainty. Notably, permit activity in Los Angeles City declines sharply, in contrast to stable trends in areas outside the city, which serve as the control group. The negative effect persists through 2024Q3.

Beginning in 2023Q2, the treatment effect steadily becomes more negative in subsequent quarters. This suggests that developers increasingly paused or scaled back projects as the tax's operational and financial implications solidified. Given the long lead times in construction planning, the results imply that policy shocks may take quarters to fully materialize. The stability of permits outside Los Angeles City, despite shared economic conditions, reinforces the tax's distinct role in reshaping local supply incentives. The sustained decline in the City of LA filings implies that the tax has meaningfully dampened construction activity over time.

Table 1.4 shows the effects before and after the mansion tax for scaled filed permits and the log number of filed permits for new constructions separately. To account for differences in baseline permit activity across ZIP codes, a weighted least squares (WLS) regression is estimated, with each ZIP code in the treatment and control groups weighted by its average number of permit filings during the pre-policy period (2018 to 2019). For the log number of filed permits, I use PPML regression (Santos Silva & Tenreyro, 2006) and report a pseudo-R-squared. The PPML regression estimates a model as follows:

$$log(y_{i,c,t}) = \alpha + \beta * (City_c \times Post_t) + X_{i,c,t} + Group_m \times \delta_t + \gamma_i + \varepsilon_{i,c,t}.$$

I find a decrease in both outcome variables in treated areas after the mansion tax. There is an average 40% decrease in scaled filings for new construction in the City compared to adjacent ZIP codes outside the City. The results are statistically significant at the 1% level.

1.4.2 Heterogeneous Responses

A crucial issue to address is how to identify the driving force behind the observed decrease in permit filings, as this has implications for understanding the underlying mechanisms at work.

Permit filings reveal stark differences in how housing segments responded to the mansion tax. Single-family residences (SFRs), including ground-up constructions and accessory dwelling units (ADUs), show statistically significant declines in new construction activity, as shown in Tables 1.5 and 1.6. Both scaled permit and log filings fell sharply post-tax. There is an average of 46% decrease in scaled filings for new SFR constructions in the City.

The divergence highlights unintended consequences. While the tax succeeded in curbing high-value SFR construction, a proxy for luxury-oriented development, it failed to meaningfully shift supply toward affordability. Instead, it exacerbated declines in housing types that, while not explicitly “affordable,” often serve as middle-income housing in high-cost regions. These results underscore the challenge of targeting luxury markets without collateral damage to broader supply.

A key question arises: are declines in permits for new single-family residences (SFRs) and accessory dwelling units (ADUs) driven by reductions in luxury properties, lower-priced segments, or both? I analyze permit trends across housing price tiers, testing whether the mansion tax disproportionately impacted high-value developments or broadly suppressed supply. Did it curb luxury construction as intended, or inadvertently stifle "missing middle" housing?

To isolate whether permit declines stemmed from luxury or lower-priced housing, I divide the extrapolated price distribution into three groups, using 2 sets of cutoffs. The first set of cutoffs is \$1 million and \$4 million. The second cutoff is \$1,081,276 and \$4,592,702, which aligns with the 50th and 95th percentiles of the extrapolated price distributions. Both frameworks split permits into three groups: low-priced, mid-priced, and high-priced.

For ground-up single-family constructions, permits in the high-priced tier fell sharply post-tax, according to Table 1.11. For example, when looking at the results with a price cutoff of \$1,081,276 and \$4,592,702, compared to the lowest price bin, mid-priced permits rose by an insignificant 27%, and the high-priced permits decreased by 14%, significant at the 5% level. These shifts suggest developers pivoted toward cheaper projects to avoid tax exposure, though it is possible that gains in lower tiers were smaller than losses at the top.

Has construction capacity been redirected to more affordable multifamily (MF) projects?

Contrary to expectations, construction projects did not shift toward multi-family housing, as in Table 1.9. After the mansion tax took effect, the scaled number of permits for new multi-family constructions in Los Angeles City declined by 27% compared to pre-tax trends, a statistically significant drop mirroring reductions in single-family permits. This parallel decline indicates that the tax suppressed development activity broadly, rather than redirecting resources toward denser, "missing-middle" housing.

Developers likely hesitated to pursue even moderately priced multi-family units due to perceived financial risks, such as cascading tax impacts on project valuations or uncertainty about future policy adjustments. Meanwhile, permit levels in areas outside the city remained stable, suggesting local policy effects rather than broader market shifts.

These results challenge the assumption that curbing luxury construction would naturally free up capacity for affordable or mid-tier supply. Instead, the tax appears to have dampened overall development incentives without achieving its redistributive goals. The missing middle remains missing.

The results underscore that the mansion tax most directly discouraged high-priced construction. However, without complementary policies (e.g., density bonuses or affordable subsidies), the tax alone fails to redirect supply toward the more affordable projects. For robustness check, as in Table 1.10, I exclude all the permits that are marked as ED1 eligible affordable projects. The results are parallel to those in Table 1.9, suggesting that when excluding affordable projects, whose application intentions may be affected by the mayor's multiple policies to expedite the permitting process, multi-family construction alone decreases in response to the mansion tax.

For the last part, I look at remodeling permits in Table 1.12. Compared to the lowest-priced bin, for mid-priced homes, permits fell by 3%, with no statistical significance, indicating muted tax sensitivity. For high-priced homes, scaled permit counts increase sharply by 46%.

This inverse relationship, which increases luxury remodels and decreases activity in the lower level, defies conventional expectations. One explanation is asymmetric financial capacity: Luxury owners absorb tax costs, or leverage remodels to offset tax liabilities (e.g., by adding high-end amenities to generate rental income). Lower-tier homeowners, however, face tighter budgets and may avoid renovations that push their home's value into taxable brackets.

The mansion tax appears to inadvertently encourage luxury upgrades while stifling moderate renovations, a dynamic that could widen housing quality gaps and reduce 'move-up' inventory for middle-income buyers. If high-priced remodels prioritize premium features (e.g., pools, smart homes), they may further entrench neighborhood inequality without expanding affordable supply.

In addition, the mansion tax appears to have reshaped the strategies of luxury homeowners, incentivizing them to retain and upgrade existing properties rather than sell. Facing potential tax liabilities on high-value sales, many opt to "sit" on their homes while investing in renovations. This behavior aligns with a "hold and refine" mechanism: avoid sales that trigger the tax, while remodeling to enhance property value and living standards (e.g., luxury kitchens, smart-home tech) or generate rental income (e.g., upscale ADUs).

Critically, this contrasts with low-priced tier homes. Owners of cheaper properties are likely to avoid renovations that risk pushing their home's value above the taxable threshold, as they lack the financial buffer to absorb the added costs.

The disparity underscores a policy paradox: the tax disincentivizes turnover in the luxury market, reducing the supply of high-end homes for sale. Fewer luxury units enter the market, and stagnant middle-tier supply persists. Without mechanisms to channel renovation activity toward affordability (e.g., density incentives for ADUs), the tax may deepen the divides it aimed to narrow.

1.4.3 Placebo Tests

A placebo test analyzing trends in electrical and plumbing permits, not related to the intended goals of the mansion tax, reveals no statistically significant changes in file rates before and after the implementation of the policy, as can be seen in Table 1.7 and Table 1.8. These permits, which reflect routine maintenance and upgrades (e.g., HVAC systems, plumbing, wiring), remained stable in both Los Angeles City and control areas outside the city. The absence of divergence between the treatment and control groups reinforces the view that the observed declines in construction-related permits (e.g., SFRs, ADUs) are likely tied to the tax itself, rather than broader economic shifts or sector-wide disruptions.

This null result strengthens the causal interpretation of the main findings. If unrelated permit categories showed parallel trends, it would suggest that unobserved confounders (e.g., recessionary pressures, labor shortages) drove the construction slowdown. Instead, the stability of the placebo permits isolating the mansion tax as the primary driver of reduced activity in targeted housing segments. Policymakers can thus infer that the tax meaningfully reshaped developer and homeowner incentives.

1.5 Implications for the Residential Real Estate Market

This section analyzes discontinuities in Los Angeles housing transaction taxes, identifying significant price bunching below tax thresholds and suppressed transaction volumes above them. The empirical approach leverages quasi-experimental variation to demonstrate how tax-induced discontinuities distort market behavior.

Figure 1.7 shows a significant decrease in units sold in the City of LA after the mansion tax. Figure 1.8 reveals parallel declines in luxury property transactions across jurisdictions following Los Angeles' mansion tax implementation, with significantly more severe contraction within the City of LA. The city's steeper decline indicates that tax-avoidance behaviors disproportionately affected urban luxury markets relative to broader countywide transactions.

Figure 1.9 reveals pronounced price bunching below the $5 million mansion tax threshold in Los Angeles City during 2022-2023. Sales count peaks dramatically in the $4.95-5 million bin, immediately preceding the tax applicability cutoff, then drops precipitously for properties priced just above $5 million. This pattern demonstrates strategic price manipulation to avoid tax liability, with 47% more transactions clustered in the final, pre-tax bin than in the first taxable bin. The distribution shows normal dispersion between \$3 and 4.5 million, followed by an abnormal accumulation near the threshold. Above \$5 million, transaction volumes remain suppressed across all price tiers, confirming the tax's chilling effect on luxury property transfers. This bunching behavior provides clear evidence that market participants are actively restructuring transactions to minimize tax exposure, distorting the natural price distribution around the policy threshold.

Figure 1.10 demonstrates significant market distortion following Los Angeles' mansion tax implementation, revealing pronounced price bunching immediately below the \$5 million threshold in the post-tax period. Pre-tax sales distribution shows relatively smooth dispersion across price bins between \$2-8 million, with normal transaction clustering around market-driven price points. Post-tax, a sharp discontinuity emerges at the \$5 million cutoff: sales volume spikes dramatically in the \$4.95-5 million bin while collapsing for properties priced just above this threshold. This bunching effect creates a visible "cliff" at the tax boundary. Above \$5 million, the post-tax distribution shows uniformly suppressed transaction volumes across all price tiers, confirming the tax's broad chilling effect on luxury property transfers. These findings provide empirical evidence of strategic price manipulation to avoid tax liability, fundamentally altering transaction patterns in high-value real estate markets.

Figure 1.11 demonstrates stark divergence in transaction patterns between the \$4-5 million and \$5-6 million price tiers following Los Angeles' mansion tax implementation. Properties priced just below the tax threshold (\$4-5M) maintained robust sales volumes throughout the period, with noticeable. Conversely, the \$5-6 million tier experienced a severe contraction immediately after tax enactment, with sales volumes dropping to near zero by mid-2023. This differential response creates a pronounced effect, with transactions clustering below the \$5 million threshold but avoiding crossing it. The \$4-5 million segment consistently recorded 3-5 times more monthly transactions than the \$5-6 million tier post-tax, revealing active price manipulation to remain below taxable values. Both tiers showed parallel pre-tax trends until policy implementation, after which the higher-priced segment collapsed while the sub-threshold market resumed.

The difference-in-differences analysis in Table 1.13 demonstrates statistically significant reductions in property transactions attributable to Los Angeles' mansion tax. The interaction term shows consistent negative coefficients across specifications, yielding a significant decrease in property transactions at the 5% level, indicating a 23.5% reduction in sales relative to control areas. This pattern confirms the tax's causal effect in suppressing luxury property transactions beyond broader market trends.

The triple-differences analysis in Table 1.14 reveals starkly divergent impacts across price tiers from Los Angeles' mansion tax. Properties in the highest price bin (Bin 3) show statistically significant declines in transactions, significant at the 1% level. This contrasts sharply with Bin 2's statistically insignificant positive coefficients. Excluding the 2022 Q4 announcement period reduces Bin 3's negative magnitude by 20-29% across both cutoff specifications, indicating anticipatory behavioral responses prior to formal implementation. The consistent significance of high-tier suppression persists despite price bin redefinitions and announcement-period exclusions, confirming robust tax avoidance at threshold levels and above.

1.6 Robustness Checks

Table 1.15 and Table 1.16 presents a series of additional robustness tests for the main regressions. For brevity, only the regression coefficient and standard errors of the interaction term are reported. Unless otherwise specified, all control variables and regression specifications are the same as those in the main regressions.

In Table 1.15, to validate the consistency of my findings on the impact of the mansion tax on the scaled number of permits, I conduct a series of robustness checks across various dimensions: first, in row (1), I exclude permits from Culver City to ensure results are not driven by areas with potentially similar transfer tax policies. In row (2), the sample excludes

ED1-eligible affordable housing permits to avoid the confound introduced by city policies specifically related to affordable housing. In rows (3) and (4), to account for potential noise from permit duplications at the same address, I apply stricter filters that remove duplicate permits issued to the same address within a 6- or 9-month window. In row (5), to test the robustness of the arbitrary ZIP code grouping, I re-estimate the main specification using alternative ZIP code groupings and verify that the estimated effects are not sensitive to the definition of geographic clusters. These robustness checks are applied across all permit categories, including total new permits, single-family, multi-family, accessory dwelling units (ADUs), and placebo permits. The results remain unchanged across specifications, confirming the robustness of the main findings.

I also conduct the same robustness checks for the impact on the log of the number of permits using PPML regressions, as shown in Table 1.16. The results across all permit types are robust.

1.7 Conclusion

This study evaluates whether Los Angeles' citywide mansion tax, a discontinuous increase in real estate transfer taxes designed to curb homelessness and fund affordable housing, achieved its housing supply goals. Leveraging plausibly exogenous variation in tax exposure across LA County, I find significant market distortions: Housing transactions are bunched below tax thresholds, with sharp declines in units sold at taxed price tiers. Overall construction across all residential types is suppressed. Luxury homeowners increasingly opted to remodel rather than sell, reducing turnover and exacerbating supply constraints.

The results challenge the policy's foundational premise: taxing high-value housing does not inherently redirect resources toward affordability. Instead, it risks broadly stifling supply. For future reforms, pairing taxation with direct, affordable housing incentives, such as streamlined permitting, density bonuses, or targeted subsidies, may better align revenue generation with affordable housing.

Figure 1.1: Bordering Zip Codes Colored by Group

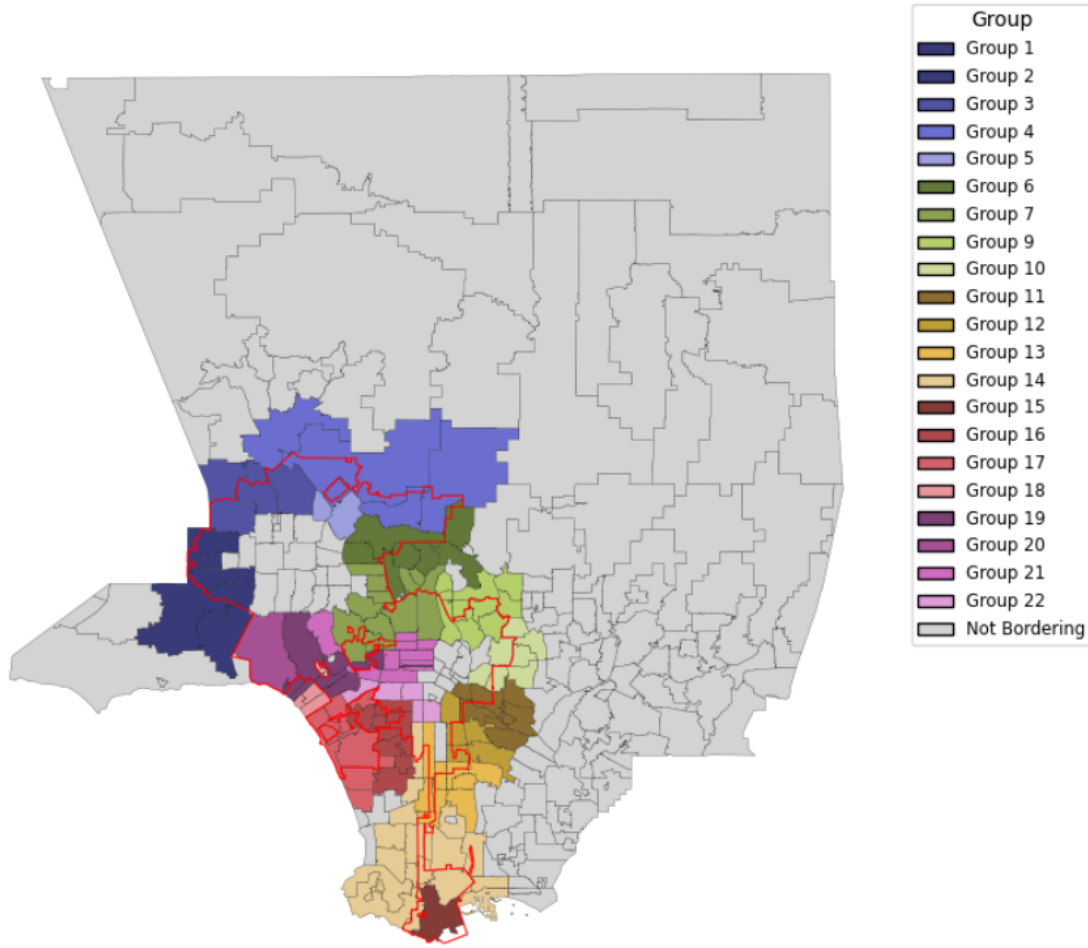

Notes: Bordering ZIP codes are color-coded by groups. The red line marks the jurisdictional boundary between the City of LA and LA County. Grey areas indicate ZIP codes excluded from the sample due to insufficient comparison data.

Figure 1.2: 95% of Sold Price

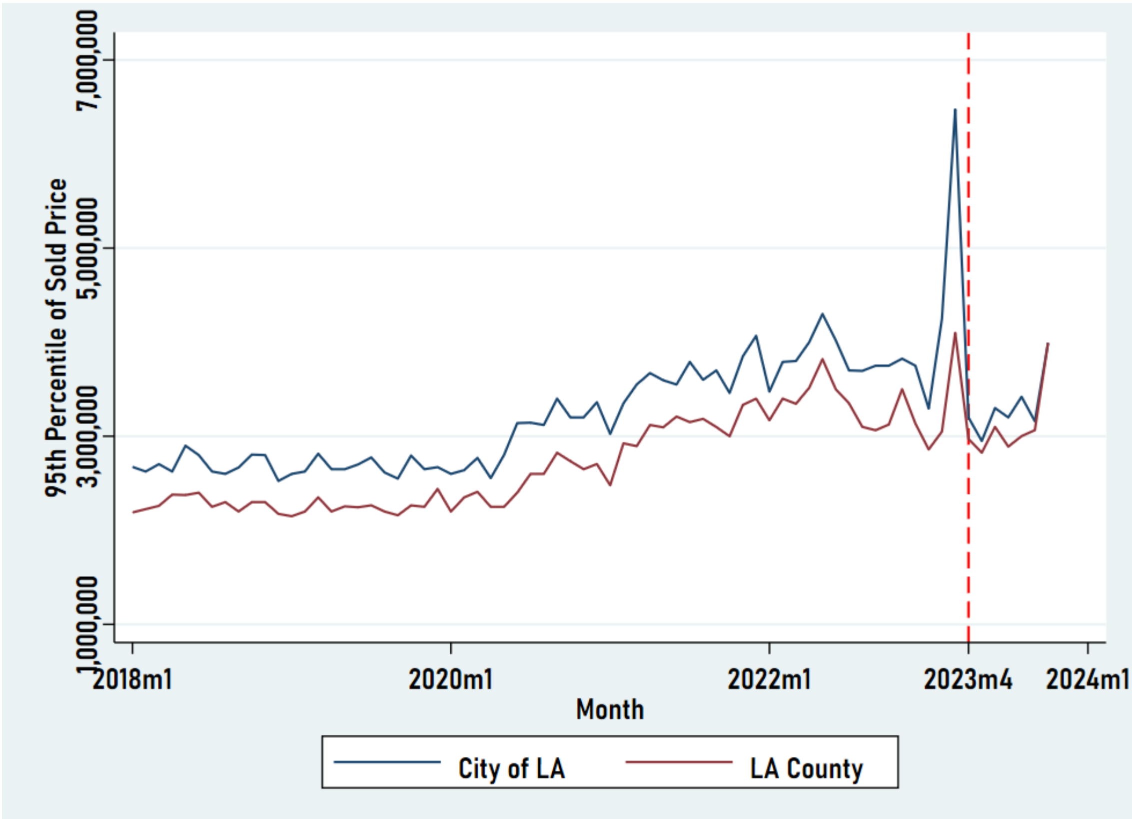

Notes: This figure reports 95% of sold price between 2018-2023 between the City of LA and LA county.
Source: MLS

Figure 1.3: 90% of Listing Price

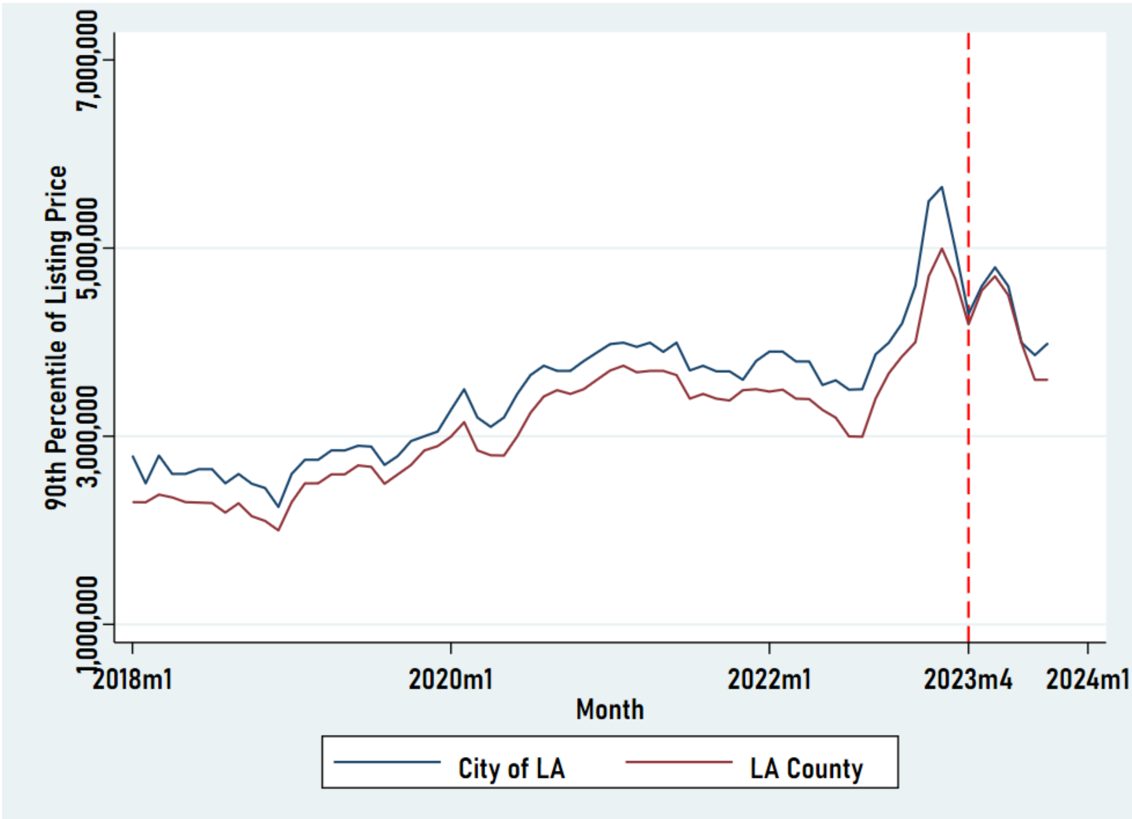

Notes: This figure reports 90% of listing price between 2018-2023 between the City of LA and LA county.
Source: MLS

Figure 1.4: Annual Permit Filing for SFR and MF

Notes: This figure shows the annual trend for permit filing of single family residence and multi-family residence in the City of LA and outside of the City of LA.
Source: ATTOM

Figure 1.5: Annual Permit Filing for ADU and Percentage Change in Permit Filing

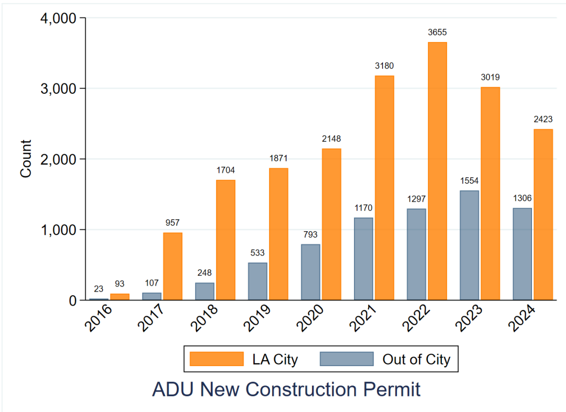

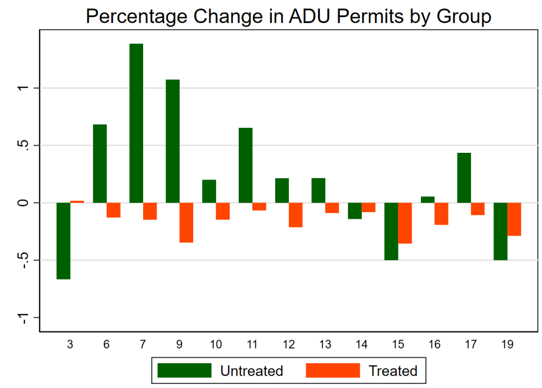

Notes: This figure shows the number of permit filing each year for ADU construction in Panel A. Panel B compares the percentage change of ADU permit filing between the City of LA and outside areas by ZIP-code-clustered groups. Some groups are omitted because there is no ADU construction within the treatment or control areas for those groups.
Source: ATTOM

Figure 1.6: Event Study on Scaled number of Permit and Log number of Permit

Notes: This figure shows the coefficients and the associated 95% confidence intervals of the interaction terms β from the baseline estimating equation for total number of new construction permits. Panel (a) shows the result for the scaled number of permits, while Panel (b) shows the result for the log number of permits. The x-axis indicates the year-quarter. The y-axis indicates the point estimates associated with the β_t estimate. The sample spans the period from 2021q1 to 2024q3.

