FHFA STAFF WORKING PAPER SERIES

Measuring Price Effects from Disasters Using Public Data: A Case Study of Hurricane Ian

Justin C. Contat
William M. Doerner
Robert N. Renner
Malcolm J. Rogers

May 2024

Working Paper 24-04

FEDERAL HOUSING FINANCE AGENCY
Division of Research and Statistics
400 7th Street SW
Washington, DC 20219, USA

Please address correspondence to Justin Contat (justin.contat@fhfa.gov). Helpful feedback was provided by Becka Brolinson, Jonah Coste, Kevin Duncan, Shan Ge, Caroline Hopkins, Alex Marr, Jesica Shui, Albert Yoon, Velma Zahirovic-Herbert, as well as participants from the American Real Estate Society annual conference, Federal Housing Finance Agency Econ Summit, and University of Central Florida / Journal of Real Estate Research Critical Issues in Real Estate Symposium.

The latest version of this paper can be found at https://www.fhfa.gov/papers/wp2404.aspx.

Measuring Price Effects from Disasters Using Public Data:
A Case Study of Hurricane Ian

Justin C. Contat, William M. Doerner, Robert N. Renner, and Malcolm J. Rogers

FHFA Staff Working Paper 24-04
May 2024

Abstract

Natural disasters can disrupt housing markets, causing destruction to communities and distress to economic activity. To estimate the effects of disasters on home prices, publicly-available data on property damages are often used to classify "treated" properties. However, by design these data lack precise geospatial information, leading to measurement error in the treatment variable as aggregate measures must be used. We leverage leading difference-in-differences and synthetic control methodologies across various treatments and levels of geography to measure price effects with such data following Hurricane Ian's unexpected landfall in southwest Florida during September 2022, coinciding with the state's initial recovery from the COVID-19 pandemic. Empirical results suggest positive, time-varying price effects, though we place caveats on these results as there may be many mechanisms underway; our results should be interpreted as descriptive correlations and not causal effects for various reasons. Our main contribution is methodological, highlighting the importance of robustness checks, functional form, statistical techniques, and testing across different samples. Additionally, quicker access to high quality public data could enhance quantitatively-informed conversations on natural disaster effects.

Keywords: causal inference · hurricane · model selection · natural experiment · real estate valuation

JEL Classification: C52 · Q54 · R21 · R31

Justin C. Contat
Federal Housing Finance Agency
Division of Research and Statistics
Office of Research and Analysis
400 7th Street SW
Washington, DC 20219, USA
justin.contat@fhfa.gov

William M. Doerner
Federal Housing Finance Agency
Division of Research and Statistics
Office of Research and Analysis
400 7th Street SW
Washington, DC 20219, USA
william.doerner@fhfa.gov

Robert N. Renner
Federal Housing Finance Agency
Division of Research and Statistics
Office of Chief Data Officer
400 7th Street SW
Washington, DC 20219, USA
robert.renner@fhfa.gov

Malcolm J. Rogers
Federal Housing Finance Agency
Division of Research and Statistics
Office of Chief Data Officer
400 7th Street SW
Washington, DC 20219, USA
malcolm.rogers@fhfa.gov

1 Introduction

Natural disasters disrupt housing markets, causing physical damage and reducing property values. While measures of damage are useful for certain purposes, understanding changes to a home's fundamental value (price) is more relevant for economic considerations. For example, potential home buyers may likely be concerned about a previously damaged home's future value even if damages have been fully repaired. Nevertheless, measures of damages are still needed as failure to account for destruction when estimating price effects may conflate changes in market conditions with changes in a home's quality. Publicly-available, claims-level damage data exist but (by design) do not have precise geospatial information. This introduces measurement error along with other practical and econometric challenges for those using this data. This paper assesses whether current publicly-available damages data can shed light on real estate market price effects after a natural disaster. Without proprietary damages data, the publicly-available damages data are often the only source for many policymakers and researchers. This paper also highlights how leading statistical techniques cannot always overcome common data challenges.

A major storm, Hurricane Ian, struck southwest Florida in September 2022. The initial trajectory anticipated landfall in Tampa Bay before suddenly veering eastward. The altered course brought landfall to Lee County, where Ian ultimately became the costliest and most destructive storm in the state's history. This paper uses several leading approaches in the literature, namely, difference-in-differences (DiD) and synthetic control methods (SCM), to estimate causal effects of Hurricane Ian on home prices.[1] To do this, residential real estate listings are collected for transactions from counties across Florida, the coastal Gulf-Coast, and the coastal Mid-Atlantic states. The information is combined with the Federal Emergency Management Agency's (FEMA) publicly-available Individual Household Program (IHP) data containing details on property level damages resulting from Hurricane Ian. Using the IHP data, we construct several aggregate level (county or ZIP Code™) treatment definitions with both counts and dollar value damages of claims. Importantly, each aggregate treatment definition is associated with some degree of measurement error because homes affected by Ian may be misclassified as not being damaged by Ian, and vice-versa.