Figure 1.7: Number of Units Sold in the City of LA(> \$ 5m)

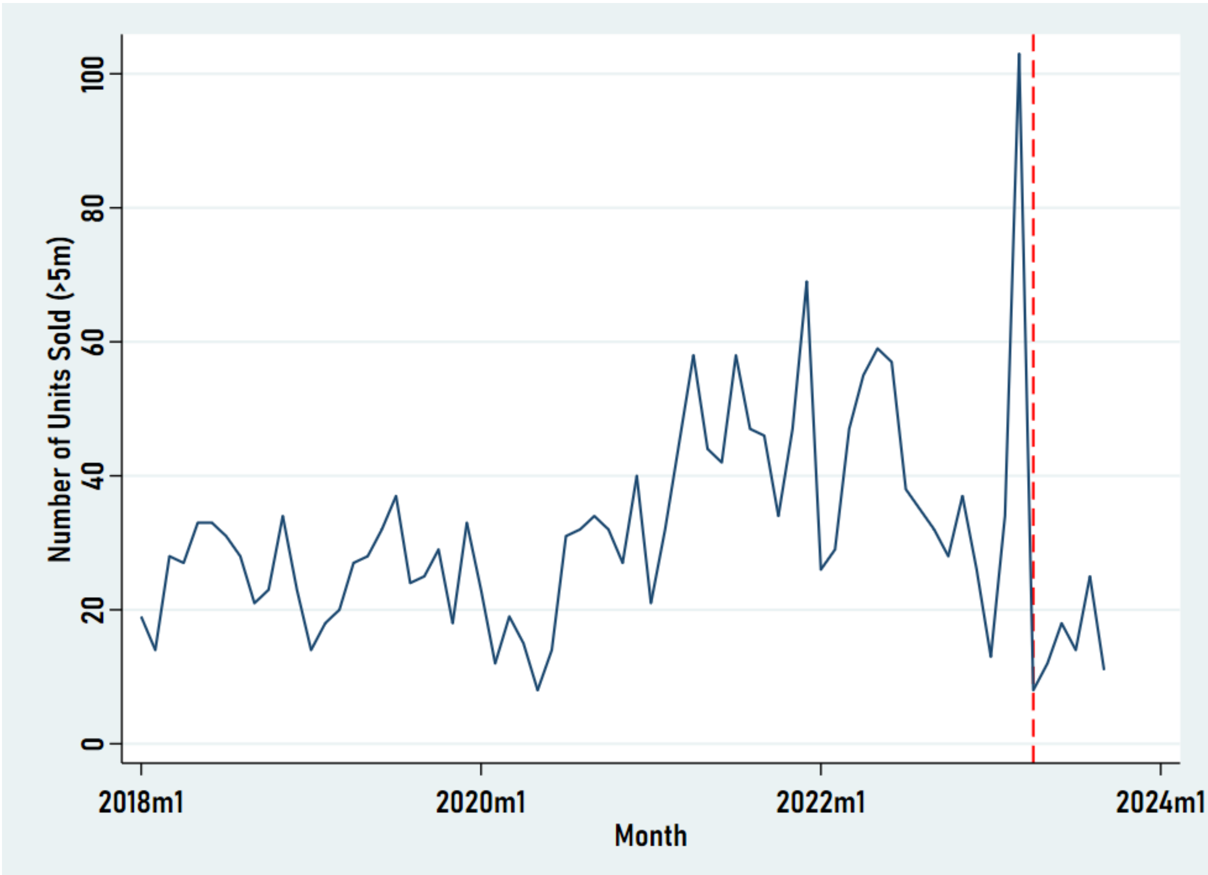

Notes: This figure reports the time series of number of units sold above \$5 million in the City of LA, from 2018-2023.

Figure 1.8: Percentage change in Units Sold(>$5m)

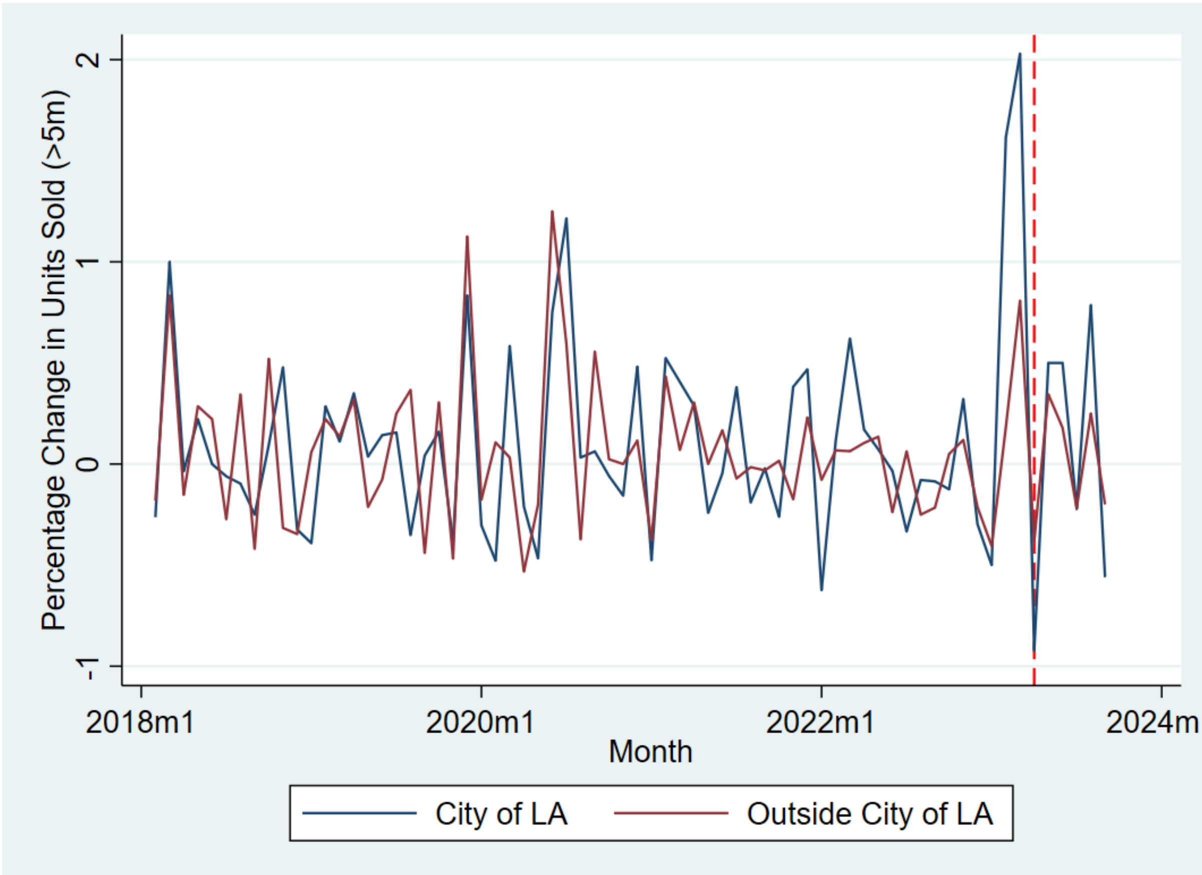

Notes: This figure reports the time series of percentage change in number of units sold above $5 million, between city of LA and outside city of LA, from 2018-2023.

Figure 1.9: Distribution of Taxable Sales

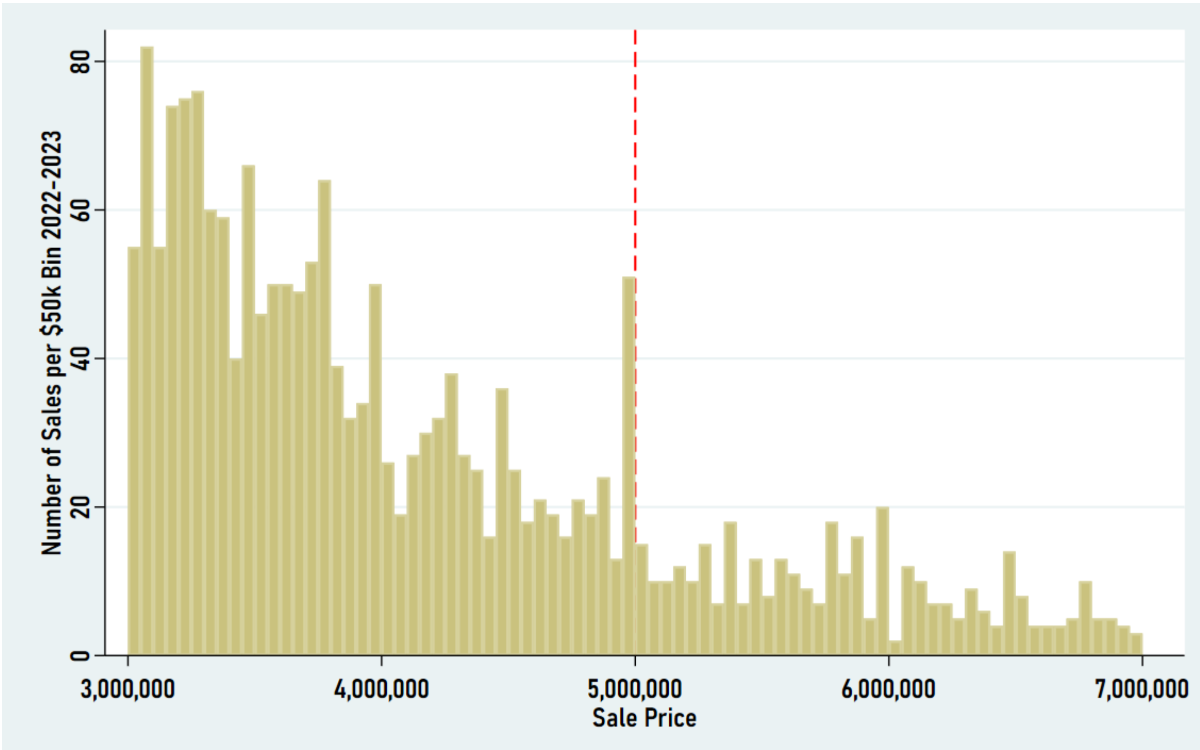

Notes: Plot of the number of mansion-tax eligible sales in each $50,000 price bin between $3,000,000 and $7,000,000 from 2022 to 2023 in the City of LA.

Figure 1.10: Distribution of Sales pre- and post- Mansion Tax

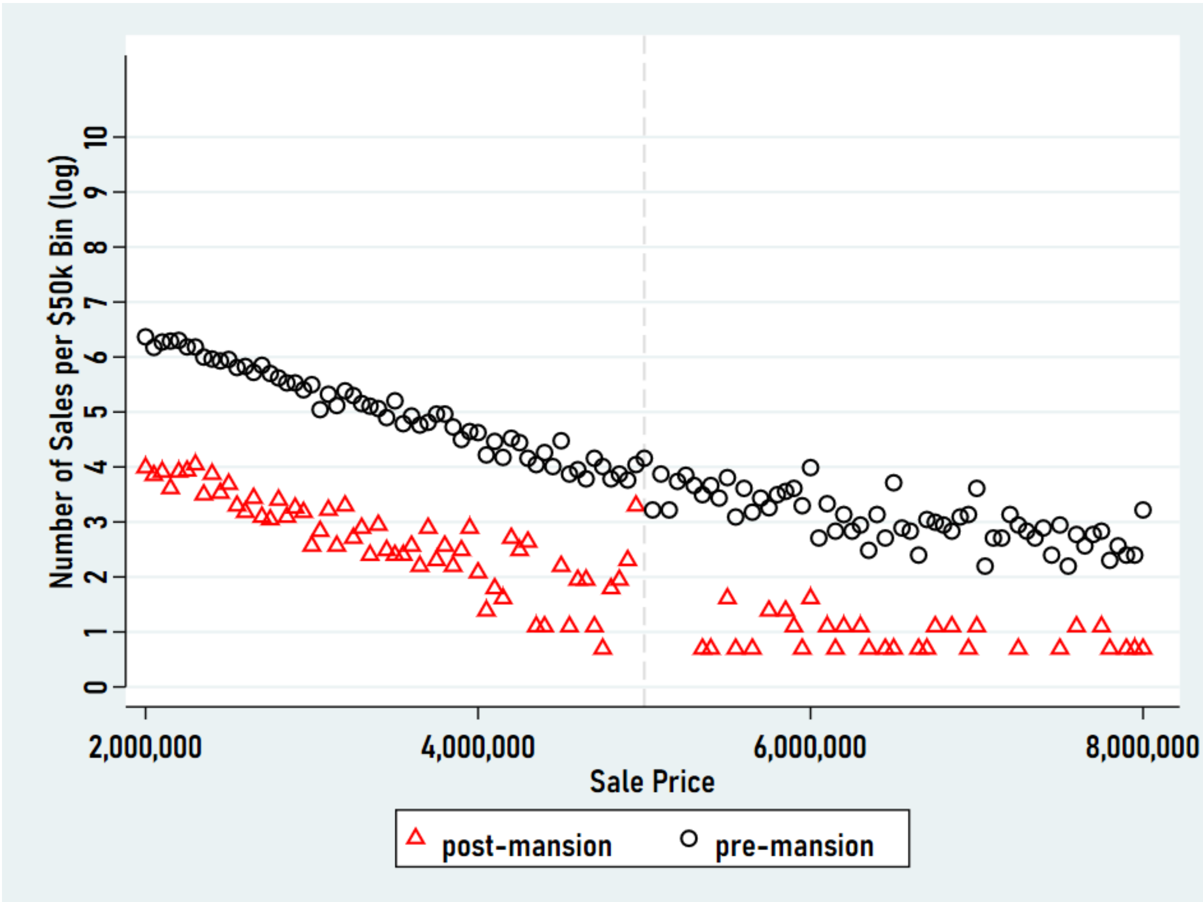

Notes: Plot of the log number of mansion-tax eligible sales in each $50,000 price bin between $2,000,000 and $8,0000,000 before and after the introduction of the tax in the City of LA.

Figure 1.11: Monthly Sales above $ 4m

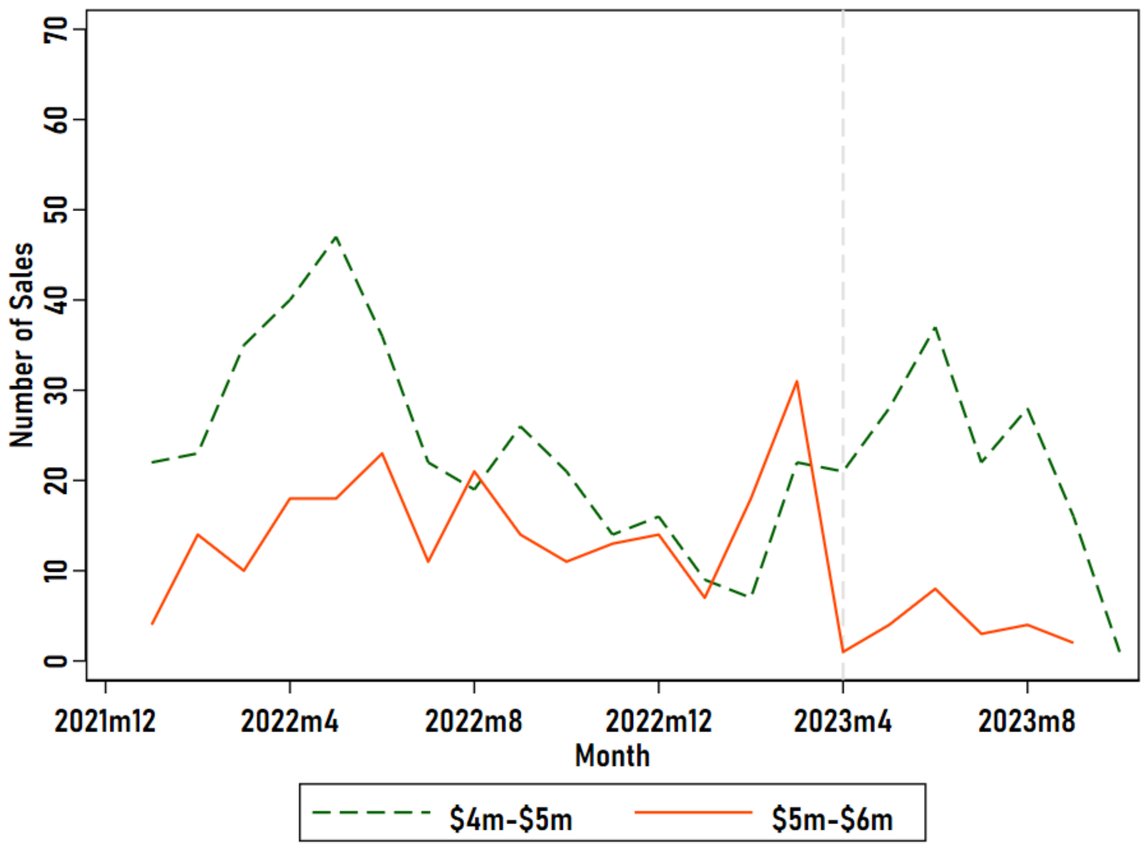

Notes: This figure shows the total number of taxable sales in the range of 4 to 5 million and 5 to 6 million by month in the City of LA.

Table 1.1: Summary Statistics

	mean	sd	min	p50	max	count
In LA City	0.55	0.50	0	1	1	1,750,225
Project Square Footage	2,322.89	19,279.33	-1,852,440	523	1,788,210	104,541
Average Project Value	136,910.80	3,888,254.12	-2,498	8,000	2,000,000,000	798,287
Building Square Footage	12,268.59	70,257.71	1	1,967	3,585,714	1,620,303
Lotsize	52,259.49	1,126,745.35	2	7,705	968,793,536	1,664,351
Population	1,849.09	742.03	34	1,916	4,016	1,064,160
Median Household Income	8,525.00	6,600.16	2,085	6,807	51,710	1,063,466
Built post 2020	48.20	35.97	6	31	155	1,064,160
HP Index	1,041,305.27	698,723.16	236,788	856,659	5,455,559	922,032
Extrapolated Price	1,766,711.46	2,789,921.14	5	1,098,401	168,958,688	422,446

Notes: This table provides summary statistics for the variables of interest. The columns provide the mean, standard deviation, min, and max for the permit data used. Population, dummy for buildings post 2020, and income are obtained from the U.S. Census Bureau. Dummy for in LA city, project square footage, project value, building square footage, and lotsize are from ATTOM. HP Index is obtained from Zillow. Extrapolated prices for houses are calculated from MLS transaction data.

Table 1.2: Summary Statistics by Area

Area	Out of LA			In LA		
	Mean	SD	Median	Mean	SD	Median
Project Square Footage	1,470	5,395	395	1,009	5,273	25
Average Project Value	119,139	5,694,599	14,630	185,974	4,028,549	7,600
Building Square Footage	7,790	46,503	1,814	14,084	80,627	2,010
Lotsize	53,684	811,446	7,835	41,591	618,773	7,551
Population	1,456	790	1,452	2,049	621	2,011
Median Household Income	9,925	6,087	8,402	7,591	5,723	5,710
Built Post 2020	36	25	28	54	39	40
HP Index	1,326,299	843,356	1,090,753	1,184,480	648,299	965,310
Extrapolated Price	2,049,007	2,100,789	1,342,625	1,865,831	1,831,632	1,279,800

Notes: This table provides summary statistics for the variables of interest. The columns provide the mean, standard deviation, min, and max for the permit data used, for inside and outside of LA city, respectively. Population, dummy for buildings post 2020, and income are obtained from the U.S. Census Bureau. Dummy for in LA city, project square footage, project value, building square footage, and lotsize are from ATTOM. HP Index is obtained from Zillow. Extrapolated prices for houses are calculated from MLS transaction data.

Table 1.3: Mansion Tax and Filed Permits

	(1)	(2)	(3)
Post Mansion	-0.213***	-0.299***	-0.203***
	(0.041)	(0.033)	(0.044)
Observations	152	152	155
Adjusted R^2	0.493	0.497	0.273
Zipcode FE	Yes	Yes	No
Control	No	Yes	Yes

Notes: This table reports the effects of mansion tax on permit filing. The dependent variable is scaled number of filed permits. Demographic controls include population, income, and housing value. Specifications include zipcode fixed effects. ***, **, and * represent statistical significance at the 1%, 5%, and 10% level, respectively.

Table 1.4: Mansion Tax and Filed Permits (Total New Constructions)

	Scaled No. of Permits		Log No. of Permits	
	(1)	(2)	(3)	(4)
LA city x Post Mansion	-0.396***	-0.415***	-0.406***	-0.361***
	(0.102)	(0.093)	(0.105)	(0.136)
Group x Mansion Tax FE	Yes	Yes	Yes	Yes
Zipcode FE	No	Yes	No	Yes
Controls	Yes	Yes	Yes	Yes
Adj. R^2	0.212	0.573		
Pseudo R^2			0.551	0.777
Obs.	272	264	272	264

Notes: Group-by-Mansion-Tax fixed effects are included to control for local trends and time-varying shocks. This table reports the effects of mansion tax on permit filing. The dependent variable is scaled number of filed permits, or the log number of filed permits. Demographic controls include population, income, and housing value. Specifications include group and zipcode fixed effects. Each ZIP code in the treatment and control groups is weighted by its average number of permit filings during the pre-policy period of 2018 to 2019. Standard errors are clustered at the group level. ***, **, and * represent statistical significance at the 1%, 5%, and 10% level, respectively.

Table 1.5: Mansion Tax and Filed Permits (Single Family Residence)

	Scaled No. of Permits		Log No. of Permits	
	(1)	(2)	(3)	(4)
LA city x Post Mansion	-0.448***	-0.462***	-0.476***	-0.451***
	(0.124)	(0.125)	(0.099)	(0.129)
Group x Mansion Tax FE	Yes	Yes	Yes	Yes
Zipcode FE	No	Yes	No	Yes
Controls	Yes	Yes	Yes	Yes
Adj. R^2	0.271	0.615		
Pseudo R^2			0.574	0.772
Obs.	257	250	257	250

Notes: Group-by-Mansion-Tax fixed effects are included to control for local trends and time-varying shocks. This table reports the effects of mansion tax on permit filing. The dependent variable is scaled number of filed permits, or the log number of filed permits. Demographic controls include population, income, and housing value. Specifications include group and zipcode fixed effects. Each ZIP code in the treatment and control groups is weighted by its average number of permit filings during the pre-policy period of 2018 to 2019. Standard errors are clustered at the group level. ***, **, and * represent statistical significance at the 1%, 5%, and 10% level, respectively.