[1] In our usage, DiD refers to both the canonical difference-in-difference framework with a single pre-treatment period and single post-treatment period, as well as event study frameworks with multiple pre- and post-treatment periods. We label these frameworks as static and dynamic treatment effects, respectively.

Despite initial summary statistics indicating aggregate price declines, traditional DiD and SCM estimates reveal *positive* price effects. While a supply channel effect—Ian reduced housing supply, driving up prices—is possible, we are hesitant to interpret the results in such a way for several reasons. First, evidence suggests the parallel trends assumption (essential for DiD) is likely unsatisfied, as hurricane-stricken areas exhibit different price trends than unaffected ones. Second, scarce housing data in some areas means balancing panels for SCM estimation reduces the sample size, yielding a weak pre-treatment fit, thus posing a problem with SCM design. Third, estimated effects are sensitive to treatment definitions, making the choice of which aggregate treatment to use non-trivial. Moreover, weights assigned to control units vary considerably when treatment definitions change, underscoring the need for additional statistical investigation. Fourth, estimated treatment effects are also sensitive to the time window used and the property type of the home; manufactured homes appear to be more affected than single-family homes, although the size of the effect varies depending upon the earliest time used in the estimation. Fifth, there are likely many mechanisms underway that influence home prices after disasters and we do not distinguish empirically the different channels, but we provide a discussion of possible channels at the end of the paper. We view the positive correlation between Ian and home prices as descriptive and not a causal effect for the various reasons previously mentioned. Finally, careful consideration of temporal frequency is crucial for analyses; using too coarse a time frequency (e.g., annual) may mask important variation within a given period. In general, despite using leading econometric approaches, the results highlight limitations of publicly available data, especially in finding suitable controls when treatment (and control) are mis-measured. Researchers investigating causal effects of disasters are likely to encounter similar problems.

This is not the first study about the effects of natural disasters on home prices. Unexpected events, like hurricanes, can shock markets with lasting consequences. As detailed in the next section, prior work most commonly considers disasters' differential impact on flood zones and non-flood zones.[2] Some papers consider the direct effect of damages on residential prices, though often with different treatment definitions. These papers typically have granular proprietary damage information, thus avoiding the measurement error issues raised here. Such granularity would be extremely valuable for policymakers reacting to a storm's potential damages and price effects. A small literature is emerging about misclassification (treatment mismeasurement) in a broader setting, but it is too early to apply to this setting.

[2]See Contat et al. (2024) for a recent survey of this literature, including methods used and findings.

This paper's contribution is twofold. First, using public data, estimates are derived for the effects of a hurricane on residential real estate prices, providing valuable information for those still grappling with a storm's aftermath. While DiD results are mixed, SCM findings suggest a time-varying premium with the largest treatment magnitude occurring about three months after the storm. In particular for Hurricane Ian, the results suggest around a 5% to 11% price *increase*, which we view with some skepticism. Second, we provide practical advice on methodological choices while highlighting the limitations of using public data. Our primary recommendations include employing matching before estimation, imputation of missing geographical areas for balancing panels, careful scrutiny of the time window being considered, sensitivity analysis with respect to treatment aggregation choice, consideration of proper form of dependent variable, and supplementing damage measures to enhance accuracy. The results underscore the importance of robustness checks with DiD and SCM methods.

The plan of the paper is as follows. After a brief review of the applied disaster literature, section 3 describes the data and presents preliminary descriptions of price effects. Section 4 uses traditional causal techniques to test if a natural disaster might disrupt real estate prices. The first empirical specification, DiD, and the resulting estimates, pay attention to pre-trends tests for identifying suitable controls groups. Dynamic treatment effects further document heterogeneity over time. Section 5 introduces the second specification, synthetic controls, as a potential solution when a suitable control cannot be found or if the aim is an aggregated treatment effect. Both specifications are performed at the county and ZIP Code levels to highlight spatial heterogeneity. Finally, section 6 provides a concluding discussion.

2 Literature Review

This paper contributes to two distinct branches of the literature: studies of natural disasters and applied experimental techniques. As commonly noted by the popular press, natural disasters can generate substantial costs. However, those figures often focus on damages or disaster relief that are totaled for individual properties. While we do estimate price effects at the individual property level, we also consider whether such events can fundamentally disrupt an entire market's valuation or its potential price trajectory. If any event could have an impact, Hurricane Ian is a good candidate. However, the exercise becomes difficult when a storm destroys the existing stock or leads to only certain kinds of properties being sold. Luckily, the federal government has been making considerable strides to release more detailed information on disasters claims and there already exists very rich residential real estate data. The subsections below highlight academic works that utilize such data sources and evolving causal methodologies that this paper will bridge together.