Table 1.6: Mansion Tax and Filed Permits (ADU)

	Scaled No. of Permits		Log No. of Permits	
	(1)	(2)	(3)	(4)
LA city x Post Mansion	-0.727***	-0.761***	-0.548***	-0.476***
	(0.221)	(0.189)	(0.127)	(0.142)
Group x Mansion Tax FE	Yes	Yes	Yes	Yes
Zipcode FE	No	Yes	No	Yes
Controls	Yes	Yes	Yes	Yes
Adj. R^2	0.303	0.549		
Pseudo R^2			0.561	0.795
Obs.	230	228	230	228

Notes: Group-by-Mansion-Tax fixed effects are included to control for local trends and time-varying shocks. This table reports the effects of mansion tax on permit filing. The dependent variable is scaled number of filed permits, or the log number of filed permits. Demographic controls include population, income, and housing value. Specifications include group and zipcode fixed effects. Each ZIP code in the treatment and control groups is weighted by its average number of permit filings during the pre-policy period of 2018 to 2019. Standard errors are clustered at the group level. ***, **, and * represent statistical significance at the 1%, 5%, and 10% level, respectively.

Table 1.7: Mansion Tax and Filed Permits (Plumbing)

	Scaled No. of Permits		Log No. of Permits	
	(1)	(2)	(3)	(4)
LA city x Post Mansion	-0.044	-0.040	-0.022	-0.020
	(0.056)	(0.059)	(0.088)	(0.110)
Group x Mansion Tax FE	Yes	Yes	Yes	Yes
Zipcode FE	No	Yes	No	Yes
Controls	Yes	Yes	Yes	Yes
Adj. R^2	0.218	0.555		
Pseudo R^2			0.471	0.803
Obs.	345	342	345	342

Notes: Group-by-Mansion-Tax fixed effects are included to control for local trends and time-varying shocks. This table reports the effects of mansion tax on permit filing. The dependent variable is scaled number of filed permits, or the log number of filed permits. Demographic controls include population, income, and housing value. Specifications include group and zipcode fixed effects. Each ZIP code in the treatment and control groups is weighted by its average number of permit filings during the pre-policy period of 2018 to 2019. Standard errors are clustered at the group level. ***, **, and * represent statistical significance at the 1%, 5%, and 10% level, respectively.

Table 1.8: Mansion Tax and Filed Permits (Electrical)

	Scaled No. of Permits		Log No. of Permits	
	(1)	(2)	(3)	(4)
LA city x Post Mansion	0.054	0.032	0.114	0.082
	(0.055)	(0.050)	(0.081)	(0.084)
Group x Mansion Tax FE	Yes	Yes	Yes	Yes
Zipcode FE	No	Yes	No	Yes
Controls	Yes	Yes	Yes	Yes
Adj. R^2	0.200	0.549		
Pseudo R^2			0.652	0.903
Obs.	339	335	339	335

Notes: Group-by-Mansion-Tax fixed effects are included to control for local trends and time-varying shocks. This table reports the effects of mansion tax on permit filing. The dependent variable is scaled number of filed permits, or the log number of filed permits. Demographic controls include population, income, and housing value. Specifications include group and zipcode fixed effects. Each ZIP code in the treatment and control groups is weighted by its average number of permit filings during the pre-policy period of 2018 to 2019. Standard errors are clustered at the group level. ***, **, and * represent statistical significance at the 1%, 5%, and 10% level, respectively.

Table 1.9: Mansion Tax and Filed Permits (Multifamily)

	Scaled No. of Permits		Log No. of Permits	
	(1)	(2)	(3)	(4)
LA city x Post Mansion	-0.300***	-0.270**	-0.331*	-0.238
	(0.099)	(0.109)	(0.182)	(0.204)
Group x Mansion Tax FE	Yes	Yes	Yes	Yes
Zipcode FE	No	Yes	No	Yes
Controls	Yes	Yes	Yes	Yes
Adj. R^2	0.180	0.508		
Pseudo R^2			0.435	0.683
Obs.	245	234	245	234

Notes: Group-by-Mansion-Tax fixed effects are included to control for local trends and time-varying shocks. This table reports the effects of mansion tax on permit filing. The dependent variable is scaled number of filed permits, or the log number of filed permits. Demographic controls include population, income, and housing value. Specifications include group and zipcode fixed effects. Each ZIP code in the treatment and control groups is weighted by its average number of permit filings during the pre-policy period of 2018 to 2019. Standard errors are clustered at the group level. ***, **, and * represent statistical significance at the 1%, 5%, and 10% level, respectively.

Table 1.10: Mansion Tax and Filed Permits (Multifamily exclud. Affordable)

	Scaled No. of Permits		Log No. of Permits	
	(1)	(2)	(3)	(4)
LA city x Post Mansion	-0.274**	-0.247*	-0.319*	-0.229
	(0.109)	(0.118)	(0.192)	(0.211)
Group x Mansion Tax FE	Yes	Yes	Yes	Yes
Zipcode FE	No	Yes	No	Yes
Controls	Yes	Yes	Yes	Yes
Adj. R^2	0.160	0.472		
Pseudo R^2			0.420	0.674
Obs.	245	234	245	234

Notes: Group-by-Mansion-Tax fixed effects are included to control for local trends and time-varying shocks. This table reports the effects of mansion tax on permit filing. The dependent variable is scaled number of filed permits, or the log number of filed permits. Demographic controls include population, income, and housing value. Specifications include group and zipcode fixed effects. Each ZIP code in the treatment and control groups is weighted by its average number of permit filings during the pre-policy period of 2018 to 2019. Standard errors are clustered at the group level. ***, **, and * represent statistical significance at the 1%, 5%, and 10% level, respectively.

Table 1.11: Mansion Tax and Filed Permits by Price Levels (Ground-up)

	(1)	(2)
LA city × Post Mansion	-0.470***	-0.426***
	(0.061)	(0.066)
LA city x Post Mansion x Bin 2	0.271	0.277
	(0.175)	(0.162)
LA city x Post Mansion x Bin 3	-0.141**	-0.189***
	(0.058)	(0.062)
Group x Mansion Tax FE	Yes	Yes
Zipcode FE	Yes	Yes
Controls	Yes	Yes
Adj. R^2	0.556	0.581
Obs.	186	186

Notes: Group-by-Mansion-Tax fixed effects are included to control for local trends and time-varying shocks. This table reports the effects of mansion tax on permit filing. The dependent variable is scaled number of filed permits. Demographic controls include population, income, and housing value. Prices are partitioned into three bins using two cutoffs. *Bin 2* is an indicator for observations in the middle price bin, and *Bin 3* is an indicator for observations in the highest price bin; the omitted category is the lowest price bin (*Bin 1*). The two sets of regressions are for two sets of price level cutoffs. Specifications include group and zipcode fixed effects. Each ZIP code in the treatment and control groups is weighted by its average number of permit filings during the pre-policy period of 2018 to 2019. Standard errors are clustered at the group level. ***, **, and * represent statistical significance at the 1%, 5%, and 10% level, respectively.

Table 1.12: Mansion Tax and Filed Permits by Price Levels (Remodel)

	(1)	(2)
LA city × Post Mansion	-0.117	-0.173
	(0.134)	(0.123)
LA city x Post Mansion x Bin 2	-0.034	-0.027
	(0.109)	(0.086)
LA city x Post Mansion x Bin 3	0.474***	0.460***
	(0.134)	(0.136)
Group x Mansion Tax FE	Yes	Yes
Zipcode FE	Yes	Yes
Controls	Yes	Yes
Adj. R^2	0.599	0.617
Obs.	278	278

Notes: Group-by-Mansion-Tax fixed effects are included to control for local trends and time-varying shocks. This table reports the effects of mansion tax on permit filing. The dependent variable is scaled number of filed permits. Demographic controls include population, income, and housing value. Prices are partitioned into three bins using two cutoffs. *Bin 2* is an indicator for observations in the middle price bin, and *Bin 3* is an indicator for observations in the highest price bin; the omitted category is the lowest price bin (*Bin 1*). The two sets of regressions are for two sets of price level cutoffs. Specifications include group and zipcode fixed effects. Each ZIP code in the treatment and control groups is weighted by its average number of permit filings during the pre-policy period of 2018 to 2019. Standard errors are clustered at the group level. ***, **, and * represent statistical significance at the 1%, 5%, and 10% level, respectively.

Table 1.13: Mansion Tax and No. of Sales

	(1)	(2)	(3)
LA city x Post Mansion	-0.180*	-0.211*	-0.235**
	(0.102)	(0.113)	(0.102)
Month FE	Yes	No	Yes
Zipcode FE	No	Yes	Yes
Pseudo R^2	0.019	0.253	0.270
Obs.	541	479	479

Notes: Month and zipcode fixed effects are included to control for local trends and time-varying shocks. This table reports the effects of mansion tax on number of sales. Demographic controls include population, income, and housing value. Standard errors are clustered at the group level. ***, **, and * represent statistical significance at the 1%, 5%, and 10% level, respectively.

Table 1.14: Mansion Tax and No. of Sales by Price Levels

	(1)	(2)	(3)	(4)
LA city × Post Mansion	0.199***	0.243***	0.178**	0.215**
	(0.064)	(0.076)	(0.064)	(0.077)
LA city x Post Mansion x Bin 2	0.064	0.103	0.092	0.137
	(0.082)	(0.096)	(0.081)	(0.097)
LA city x Post Mansion x Bin 3	-0.905***	-0.724***	-0.807***	-0.572***
	(0.298)	(0.234)	(0.281)	(0.181)
Month FE	Yes	Yes	Yes	Yes
Zipcode FE	Yes	Yes	Yes	Yes
Notch Omitted	No	Yes	No	Yes
Pseudo R^2	0.501	0.517	0.501	0.517
Obs.	4250	2581	4250	2581

Notes: Month and zipcode fixed effects are included to control for local trends and time-varying shocks. This table reports the effects of mansion tax on number of sales. The dependent variable is log number of sales. Demographic controls include population, income, and housing value. The (1) and (2) columns, and (3) and (4) columns, are for two sets of price level cutoffs. Column (2) and (4) exclude the observations in the annoucing period. Standard errors are clustered at the group level. ***, **, and * represent statistical significance at the 1%, 5%, and 10% level, respectively.

Table 1.15: Robustness Checks for Impact on Scaled No. of Permits

DepVar =	Scaled No. of Permits				
	Total New	SFR	MF	ADU	Placebo
	(1)	(2)	(3)	(4)	(5)
Subsample analyses					
(1) Exclude Culver City	-0.424***	-0.446***	-0.271**	-0.760***	-0.039
	(0.097)	(0.120)	(0.109)	(0.189)	(0.059)
(2) Exclude ED1 projects	-0.417***	-0.463***	-0.247*	-0.761***	-0.039
	(0.097)	(0.125)	(0.118)	(0.188)	(0.060)
Alternative Sample Selection					
(3) Drop duplicates within 6 months	-0.437***	-0.476***	-0.298**	-0.802***	-0.035
	(0.093)	(0.120)	(0.111)	(0.181)	(0.059)
(4) Drop duplicates within 9 months	-0.418***	-0.458***	-0.269**	-0.763***	-0.042
	(0.095)	(0.128)	(0.112)	(0.186)	(0.060)
Alternative Group Definition					
(5) Alternative groups	-0.473***	-0.506***	-0.342**	-0.842***	-0.027
	(0.097)	(0.128)	(0.146)	(0.194)	(0.061)

Notes: Group-by-Mansion-Tax fixed effects are included to control for local trends and time-varying shocks. This table reports the robustness checks for the effects of mansion tax on permit filing. Each column represents the types of permit being affected. The dependent variable is scaled number of filed permits. Demographic controls include population, income, and housing value. Specifications include group-mansion-tax and zipcode fixed effects. Each ZIP code in the treatment and control groups is weighted by its average number of permit filings during the pre-policy period of 2018 to 2019. Standard errors are clustered at the group level. ***, **, and * represent statistical significance at the 1%, 5%, and 10% level, respectively.

Table 1.16: Robustness Checks for Impact on Log No. of Permits

DepVar =	Log No. of Permits				
	Total New	SFR	MF	ADU	Placebo
	(1)	(2)	(3)	(4)	(5)
Subsample analyses					
(1) Exclude Culver City	-0.359***	-0.454***	-0.240	-0.475***	-0.020
	(0.135)	(0.130)	(0.203)	(0.142)	(0.110)
(2) Exclude ED1 projects	-0.367***	-0.454***	-0.229	-0.477***	-0.021
	(0.138)	(0.129)	(0.211)	(0.141)	(0.110)
Alternative Sample Selection					
(3) Drop duplicates within 6 months	-0.357***	-0.457***	-0.250	-0.476***	-0.021
	(0.133)	(0.127)	(0.205)	(0.126)	(0.109)
(4) Drop duplicates within 9 months	-0.354***	-0.442***	-0.232	-0.458***	-0.025
	(0.137)	(0.132)	(0.202)	(0.135)	(0.109)
Alternative Group Definition					
(5) Alternative groups	-0.403***	-0.471***	-0.395**	-0.500***	-0.002
	(0.125)	(0.121)	(0.195)	(0.116)	(0.120)

Notes: Group-by-Mansion-Tax fixed effects are included to control for local trends and time-varying shocks. This table reports the robustness checks for the effects of mansion tax on permit filing. Each column represents the types of permit being affected. The dependent variable is log number of filed permits. Demographic controls include population, income, and housing value. Specifications include group-mansion-tax and zipcode fixed effects. Each ZIP code in the treatment and control groups is weighted by its average number of permit filings during the pre-policy period of 2018 to 2019. Standard errors are clustered at the group level. ***, **, and * represent statistical significance at the 1%, 5%, and 10% level, respectively.

Chapter 2

Additive Effect of State Policies on House Prices: Evidence from COVID-19

On March 19, 2020, California became the first state to issue a statewide stay-at-home order, directing nearly 40 million residents to remain at home except for essential activities. Within two weeks, more than 40 states had followed suit, though with strikingly different timing, scope, and stringency. A household in El Paso, Texas, could walk across the state line into Las Cruces, New Mexico, and find itself under a different set of restrictions entirely. This patchwork of state-level responses created a natural experiment: neighboring communities that shared labor markets, school districts, and shopping centers were suddenly subject to different regulatory regimes, purely as an accident of which side of a state border they happened to occupy.

Against this backdrop, the U.S. housing market exhibited behavior that puzzled many observers. After an initial contraction in the spring of 2020, existing home sales rebounded to their highest seasonally adjusted annual rate since the mid-2000s boom, and year-over-year median price growth accelerated into double digits, as shown in Figure 2.1. How much of this housing market upheaval can be attributed to the policies themselves, as distinct from the pandemic that occasioned them?

This paper takes up that question by investigating the *additive* effect of COVID-19 state policies on house prices. The term "additive" is deliberate. Rather than asking whether a single lockdown raised or lowered prices, I examine whether the cumulative layering of successive policy interventions, each tightening restrictions, produced compounding effects on the housing market. The distinction matters because housing decisions are forward-looking: a household contemplating a purchase or relocation responds not only to the current policy environment but also to the trajectory of policy changes and the beliefs those changes engender about future conditions. As Kozlowski et al. (2020) argue in a related context, rapidly evolving policy regimes force continuous updating of beliefs, and the economic consequences of such updating can persist well beyond the duration of any individual measure.

The additive dimension is particularly relevant in the housing context because real estate transactions involve large, illiquid, and long-lived commitments. A household that buys a home in response to one workplace closure is unlikely to reverse course when the next closure is lifted, especially if the purchase has already been made and the household has relocated. Each successive tightening reinforces the expectation that remote work will persist and that the pre-pandemic geographic equilibrium will not return. If these expectations are capitalized into prices, then the cumulative stock of prior policy interventions should have a positive and growing effect on house values, over and above the effect of the most recent intervention. This is precisely the additive effect that the empirical strategy is designed to detect.

Identifying the causal effect of state policies on house prices is complicated by the obvious endogeneity of policy adoption. States with more severe outbreaks, stronger political preferences for intervention, or different underlying economic conditions may have adopted stricter policies, and these factors could also affect housing demand independently. To address this challenge, I adopt a border-discontinuity design. The unit of analysis is a pair of adjacent counties that share a state border but lie in different states, and therefore face different policy regimes. Within each county pair and month, I compare house price changes in the county that received an additional workplace-closure policy to those in the county that did not. The design absorbs county-pair-by-month fixed effects, which sweep out any time-varying shock common to both sides of the border, including local labor market fluctuations, regional infection trends, seasonal patterns, and so forth. Identification rests on the residual within-pair policy variation, which owes entirely to the accident of the state line. Figure 2.2 illustrates the geography of the county pairs that constitute the estimation sample.

I implement this strategy using zipcode-month-level house price data from the Zillow House Value Index (ZHVI), which tracks the typical home value for the middle tier of the distribution; state-daily policy indicators from the Oxford COVID-19 Government Response Tracker (OxCGRT), aggregated to the state-month level; and zipcode-level demographic characteristics from the 2016–2020 American Community Survey. For each county-pair and treatment month t, I estimate a sequence of regressions indexed by the horizon $j \in \{1, \ldots, 12\}$, tracing the effect of the additional policy on house price changes over the following year.

Three findings emerge. First, the immediate impact of an additional workplace-closure policy is contractionary. At horizons $t + 1$ and $t + 2$, the treated county's monthly house price change falls by roughly 0.05 percentage points relative to its cross-border neighbor, a statistically significant decline (Table 2.2). Second, this effect reverses. Beginning around month $t + 4$, the treatment coefficient turns positive and grows in magnitude, reaching statistical significance from $t + 9$ onward. Third, cumulative house price effects, obtained by summing the monthly coefficients or, equivalently, by regressing the cumulative price change on treatment, display a characteristic V-shape: negative for the first three to four months, then steadily positive, with cumulative gains of approximately 0.5 percent by the end of the year (Figures 2.3 and 2.4). This pattern is not an artifact of the baseline sample construction. Restricting the analysis to one-to-one matched county pairs, so that each county enters the sample exactly once, yields qualitatively identical dynamics, though the initial trough is somewhat deeper and persists for one or two additional months (Table 2.3; Figures 2.5 and 2.6).

The reversal from short-run decline to medium-run appreciation is consistent with several channels emphasized in the existing literature. Workplace closures compelled a sudden and widespread shift to remote work (Barrero et al., 2021; Bick et al., 2021), thereby relaxing the geographic constraint that ties workers to their employers' vicinity. The resulting expansion in locational choice increased demand for housing in areas that had previously been marginal commuting locations, precisely the kinds of border counties that populate the estimation sample. At the same time, supply responded sluggishly: construction activity slowed during lockdown months, and the existing stock of homes for sale thinned as owners pulled listings amid uncertainty. The combination of delayed demand and constrained supply produced the price appreciation observed at longer horizons. The additive dimension of the result, captured by the positive and growing coefficient on cumulative prior policies, suggests that each successive tightening reinforced these dynamics, deepening the shift toward remote work and further tightening housing supply.

The housing market is a natural setting for studying the effects of COVID-19 policies for several reasons. Housing is the single largest asset in most American households' portfolios, and changes in house prices have direct consequences for household wealth, consumption, and borrowing capacity. The decision to buy, sell, or relocate is sensitive to expectations about future income, future commuting patterns, and the future trajectory of local economic conditions, all of which are shaped by the policy environment. Moreover, housing transactions are recorded with geographic precision and at high frequency, providing a rich data source for detecting localized treatment effects. And because homes are immobile, prices observed in a given location reflect the willingness to pay of the marginal buyer in that location, making house prices a particularly informative indicator of how policies reshape the spatial distribution of economic activity.

The paper proceeds as follows. Section 3.1 situates the contribution within the existing literature. Section 3.2 describes the data and reports summary statistics. Section 3.3 develops the identification strategy and econometric specification. Section 3.4 presents the baseline results. Section 2.5 reports the robustness exercise with one-to-one matched pairs. Section 3.5

concludes.

2.1 Related Literature

This paper engages with five bodies of work.

The pandemic and residential real estate. Davis et al. (2021) develop a spatial equilibrium model showing that a permanent increase in remote work shifts housing demand from city centers to suburbs. Brueckner et al. (2021) reach a similar conclusion using a monocentric city framework, finding that remote work flattens the bid-rent gradient. Delventhal et al. (2021) calibrate a model of Los Angeles and estimate that full-time remote work would reduce central-city rents by 10 to 15 percent while raising suburban rents comparably. Gupta et al. (2022) provide evidence that the pandemic flattened the bid-rent curve in major U.S. cities, a pattern they attribute to the differential flight of higher-income workers from the urban core. Ramani and Bloom (2021) document a "donut effect" in U.S. cities, with population and business activity moving from dense urban cores to suburban rings. Liu and Su (2021) show that housing demand shifted toward lower-density areas, with the magnitude of the shift proportional to the local share of jobs amenable to remote work. Mondragon and Wieland (2022) estimate that the transition to remote work accounts for more than half of the 24-percent increase in national house prices between November 2019 and November 2021. Ling et al. (2020) show that the pandemic depressed commercial real estate prices, particularly for retail and hospitality properties. The present paper moves beyond the pandemic's aggregate effect on real estate and instead identifies the marginal and cumulative price impacts of specific policy interventions, exploiting within-county-pair variation to isolate the policy channel from the broader pandemic shock.

State policies and economic behavior. Baker et al. (2020) show that consumer spending responded sharply to stay-at-home orders, with restaurant and retail spending falling while grocery and home-improvement spending surged. Choe et al. (2021) document a tight link between stay-at-home orders and reductions in cell-phone-based mobility measures. Coibion et al. (2020) find that lockdown measures account for a substantial share of the pandemic's labor market displacement. Goolsbee and Syverson (2021) separate the effects of government-imposed restrictions from voluntary behavioral changes, finding that fear of infection drove much of the decline in consumer activity, but that state orders had additional binding effects. Chetty et al. (2020) construct high-frequency economic indicators from private-sector data and show that state reopening policies produced only modest recoveries in spending and employment, suggesting that the initial closures had effects that outlasted their formal duration. The present paper extends this line of inquiry to the housing market, where the forward-looking nature of prices means that the dynamic and cumulative effects of policy may differ substantially from the contemporaneous effects observed in spending or employment data.

The rise of remote work. Bick et al. (2021) estimate that roughly 35 percent of the U.S. workforce was working from home full-time by May 2020, a figure that stabilized well above pre-pandemic levels. Barrero et al. (2021) argue that the pandemic permanently raised the equilibrium share of remote work, driven by sunk investments in home offices and demonstrated feasibility. Emanuel and Harrington (2020) provide micro-level evidence that the transition to remote work was associated with modest productivity gains and significant worker preference for continued flexibility. Bloom et al. (2021) document that firms invested heavily in remote-work technologies during 2020, and that these investments are largely irreversible, making a full return to office unlikely. Althoff et al. (2022) map the geography of remote work across U.S. cities and find that it is concentrated in high-skill, high-wage occupations, producing a spatial pattern that amplifies demand for housing in amenity-rich locations outside traditional employment centers. These findings suggest a mechanism through which workplace closures affected housing demand: by untethering workers from their offices, closures expanded the set of locations where households could viably live.

Pandemic-era migration and spatial reallocation. The housing market response to workplace closures is mediated by the geographic mobility of households. Whitaker (2021) use USPS change-of-address data to document that the pandemic accelerated out-migration from dense urban counties, with the largest outflows concentrated in the highest-cost cities. Haslag and Weagley (2021) find that net migration flows reversed sign in many metropolitan areas during 2020, with suburban and exurban counties gaining population at the expense of central cities. Howard (2021) show that the pandemic reduced within-city mobility but increased cross-city migration, consistent with a model in which workplace closures relaxed the geographic constraint on residential location while simultaneously reducing the value of urban amenities. Brueckner and Sayantani (2023) extend the monocentric framework to incorporate inter-city migration, showing that remote work not only flattens the bid-rent gradient within cities but also compresses house price differentials across cities of different productivity levels. These migration patterns are directly relevant to the border-discontinuity design because they suggest that the treatment effect may operate in part through the relocation of households across state borders in response to differential policy regimes.