2.1 Studies of Natural Disasters

A large applied climate literature estimates the disruptive effects of disasters on residential prices.[3] Difference-in-differences (DiD) and synthetic control methods (SCM) are the dominant methods, with DiD being much more prevalent, although there are papers that use other estimators to answer particular questions.[4] For DiD papers, the crucial statistical requirement is satisfying the parallel trends assumption, without which there is not a suitable control group.[5] This paper builds on the literature by explicitly considering how to find a suitable control in the presence of a misclassified treatment. While there are econometric tests (e.g., pre-trends tests for DiD, pre-treatment fit for SCM) that can determine whether a proposed control group is suitable, this paper builds off Miller (2023) and offers practical advice about choosing suitable controls using a variety of statistical tests. For SCM papers, although these methods were developed to deal with treatment at aggregate levels when there is no obvious control group (Abadie and Gardeazabal, 2003; Abadie, Diamond, and Hainmueller, 2010), the highly localized nature of real estate prices may introduce bias when assigning treatment. For the few papers that use SCM, a wide variety of treatment aggregations are assigned at the Census tract (Keys and Mulder, 2020), city block (Ho et al., 2023), and city levels (Kim and Lee, 2023). We build off of this literature and consider treatment at two different geographic levels (the county and ZIP Code) to show whether an appropriate aggregation level can be established. Moreover, we believe this is the first paper to consider how estimated price effects are affected by different aggregating procedures (i.e., treatment definitions) with a certain geography.

Restricting attention to hurricanes and flooding events, the definition of a "treated home" or a "treated area" is typically expressed in terms of flood zone status (Fang, Li, and Yavas, 2023; Muller and Hopkins, 2019; Hino and Burke, 2021; Harrison, T. Smersh, and Schwartz, 2001), actual damages sustained (Zivin, Liao, and Panassie, 2023; Strobl, 2011; Fisher and

[3] For an alternative dependent variable, see Turnbull, Zahirovic-Herbert, and Mothorpe (2013) and Salter and King (2009), each of whom consider the effect of flood risk on price and liquidity.

[4] For example, Zhang (2016) uses a spatial quantile regression to study the effects of flood hazards on residential values across the price distribution and finds lower-priced properties are impacted more.

[5] Recent works are becoming more likely to explicitly mention or even test for parallel (common) trends; see, for example, the papers surveyed in Contat et al. (2024). Roth (2022) points out that even conditional on passing pre-trends tests one may still have biased results.

Rutledge, 2021; Gould Ellen and Meltzer, 2024), or both (Atreya and Ferreira, 2015; Hennighausen and Suter, 2020; Yi and Choi, 2020; Gibson and Mullins, 2020). Building on this literature, the present study focuses solely on damages and asks: Do affected areas experience lasting changes in fundamental real estate values? Unsurprisingly, the literature generally finds evidence that damages negatively impact property prices, but there is a lack of consensus about the right magnitude due to differences in definitions of damages, empirical approaches, and other factors. There are exceptions where some papers do find positive price effects, which is justified as a correlation with amenities (Muller and Hopkins, 2019) or even rebuilding to stronger codes (Dumm, Sirmans, and Smersh, 2012).

Further restricting attention to hurricane and flood papers that use damages data, most studies have fairly granular measures of data. Exceptions to this include Fisher and Rutledge (2021), who focus on commercial real estate and define a property to be treated if it is located in a core-based statistical area (CBSA) that was affected by a hurricane. They find evidence of significant and heterogeneous price effects due to damage from hurricanes. Similarly, Gould Ellen and Meltzer (2024) aggregate damages to the Census block and tract levels using highly granular storm surge data, indicating heterogeneous price effects with respect to income. The present study builds on these works by explicitly considering the consequences of deliberately (and necessarily) introducing measurement error.[6] In terms of a similar sample area, Hallstrom and Smith (2005) study the effects of the "near-miss" Hurricane Andrew (1992) on the Lee County area, which was not damaged by the storm. Using a DiD approach, they find that residences in flood zones experienced price discounts of about 19%. Importantly, their treatment definitions correspond to flood zone status, not direct damage, making it difficult to compare price effects with our study.