Persistent belief effects of shocks. Kozlowski et al. (2020) model how a rare disaster such as a pandemic shifts agents' beliefs about the probability of future tail events, permanently altering consumption and investment behavior. Applied to the housing market, this logic implies that each successive policy tightening, each new closure order, each escalation in stringency, reinforces the belief that remote work is the new normal, thereby cumulatively strengthening the incentive to relocate. Bordalo et al. (2020) develop a model of diagnostic expectations in which agents overweight recent experiences when forming beliefs about the future, implying that a sequence of workplace closures would produce a disproportionately large revision in expectations relative to a single closure of the same total duration. Gennaioli and Shleifer (2010) provide a more general framework for understanding how salient events reshape risk perceptions and economic behavior, a mechanism that may help explain why cumulative policies have effects that exceed the sum of their individual parts.

2.2 Data

The empirical analysis brings together three datasets: a comprehensive measure of local house prices, a detailed record of state-level COVID-19 policy interventions, and a rich set of ZIP code-level demographic characteristics. I describe each in turn, then present summary statistics.

2.2.1 House Price Data

House prices are measured using the Zillow House Value Index (ZHVI), a smoothed, seasonally adjusted index of the typical home value in a given geography. Zillow defines "typical" as the 35th-to-65th percentile of the local value distribution, a range intended to represent the broad middle of the market while excluding outliers at both tails. The smoothing algorithm filters out short-term noise in transaction prices, producing a measure that tracks the underlying trend in local market conditions rather than the idiosyncratic variation in individual sales. The ZHVI is available at the zipcode-month level and is disaggregated by property type: all single-family residences, condominiums and co-ops, and homes classified by bedroom count. Throughout the analysis, I work with the broadest category, all homes, although the results are qualitatively similar for single-family residences alone.

The dependent variable is the log first difference of ZHVI, denoted $\Delta\text{House Price}_{i,t}$, which approximates the percentage change in the typical home value for zipcode i from month $t-1$ to month t. This first-differenced specification removes zipcode-level fixed differences in price levels and focuses on the month-to-month growth rate, which is the margin along which the policy treatment is expected to operate. Table 2.1 reports that the sample mean of this variable is 0.0067, or roughly 0.67 percent per month, with a standard deviation of 0.0084.

The 5th percentile stands at -0.0044, and the 95th at 0.0218, spanning a range from modest price declines to brisk appreciation. The dispersion reflects the heterogeneity of local housing market conditions during the pandemic, with some zip codes experiencing falling values in the early months and others seeing rapid appreciation throughout. The estimation sample contains approximately 1.71 million ZIP code-month observations.

An important advantage of the ZHVI relative to transaction-based measures such as repeat-sale indices is that it is available for every ZIP code in every month, regardless of whether any homes actually sold. Transaction-based indices are undefined in zipcode-months with no sales, which can introduce selection bias if the decision to transact is itself affected by the policy treatment. During the early months of the pandemic, transaction volumes fell sharply in many areas, precisely the period in which the treatment is most likely to be active. The ZHVI, because it is model-based rather than transaction-based, avoids this problem and provides a continuous measure of housing market conditions even in thin markets.

2.2.2 State Policy Data

Information on state-level containment measures comes from the Oxford COVID-19 Government Response Tracker (OxCGRT). The OxCGRT records daily observations on 21 ordinal indicators for each U.S. state. These indicators span four broad categories: containment and closure policies, including workplace shutdowns, school closures, public gathering restrictions, and stay-at-home requirements; economic support measures, such as income support and debt relief programs; health system interventions, including testing regimes, contact tracing protocols, and mask mandates; and vaccination campaigns. Each indicator is coded on an ordinal scale, typically ranging from 0 (no policy) to 2 or 3 (the most stringent level). The indicators are further distinguished by whether the policy is "targeted," applying to a specific region or population subgroup, or "general," applying statewide.

The daily frequency of the OxCGRT data is important because it captures the rapid pace at which state governments adjusted their policies during the pandemic. A state might impose a workplace closure on a Monday, relax it the following Friday, and reimpose a more stringent version two weeks later. I aggregate the daily indicators to the state-month level to align with the monthly frequency of the house price data. Within each month, I record the maximum stringency level attained by each indicator, so that a month in which a policy was active for even a few days is coded as having that policy in force.

A central variable in the analysis is the count of "additional policies" that a state has imposed. I define an additional policy as the occurrence of any of the following events within a given month: (i) a previously inactive indicator becomes active; (ii) the stringency of an active indicator increases; or (iii) an existing policy shifts from targeted to general scope. Each such event represents a discrete tightening of the regulatory environment that households and market participants must adjust to. The treatment variable in the regression is an indicator for whether the county in question experienced one more such tightening event than its cross-border neighbor in a given month. This definition focuses on the margin of policy change rather than the level, capturing the incremental effect of each new restriction.

The focus on workplace closures specifically, rather than on the full bundle of containment policies, is motivated by the theoretical mechanism. Workplace closures are the policy instrument most directly linked to the remote-work channel emphasized in the economic interpretation. A school closure or a gathering restriction may affect household behavior in important ways, but it does not directly alter the geographic constraint on where a worker can live. A workplace closure does so by forcing workers to perform job tasks from home, thereby demonstrating that the job can be done remotely. The OxCGRT coding distinguishes between "recommended" and "required" closures, and between closures that apply to "some sectors" versus "all but essential sectors." Each escalation along either dimension constitutes a separate tightening event in the data.

The pace at which states adopted and revised these policies was rapid. During the initial wave in March and April 2020, most states moved from no workplace restrictions to full closures within two to three weeks. The subsequent reopening was more gradual and unevenly paced, with some states lifting restrictions by May while others maintained closures into the summer. A second wave of tightening occurred in late 2020 and early 2021 as infection rates surged. This temporal variation is central to the identification strategy: it means that at any given month, the two counties in a border pair may face different policy regimes because their respective states are at different points in the tightening-relaxation cycle.

2.2.3 Demographic Controls

I control for cross-sectional differences in ZIP code characteristics using data from the U.S. Census Bureau's American Community Survey (ACS), 2016–2020 five-year estimates. The control vector includes log population density, median age, racial composition measured as the shares of the population that are white and Black, the share of the population that is married, log median household income, average household size, the share of adults with a high school diploma, the share with a four-year college degree, the homeownership rate, log median house value, the share of workers who work from home, and a set of commuting-time shares covering the categories under 15 minutes, 15 to 29 minutes, 30 to 44 minutes, 45 to 59 minutes, and 60 to 89 minutes.

These variables serve two purposes. First, they account for observable differences in the socioeconomic and housing-market characteristics of zip codes within a county pair that might independently predict price changes. Two zip codes on opposite sides of a state border may differ in income, education, or homeownership even if they share the same labor market and regional economy. Failing to control for these differences could bias the treatment estimate if, for example, higher-income zipcodes tend to be located in states with more stringent policies and also experience different price dynamics for reasons unrelated to the policy. Second, several of the control variables speak directly to the plausibility of the remote-work mechanism that underlies the economic interpretation of the results. The work-from-home share captures the pre-existing propensity for remote work in each ZIP code. Areas with a higher baseline share of remote workers may respond differently to workplace closures than areas where most workers commute to a physical location. Similarly, the commuting-time distribution captures the geographic relationship between each ZIP code and its associated employment centers. Zip codes with long average commutes are likely farther from urban cores and may stand to gain more from the relaxation of commuting requirements resulting from workplace closures.

2.2.4 Summary Statistics

Table 2.1 reports descriptive statistics for the full set of variables. The average log population density of 4.68 corresponds to roughly 108 persons per square mile. The wide standard deviation of 2.43 confirms that the sample spans a broad cross-section of U.S. geographies, from sparsely populated rural areas at the 5th percentile to dense urban cores at the 95th. The mean homeownership rate is 60 percent, and the median house value averages about $165,000 in levels, consistent with a sample that is tilted toward the middle of the national housing market rather than dominated by expensive coastal metros.

The work-from-home share averages 26 percent, elevated relative to pre-pandemic norms and reflecting the ACS survey's coverage of the 2020 calendar year. The commuting-time distribution is right-skewed, with the largest share of workers, 33 percent, commuting 15 to 29 minutes. Approximately 28 percent commute less than 15 minutes, 20 percent commute 30 to 44 minutes, and progressively smaller shares commute at longer distances. The concentration of commuting times in the 15- to 29-minute range suggests that the typical border-county worker lives within a moderate distance of an employment center, consistent with the suburban and exurban character of many border counties.

The educational attainment variables indicate that the sample ZIP codes are broadly representative of the national population. About 89 percent of adults hold a high school diploma, and 25 percent hold a four-year college degree. The median age is 43, with a range from 29 at the 5th percentile to 60 at the 95th, capturing both young, fast-growing communities and older, more established ones. These patterns are broadly representative of the universe of U.S. zipcodes located in border counties and provide assurance that the results are not driven by an unusual or unrepresentative subsample of the national housing market.

2.3 Identification Strategy

2.3.1 The Cross-Border Design

The empirical strategy rests on a simple geographic insight: counties that share a state border are, in most economically relevant respects, very similar to one another. They draw on the same local labor markets, face similar weather patterns and geographic constraints, and are exposed to common regional shocks. Yet because they lie on opposite sides of a state line, they may be subject to different COVID-19 policies, policies determined at the state capital, often hundreds of miles away, in response to statewide conditions that need not mirror those at the border. This generates quasi-random variation in policy exposure between county pairs, which I exploit for identification.

The advantage of this design over a simple cross-state comparison is that it holds the local economic environment constant. A regression that compares house prices in New York to house prices in Florida, for example, would confound the effect of different COVID-19 policies with the many other ways in which New York and Florida differ: climate, industry composition, tax regimes, demographics, and so forth. The border-discontinuity design narrows the comparison to pairs of counties that are geographically adjacent and economically integrated. The only systematic difference between the two members of a pair is the state in which they are located, and therefore the state policy regime under which they operate.

The approach is a close cousin of the cross-border designs that have become standard in the empirical public finance and labor economics literatures. Dube et al. (2010) use contiguous county pairs to study the employment effects of minimum wage differences, arguing that adjacent counties share similar economic trends and therefore provide a compelling counterfactual for one another. Holmes (1998) exploits the sharp change in state business climate policies at the borders to study the effect of regulation on manufacturing location. The logic here is the same: conditional on pair-by-time effects, the only systematic difference between the two counties in a pair is the policy regime imposed by their respective states.

Figure 2.2 maps the continental United States at the county level, with shading indicating the border counties included in the estimation sample. The sample is geographically dispersed, spanning all major regions of the country. This geographic breadth is important because it means the estimates are not driven by conditions at any single border or in any single region. The treatment arises from policy variation at dozens of different state boundaries, each contributing independent identifying information.

A natural question is whether border counties are representative of the broader U.S. housing market or whether they constitute a peculiar subsample. Border counties tend to be more rural and less densely populated than the national average, though the sample includes several metropolitan border pairs, such as those straddling the Kansas-Missouri border in the Kansas City area or the Pennsylvania-New Jersey border in the Philadelphia-Trenton corridor. The summary statistics reported in Table 2.1 confirm that the sample spans a wide range of demographic and housing market characteristics. The key identifying assumption does not require that border counties be nationally representative; it requires only that, within each pair, the two counties be comparable to each other. The geographic proximity of paired counties makes this a plausible assumption.

2.3.2 Specification

I estimate the following regression for each horizon $j \in \{1, 2, \ldots, 12\}$:

$$Y_{in,t+j} \;=\; \beta\,\text{Treatment}_{in,t} \;+\; \gamma' X_i \;+\; \phi_c \;+\; \alpha_{nt} \;+\; \varepsilon_{int}, \tag{2.1}$$

where i indexes zipcodes, n indexes county-pairs, c indexes counties, and t indexes months.

The dependent variable $Y_{in,t+j}$ is the house price change of zipcode i in county-pair n from month $t+j-1$ to $t+j$: in words, the price growth observed j months after the treatment event. The key regressor, $\text{Treatment}_{in,t}$, equals one if the county containing zipcode i has one more workplace-closure policy than the opposing county in pair n at time t, and zero otherwise. The vector X_i collects the time-invariant zipcode-level demographic controls described above. ϕ_c denotes county fixed effects and α_{nt} denotes county-pair-by-month fixed effects.

The county fixed effects absorb any county-level permanent characteristics that affect its average rate of house price appreciation. Counties with historically faster-growing housing markets, better school districts, or more desirable amenities will have these features absorbed by ϕ_c, so they do not confound the treatment estimate. The county-pair-by-month effects are considerably more demanding: they absorb every time-varying shock that is common to both members of a pair in a given month. This includes local infection rates, which might independently affect housing demand through fear of contagion or hospital overcrowding. It includes labor market fluctuations, which affect household income and the ability to service a mortgage. It includes seasonal housing market cycles, which are well documented in the real estate literature. And it includes national monetary policy, which during 2020 and 2021 pushed mortgage rates to historic lows and stimulated housing demand nationwide. All of these forces affect both counties in a pair equally and are therefore absorbed by α_{nt}. What remains after conditioning on the pair-by-month effects is the within-pair difference in policy exposure and its correlation with the within-pair difference in house price changes. The coefficient β captures precisely this relationship.

By estimating Equation (3.1) separately for $j = 1$ through $j = 12$, I trace out an impulse-response function describing how an additional workplace-closure policy in month t affects house price dynamics over the ensuing year. The horizon structure is essential to the analysis because the hypothesis predicts different effects at different lags. A specification that pooled all twelve horizons into a single regression would average across the negative short-run effect and the positive medium-run effect, producing an attenuated estimate that misses the dynamic pattern entirely. I also examine the cumulative treatment effect, obtained either by summing the point estimates $\sum_{k=1}^{j} \hat{\beta}_k$ or, more directly, by replacing the dependent variable with the cumulative price change from t to $t+j$. The two approaches yield essentially identical results.

Standard errors are clustered at the county-pair level to account for serial correlation within pairs over time and spatial correlation across zipcodes within the same county. Clustering at a coarser level, such as the state, would be inappropriate in this setting because the variation of interest is within county pairs, not across states. The county-pair cluster allows arbitrary correlation in the error term across all zip codes and time periods within a given pair, which is the most conservative level of clustering consistent with the structure of the identifying variation.

It is worth noting how this design differs from a standard difference-in-differences approach. In a typical difference-in-differences setting, the researcher compares treated and untreated units before and after a single treatment event. Here, the treatment is not a single event but a continuous sequence of policy tightenings, each of which constitutes a separate treatment. The regression is estimated separately at each horizon, with the treatment month t varying across observations. This means that the identifying variation comes from many different treatment events occurring at different times and at different borders throughout the sample period. The staggered nature of the treatment is a strength of the design because it reduces the likelihood that any single confounding event could drive the results.

2.3.3 Identifying Assumptions and Threats

The key identifying assumption is that, conditional on county-pair-by-month fixed effects and demographic controls, the timing and incidence of one county receiving an additional policy, relative to its cross-border neighbor, are uncorrelated with unobserved determinants of that county's house price changes. This assumption would be violated if, for example, state governments systematically tightened restrictions in response to county-level housing market conditions near the border. Such feedback is implausible: state COVID-19 policies were driven by statewide infection trends, hospital capacity, and political considerations, not by real estate prices in a handful of border counties.

A second concern is that the treatment might be correlated with unobserved county-level characteristics that vary over time. Suppose, for instance, that counties on the stricter side of the border tend to have higher-income residents, and that these residents' housing demand is more sensitive to macroeconomic conditions. If macroeconomic conditions fluctuate over the sample period, the estimated treatment effect could partly reflect the interaction of income levels and aggregate shocks rather than the causal effect of the policy. The county-pair-by-month fixed effects address this concern by absorbing any shock that is common to both counties in the pair. What remains is the within-pair difference, which by construction cannot be driven by aggregate conditions.

A subtler concern is that differential sorting or migration across the border could both respond to policy differences and affect house prices. If households move from the more restrictive county to the less restrictive one, this would increase demand on the less restrictive side and depress it on the more restrictive side, generating a pattern that is the opposite of the one I find at medium horizons. Alternatively, if workplace closures induce remote workers to relocate to the more restrictive side because it offers lower housing costs, this migration-induced demand would be part of the mechanism of interest rather than a confound. I note that the time horizon of the analysis, twelve months, is long enough for some migration to have occurred but short enough that the housing stock is essentially fixed, so any price effect attributable to migration-induced demand shifts is itself part of the treatment effect of interest. In the concluding section, I discuss the potential to disentangle the migration channel from the remote-work channel using administrative data on household moves.

2.4 Results

2.4.1 Main Estimates

Table 2.2 reports the estimates of Equation (3.1) for all twelve horizons. Each column corresponds to a separate regression; the first row reports the coefficient on Treatment, and the second row reports the coefficient on the count of cumulative policies in the county as of month $t-1$.

The treatment effect at horizon $t+1$ is -0.00054, significant at the one-percent level. In economic terms, a county that receives one additional workplace-closure policy relative to its cross-border neighbor experiences a reduction of 0.054 percentage points in its monthly house price growth, a roughly eight-percent decline relative to the sample mean of 0.67 percent per month. The effect at $t+2$ is nearly identical in magnitude at -0.00056, also significant at the one-percent level. By $t+3$, the coefficient has attenuated to -0.00030 and is no longer statistically distinguishable from zero, signaling a rapid dissipation of the initial contractionary impulse.

The reversal begins at $t+4$, where the treatment coefficient turns positive at 0.00037, significant at the five-percent level. Positive point estimates persist from this horizon onward, though they are imprecisely estimated at some intermediate horizons. The coefficients tighten and become strongly significant at $t+9$, $t+10$, and $t+11$, with the $t+11$ estimate of 0.00105 being the largest in magnitude. At $t+12$, the point estimate remains positive but is no longer significant, possibly reflecting the attenuation of the policy-specific effect as the housing market absorbs the shock.

The second variable of interest, the cumulative policy count at $t-1$, provides direct evidence of the additive effect. This coefficient is near zero at $t+1$ and $t+2$ but rises steadily over subsequent horizons, reaching 0.00109 at $t+11$, significant at the one-percent level. The interpretation is that, controlling for the current period's treatment, the stock of prior policy tightenings exerts an independent and growing positive influence on house price growth. The finding is difficult to reconcile with a purely transitory view of policy effects and instead supports the hypothesis that successive interventions cumulatively reshaped housing demand.

All specifications include the full set of zipcode-level demographic controls, county fixed effects, and county-pair-by-month fixed effects.

2.4.2 Dynamic Treatment Effects

Figure 2.3 provides a graphical summary of the dynamic pattern. The colored dots plot the point estimates $\hat{\beta}_j$ at each horizon, with vertical bars spanning the 95-percent confidence interval. The connected line traces the cumulative sum $\sum_{k=1}^{j} \hat{\beta}_k$, which represents the implied total effect of the treatment on the house price *level* through horizon j.

The picture tells a clear story. The cumulative effect descends into negative territory at $t+1$, deepens through $t+3$ where it reaches a trough of approximately -0.0019, and then begins a steady ascent. The zero line is crossed around $t+8$, and the cumulative effect climbs sharply thereafter, peaking near $+0.0019$ at $t+11$. The magnitude of the swing, roughly 0.38 percentage points from trough to peak, is economically meaningful. It implies that, over a one-year horizon, the additional policy generates a net cumulative price appreciation of approximately 0.2 percent relative to the counterfactual of no additional policy.

To put this number in perspective, the median home value in the sample is approximately \$165,000. A cumulative price effect of 0.2 percent translates to roughly \$330 in additional home equity for the median homeowner. While modest for any individual household, the aggregate effect across the roughly 33,000 zipcodes in the estimation sample is substantial. Moreover, the effect is measured relative to the cross-border neighbor rather than to a national average. The absolute price change may be considerably larger if both counties in the pair experienced positive appreciation, but the treated county appreciated more.

The confidence intervals for the monthly coefficients provide insight into the precision of the estimates. At horizons $t+1$ and $t+2$, the intervals are tight and exclude zero, indicating that the negative short-run effect is precisely estimated. At intermediate horizons from $t+4$

through $t+8$, the intervals are wider and typically include zero, reflecting the transitional period in which the effect is shifting sign and the cross-sectional heterogeneity in response timing is greatest. At $t+9$ through $t+11$, the intervals again tighten and exclude zero, confirming that the positive medium-run effect is estimated with reasonable precision.

2.4.3 Cumulative House Price Change

Figure 2.4 presents an alternative visualization that reinforces the same conclusion. Here, the dependent variable in each horizon's regression is the *cumulative* house price change from month t to month $t+j$, rather than the month-on-month change. Each dot therefore represents the total price impact of the treatment over the entire interval from t to $t+j$.

The resulting trajectory is a smooth, monotonically increasing curve that starts below zero and crosses into positive territory around $t+5$. The point estimates grow to approximately 0.003 by $t+9$, 0.005 by $t+11$, and 0.006 by $t+12$, with confidence intervals that exclude zero from $t+6$ onward.

The cumulative specification complements the month-by-month estimates in Table 2.2 by providing a direct answer to the question that a policymaker or homeowner would ask: what is the total effect on the price level of having one additional workplace-closure policy imposed in month t? The answer is that prices are lower for the first four to five months but higher by the end of the year, with the positive cumulative gain more than offsetting the initial decline. This V-shaped pattern is important because it suggests that the short-run and medium-run effects operate through different mechanisms. The short-run decline reflects the immediate contraction in housing demand that accompanies the announcement of a new restriction. The medium-run appreciation reflects the structural reallocation of demand that the restriction sets in motion, operating through the remote-work channel and the supply-side tightening described below.

2.4.4 Economic Interpretation

The two-phase dynamic response, contraction followed by appreciation, maps onto a straightforward economic narrative.

In the short run, an additional workplace closure acts as a demand shock. Prospective buyers, confronted with new uncertainty about their employment, income, and the pandemic's trajectory in their area, postpone purchase decisions. Real estate agents report fewer showings, and transaction volumes thin. Sellers, unwilling to accept lower offers in a market they perceive as temporarily disrupted, may pull listings rather than reduce asking prices. The net effect is a decline in the rate of price appreciation, consistent with the negative coefficients at $t + 1$ and $t + 2$. The speed with which this effect appears, within a single month of the policy announcement, suggests that the demand response is driven by uncertainty and expectations rather than by any material change in economic fundamentals. A workplace closure announced in the first week of a month can affect buyer behavior before any economic impact is felt.

In the medium run, several countervailing forces emerge. The most prominent is the remote-work channel. Workplace closures compel firms and workers to adopt remote arrangements, and Barrero et al. (2021) document that a large fraction of these arrangements prove durable. Workers who discover that they can perform their jobs from home are no longer tethered to a commuting radius around the office. For those working in expensive urban centers, this opens the possibility of relocating to less costly areas, including the kinds of border-county locations that make up the estimation sample. The resulting inflow of demand pushes prices upward. The effect is amplified by the fact that remote workers relocating from high-cost cities to lower-cost border counties bring with them purchasing power that is large relative to local price levels, allowing them to bid aggressively for housing.

A second channel operates through supply. Housing construction is slow to initiate and even slower to complete. If the initial policy shock reduces new construction starts or delays projects in progress, the supply of available homes will be tighter in subsequent months, amplifying the price effect of any demand recovery. The pandemic period saw widespread reports of construction delays attributable to labor shortages, supply chain disruptions in building materials, and regulatory uncertainty. These supply-side constraints were not confined to the counties that experienced workplace closures, but were likely more severe there, reinforcing upward pressure on prices.

Third, the accumulation of policies may itself be informative. Each new tightening reinforces the perception that remote work is not a fleeting arrangement but a structural shift. To the extent that house prices capitalize expectations about future living patterns, the additive signal embedded in repeated closures should raise valuations in areas that stand to benefit from decentralization. A single workplace closure might be dismissed as a temporary measure. A third or fourth closure in the same state, each more stringent than the last, sends a much stronger signal about the permanence of the new equilibrium.

Disentangling these channels is beyond the scope of the present paper, but the pattern of coefficients, and particularly the growing positive effect of cumulative prior policies, is most naturally read as reflecting a permanent rather than temporary reallocation of housing demand.

A fourth channel, related to the third but distinct from it, operates through the mortgage market. Households that decide to relocate in response to workplace closures must obtain financing, and the mortgage application and approval process typically takes 30 to 60 days. This processing lag means that demand generated by a workplace closure in month t may not translate into a completed transaction and a recorded price change until $t + 2$ or $t + 3$, thereby delaying the emergence of the positive effect. Moreover, the historically low mortgage rates that prevailed throughout 2020 and 2021 reduced the carrying cost of homeownership, amplifying the demand response to any given relocation incentive. While the county-pair-by-month fixed effects absorb the common component of mortgage rate variation, the interaction between low rates and differential policy exposure may have produced a larger treatment effect than would have obtained in a higher-rate environment.

Finally, the results are consistent with a behavioral channel whereby the salience of workplace closures, rather than their economic content, shapes household expectations. A workplace closure is a highly visible, widely reported event that dominates local news coverage and social media. Its salience may cause households to overweight the probability of future closures and to overinvest in the flexibility provided by a housing location that does not depend on daily commuting. This is the mechanism formalized by Bordalo et al. (2020) in the context of diagnostic expectations. If households are diagnostic rather than rational, the cumulative effect of a sequence of closures on housing demand will exceed the rational-expectations benchmark, which may help explain the large magnitudes observed at longer horizons.

2.5 Robustness: One-to-One Matched County Pairs

A potential concern with the baseline specification is that some counties participate in multiple county pairs. A large county that borders several states, for example, will appear once in each pair, and its house price data will be used repeatedly. If such counties have unusual price dynamics, their over-representation could distort the estimates. This concern is particularly relevant for counties at the intersection of three or more state borders, which may appear in the sample multiple times and receive disproportionate weight in the estimation.

To address this, I construct a one-to-one matched sample in which each county is paired with exactly one cross-border neighbor. The matching eliminates repeated use of any county and produces a cleaner, though smaller, estimation sample. The one-to-one restriction reduces the number of observations by roughly 30 percent relative to the baseline, but it ensures that no county's price dynamics unduly influence the results.

Table 2.3 reports the results. The broad contours of the dynamic response are preserved. At $t + 1$, the treatment effect is -0.00047, significant at the five-percent level. At $t + 2$, it deepens to -0.00062, significant at the one-percent level. At $t + 3$, it remains negative and significant at -0.00060. Relative to the baseline, the initial contraction is somewhat more persistent: the coefficients at $t+4$ and $t+5$ remain negative but insignificant, whereas in the baseline the reversal was already underway by $t+4$. This persistence is consistent with the matched sample placing greater weight on county pairs where the policy asymmetry is particularly stark, generating a stronger and more prolonged short-run demand response.

The positive phase, when it arrives, is comparable in magnitude to the baseline. The $t+9$ coefficient is 0.00066, significant at the five-percent level, and the $t+11$ coefficient is 0.00101, significant at the one-percent level, both close to their baseline counterparts. The cumulative policy variable again exhibits a positive and strengthening association with price growth at longer horizons, confirming that the additive effect is not an artifact of the baseline sample construction.

Figure 2.5 overlays the monthly coefficients and the cumulative treatment effect for the matched sample. The cumulative line traces a deeper trough than in the baseline, reaching roughly -0.002 near $t+6$, before climbing through zero around $t+9$ and into positive territory by $t+10$.

Figure 2.6 shows the treatment effects on cumulative house price change in the matched sample. The qualitative pattern is unchanged: negative and declining through approximately $t+4$, crossing zero around $t+5$ to $t+6$, and growing steadily positive thereafter. By $t+11$, the cumulative effect reaches roughly $+0.004$, and by $t+12$ it exceeds $+0.005$. These magnitudes are modestly smaller than in the baseline but remain economically and statistically significant.

Taken together, the baseline and matched-sample results tell a consistent story. The precise timing of the reversal and the depth of the initial trough differ slightly across specifications, as one would expect when the sample composition changes. But the fundamental pattern, namely short-run contraction, medium-run appreciation, and a cumulative effect that is positive over a one-year horizon, is robust. The fact that the one-to-one matched sample, which eliminates any mechanical correlation introduced by repeated county observations, produces nearly identical results strengthens the case for a causal interpretation.

It is also instructive to compare the magnitude of the treatment effects across the two samples. The $t+1$ coefficient is -0.00047 in the matched sample versus -0.00054 in the baseline, a difference of roughly 13 percent. At $t+11$, the coefficient is 0.00101 in the matched sample versus 0.00105 in the baseline, a difference of less than 4 percent. The convergence of the estimates at longer horizons is consistent with the view that the short-run effect is somewhat sensitive to sample composition, while the medium-run effect, which reflects deeper structural forces, is not. The observation counts in the matched sample are approximately 30 percent smaller than in the baseline, reflecting the loss of county pairs when the one-to-one restriction is imposed. Despite this substantial reduction in sample size, the standard errors increase only modestly, and the key coefficients remain significant at conventional levels.

One might wonder whether the results are sensitive to the definition of "additional policy" used to construct the treatment variable. The baseline definition codes a treatment event whenever any of the 21 OxCGRT indicators increases in stringency, shifts from targeted to general, or is newly activated. An alternative would be to restrict attention to workplace-closure policies specifically, ignoring changes in school closures, gathering restrictions, or other containment measures. I focus on workplace closures because they are the policy most directly relevant to the remote-work mechanism emphasized in the economic interpretation. The results are qualitatively similar when the treatment is defined more broadly to include all types of policy tightenings, though the magnitudes are somewhat smaller, consistent with the expectation that non-workplace policies have a weaker effect on housing demand than workplace closures per se.

2.6 Conclusion

This paper has asked whether the sequential tightening of COVID-19 state policies produced an additive effect on U.S. house prices. Exploiting the discontinuous change in policy regimes at state borders, I compare house price trajectories in adjacent county-pairs where one county experienced an additional workplace-closure policy, and its cross-border neighbor did not.

The answer is yes, and the effect unfolds in two phases. In the first two months, the additional policy depresses house price growth, consistent with a short-run demand contraction driven by uncertainty and reduced market activity. Over the following quarters, the effect reverses, and prices in the treated county begin to outpace those in the control county. By the end of one year, the cumulative house price effect is unambiguously positive. Moreover, the coefficient on the stock of prior policies is positive and growing over time, indicating that successive tightenings do not merely repeat the same impulse but compound it, each reinforcing the shift in residential demand that workplace closures set in motion.

These findings carry several broader implications. First, they suggest that the pandemic-era run-up in house prices was not solely a consequence of low interest rates or pent-up demand, but was partly driven by the policy responses to COVID-19 themselves. The border-discontinuity design isolates the policy channel from the aggregate monetary and fiscal forces that simultaneously affected all housing markets. The fact that the policy effect is positive and growing at longer horizons implies that government containment measures played a direct role in reshaping the geography of housing demand.

Second, the results underscore the importance of studying the dynamic and cumulative effects of policy, rather than focusing on the contemporaneous impact of a single intervention. In a rapidly evolving policy environment, the steady accumulation of restrictions can reshape economic behavior in ways that no individual measure, taken in isolation, would predict. The additive structure documented here suggests that the housing market response to COVID-19 policies was not simply a sum of independent shocks but a compounding process in which each new restriction reinforced the behavioral changes initiated by prior ones.

Third, the findings have implications for the welfare consequences of containment policies. If workplace closures durably increase house prices in certain areas, the distributional effects depend on who lives there. Existing homeowners in treated counties experienced capital gains. Prospective first-time buyers in the same counties faced higher barriers to entry. And renters may have seen their housing costs increase without any corresponding gain in wealth. A full welfare analysis would need to account for these heterogeneous effects, as well as the public health benefits of containment.

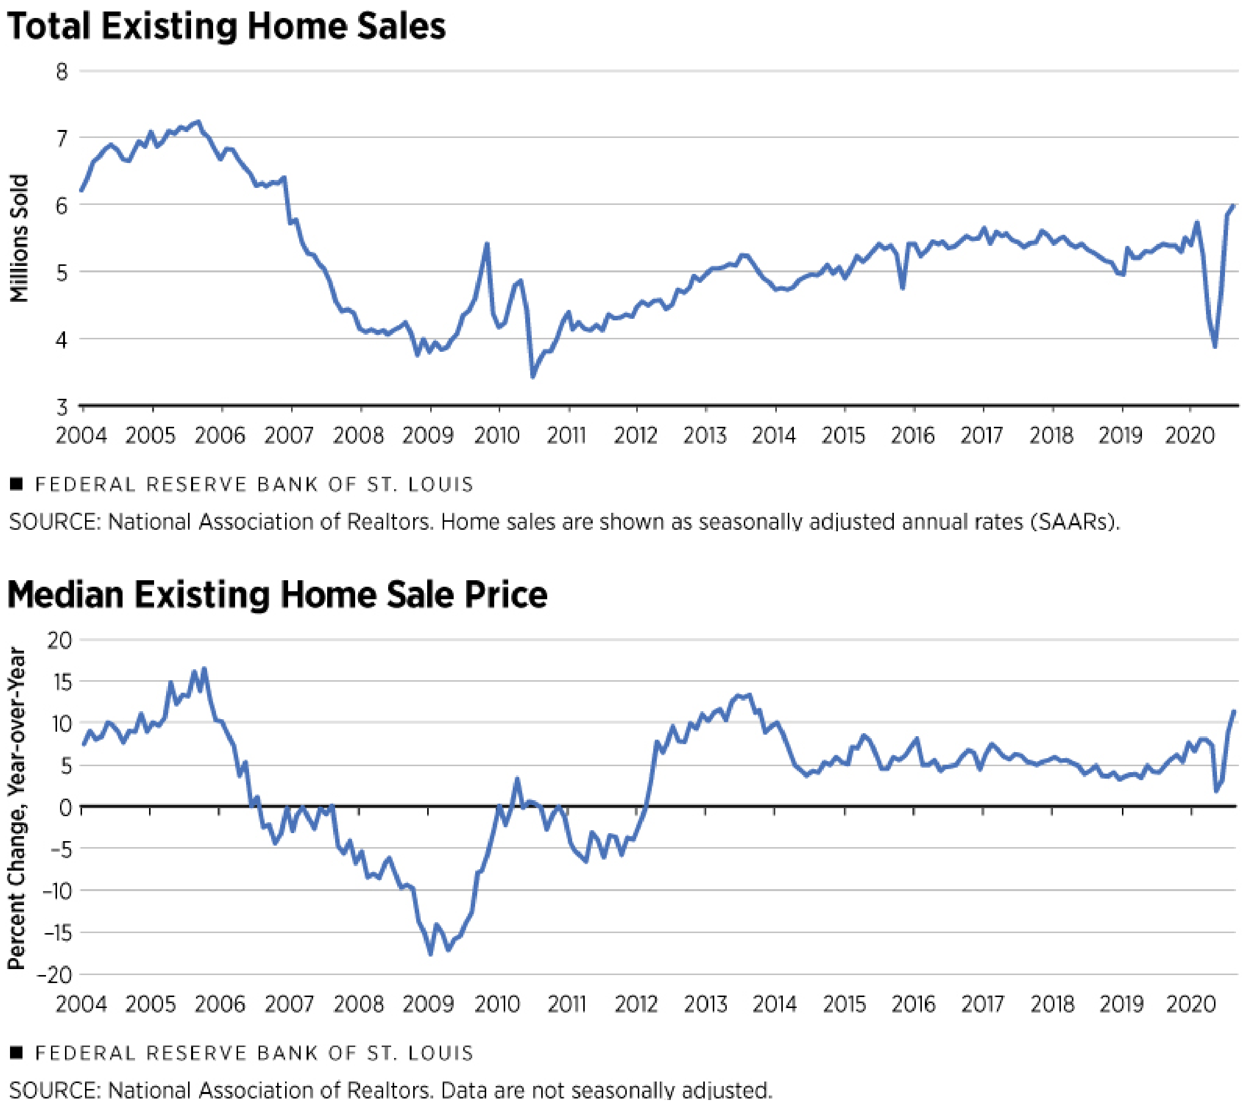

Figure 2.1: U.S. Housing Market Trends, 2004–2020

Notes: Panel (a) plots total existing home sales in millions at a seasonally adjusted annual rate (SAAR) from 2004 to 2020. The vertical axis is in millions of units. Panel (b) plots the year-over-year percent change in the median existing home sale price over the same period, not seasonally adjusted. The vertical axis is in percentage points. Source: National Association of Realtors via the Federal Reserve Bank of St. Louis (FRED).

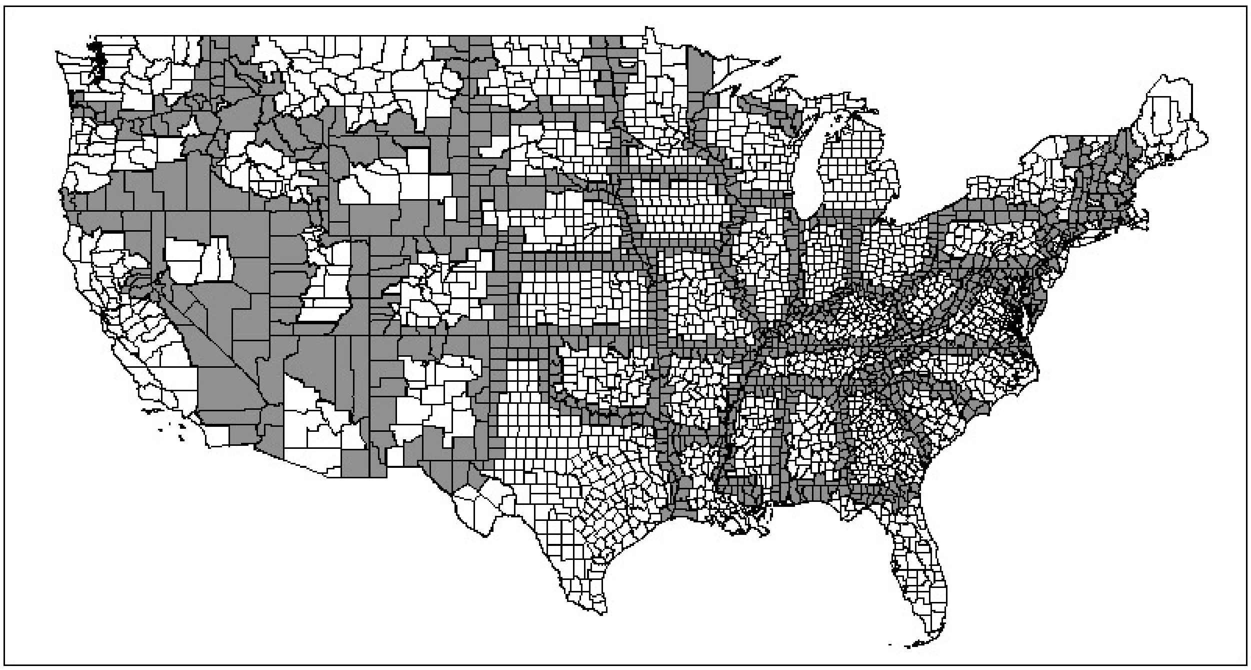

Figure 2.2: Geography of Cross-Border County Pairs

Notes: Each shaded area represents a county included in the estimation sample. A county pair consists of two adjacent counties that share a state border and therefore face different state-level policy regimes. Gray and white shading distinguish the two sides of each pair. The sample covers state borders in the West, Midwest, South, and Northeast.

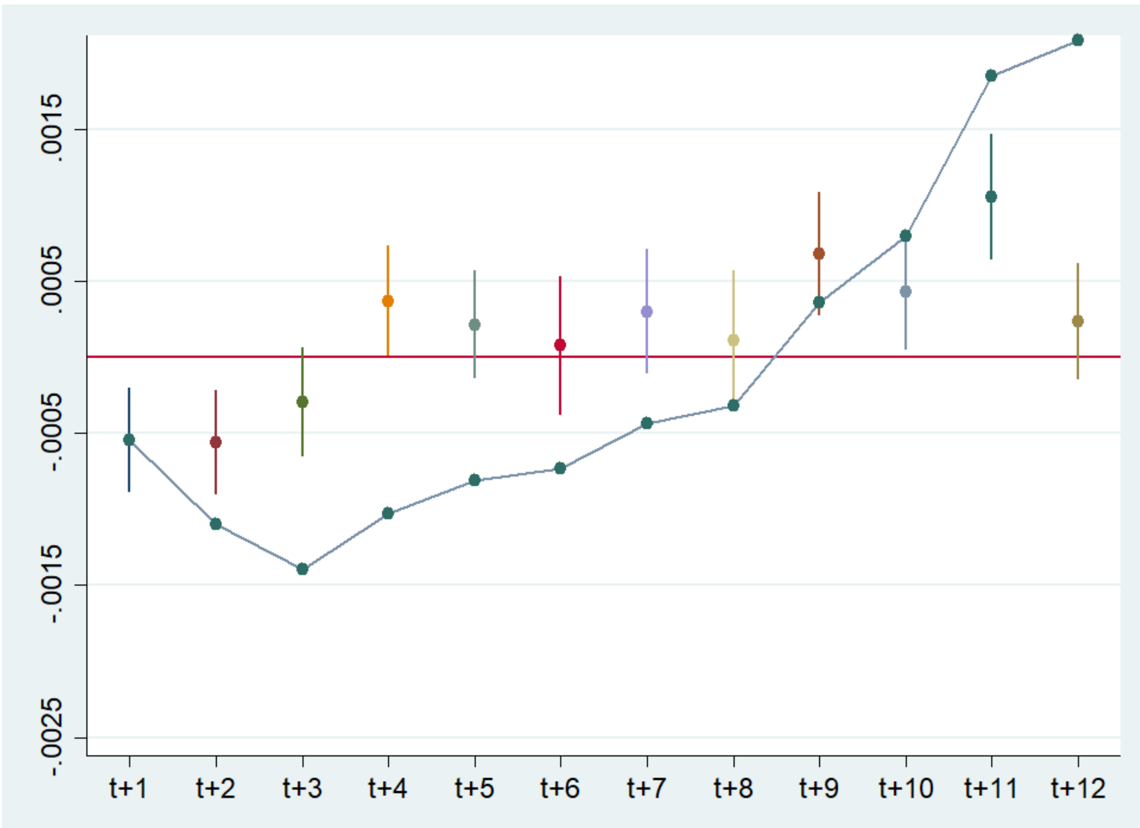

Figure 2.3: Dynamic Treatment Effects: Baseline Sample

Notes: Each dot plots the estimated coefficient $\hat{\beta}_j$ from Equation (3.1) at horizon $j = 1, \ldots, 12$, where Treatment equals one if the zipcode's county received one more workplace-closure policy than its cross-border neighbor in month t. Vertical bars indicate 95-percent confidence intervals. The connected line traces the cumulative treatment effect $\sum_{k=1}^{j} \hat{\beta}_k$. The horizontal line at zero is shown for reference. All regressions include zipcode-level demographic controls, county fixed effects, and county-pair-by-month fixed effects.

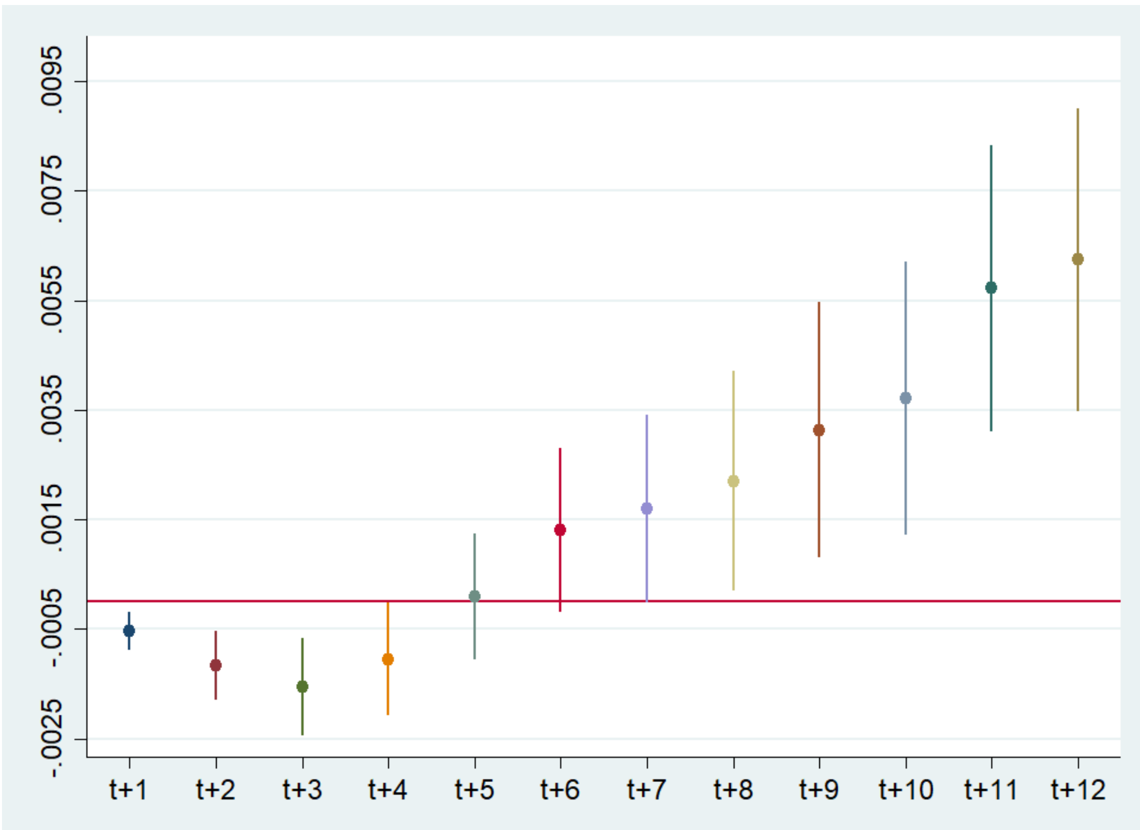

Figure 2.4: Treatment Effects on Cumulative House Price Change: Baseline Sample

Notes: Each dot reports the coefficient on Treatment from a separate regression in which the dependent variable is the cumulative house price change from month t to month $t+j$, for $j = 1, \ldots, 12$. Treatment equals one if the zipcode's county received one more workplace-closure policy than its cross-border neighbor in month t. Vertical bars denote 95-percent confidence intervals. The horizontal line at zero is shown for reference. All regressions include zipcode-level demographic controls, county fixed effects, and county-pair-by-month fixed effects.

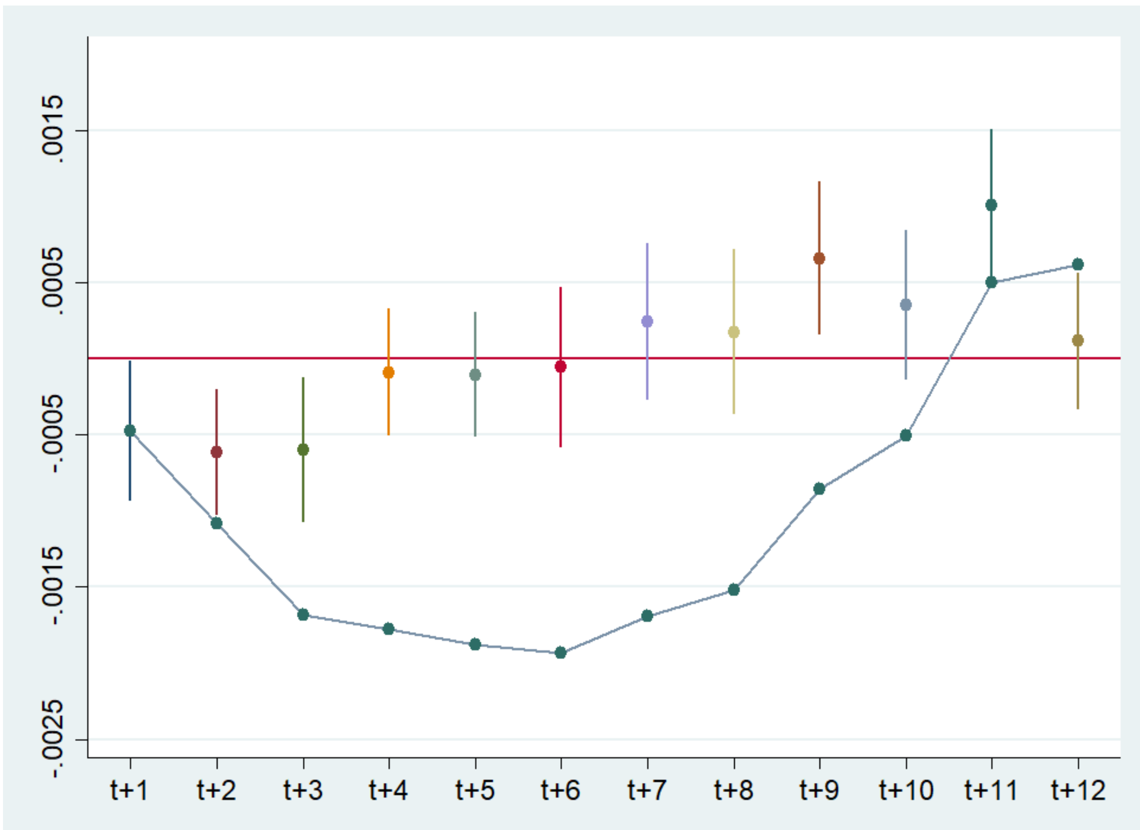

Figure 2.5: Dynamic Treatment Effects: One-to-One Matched County Pairs

Notes: Replication of Figure 2.3 using the one-to-one matched sample, in which each county is paired with exactly one cross-border neighbor. Each dot plots the estimated coefficient $\hat{\beta}_j$ from Equation (3.1) at horizon $j = 1, \ldots, 12$, where Treatment equals one if the zipcode's county received one more workplace-closure policy than its cross-border neighbor in month t. Vertical bars indicate 95-percent confidence intervals. The connected line traces the cumulative treatment effect $\sum_{k=1}^{j} \hat{\beta}_k$. The horizontal line at zero is shown for reference. All regressions include zipcode-level demographic controls, county fixed effects, and county-pair-by-month fixed effects.