One aspect of the literature not often explicitly considered is the time dimension, which is often dictated by a data constraint. However, when a researcher has the luxury of making a free choice, it would be useful to know when and what matters to the dependent variable (price). The literature contains a range of time frequencies. Gibson and Mullins (2020) use a weekly frequency. Other papers rely on months (Bin and Landry, 2013; Zivin, Liao, and Panassie, 2023), quarters (Fang, Li, and Yavas, 2023; Gould Ellen and Meltzer, 2024), and years (Yi and Choi, 2020; Atreya and Ferreira, 2015). The present study is agnostic

[6] Interestingly Zivin, Liao, and Panassie (2023) note another measurement error when they assume treatment (exposure) is symmetric about the eye of the storm. This is not a concern here because the starting point is a property level measure of damages and not scientific measures like wind speed or the storm path.

about the most appropriate choice and is more interested in discerning whether the choice matters for finding a suitable control. For example, if a DiD specification has 4 years of pre-treatment data then (ignoring normalization of coefficients) four coefficients are needed to jointly be equal to zero for an annual time frequency whereas 48 coefficients are needed for a monthly frequency. Additional complications arise if the larger number of coefficients also entail larger standard errors because of the increasing offset between precision and noise when increasing granularity. The lack of a clear ex-ante prediction about directionality is discussed later in more detail. In our dynamic specifications we see some evidence of price effects that change over time, suggesting that a coarser frequency on a quarterly or annual basis may miss important time heterogeneity.

2.2 Applied Experimental Techniques

This paper relates to an old and extensive literature on measurement (see Schennach, 2016, for a recent survey). A relatively new strand of work analyzes misclassification of treatment in the DiD setting (Denteh and Kédagni, 2022; Negi and Negi, 2022). Misclassified treatment leads to biased and inconsistent estimation of the average treatment effect on the treated (ATT) as the estimated parameter can be written as a weighted average of the correctly measured and mismeasured groups.[7] Our paper also relates to a growing DiD literature that points out threats to the parallel trends assumption (Roth et al., 2023).[8] Several papers have pointed out problems with common "pre-trends" tests that are used to evaluate the validity of the parallel trends assumption (Roth, 2022; Rambachan and Roth, 2023). In particular, pre-trends tests have low power (failing to reject the presence of differing trends between treated and control groups) and, conditional on passing these tests, the resulting estimates of treatment effects are likely to be biased (Roth, 2022). Can the parallel trends assumption be satisfied under misclassification in an applied setting? Before moving on to SCM, it may be useful to point out this current paper is conceptually related to the "fuzzy" difference-

[7] Alternatively, the problem of measurement is eliminated, or at least reduced, if one is interested in a more market-level (aggregate) effect rather than the effect of the disaster on individual home prices.

[8] This paper only considers a single event. A reviewer pointed out that multiple events or varying degrees of treatment could be happening in the pre-treatment sample. Unfortunatley, if using staggered or stacked DiD methods, any mismeasurement could be compounded over multiple disasters. This paper already points out applied limitations with estimating the effects from one disaster. Realistically, the problem exists because there is enough time for multiple events to affect treated and controlled locations. A partial solution is to limit pre- and post-treatment windows. Based on feedback, we run tests with shorter pre-treament lengths back to 2019 and 2021. Results remain similar. This approach isolates the disaster by itself (i.e. no other major hurricanes happen within that window), but it complicates the selection of control locations. We defer to future work for applying empirical techniques to the case of multiple disasters. See Marcus and Sant'Anna (2021) for a comparison of recently proposed estimators with staggered treatment in an environmental setting.

in-difference treatment definition developed by De Chaisemartin and d'Haultfoeuille (2018), where units in treatment group may be allowed to be untreated and units in the control may be allowed to be treated. The authors show that with certain assumptions and fuzzy treatment, alternative estimands can identify a local average treatment effect (LATE) of the units that actually become treated.

SCM estimation has a rich academic basis.[9] Two notable sub-threads relate to covariate selection and goodness of fit. For covariate selection, Kaul et al. (2015) show that using all pre-treatment outcomes as controls will maximize pre-treatment fit while rendering other factors unusable, which invokes a fit versus bias trade-off. We follow the advice in Ferman, Pinto, and Possebom (2020) who use all pre-treatment lags as controls to maximize fit and avoid criticisms about selected controls. For goodness of fit, papers propose solutions for dealing with the problem of a weak pre-treatment fit, which can cause bias in the estimates (Ben-Michael, Feller, and Rothstein, 2021; Ferman and Pinto, 2021). In particular, Arkhangelsky et al. (2021) propose a combination of DiD with SCM ("synthetic differences"), where a synthetic control establishes a baseline comparison in a DiD estimation rather than in direct calculation of the desired treatment effect from the relevant formula. Finally, for many treated units, Abadie and L'Hour (2021) propose a method for selecting a suitable synthetic control among disaggregated data if the standard approach does not yield a unique solution. This technique is tested and presented later.

3 Data and Background

This section begins with Hurricane Ian background information, followed by a summary of utilized datasets, including real estate transactions and disaster claims. It defines various aggregate treatments and presents preliminary evidence of Ian's impact on market valuation metrics. The findings suggest an Ian-related disruption, with a more substantial shock observed in the two years preceding it due to COVID-19.[10]

[9]For a recent survey of synthetic control methods, see Abadie (2021). Work relating DiD and SCM is in relatively early stages. Several important contributions have already been made by Doudchenko and Imbens (2016) and Arkhangelsky et al. (2021).