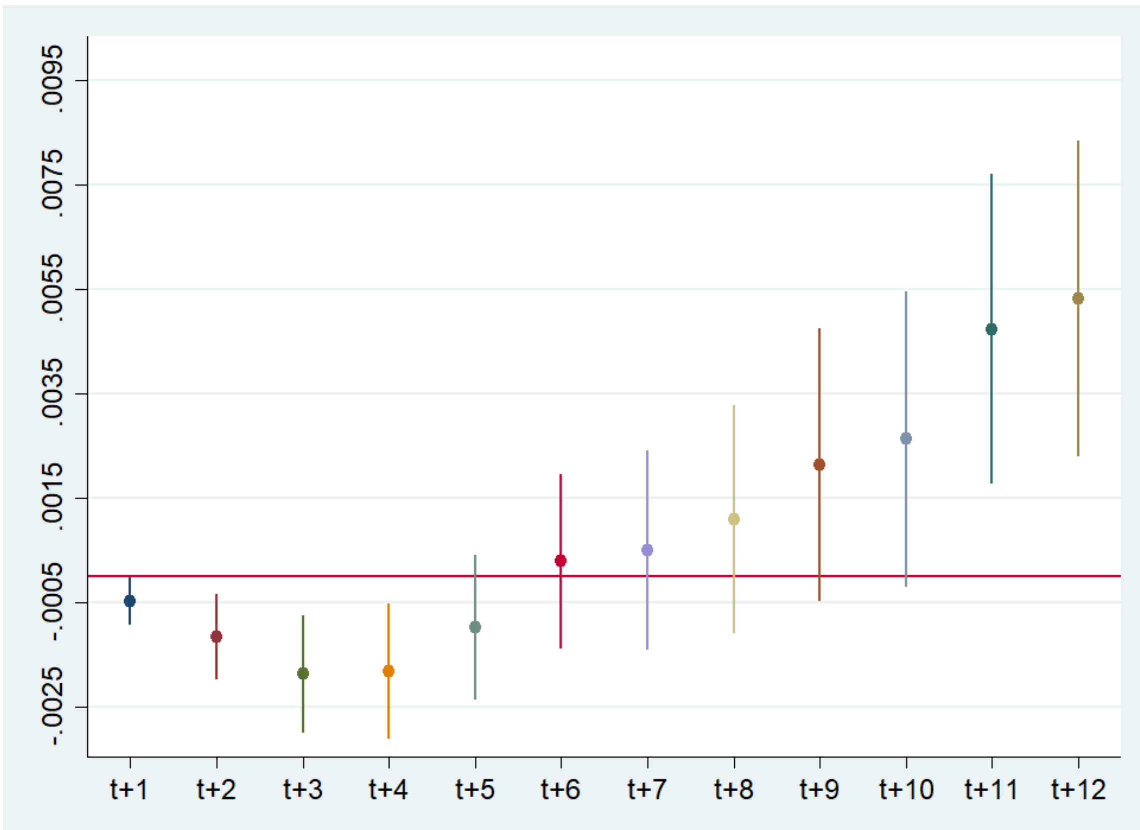

Figure 2.6: Treatment Effects on Cumulative House Price Change: One-to-One Matched County Pairs

Notes: Replication of Figure 2.4 using the one-to-one matched sample, in which each county is paired with exactly one cross-border neighbor. Each dot reports the coefficient on Treatment from a separate regression in which the dependent variable is the cumulative house price change from month t to month $t + j$, for $j = 1, \ldots, 12$. Treatment equals one if the zipcode's county received one more workplace-closure policy than its cross-border neighbor in month t. Vertical bars denote 95-percent confidence intervals. The horizontal line at zero is shown for reference. All regressions include zipcode-level demographic controls, county fixed effects, and county-pair-by-month fixed effects.

Table 2.1: Summary Statistics

	Mean	Std	p5	p50	p95	Obs
Dependent variable						
Δ House price	0.0067	0.0084	−0.0044	0.0056	0.0218	1,708,430
Control variables						
Log population density	4.68	2.43	0.86	4.39	8.66	32,620
Median age	43	9.44	29	42	60	32,373
Percentage of white	0.76	0.26	0.17	0.86	1	32,620
Percentage of black	0.08	0.16	0	0.01	0.42	32,620
Percentage of married	0.54	0.14	0.28	0.55	0.73	32,619
Log median household income	10.98	0.39	10.35	10.96	11.64	30,439
Average household size	2.50	0.65	1.78	2.48	3.36	32,620
Pct. high school graduates	0.89	0.09	0.71	0.91	0.99	32,584
Pct. 4-year college graduates	0.25	0.18	0.04	0.21	0.61	32,584
Pct. housing owned	0.60	0.19	0.24	0.62	0.86	32,389
Log median house value	12.02	0.68	11.07	11.97	13.26	30,259
Pct. work-at-home	0.26	0.07	0.15	0.36	0.39	30,326
Pct. commuting $<$ 15 min	0.28	0.18	0.04	0.24	0.64	30,320
Pct. commuting 15–29 min	0.33	0.15	0.08	0.33	0.57	30,320
Pct. commuting 30–44 min	0.20	0.13	0.00	0.19	0.41	30,320
Pct. commuting 45–59 min	0.08	0.08	0.00	0.07	0.22	30,320
Pct. commuting 60–89 min	0.06	0.08	0.00	0.04	0.18	30,320

Notes: p5, p50, and p95 denote the 5th, 50th, and 95th percentiles, respectively. Obs denotes the number of observations. Δ House price is the month-over-month log change in the Zillow House Value Index (ZHVI) at the zipcode level. Log population density is the natural log of population per square mile. Percentage of white, Percentage of black, and Percentage of married are expressed as shares of the total population. Log median household income is the natural log of median household income in dollars. Log median house value is the natural log of the median self-reported home value in dollars. Pct. work-at-home is the share of workers whose primary mode of commuting is working from home. Commuting shares are defined over workers who do not work from home and sum to less than one because the 90-plus-minute category is omitted. Source: ZHVI from Zillow; control variables from the U.S. Census Bureau American Community Survey (ACS) 2016–2020 five-year estimates.

Table 2.2: Workplace Closing and House Price Change in Subsequent Months

Δ House Price	t+1	t+2	t+3	t+4	t+5	t+6	t+7	t+8	t+9	t+10	t+11	t+12
Treatment	-0.00054^{***}	-0.00056^{***}	-0.00030	0.00037^{**}	0.00022	0.00008	0.00030	0.00012	0.00068^{***}	0.00043^{**}	0.00105^{***}	0.00024
	(0.00018)	(0.00018)	(0.00018)	(0.00018)	(0.00018)	(0.00023)	(0.00021)	(0.00023)	(0.00021)	(0.00019)	(0.00021)	(0.00019)
# Cumulative Policies ($t-1$)	0.00005	0.00004	0.00050^{**}	0.00068^{***}	0.00076^{***}	0.00067^{**}	0.00062^{***}	0.00062^{***}	0.00076^{***}	0.00100^{***}	0.00109^{***}	0.00072^{***}
	(0.00020)	(0.00018)	(0.00020)	(0.00023)	(0.00027)	(0.00026)	(0.00023)	(0.00021)	(0.00028)	(0.00026)	(0.00022)	(0.00021)
Observations	130,330	130,160	130,623	128,676	122,517	115,856	112,969	112,756	110,119	109,621	106,184	104,769
Control	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
County-pair-month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Each column reports estimates from a separate regression of Equation (3.1) at horizon j. The dependent variable Δ House Price is the monthly log change in the Zillow House Value Index at the zipcode level. Treatment is an indicator equal to one if the zipcode's county received one more workplace-closure policy than its cross-border neighbor in month t, and zero otherwise. # Cumulative Policies ($t-1$) is the running count of all policy tightenings in the county through month $t-1$. "Control" indicates the inclusion of zipcode-level demographic controls from the ACS: log population density, median age, racial composition, marital share, log median household income, average household size, educational attainment, homeownership rate, log median house value, work-from-home share, and commuting-time shares. County FE denotes county fixed effects. County-pair-month FE denotes county-pair-by-month fixed effects. Standard errors are clustered at the county-pair level and reported in parentheses. $^{***}\ p < 0.01$; $^{**}\ p < 0.05$; $^{*}\ p < 0.10$.

Table 2.3: Workplace Closing and House Price Change in Subsequent Months: One-to-One Matched County Pairs

Δ House Price	t+1	t+2	t+3	t+4	t+5	t+6	t+7	t+8	t+9	t+10	t+11	t+12
Treatment	-0.00047^{**}	-0.00062^{***}	-0.00060^{**}	-0.00009	-0.00011	-0.00006	0.00024	0.00018	0.00066^{**}	0.00035	0.00101^{***}	0.00012
	(0.00023)	(0.00021)	(0.00024)	(0.00021)	(0.00021)	(0.00027)	(0.00026)	(0.00027)	(0.00026)	(0.00025)	(0.00026)	(0.00023)
# Cumulative Policies ($t-1$)	-0.00002	-0.00005	0.00038	0.00057^{**}	0.00055	0.00045	0.00042	0.00061^{**}	0.00065^{*}	0.00080^{**}	0.00089^{***}	0.00057^{**}
	(0.00027)	(0.00021)	(0.00025)	(0.00027)	(0.00034)	(0.00034)	(0.00027)	(0.00025)	(0.00034)	(0.00034)	(0.00024)	(0.00025)
Observations	93,634	93,560	93,840	92,226	87,907	83,389	81,383	81,128	79,071	78,767	76,305	75,185
Control	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
County-pair-month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Replication of Table 2.2 using a one-to-one matched sample in which each county is paired with exactly one cross-border neighbor, eliminating repeated use of counties that border multiple states. Each column reports estimates from a separate regression of Equation (3.1) at horizon j. The dependent variable Δ House Price is the monthly log change in the Zillow House Value Index at the zipcode level. Treatment is an indicator equal to one if the zipcode's county received one more workplace-closure policy than its cross-border neighbor in month t, and zero otherwise. # Cumulative Policies ($t-1$) is the running count of all policy tightenings in the county through month $t-1$. "Control" indicates the inclusion of zipcode-level demographic controls from the ACS: log population density, median age, racial composition, marital share, log median household income, average household size, educational attainment, homeownership rate, log median house value, work-from-home share, and commuting-time shares. County FE denotes county fixed effects. County-pair-month FE denotes county-pair-by-month fixed effects. Standard errors are clustered at the county-pair level and reported in parentheses. $^{***}\, p < 0.01$; $^{**}\, p < 0.05$; $^{*}\, p < 0.10$.

Chapter 3

Supply Chain Network Disruption by Natural Disasters

In August 2017, Hurricane Harvey stalled over the Texas Gulf Coast for four days, dumping more than 60 inches of rain on parts of Houston and inundating one of the world's largest petrochemical complexes. Refineries, polymer plants, and chemical facilities shut down across the region. Within days, automakers in Michigan reported parts shortages. Packaging companies in the Midwest scrambled to find substitute resins. Retailers on the East Coast warned of delivery delays for products they had never associated with hurricanes in Texas. The episode illustrated, in concentrated form, a phenomenon that plays out more quietly after hundreds of smaller disasters each year: the economic damage from a natural disaster does not stay where it lands.

The aggregate scale of this damage is staggering. Between 1960 and 2019, natural disasters in the United States caused a cumulative $1.14 trillion in direct damage, claimed nearly 35,000 lives, and injured over 250,000 people.[1] The pace of destruction has accelerated markedly over the past three decades. Weather-related losses alone exceeded $130 billion in 2017, and the year 2005 still holds the record at $133.6 billion following Hurricane Katrina.

[1] All aggregate disaster statistics in this paragraph are drawn from Arizona State University's Center for Emergency Management and Homeland Security. See Figure 3.1 for a detailed breakdown.

In an average year, the United States sustains $19.1 billion in direct damage, 583 fatalities, and 4,206 injuries from natural hazards. These are not isolated events. They are a recurring feature of economic life across virtually every region of the country.

The composition of this damage is heavily concentrated in hydrological and meteorological events. As Figure 3.2 documents, flooding accounts for the single largest share of total losses at $335 billion, or 29 percent of the cumulative total, followed closely by hurricanes and tropical storms at $307 billion, or 27 percent. Severe weather events contribute $162 billion, while geophysical events, drought and heat, tornadoes, winter weather, and wildfires account for the remainder. The geographic footprint is correspondingly broad. Figure 3.3 displays the county-level distribution of losses for a single year, 2013, revealing that damage exceeding $1 million occurs in counties spanning the Gulf Coast, the Great Plains, the Pacific Northwest, and portions of the Northeast.

These aggregate figures, however, measure only the damage that occurs where the disaster physically strikes. A factory destroyed by a tornado in Oklahoma, a warehouse flooded in Houston, or a port shut down by a hurricane in Florida incur losses recorded at the point of impact. What these figures do not capture is the economic disruption that radiates outward through the web of commercial relationships connecting the affected facility to its customers and, in turn, to their customers further down the production network. Modern supply chains are long and specialized. A single finished good may depend on dozens of intermediate inputs, each sourced from a different firm, often in a different state or country. The failure of any one link in this chain can interrupt the entire production process. When a natural disaster halts production at a supplier, the consequences do not remain at the supplier's factory gate. Downstream customers that depend on that supplier for critical inputs face a sudden interruption in the flow of materials, components, or services. Production lines may slow, delivery schedules slip, and contractual obligations go unfulfilled. Workers at the customer firm may idle while waiting for inputs that do not arrive. And because the customer's own customers, in turn, depend on the customer's output, the disruption can cascade through multiple tiers of the production network, amplifying the initial shock and transmitting it to firms located hundreds or thousands of miles from the disaster zone.

The mechanisms through which a supplier disaster disrupts downstream customers are multiple and reinforcing. A *halt in production* is the most direct channel: if a factory is physically damaged or workers cannot reach the workplace due to flooding or evacuation orders, the supplier's output drops to zero. For customers that rely on just-in-time delivery and maintain thin inventories, even a brief interruption can ripple into production delays. *Logistics disruption* compounds the problem, as natural disasters frequently damage roads, bridges, rail lines, and port facilities, making it difficult to ship goods even from undamaged facilities nearby. *Workforce displacement* may prevent the supplier from resuming operations even after physical repairs are complete, as employees who have been evacuated or whose homes have been destroyed may be unable to return for weeks or months. And *local property damage* may destroy not only the supplier's plant and equipment but also its finished goods awaiting shipment. Each of these channels reinforces the others. A supplier whose factory is damaged, whose workers are displaced, and whose local transportation links are severed faces a far longer recovery than one suffering from a single disruption.

Faced with such disruption, how should a rational downstream customer respond? The hypothesis motivating this paper posits that customers will react along two dimensions with distinct time profiles. In the short run, the customer's most urgent need is to secure continuity of supply. This may involve procuring emergency inventory from alternative suppliers, possibly at premium prices, or investing in new equipment to bring previously outsourced processes in-house. The customer may also search for and establish new supplier relationships to replace the damaged source. Firms may expedite existing orders from unaffected suppliers, pay for faster shipping, or expand capacity at their own plants to reduce dependence on the disrupted input. All of these responses involve capital outlays that would register as increases in the customer's investment rate, and they should appear quickly, within the first quarter following the disaster.

In the medium run, however, the consequences for revenue are less favorable. If the supplier's disruption is prolonged and the customer's inventory buffers are eventually exhausted, the customer's own output will fall, translating into lower sales. The key insight is that inventories act as a buffer, absorbing the initial shock. In the weeks and months immediately following the supplier's disaster, the customer can continue to produce by drawing on safety stock, completing work in process already on the factory floor, and fulfilling orders from finished-goods inventory. From the outside, the firm's revenue looks normal during this period. But inventories are finite. Once they are depleted, the customer must either slow production or stop altogether, and the revenue consequences become visible. Because this depletion takes time, the sales decline should materialize only with a lag of several quarters, well after the initial investment spike. The hypothesis, in short, predicts an immediate increase in investment and a delayed decrease in sales growth.

To test this hypothesis, I assemble a dataset that links four sources. Supply chain relationships are identified using the Compustat segment file, which records customers to whom a firm reports sales exceeding 10 percent of total revenue, as mandated by SFAS 131. Disaster exposure is measured at the county-month level using the Spatial Hazard Events and Losses Database for the United States (SHELDUS), maintained by Arizona State University. Firm-level financial data, including quarterly sales, capital expenditure, total assets, and investment rates, come from the merged CRSP-Compustat database. The final sample is constructed by mapping each supplier to its headquarters county, flagging quarters in which that county experienced significant property damage, and tracing the subsequent financial outcomes of the supplier's listed customers. The merge links these datasets through GVKEY identifiers for Compustat, PERMNO identifiers for CRSP, and county FIPS codes for SHELDUS.

The estimation strategy is a panel regression of customer-level outcomes on an indicator for whether the customer's supplier was hit by a natural disaster in the current or preceding quarter. The regression includes controls for firm size, profitability, and the number of supplier relationships, as well as firm fixed effects, year-quarter fixed effects, customer industry fixed effects, and supplier ZIP code fixed effects. In more demanding specifications, I additionally include customer ZIP code fixed effects. The identifying variation comes from the staggered timing of natural disasters across counties and the cross-sectional variation in which firms happen to be linked to affected suppliers through pre-existing supply chain relationships.

Three sets of results emerge. First, the effect of a supplier disaster on customer sales growth is negligible in the first quarter but grows steadily negative over subsequent quarters. At the four-quarter horizon, the treated customer's annual sales growth declines by 2.3 percentage points relative to an unaffected customer, a significant effect at the one-percent level. The gradual emergence of this effect is consistent with the depletion of inventory buffers. The result survives sample restrictions that exclude customers who were themselves located in a disaster area or who were never hit by any disaster over the entire sample period, confirming that the channel runs through the supply chain rather than through geographic proximity.

Second, the investment response is immediate and large. Customers whose suppliers are hit increase their investment rate by approximately 11 percentage points in the quarter following the event, significant at the one-percent level. The sample mean investment rate is 0.16, indicating a near doubling of the normal rate. The coefficient is stable across all specifications.

Third, heterogeneity analysis reveals that the investment response is moderated by the structure of the customer's supply base. When the disaster indicator is interacted with measures of supplier-industry diversification, the interaction terms are negative and significant, indicating that customers whose inputs span a broader set of industries face a smaller investment spike following an upstream disruption. This pattern is consistent with the view that diversification across input markets provides a form of natural insurance.

The paper proceeds as follows. Section 3.1 discusses the related literature. Section 3.2 describes the data and reports summary statistics. Section 3.3 develops the identification strategy. Section 3.4 presents the results. Section 3.5 concludes.

3.1 Related Literature

This paper sits at the intersection of two literatures: one on the direct firm-level effects of natural disasters, and another on the propagation of economic shocks through production networks.

Direct effects of natural disasters on firms. Seetharam (2017) quantifies the indirect spatial impact of hurricanes through firms' internal plant networks, estimating that for every manufacturing job lost at the point of impact, an additional 0.19 to 0.25 jobs are lost in distant, undamaged regions connected through multi-establishment firms. Bourdeau-Brien (2017) examines a broader set of weather disasters and finds that hurricanes and major floods produce the largest and most persistent damage to operating income. Basker and Miranda (2018) study establishment-level outcomes after Hurricane Katrina, documenting elevated closure rates and the critical role of financing in firm survival. Dessaint and Matray (2017) show that managers of firms located near recent hurricane strikes increase cash holdings as a precautionary measure, even when their own operations were undamaged. Giroud and Mueller (2019) find that multi-plant firms transmit local economic shocks to distant establishments through internal reallocation of capital and labor.

These studies focus on firms that are directly exposed to disaster, or on propagation within the boundaries of a single corporate entity. The present paper shifts attention to propagation across firm boundaries through commercial supplier-customer relationships.

Production networks and shock propagation. The theoretical foundation is Acemoglu et al. (2012), who show that idiosyncratic shocks to individual sectors can generate aggregate fluctuations when the input-output network is sufficiently asymmetric. Henriet et al. (2012) extend this framework to natural disasters, demonstrating that the network position of affected firms can multiply direct losses by a factor of two or more. Fujiwara and Aoyama (2010) document that the Japanese firm-level production network exhibits scale-free properties, implying that disruption to a small number of hub firms can cascade widely.

The most directly related empirical work is Barrot and Sauvagnat (2016), who use the Compustat segment file and natural disasters to show that affected suppliers impose substantial output losses on their customers, with propagation concentrated among suppliers producing specific, hard-to-substitute inputs. Carvalho et al. (2021) exploit the 2011 Great East Japan Earthquake to document that disruption propagated both upstream and downstream along supply chains, reducing Japan's real GDP growth by 0.47 percentage points in the following year. Todo et al. (2015) study the same earthquake and find that firms with more diversified supply networks recovered more quickly.

The present paper contributes along three dimensions. First, I document the contrasting time profiles of two outcome variables, an immediate investment increase and a delayed sales decline, a margin that has not been examined in the propagation literature. Second, I show that the magnitude of propagation depends on the industrial diversification of the customer's supplier base, complementing the product-level specificity channel in Barrot and Sauvagnat (2016). Third, the dynamic structure of the analysis provides indirect evidence on the depth of inventory buffers and the typical duration of disaster-induced supply disruptions.

3.2 Data

The empirical analysis brings together four datasets: a record of firm-to-firm supply chain relationships, a catalog of U.S. natural disaster events and their associated damage, a database of firm-level financial fundamentals, and stock return data for firm identification and matching. I describe each source in turn.

3.2.1 Supply Chain Relationships

Supply chain links are identified from the Compustat Segment File. Under SFAS 131, publicly traded firms are required to disclose the identity of any customer that accounts for more than 10 percent of total sales. The data are reported at the annual level. For each fiscal year, the supplying firm reports the name of each qualifying customer and the dollar amount of sales to that customer. I match the reported customer names to Compustat GVKEYs to create a directed network of supplier-customer pairs.

The 10-percent disclosure threshold means that the observed network is weighted toward the most economically significant commercial links. Smaller or more diffuse relationships are not captured. While this introduces a truncation, it ensures that every link in the sample represents a relationship large enough to plausibly transmit a meaningful economic shock. If a supplier loses the ability to deliver to a customer that accounts for more than 10% of its revenue, the disruption is material to the supplier. Symmetrically, if a customer loses access to a supplier whose deliveries constitute a significant fraction of the customer's input requirements, the disruption is material for the customer as well. The links that fall below the threshold are, by construction, those for which the bilateral exposure is relatively small, and for which the propagation effect, if any, would be harder to detect.