[10]A differential empirical framework could attempt to isolate the COVID-19 disruption from the Hurricane Ian disruption using a staggered treatment design. Such a design is beyond the current scope of this paper because it suggests testing for a mismeasured treatment. However, in some robustness checks, we partially address the issue by varying lengths of treatment windows in the estimation.

Figure 1: Mapping Out Coastal Areas, Hurricane Ian's Path, and Real Estate Listings and Damages

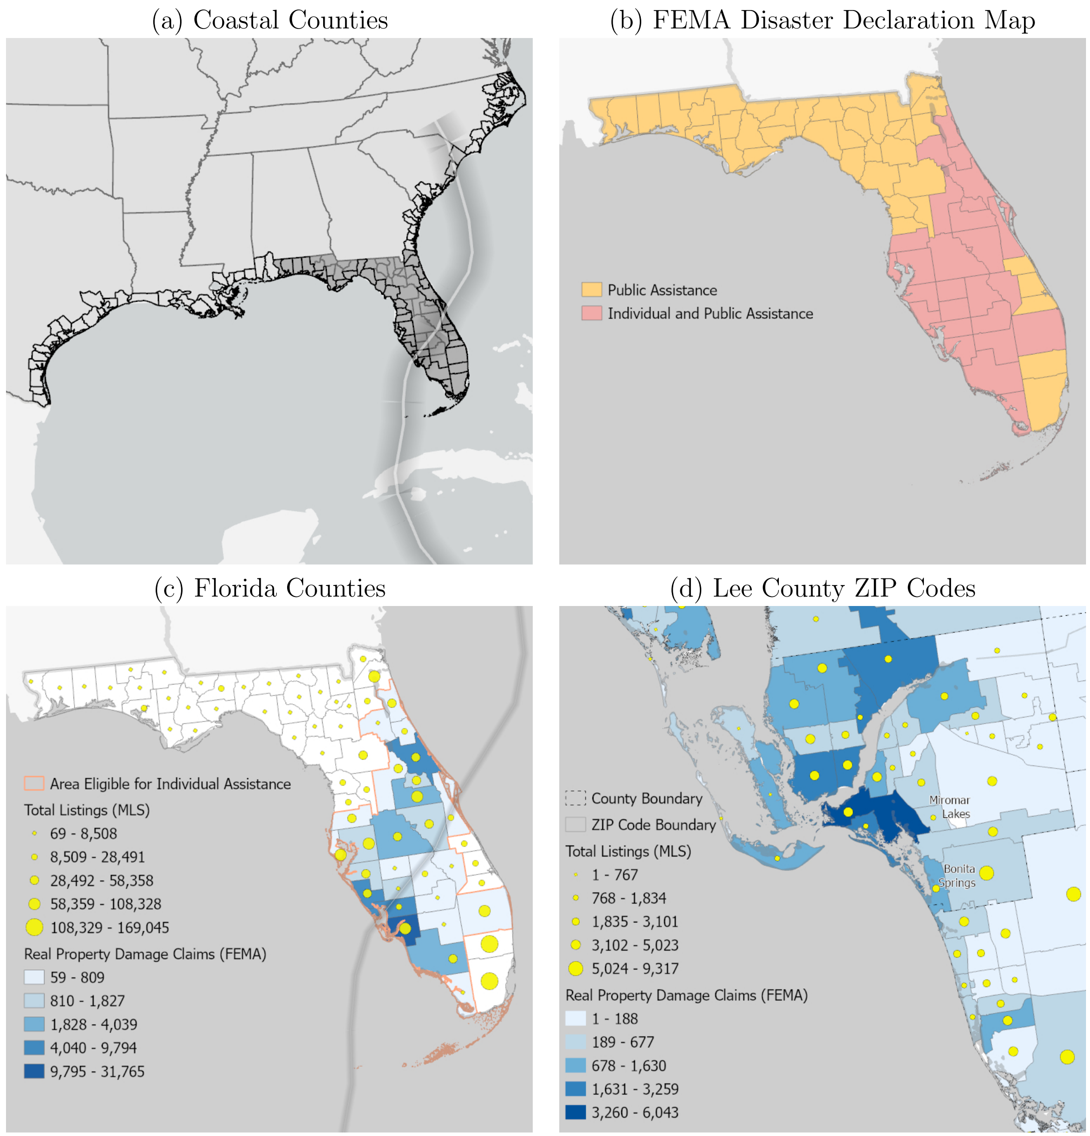

Notes: Panel (a) was created by the authors using ArcGIS while Panel (b) is available online at https://gis.fema.gov/maps/dec_4673.pdf. Listings data come from MLS and damages from FEMA.

Table 1: Recent Declared Disasters Affecting Lee County, Florida

Disaster	Date	Declaration	Disaster Severity	IH	IA	PA	HM
Hurricane Irma	9/4/2017	DR-4337	Cat 3, $60.13 billion	1	0	1	1
COVID-19 Pandemic	1/20/2020	DR-4486	-	1	0	1	1
Hurricane Ian	9/23/2022	DR-4673	Cat 4, $112 billion	1	0	1	1
Hurricane Nicole	11/7/2022	DR-4680	Cat 1, $1 billion	0	1	1	1

Notes: This paper focuses on the period from March 2015 to June 2023 due to coverage of privately licensed real estate information that is combined with publicly available data on disasters and damages. Information shown above comes from FEMA (https://www.fema.gov/openfema-data-page/fema-web-disaster-declarations-v1) and the National Hurricane Center Tropical Cyclone Reports. Cat is the Category of the storm as it made landfall in Florida and the dollar amount is the estimated total damages. All dollars are 2022 dollars; Irma's dollars adjusted using BLS CPI Inflator with Sept 2017 and Sept 2022 dates, $50 billion as 2017 estimate of damages. The two letter acronyms in the title indicate whether the associated programs were declared for the particular disaster: IH = Individuals and Households, IA = Individual Assistance, PA = Public Assistance, HM = Hazard Mitigation. See OPENFEMA for more details on these programs.

3.1 Hurricane Ian and Southwest Florida

Hurricane Ian made landfall in southwest Florida on September 23, 2022, near Fort Myers (Lee County), then traversed northeast through the state until reaching Georgetown, South Carolina on September 30, 2022. This Category 4 storm was the most expensive in Florida's history and the third costliest in national history, resulting in $109.5 billion in damages and claiming at least 156 lives in the state.[11] In the top row of Figure 1, Ian's trajectory is depicted in panel (a) while FEMA's official map of disaster declarations in Florida is showcased in panel (b). Ian's impact was widespread, affecting the entire state and resulting in a significant number and extensive area of damaged (treated) units. All counties qualified for public assistance and the figure shows (in pink) the areas hit the hardest which qualified for the Individual Assistance program.[12] The governor of Florida declared a state of emergency for the entire state, making every county eligible for some type of public assistance.

Hurricane Ian is not the first of its kind to impact the area; unfortunately, southwest Florida has witnessed several hurricanes in recent years. Table 1 lists major disasters affecting Lee County since 2015 and details the relative severity of each disaster, categorized by storm strength and the 2022-dollar value of damages. The table also outlines the types of aid available for each disaster, with more information on these programs accessible on FEMA's website. It is noteworthy that a variety of disaster aid options are typically accessible post-

[11] For additional details on Hurricane Ian, refer to: https://www.nhc.noaa.gov/data/tcr/AL092022_Ian.pdf.

[12] For example, all municipalities in the State of Florida were eligible to apply for the Hazard Mitigation Grant Program See `https://gis.fema.gov/maps/dec_4673.pdf` and `https://www.fema.gov/sites/default/files/2020-07/fema_public-assistance-fact-sheet_10-2019.pdf` for more details.

Table 2: Sample and Observation Counts

	MLS		IHP		NFIP	
State	Counties	Sold Listings	Counties	Claims	Counties	Claims
FL	67	1,355,679	28	467,273	42	38,554
AL	4	49,037	0	0	0	0
MS	4	16,194	0	0	0	0
LA	15	41,737	0	0	0	0
TX	22	206,017	0	0	0	0
GA	9	36,410	0	0	4	4
SC	13	75,013	3	3,018	11	780
NC	23	75,436	0	0	15	136
N	157	1,855,523	31	470,291	72	39,474

Notes: The acronyms in the column titles are defined as MLS = Multiple Listing Services, IHP = Individuals and Households Program, and NFIP = National Flood Insurance Program. The MLS data capture real estate prices for transacted properties. The IHP and NFIP data are used to calculate damages that are classified into empirical treatments for places in Florida. For IHP data, the only counties kept are those that are in Florida and South Carolina that are affected by Hurricane Ian.

disaster. Additionally, it is important to acknowledge there may be confounding impacts of both COVID-19 and Hurricane Irma in the analysis. As we discuss in more detail later, we attempt to mitigate this by changing the pre-treatment window size.