The data reveal a pronounced size asymmetry between the two sides of the relationship. Because the threshold is defined relative to the supplier's revenue, the observed network disproportionately captures relationships between smaller suppliers and larger customers. A small supplier selling specialized components to a large manufacturer is likely to have that manufacturer represent more than ten percent of its sales, even if the manufacturer sources only a tiny fraction of its total inputs from that supplier. This asymmetry, confirmed by the summary statistics discussed below, is useful for the research design because it means that the disaster shock originates at the smaller firm, but the outcome of interest is measured at the larger firm, reducing the likelihood that reverse causality drives the results.

3.2.2 Natural Disaster Data

Information on disaster exposure comes from SHELDUS, the Spatial Hazard Events and Losses Database for the United States, maintained by Arizona State University. SHELDUS records property and crop damage for 18 categories of natural hazards at the county-month level, with damage values adjusted to constant 2015 U.S. dollars. The 18 categories include floods, hurricanes, tropical storms, tornadoes, severe thunderstorms, hail, wildfires, earthquakes, landslides, winter storms, drought, heat waves, and several others. The database covers the period from 1960 to the present, though the analysis in this paper focuses on the years when Compustat segment data and quarterly financial data are jointly available.

Table 3.1 reports the ten states with the highest per capita property damage from 1976 to 2013. Hawaii leads the ranking at \$539.55 per capita, driven by its exposure to volcanic activity, tropical storms, and tsunamis. Washington ranks second at \$473.53, Louisiana third at \$236.56, and Florida fourth at \$158.40. The remainder of the top ten includes Mississippi, Alaska, California, New Mexico, North Dakota, and Texas. The list spans Pacific islands, the Pacific Northwest, the Gulf Coast, the Great Plains, and the Southwest, confirming that disaster risk in the United States is not confined to any single region or hazard type. This geographic breadth is important for the research design because it ensures that identification is not driven by a handful of events in a single area of the country.

I define a county-month as "disaster-hit" if total property damage exceeds a threshold representing economically significant destruction. A supplier is coded as hit if its headquarters county experiences such an event.[2]

[2]Headquarters-based coding may not coincide with the location of the supplier's production facilities, but it is standard in the literature and has the advantage of being available for all firms in Compustat. To the extent that some suppliers' operations are located in counties different from their headquarters, the coding introduces measurement error that would bias the estimates toward zero, making the results conservative.

3.2.3 Firm Financial Data

Firm-level financial variables come from the CRSP-Compustat Merged Database. For each firm-quarter, I observe total assets, quarterly net sales, capital expenditure, and operating income. The investment rate is defined as capital expenditure divided by one-quarter-lagged total assets, a standard normalization that allows comparison across firms of different sizes. Sales growth is defined as the percentage change in quarterly sales over the relevant horizon. I measure cumulative sales growth over one, two, three, and four quarters following the supplier disaster event, which allows me to trace the dynamic profile of the revenue response. The CRSP component provides stock return data and permanent firm identifiers that facilitate the merge.

3.2.4 Sample Construction

The final sample is constructed through a multi-step merge. First, the Compustat segment file identifies all supplier-customer pairs in each fiscal year. Because the segment data are annual while the outcome data are quarterly, each annual supplier-customer link is assumed to hold for the four quarters of the corresponding fiscal year. Second, each supplier is mapped to its headquarters county using FIPS codes from the Compustat location variables. Third, SHELDUS disaster data are merged to the supplier's county at the monthly level, and the monthly disaster indicators are aggregated to the quarterly frequency. Fourth, the resulting supplier-disaster dataset is linked to the customer's quarterly financial data through GVKEY identifiers. Finally, CRSP PERMNO identifiers are used to bring in stock return data and confirm the match.

For each customer-quarter, I construct the key treatment variable, Supplier Hit_{it}, equal to one if any of customer i's reported suppliers was located in a county that experienced a qualifying disaster in the current or preceding quarter, and zero otherwise. The preceding-quarter clause accounts for the fact that a disaster occurring late in one quarter may not affect supplier output until the next. A customer with multiple suppliers is coded as treated if at least one supplier was hit. This is a conservative definition in that it does not distinguish between a customer with a single affected supplier and one with multiple affected suppliers. To the extent that having more affected suppliers leads to greater disruption, the binary treatment indicator will underestimate the marginal effect.

3.2.5 Summary Statistics

Tables 3.2 and 3.3 report descriptive statistics separately for customers and suppliers.

Customer firms are large. The median customer has total assets of roughly \$3.2 billion and quarterly sales of about \$808 million. Even at the 5th percentile, customer assets are \$77 million and sales are \$16 million, indicating that the customer sample consists overwhelmingly of mid-cap to large-cap firms. Capital expenditure at the median is \$84 million per quarter. The mean investment rate is 0.16, with a standard deviation of 0.15 and a 95th percentile of 0.47, reflecting the right-skewed distribution of capital intensity across the sample. The average customer maintains relationships with about 18 reported suppliers, though there is substantial heterogeneity. The median is only 6, indicating that most customers have relatively few major suppliers, while the 95th percentile is 79, reflecting a tail of firms with complex, extensive supply bases. The disaster-hit indicator averages 0.01, confirming that direct disaster exposure is rare in any given quarter.

Suppliers are an order of magnitude smaller. The median supplier has total assets of \$104 million and quarterly sales of \$23 million, roughly one-thirtieth the size of the median customer on both dimensions. Median capital expenditure is just \$2.15 million per quarter. The mean investment rate for suppliers is 0.21, higher than the customer mean of 0.16, consistent with the tendency of smaller, higher-growth firms to invest a larger share of their asset base. Suppliers report an average of 2.57 major customers, with a median of 2 and a 95th percentile of 6. The relatively small number of major customers per supplier reflects the concentrated nature of the supplier's revenue base and underscores why the loss of even one major customer, or the inability to deliver to that customer, represents a material event.

The supplier disaster-hit rate is 0.02, modestly higher than for customers, consistent with suppliers being more numerous and more geographically dispersed.

3.3 Identification Strategy

3.3.1 Empirical Design

The empirical strategy exploits the interaction of two sources of variation: the geographic incidence of natural disasters, which determines which suppliers are disrupted and when, and the pre-existing structure of supply chain networks, which determines which customers are exposed to those disrupted suppliers. A customer's treatment status is determined by the coincidence of where its supplier happens to be headquartered and where a natural disaster happens to strike, a combination that is plausibly exogenous to the customer's own financial condition, conditional on the fixed effects described below.

The unit of observation is a customer-firm-quarter. The baseline specification is:

$$Y_{it} \;=\; \alpha_t \;+\; \alpha_i \;+\; \beta\,\text{Supplier Hit}_{it} \;+\; \gamma'\,X_{it} \;+\; \varepsilon_{it}, \tag{3.1}$$

where Y_{it} is an outcome variable for customer firm i in quarter t. In the profitability analysis, Y_{it} is cumulative sales growth from quarter t to quarter $t+k$, for $k \in \{1, 2, 3, 4\}$. In the investment analysis, Y_{it} is the one-quarter change in the investment rate. Supplier Hit_{it} equals one if at least one of firm i's reported suppliers was located in a disaster-hit county, and zero otherwise. X_{it} includes log total assets, which controls for the well-documented relationship between firm size and growth; operating profitability, which captures the firm's current financial health and its capacity to absorb shocks; and the number of supplier relationships the customer maintains, which proxies for the complexity of the customer's supply chain and its ex ante diversification. α_t and α_i denote year-quarter and firm fixed effects, respectively.

The firm fixed effects absorb all permanent characteristics of the customer, including its industry, managerial quality, baseline growth trajectory, geographic location, historical investment patterns, and any other time-invariant attributes. A customer that happens to be located in a disaster-prone area, or that happens to source inputs from suppliers in disaster-prone areas, will have these permanent features absorbed by the firm effect. What remains is the within-firm variation over time. The year-quarter effects absorb aggregate conditions common to all firms in a given period, such as macroeconomic fluctuations, Federal Reserve policy, commodity price movements, and seasonal patterns.

In augmented specifications, I additionally include customer industry fixed effects, supplier ZIP code fixed effects, and customer ZIP code fixed effects. Customer industry effects reflect sector-specific investment and sales trends that may coincide with disaster timing. For example, if construction firms tend to experience higher investment around the same time that natural disasters strike, this correlation would be absorbed by the industry effect rather than attributed to the supply chain channel. Supplier ZIP code effects absorb all time-invariant geographic characteristics of the supplier's location, including baseline disaster risk, local economic conditions, and proximity to transportation infrastructure. Customer ZIP code effects apply to the customer's location as well.

3.3.2 Identifying Assumption

The key identifying assumption is that, conditional on firm and time effects, the timing of a natural disaster in a supplier's county is uncorrelated with unobserved determinants of the customer's outcomes. This assumption requires that the timing of natural disasters is exogenous to the customer's financial trajectory. Given that hurricanes, floods, tornadoes, and earthquakes are determined by meteorological and geophysical forces rather than by firm-level economic conditions, this is a reasonable starting point.

The most obvious threat, that customers with disaster-exposed suppliers are permanently different from other customers, is addressed by the firm fixed effects. A customer that selects into relationships with suppliers in flood-prone coastal areas may differ from one that sources from inland suppliers, but these differences are time-invariant and are absorbed by the firm effect. The remaining concern is a time-varying confounder that simultaneously triggers a disaster in the supplier's county and affects the customer's outcomes through a non-supply-chain channel. One might worry, for example, that a regional economic downturn both increases the severity of a disaster's impact and depresses the sales of firms connected to that region through demand linkages. This scenario is unlikely for several reasons. The customer and supplier are typically in different counties and often in different states, so a local demand shock at the supplier's location would not reach the customer directly. The year-quarter effects absorb any aggregate conditions that might coincide with the timing of disasters. And the supplier ZIP code effects absorb any persistent local economic characteristics of the supplier's area.

Most importantly, I estimate the regressions on restricted samples that progressively exclude customers with direct disaster exposure. The "Customer Not Hit at t" restriction removes customer-quarters in which the customer's own county experienced a disaster at the same time as the supplier's county. This ensures that the estimated effect is not driven by situations in which both the supplier and the customer were hit by the same regional event. The "Customer Never Hit" restriction is more stringent: it excludes any customer that was ever located in a disaster-hit county at any point during the entire sample period. This addresses the deeper concern that firms located in chronically disaster-prone areas might differ from those in safer areas in unobservable, time-varying ways that correlate with disaster timing. If the coefficient on Supplier Hit survives both restrictions, the supply chain channel is the most plausible remaining explanation.

3.4 Results

3.4.1 Supplier Disasters and Customer Sales Growth

Table 3.4 reports the central results on customer profitability. Each column corresponds to a different horizon, from one to four quarters. All specifications include firm-level controls, firm fixed effects, customer industry fixed effects, year-quarter fixed effects, and supplier ZIP code fixed effects.

At the one-quarter horizon, the coefficient on Supplier Hit is 0.002, statistically indistinguishable from zero. In the quarter immediately following an upstream disaster, customer sales are essentially unaffected. This is consistent with the role of safety stock. Most manufacturing and retail firms hold finished goods, work-in-process, and raw materials in reserve to buffer against supply interruptions. In the first quarter after a supplier is hit, the customer can draw on these inventories to sustain production and fulfill existing orders. To an outside observer examining the customer's quarterly income statement, nothing appears to have changed.

At the two-quarter horizon, the point estimate turns negative at -0.010 but remains insignificant. The inventory buffer is beginning to erode, but the average effect across the full sample is still within the range of normal variation. This is the transitional period during which some customers, particularly those with thinner inventories or greater dependence on the affected supplier, are starting to experience output declines, while others with deeper reserves continue to produce at normal rates. The average across these heterogeneous responses is negative but imprecisely estimated.

By the three-quarter horizon, the coefficient reaches -0.021 and is significant at the ten-percent level. For most affected customers, the inventory cushion has been largely exhausted by this point, and the lost or delayed inputs are translating into measurable declines in output and revenue. The effect is now large enough and sufficiently concentrated across the sample to exceed the conventional threshold for statistical significance.

At the four-quarter horizon, the coefficient stands at -0.023, significant at the one-percent level. This is the central estimate of the paper. A customer whose supplier is hit by a major natural disaster experiences a 2.3 percentage-point decline in annual sales growth relative to an otherwise identical customer whose supplier was unaffected. The precision of the estimate improves at the four-quarter horizon because the cumulative nature of the dependent variable smooths out quarter-to-quarter noise in sales, and because the inventory buffer has fully unwound, so the disruption effect is present across a broader cross-section of treated customers.

The monotonic progression from zero at quarter one to -0.023 at quarter four is exactly the pattern one would expect if disruption operates through inventory depletion. It also provides indirect evidence on the typical duration of the underlying supply disruption. If disasters were short-lived and suppliers could resume full production within a few weeks, the inventory buffer would suffice to absorb the shock, and no sales effect would appear at any horizon. The fact that the effect materializes only by the third quarter and deepens through the fourth suggests that the typical disaster-induced disruption lasts several months, long enough to exhaust the customer's reserves. This is consistent with evidence from the disaster recovery literature indicating that reconstruction of damaged industrial facilities routinely takes six to twelve months.

3.4.2 Robustness of the Sales Effect

Table 3.5 subjects the four-quarter result to a sequence of robustness tests. Each test addresses a specific identification concern by altering the sample or the fixed-effect structure.

The first column reproduces the baseline without customer industry fixed effects. The coefficient is -0.023, significant at the one-percent level. Adding customer industry effects in the second column leaves the estimate unchanged. The fact that the coefficient is invariant to the inclusion of industry effects is reassuring: it rules out the possibility that the result is driven by sector-level trends in sales growth that happen to coincide with the timing of supplier disasters.

The third column restricts to customer-quarters in which the customer was not itself in a disaster-hit county at the time of the supplier event. This addresses the concern that some customers may be geographically close to their affected suppliers, and that the estimated effect might partly reflect the customer's own direct exposure to the same disaster rather than propagation through the supply chain. Removing these observations leaves the coefficient at -0.022, virtually identical to the baseline. The result confirms that the sales decline is not an artifact of co-location.

The fourth column imposes the most stringent restriction, excluding any customer who was ever located in a disaster-hit county at any point during the sample period. This eliminates roughly 4 percent of observations and addresses the concern that firms in chronically disaster-prone areas might differ from those in safer areas in unobservable ways that are correlated with disaster timing. The coefficient is -0.023, significant at the five-percent level. The stability of the coefficient across all four specifications, each of which narrows the comparison group in a different way, is the strongest evidence that the sales decline operates through the supply chain linkage rather than through geographic proximity or the customer's own disaster exposure.

3.4.3 Supplier Disasters and Customer Investment

Table 3.6 reports the investment results. The dependent variable is the one-quarter change in the investment rate.

In the baseline specification where customers are not hit, the coefficient is 0.110 and significant at the 1% level. Customers whose suppliers are hit by a disaster increase their investment rate by approximately 11 percentage points in the quarter following the event. Given the sample mean investment rate of 0.16, this represents a near-doubling of normal capital spending. For the median customer with quarterly CAPEX of $84 million, the implied increase in spending is on the order of $350 million.[3]

The magnitude reflects the nature of the actions the customer must undertake. When a critical input supplier is suddenly unable to deliver, the customer can either accept the disruption passively or take costly action. The actions available include procuring emergency inventory from alternative sources, often at premium prices and with expedited shipping charges. Customers may invest in equipment to bring outsourced processes in-house on a temporary or permanent basis. They may expand capacity at existing plants to reduce dependence on the damaged supplier. They may establish new supplier relationships, which involve site visits, quality audits, and upfront tooling costs. And they may build up safety stock in anticipation of a prolonged disruption. Each of these responses involves substantial capital outlays that register as increases in the investment rate.

The coefficient is stable across all five specifications in Table 3.6. Adding customer industry effects moves the estimate to 0.111. Adding customer ZIP effects and restricting to customers never hit by any disaster yields 0.109. Restricting to customers who never hit after the supplier event produces 0.104, and restricting to customers who never hit before the event produces 0.116. All five estimates are significant at the one-percent level. The result does not depend on any particular sample definition or fixed-effect structure.

The contrast between the investment and sales results is striking and informative. Investment responds within a single quarter; sales do not respond for three to four quarters. This asymmetry is consistent with a model in which the customer responds to the disaster in two stages. In the first stage, the customer takes immediate action to mitigate the disruption, spending aggressively to secure alternative supply and build buffer stock. These expenditures show up in the investment data right away. In the second stage, the disruption works its way through the production process as inventories are depleted and the customer's own output begins to fall. These consequences show up in the sales data with a lag. The two-stage pattern rules out alternative interpretations in which the sales decline and the investment

[3]Calculated as $0.11\times$ median lagged assets of approximately $3.2 billion.

increase are driven by different shocks or different mechanisms. They are two faces of the same underlying event: the supplier disaster.

3.4.4 Heterogeneity by Supplier-Industry Diversification

The baseline results establish that customers experience both an investment spike and a sales decline when their suppliers are hit. A natural follow-up question is whether the magnitude of the response depends on the structure of the customer's supply base.

Table 3.7 tests this by interacting the Supplier Hit indicator with measures of the customer's supplier-industry diversification, constructed by counting the number of distinct industry groups represented among the customer's reported suppliers. Customers whose suppliers all belong to a single industry group are classified as having low diversification. Those whose suppliers span multiple groups are classified as having higher diversification.

The main effect of Supplier Hit for the least-diversified customers is approximately 10 percentage points, significant at the 1% level. The interaction with the first diversification category is −0.106, significant at the five-percent level, and the interaction with the second, more diversified category is −0.119, significant at the one-percent level. The pattern is monotonic: the broader the customer's input portfolio, the smaller the emergency investment response. These results hold across all three sample specifications reported in the table.

The economic logic is intuitive. A customer whose inputs all come from suppliers within a single industry has few substitution options when one of those suppliers is disrupted and must invest heavily to secure alternatives. A customer sourcing from many distinct industries can redirect orders to suppliers in unaffected sectors, substitute alternative inputs, or reallocate production across its existing network. This flexibility dampens the urgency of the response. The finding is consistent with Acemoglu et al. (2012), who show that the aggregate impact of idiosyncratic shocks depends on network concentration, and with Todo et al. (2015), who find that more diversified supply networks are associated with faster recovery from the 2011 Japan earthquake.

The result also has practical implications. It suggests that firms can partially insure themselves against upstream disruption by maintaining relationships with suppliers across a diverse set of input industries. While such diversification may entail higher procurement costs or reduced economies of scale in normal times, the insurance value in a crisis may justify the premium.

3.5 Conclusion

This paper has documented that the damage from natural disasters does not stay where it lands. When a supplier is struck by a flood, hurricane, or other catastrophic event, the consequences propagate downstream to its customers through the supply chain. The propagation takes two forms that operate on different timescales.

The investment response is fast. Within a single quarter, customers increase their capital expenditure rates by roughly 11 percentage points, reflecting the need to secure alternative supply, build emergency inventory, and invest in backup capacity. The sales response is slow. Customer revenue is unaffected in the first quarter, declines modestly in the second, and falls by 2.3 percentage points in annual terms by the fourth quarter. The lag is consistent with the inventory-buffer mechanism: firms sustain output from existing stock for a time, but once those reserves are exhausted, the input shortage translates into lost output and lower sales. The fact that the sales effect does not appear immediately but rather builds over time provides indirect evidence on the depth of the typical firm's inventory buffer and the duration of disaster-induced supply disruptions.

Two additional findings sharpen the picture. First, the effects survive sample restrictions that exclude customers with any direct disaster exposure, providing strong evidence that propagation runs through the supply chain rather than geographic proximity. The coefficient on the supplier disaster indicator is virtually identical whether I include all customers, exclude those co-located with the disaster, or exclude those ever located in a disaster-hit county. This stability is difficult to reconcile with any explanation other than supply chain transmission. Second, customers with more diversified supplier bases exhibit a smaller investment response, consistent with the view that diversification across input markets provides a natural hedge against upstream disruption. The insurance value of diversification is monotonically increasing in the number of distinct supplier industries.

Several questions merit further investigation. First, do customers search for permanent replacement suppliers after a disaster, and do those new relationships persist once the original supplier recovers? If the new relationships prove durable, disasters may permanently rewire portions of the production network. The recovering supplier may lose market share to competitors who filled the gap during the disruption, while those competitors may gain a foothold in a market they would not otherwise have entered. The long-run reallocation of market share is an important but largely unexplored consequence of supply chain shocks.

Second, do repeated experiences with supply disruption alter firms' long-run behavior? A customer that has been burned once by a supplier disaster may choose to maintain higher inventory levels, invest more in supply chain redundancy, dual-source critical inputs, or shift procurement toward suppliers in less disaster-prone regions. If firms do learn from experience, the aggregate economy may become more resilient to supply shocks over time, even as the frequency and severity of natural disasters increase with climate change. Conversely, if adjustment is slow or incomplete, repeated disruptions may impose cumulative costs that compound over the years.

Third, the present analysis is limited to the first tier of the supply chain, that is, the direct customer of the affected supplier. The theoretical literature on network propagation suggests that the effects should cascade further downstream, attenuating with each tier but potentially remaining significant at two or three tiers of separation. Extending the empirical analysis to second- and third-tier effects would require richer network data than the Compustat segment file provides, but it is an important direction for future work. The recent emergence of proprietary datasets that map multi-tier supply chains in greater detail offers a promising avenue for addressing this question.

SUMMARY OF LOSSES

Year	Total Direct Damage	Fatalities	Injuries	Property ($2019b)	Crop ($2019b)
1960 - 2019	$1,143.9	34,933	252,361	$946.3	$197.6
Annual Average	$19.1	583	4,206	$15.8	$3.3

ANNUAL LOSSES

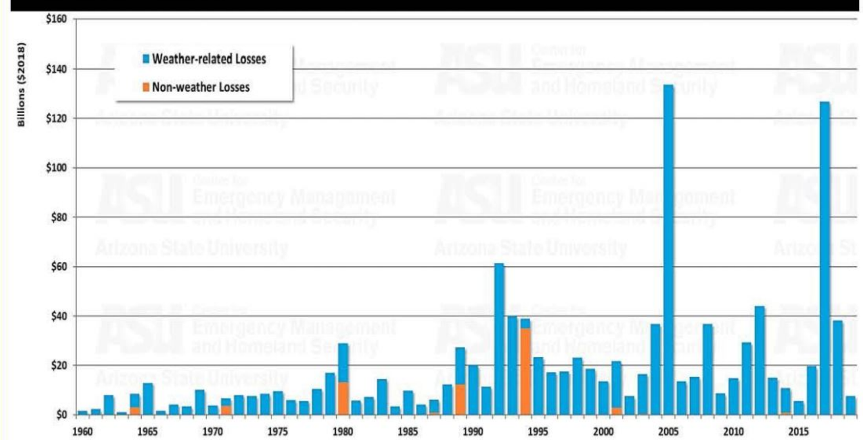

RECORD YEARS

Year	*Costliest*	*Deadliest*	*Year*
2005	$133.6	1,638	1995
2017	$126.8	1,449	2005
1992	$61.4	1,156	2011
2012	$44.0	1,054	1980
1993	$39.8	998	1972
1994	$38.9	917	1999
2018	$38.3	871	1974
2008	$36.8	816	1969
2004	$36.7	748	1971
2011	$29.4	718	1960
1980	$29.0	707	1965
1989	$27.3	671	1998
1995	$23.3	669	1983
1998	$23.2	623	1985
2001	$21.8	618	1989

Figure 3.1: Natural Disaster Losses in the United States, 1960–2019

Notes: The upper panel reports cumulative and annual average losses from 1960 to 2019 in 2019 dollars. The lower-left panel plots annual losses in billions of 2018 dollars, with weather-related losses in blue and non-weather losses in orange. The lower-right panel lists record years ranked by total cost and by fatalities. Source: Arizona State University Center for Emergency Management and Homeland Security.