3.2 Real Estate Data

Our real estate data are a composite of multiple listing services (MLS) data obtained from CoreLogic. Each MLS represents a regional database containing property characteristics input by real estate agents, including list date, contract date, list price, beds, baths, square footage, and address. As our focus is on residential properties, we keep only single-family homes, condominiums, townhouses, and manufactured homes, thereby excluding rental properties, commercial properties, and multifamily properties. After conducting basic data cleaning procedures, the sample comprises more than 1.8 million listings sold between March 2015 and June 2023.[13] Table 2 presents observation counts of counties by state for several data sources. Figure 1 maps out coverage in different ways. Panel (a) shows the MLS sample encompasses 20 major regional real estate listing services in Florida, as well as coastal counties on the Gulf Coast and Atlantic Coast (including as far north as North Carolina). Panel (b) indicates how FEMA may classify areas in different ways that affect eligible aid. The state is broadly shown in panel (c) compared to a zoomed view of ZIP Codes in Lee County in panel (d). Both panels have visual representations of the number of real estate listings (larger yellow dots indicate more listings) and claims (darker blue colors reflect more claims) throughout the sample period for the affected regions in Florida.

A potential concern is the changing composition of transacted homes before and after the disaster. For example, perhaps only those homes with higher quality construction survive the storm and are listed on the market. Alternatively, owners may want to move on quickly by selling damaged homes to avoid dealing with repairs and restorations. Therefore, estimated price effects could pick up sales of higher or lower quality homes, and not a true price effect. Table 3 contrasts transacted home characteristics after Ian against homes sold in different periods leading up to the storm. Generally, the characteristics are similar for homes sold before and after Ian. An exception is that homes sold after Ian are slightly smaller (40 sqft), but the amount is less than a small bedroom size. After the storm, homes also sell faster (16 days less), although this is a common trend noted across real estate markets. Statistical estimations can eliminate these issues by introducing proper controls, which is demonstrated later in Table 11 of the Appendix.

3.3 Public Disaster Damages Data

Real estate data are combined with disaster damages data from the OPEN Federal Emergency Management Agency (FEMA) Dataset, specifically the Individuals and Households Program (IHP) Valid Registrations data. As indicated on FEMA's website, the IHP is de-

[13] The wide sample range attempts to capture enough history while balancing the panel at a granular level for empirical estimations. A property can appear multiple times. Care is taken to distinguish between independent repeat listings versus deliberate attempts to edit a listing to reset the recorded time-on-market. The cleaning process reveals nested and overlapping listings. Nested listings are eliminated because they are entirely contained within another listing. Overlapping listings are adjusted because they have either a start date or an end date (but not both) within the range of another listing. Unsold listings that are relisted within 60 days are reclassified as the same listing, with adjustments made to the time-on-market and list date. Listings with negative or zero list prices, those with list dates after the data delivery date, listings for rental properties, and those without list dates or sale prices are removed. Additionally, the top and bottom 1% of homes with respect to price are excluded for each county(ZIP Code)-year pair to trim outliers.

Table 3: Comparing Housing Characteristics with Different Pre-Treatment Windows

	Before Ian				After Ian
	Since 2015	Since 2019	Since 2021	9 Months Prior	9 Months Later
beds	3.1	3.1	3.1	3.1	3.1
	(1.0)	(1.0)	(1.0)	(1.0)	(0.9)
baths	2.3	2.4	2.3	2.3	2.3
	(1.7)	(1.8)	(2.0)	(0.8)	(3.9)
home size (sqft)	1,879	1,897	1,886	1,826	1,781
	(27,542)	(37,643)	(43,251)	(4,313)	(771)
home age (years)	27	27	28	28	28
	(22)	(23)	(23)	(23)	(24)
cooling (yes/no)	99.6%	99.6%	99.5%	99.4%	99.2%
	(6.2%)	(6.0%)	(7.0%)	(7.9%)	(9.0%)
time-on-market	121	101	66	58	42
	(159)	(144)	(106)	(80)	(44)
sale price (logged)	12.37	12.53	12.70	12.76	12.70
	(0.79)	(0.77)	(0.74)	(0.72)	(0.73)
N	1,801,257	819,616	340,341	148,028	54,266

Notes: Based on author calculations using MLS data licensed from CoreLogic. The table compares summary statistic values before Hurricane Ian (for several period lengths) and after the storm (only for a nine month period based on data availability). Variables are selected based on being used later as controls for difference-in-difference estimations.[14] Each variable presents mean values with standard deviations below in parentheses. Beds and baths are the number of units reported in the home. Home size is the structure's total livable space in square feet (sqft). The pre-treatment standard deviations are large, but acceptable because values at the 99th percentiles do not exceed 5,000 sqft. Home age is defined as the difference between the closing date and year built. Cooling indicates (yes/no) if a cooling system is present. Time-on-market is the number of days from list date to contract date. Sale price is the natural log of the close price (nominal dollars) of the home.

signed to assist individuals in "meeting basic needs and supplementing disaster recovery efforts." Following a disaster, individuals can apply for funds if they are under-insured or lack insurance and face significant financial challenges due to the disaster. These funds can be used for temporary housing (like rental housing or a hotel), housing repairs for owner-occupied homes, hazard mitigation assistance (rebuilding for higher durability), and various other miscellaneous expenses. Sample selection might pose a concern in this particular dataset, as households and homes without flood insurance may differ from those with flood insurance. A more detailed discussion on this matter is provided in the conclusion of the paper, where we suggest other potential complementary data sources. A benefit of the IHP is that has a relatively short data lag and is available soon after a disaster.