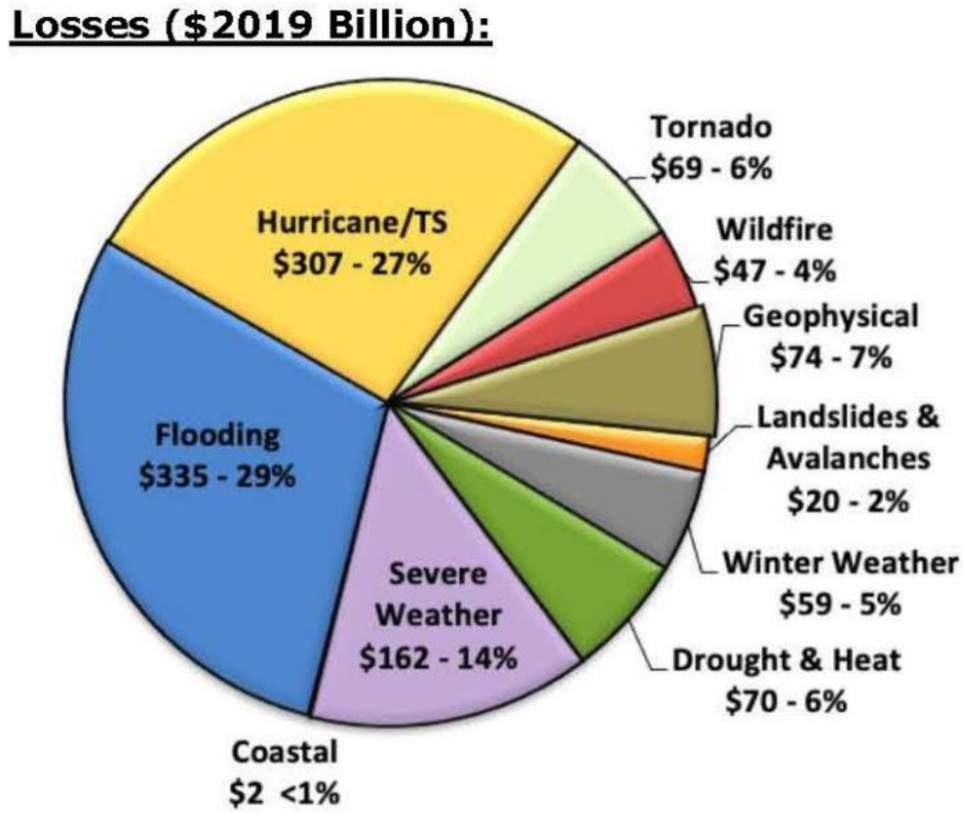

Figure 3.2: Composition of Natural Disaster Losses by Hazard Type, 1960–2019

Notes: Dollar amounts are in 2019 billions. Each segment reports the hazard type, total losses, and percentage share of the cumulative total. Source: Arizona State University Center for Emergency Management and Homeland Security.

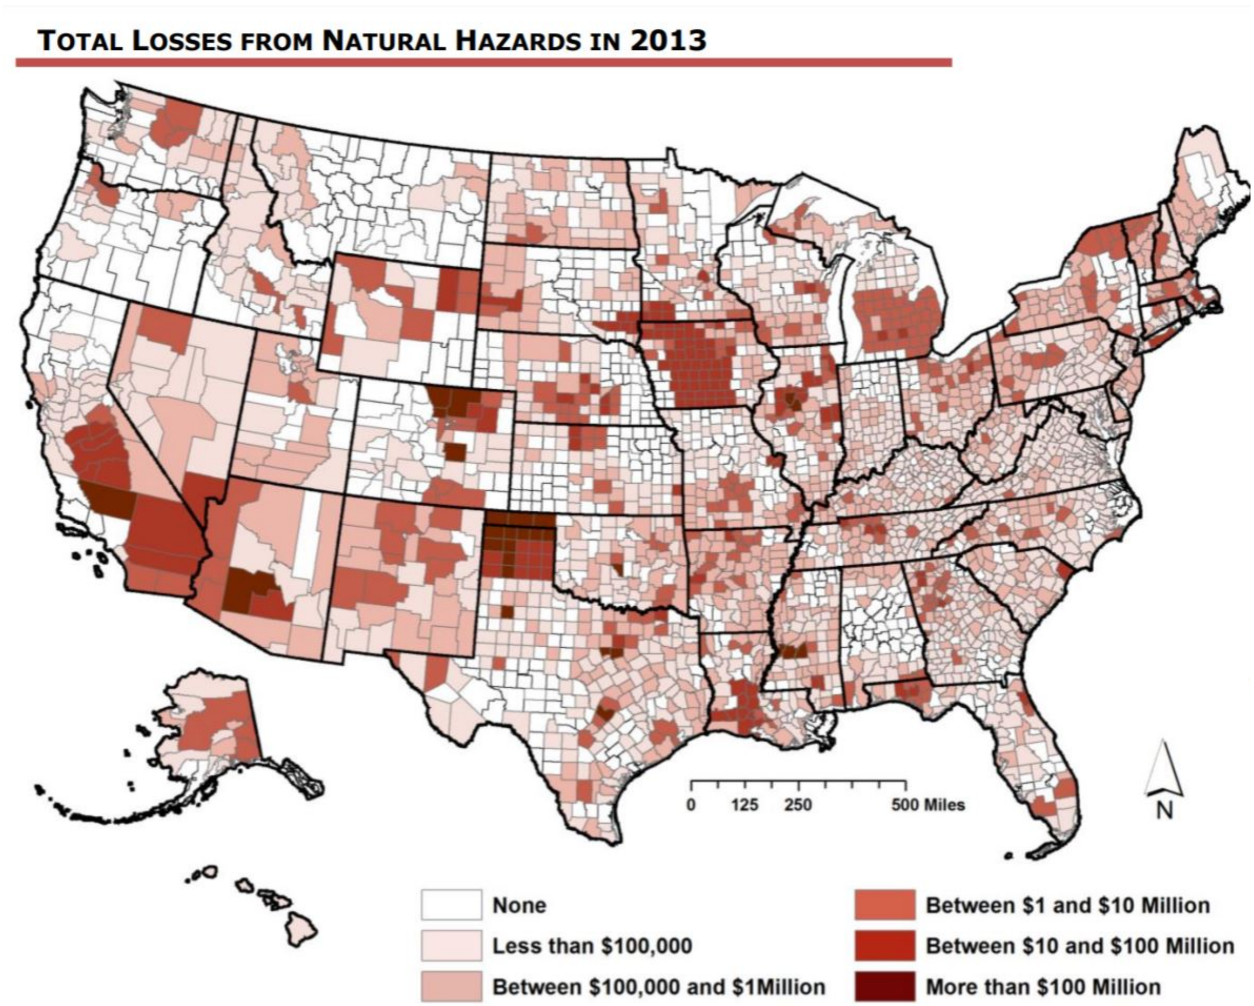

Figure 3.3: Geographic Distribution of Natural Hazard Losses by County, 2013

Notes: Counties are shaded by total property damage: white (none), light pink (less than $100,000), pink ($100,000 to $1 million), light red ($1 to $10 million), red ($10 to $100 million), and dark red (more than $100 million). Source: SHELDUS.

Table 3.1: Top 10 States Ranked by Property Damage per Capita, 1976–2013

State	Prop Damage	Prop Damage p.c.	Crop Damage	Crop Damage p.c.	Population
Hawaii	7,068,726	539.55	179,912	80.16	13,101
Washington	7,180,968	473.53	374,083	8.79	15,164
Louisiana	11,329,277	236.56	474,897	13.93	47,891
Florida	14,050,150	158.40	1,093,147	9.50	88,702
Mississippi	3,904,299	148.24	478,917	27.47	26,337
Alaska	691,219	124.01	317,925	0.07	5,573
California	16,003,987	104.92	4,979,787	150.90	152,537
New Mexico	1,833,467	101.70	107,571	5.66	18,027
North Dakota	1,642,857	97.77	317,825	48.09	16,802
Texas	1,867,622	66.94	830,874	129.59	27,899

Notes: Prop Damage and Crop Damage are cumulative totals in constant 2015 U.S. dollars (thousands). Prop Damage p.c. and Crop Damage p.c. are per capita figures. Population is the average number of people in a disaster-exposed county within each state. States are ranked by property damage per capita. Source: SHELDUS.

Table 3.2: Summary Statistics: Customers

	Mean	Std	p5	p50	p95	Obs
Assets	19,934.59	59,226.07	77.32	3,198.52	87,099	72,934
Sales	2,838.02	5,652.98	16.19	808.48	13,374.23	73,804
Capx	488.99	1,199.35	0.71	84	2,379	58,792
Investment	0.16	0.15	0.02	0.12	0.47	53,422
Number of Suppliers	17.95	26.72	1	6	79	68,319
Disaster-hit	0.01	0.11	0	0	0	48,236

Notes: Assets, Sales, and Capx (capital expenditure) are in millions of U.S. dollars. Investment is capital expenditure divided by one-quarter-lagged total assets. Number of Suppliers is the count of supplier relationships reported in the Compustat segment file. Disaster-hit is an indicator equal to one if the customer's headquarters county experienced a qualifying natural disaster in the given quarter. Source: CRSP-Compustat, Compustat Segment File, SHELDUS.

Table 3.3: Summary Statistics: Suppliers

	Mean	Std	p5	p50	p95	Obs
Assets	1,031.21	3,003.62	4.58	103.87	5,129.40	147,931
Sales	203.80	597.96	0.54	23.21	1,010.11	150,811
Capx	36.56	123.84	0.01	2.15	178.86	130,617
Investment	0.21	0.23	0.01	0.13	0.69	119,467
Number of Customers	2.57	2.34	1	2	6	68,319
Disaster-hit	0.02	0.14	0	0	0	71,625

Notes: Variables are defined as in Table 3.2. Number of Customers is the count of major customer relationships reported in the Compustat segment file. Source: CRSP-Compustat, Compustat Segment File, SHELDUS.

Table 3.4: Customer Sales Growth in 1 to 4 Quarters Following Supplier Disaster

	1-qtr $\Delta Sales$	2-qtr $\Delta Sales$	3-qtr $\Delta Sales$	4-qtr $\Delta Sales$
Supplier hit by disaster	0.002	-0.010	-0.021^{*}	-0.023^{***}
	(0.008)	(0.011)	(0.012)	(0.009)
Observations	38,004	37,050	36,312	35,612
Control	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Customer Industry FE	Yes	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes	Yes
Supplier ZIP FE	Yes	Yes	Yes	Yes

Notes: The dependent variable is cumulative customer sales growth over the indicated horizon. Supplier Hit is an indicator equal to one if at least one of the customer's reported suppliers was located in a disaster-hit county. Controls include log total assets, operating profitability, and the number of supplier relationships. Standard errors in parentheses. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.10$.

Table 3.5: One-Year Sales Growth When Supplier Hit by Disaster: Robustness

	Customer Not Hit at t		Customer Never Hit	
Supplier hit by disaster	-0.023^{***}	-0.023^{***}	-0.022^{***}	-0.023^{**}
	(0.009)	(0.009)	(0.009)	(0.009)
Observations	35,612	35,612	35,188	34,073
Control	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Customer Industry FE	No	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes	Yes
Supplier ZIP FE	Yes	Yes	Yes	Yes
Customer ZIP FE	No	No	Yes	Yes

Notes: The dependent variable is four-quarter cumulative customer sales growth. "Customer Not Hit at t" excludes customer-quarters in which the customer's own county experienced a concurrent disaster. "Customer Never Hit" restricts to customers never located in a disaster-hit county at any point in the sample. All specifications include controls for log total assets, operating profitability, and the number of supplier relationships. Standard errors in parentheses. $^{***}\,p < 0.01$; $^{**}\,p < 0.05$; $^{*}\,p < 0.10$.

Table 3.6: One-Quarter Change in Investment When Supplier Hit by Disaster

	Cust. Not Hit at t		Never Hit	Never Hit After	Never Hit Before
Supplier hit by disaster	0.110***	0.111***	0.109***	0.104***	0.116***
	(0.026)	(0.026)	(0.027)	(0.027)	(0.027)
Observations	31,951	31,951	30,717	31,280	31,219
Control	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes
Customer Industry FE	No	Yes	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes	Yes	Yes
Supplier ZIP FE	Yes	Yes	Yes	Yes	Yes
Customer ZIP FE	No	No	Yes	Yes	Yes

Notes: The dependent variable is the one-quarter change in the customer's investment rate, defined as capital expenditure divided by one-quarter-lagged total assets. "Cust. Not Hit at t" excludes customer-quarters in which the customer's own county experienced a concurrent disaster. "Never Hit" restricts to customers never located in a disaster-hit county. "Never Hit After" and "Never Hit Before" restrict to customers never hit after or before the supplier event, respectively. All specifications include controls for log total assets, operating profitability, and the number of supplier relationships. Standard errors in parentheses. $^{***}\, p < 0.01$; $^{**}\, p < 0.05$; $^{*}\, p < 0.10$.

Table 3.7: One-Quarter Investment Change: Heterogeneity by Supplier-Industry Diversification

	Cust. Not Hit at t	Never Hit	
Supplier hit by disaster	0.100***	0.102***	0.082***
	(0.032)	(0.032)	(0.027)
Supplier hit × No. Industry Group 1	−0.106**	−0.109**	−0.092*
	(0.051)	(0.051)	(0.048)
Supplier hit × No. Industry Group 2	−0.119***	−0.118***	−0.095**
	(0.043)	(0.044)	(0.041)
Observations	31,667	31,510	30,458

Notes: The dependent variable is the one-quarter change in the customer's investment rate. "No. Industry Group 1" and "No. Industry Group 2" are indicators for progressively higher levels of supplier-industry diversification, defined by the number of distinct industry groups among the customer's reported suppliers. Sample definitions follow those in Table 3.6. All specifications include controls for log total assets, operating profitability, and the number of supplier relationships, as well as firm FE, customer industry FE, year-quarter FE, supplier ZIP FE, and customer ZIP FE. Standard errors in parentheses. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.

Bibliography

Baker, S. R., Farrokhnia, R. A., Meyer, S., Pagel, M., & Yannelis, C. (2020). How does household spending respond to an epidemic? Consumption during the 2020 COVID-19 pandemic. *Review of Asset Pricing Studies*, 10(4), 834–862.

Barrero, J. M., Bloom, N., & Davis, S. J. (2021). Why working from home will stick. *NBER Working Paper No. 28731*.

Bick, A., Blandin, A., & Mertens, K. (2021). Work from home before and after the COVID-19 outbreak. *American Economic Journal: Macroeconomics*, 15(4), 1–39.

Brueckner, J. K., Kahn, M. E., & Lin, G. C. (2021). A new spatial hedonic equilibrium in the emerging work-from-home economy? *NBER Working Paper No. 28526*.

Choe, C., Flores-Lagunes, A., & Lee, M. (2021). Stay-at-home orders, mobility, and COVID-19 outcomes. *Working Paper*.

Coibion, O., Gorodnichenko, Y., & Weber, M. (2020). Labor markets during the COVID-19 crisis: A preliminary view. *NBER Working Paper No. 27017*.

Davis, M. A., Ghent, A. C., & Gregory, J. M. (2021). The work-from-home technology boon and its consequences. *NBER Working Paper No. 28461*.

Delventhal, M., Kwon, E., & Parkhomenko, A. (2021). JUE Insight: How do cities change when we work from home? *Journal of Urban Economics*, 127, 103331.

Dube, A., Lester, T. W., & Reich, M. (2010). Minimum wage effects across state borders. *Review of Economics and Statistics*, 92(4), 945–964.

Emanuel, N. & Harrington, E. (2020). Working remotely? Selection, treatment, and the market provision of remote work. *Working Paper*.

Gupta, A., Mittal, V., Peeters, J., & Van Nieuwerburgh, S. (2022). Flattening the curve: Pandemic-induced revaluation of urban real estate. *Journal of Financial Economics*, 146(2), 594–636.

Holmes, T. J. (1998). The effect of state policies on the location of manufacturing: Evidence from state borders. *Journal of Political Economy*, 106(4), 667–705.

Kozlowski, J., Veldkamp, L., & Venkateswaran, V. (2020). Scarring body and mind: The long-term belief-scarring effects of COVID-19. *NBER Working Paper No. 27439*.

Ling, D. C., Wang, C., & Zhou, T. (2020). A first look at the impact of COVID-19 on commercial real estate prices: Asset-level evidence. *Review of Asset Pricing Studies*, 10(4), 669–704.

Liu, S. & Su, Y. (2021). The impact of the COVID-19 pandemic on the demand for density: Evidence from the U.S. housing market. *Economics Letters*, 207, 110010.

Mondragon, J. & Wieland, J. (2022). Housing demand and remote work. *Federal Reserve Bank of San Francisco Working Paper 2022-11*.

Ramani, A. & Bloom, N. (2021). The donut effect of COVID-19 on cities. *NBER Working Paper No. 28876*.

Goolsbee, A. & Syverson, C. (2021). Fear, lockdown, and diversion: Comparing drivers of pandemic economic decline 2020. *Journal of Public Economics*, 193, 104311.

Chetty, R., Friedman, J. N., Hendren, N., Stepner, M., & The Opportunity Insights Team (2020). The economic impacts of COVID-19: Evidence from a new public database built using private sector data. *NBER Working Paper No. 27431.*

Bloom, N., Davis, S. J., & Zhestkova, Y. (2021). COVID-19 shifted patent applications toward technologies that support working from home. *AEA Papers and Proceedings*, 111, 263–266.

Althoff, L., Eckert, F., Ganapati, S., & Walsh, C. (2022). The geography of remote work. *Regional Science and Urban Economics*, 93, 103770.

Whitaker, S. D. (2021). Did the COVID-19 pandemic cause an urban exodus? *Federal Reserve Bank of Cleveland, District Data Brief*, February.

Haslag, P. H. & Weagley, D. (2021). From L.A. to Boise: How migration to small cities is reshaping the housing market. *Working Paper, Vanderbilt University.*

Howard, G. (2021). The migration accelerator: Labor mobility, housing, and demand. *American Economic Journal: Macroeconomics*, 12(4), 147–179.

Brueckner, J. K. & Sayantani, S. (2023). Intercity impacts of work-from-home with both remote and non-remote workers. *Journal of Housing Economics*, 59, 101905.

Bordalo, P., Gennaioli, N., & Shleifer, A. (2020). Memory, attention, and choice. *Quarterly Journal of Economics*, 135(3), 1399–1442.

Gennaioli, N. & Shleifer, A. (2010). What comes to mind. *Quarterly Journal of Economics*, 125(4), 1399–1433.

Autor, D. H., Palmer, C. J., & Pathak, P. A. (2014). Housing Market Spillovers: Evidence from the End of Rent Control in Cambridge, Massachusetts. *American Economic Review*, 104(10), 3317–3347.

Benjamin, J. D., Coulson, N. E., & Yang, S. X. (1993). Real estate transfer taxes and property values: The Philadelphia story. *The Journal of Real Estate Finance and Economics*, 7, 151–157.

Best, M. C., & Kleven, H. J. (2018). Housing market responses to transaction taxes: Evidence from notches and stimulus in the UK. *The Review of Economic Studies*, 85(1), 157–193.

Chapple, K., Loukaitou-Sideris, A., Miller, A., & Zeger, C. (2023). The role of local housing policies in preventing displacement: A literature review. *Journal of Planning Literature*, 38(2), 200–214.

Chen, Y., Ma, L., & Orazem, P. F. (2025). Capitalization of state tax rates in housing values at state borders, 2000–2017. *Real Estate Economics*.

Chen, Z., Liu, Z., Suárez Serrato, J. C., & Xu, D. Y. (2021). Notching R&D Investment with Corporate Income Tax Cuts in China. *American Economic Review*, 111(7), 2065–2100.

Dachis, B., Duranton, G., & Turner, M. A. (2012). The effects of land transfer taxes on real estate markets: Evidence from a natural experiment in Toronto. *Journal of Economic Geography*, 12(2), 327–354.

de Silva, T. (2023). Insurance versus Moral Hazard in Income-Contingent Student Loan Repayment. *Forthcoming*.

Diamond, R., & McQuade, T. (2019). Who Wants Affordable Housing in Their Backyard? An Equilibrium Analysis of Affordable Housing Impact on Property Values. *American Economic Review*, 109(11), 3621–3652.

Favilukis, J., Mabille, P., & Van Nieuwerburgh, S. (2023). Affordable housing and city welfare. *The Review of Economic Studies*, 90(1), 293–330.

Gilgoff, J. (2023). The California SHIMBY Movement: Social Housing in my Backyard. *California Western Law Review*, 60, 1.

Glaeser, E. L., & Gyourko, J. (2002). The impact of zoning on housing affordability. National Bureau of Economic Research.

Glaeser, E. L., & Luttmer, E. F. P. (2003). The Misallocation of Housing Under Rent Control. *American Economic Review*, 93(4), 1027–1046.

Green, D., Jambulapati, V., Liebersohn, J., & Velayudhan, T. (2025). Fiscal Externalities of Transaction Taxes: Evidence from the Los Angeles Mansion Tax. *Available at SSRN 5273034*.

Han, L., Ngai, L. R., & Sheedy, K. D. (2022). *To Own Or to Rent?: The Effects of Transaction Taxes on Housing Markets*. Centre for Economic Policy Research.

Horton, E., LaPoint, C., Lutz, B., Seegert, N., & Walczak, J. (2024). Property Tax Policy and Housing Affordability. *National Tax Journal*, 77(4), 861–901.

Kopczuk, W., & Munroe, D. (2015). Mansion Tax: The Effect of Transfer Taxes on the Residential Real Estate Market. *NBER Working Paper* 21684.

Londoño-Vélez, J., & Avila-Mahecha, J. (2024). Behavioral Responses to Wealth Taxation: Evidence from Colombia. *NBER Working Paper* w32134.

Manville, M., & Smith, M. (2025). The Unintended Consequences of Measure ULA.

Mense, C., Michelsen, C., & Berlemann, M. (2019). The German Housing Market: Policy Proposals to Address an Overvalued Market. *Jahrbücher für Nationalökonomie und Statistik / Journal of Economics and Statistics*, 239(2), 295–334.

Phillips, S., & Ofek, M. (2022). How Will the Measure ULA Transfer Tax Initiative Impact Housing Production in Los Angeles?

Santos Silva, J. M. C., & Tenreyro, S. (2006). The log of gravity. *The Review of Economics and Statistics*, 641–658.

Simpson, A., Filipe, L., Benedetto, V., & Hill, J. (2024). The impacts of housing conditions on physical and mental health: A critical mini-review. *Frontiers in Environmental Health*, 3, 1352580.

Van Ommeren, J., & Van Leuvensteijn, M. (2005). New evidence of the effect of transaction costs on residential mobility. *Journal of Regional Science*, 45(4), 681–702.

Ward, J., Phillips, S., & Manville, M. (2025). L.A.'s 'Mansion Tax' Needs a Remodel. Here's How to Fix It. *RAND Corporation Commentary.*

Whamond, C. (2024). State Affordable Housing Tax Credits and Renter Outcomes: Evidence from Illinois.

Acemoglu, D., Carvalho, V. M., Ozdaglar, A., & Tahbaz-Salehi, A. (2012). The network origins of aggregate fluctuations. *Econometrica*, 80(5), 1977–2016.

Barrot, J.-N. & Sauvagnat, J. (2016). Input specificity and the propagation of idiosyncratic shocks in production networks. *Quarterly Journal of Economics*, 131(3), 1543–1592.

Basker, E. & Miranda, J. (2018). Taken by storm: Business financing and survival in the aftermath of Hurricane Katrina. *Journal of Economic Geography*, 18(6), 1285–1313.

Bourdeau-Brien, M. (2017). Natural disasters and firm outcomes. *Working Paper.*

Carvalho, V. M., Nirei, M., Saito, Y. U., & Tahbaz-Salehi, A. (2021). Supply chain disruptions: Evidence from the Great East Japan Earthquake. *Quarterly Journal of Economics*, 136(2), 1255–1321.

Dessaint, O. & Matray, A. (2017). Do managers overreact to salient risks? Evidence from hurricane strikes. *Journal of Financial Economics*, 126(1), 97–121.

Fujiwara, Y. & Aoyama, H. (2010). Large-scale structure of a nation-wide production network. *The European Physical Journal B*, 77(4), 565–580.

Giroud, X. & Mueller, H. M. (2019). Firms' internal networks and local economic shocks. *American Economic Review*, 109(10), 3617–3649.

Henriet, F., Hallegatte, S., & Tabourier, L. (2012). Firm-network characteristics and economic robustness to natural disasters. *Journal of Economic Dynamics and Control*, 36(1), 150–167.

Seetharam, I. (2017). The indirect effects of hurricanes: Evidence from firm internal networks. *Working Paper*.

Todo, Y., Nakajima, K., & Matous, P. (2015). How do supply chain networks affect the resilience of firms to natural disasters? Evidence from the Great East Japan Earthquake. *Journal of Regional Science*, 55(2), 209–229.