The IHP data include property level, FEMA-determined estimates of disaster-induced damages to real property (such as floors, walls, access roads, bridges, electrical, plumbing, and HVAC) and personal property (i.e., appliances, furniture). IHP also provides a measure of damages to real and personal property solely due to flooding. However, since it is not intended to restore disaster-damaged property to pre-disaster conditions and does not consider land damage, total damange is likely to be underestimated.[15] Our primary measure is real property damages because personal property damages are unlikely to impact property prices if the home was intact. Although we are agnostic about the damage source(s), including a measure of flood damage would make IHP more comparable to publicly available National Flood Insurance Program (NFIP) data, which cover claims from flood-based damages.

The IHP data encompass both owners and renters, potentially giving rise to concerns of double counting (e.g., when a tenant files a claim for personal property and the owner files a claim for real property). To address this issue, data are limited to single-family detached homes, mobile homes, townhouses, and condominiums—effectively excluding non-housing claims. The most detailed location information available is only at the county and ZIP Code levels.[16] As a result, it is not possible to merge damages data with the listings data at the property level. Instead, aggregate damages are computed at the county and ZIP Code levels and presented both as counts of claims and dollar values of damages. These aggregate measures serve as the basis for treatment definitions.

For Table 4, selected summary statistics have been computed for the raw, individual level dataset broken down by all claims, only those where real property was damaged, and where there was flood damage. In the state of Florida, there are 470,291 claims associated with Hurricane Ian in the IHP data set (the same counts as shown in Table 2).[17]

Among these claims, per FEMA, about 83% are not associated with any real or personal property damages. In other words, most claims seek funds for shelter, temporary housing, and other miscellaneous expenses, rather than for repairing or replacing damaged property. Another explanation is that for FEMA to assess damages, an inspection must be issued and carried out ("returned"), and this would not be done for (possibly damaged) homes that were ineligible for aid. A concession with using only claims where FEMA determines there is real property damage is that could limit the sample to only residences that have experienced damage that might plausibly lower their values. On the other hand, it might restrict the sample too much and exclude damaged homes simply because they are declared ineligible

[15]FEMA explicitly states that the dataset is not an official federal report.

[16]ZIP Codes are not defined to align with county boundaries and may span multiple counties.

[17]Hurricane Ian has disaster number identifiers 4673 and 4677 in the states of Florida and South Carolina, respectively. A declaration number was issued for the Seminole Tribe of Florida, but it is excluded from the sample due to the lack of housing data for this region, and this subsample corresponds to only 3,874 claims after applying other filters, which is less than 1% of all claims.

Table 4: IHP Summary Statistics

	All Claims	Real Property Damage > 0	Flood Damage > 0
real property damage	2,950.12	17,426.77	26,092.47
	(14,813.720)	(32,311.440)	(42,555.900)
personal property damage	491.59	2,771.82	4,308.22
	(1,532.087)	(2,706.286)	(2,511.413)
flood damage amount	2,299.99	13,570.90	27,438.79
	(14,749.32)	(33,646.40)	(43,652.45)
primary residence	0.957	1.000	1.000
	(0.203)	(0.024)	(0.029)
homeowners insurance	0.6849	0.612	0.690
	(0.465)	(0.487)	(0.462)
flood insurance	0.163	0.241	0.363
	(0.369)	(0.427)	(0.481)
home damage	0.827	0.996	0.997
	(0.379)	(0.061)	(0.058)
auto damage	0.02	0.052	0.076
	(0.153)	(0.223)	(0.265)
emergency needs	0.464	0.436	0.443
	(0.499)	(0.496)	(0.497)
inspection returned	0.310	1.000	1.000
	(0.462)	(0.000)	(0.000)
rental assistance eligible	0.075	0.432	0.625
	(0.264)	(0.495)	(0.484)
repair assistance eligible	0.099	0.584	0.595
	(0.299)	(0.493)	(0.491)
replacement assistance eligible	0.003	0.018	0.031
	(0.056)	(0.134)	(0.174)
personal property eligible	0.025	0.140	0.170
	(0.155)	(0.347)	(0.376)
N	470,291	79,614	39,421

Notes: Data are from IHP, sourced from OpenFEMA. The variables real property damage, personal property damage, and flood damage amount (listed as rpfvl, ppfvl, and flood-damageamount in source data) are FEMA-determined amounts of real property, personal property and flood damages, respectively, in dollars. All other are indicator variables. In order for such a determination to take place, an inspection would be required. The remaining variables are dummy variables equal to one if the condition is true. Numbers given in tables are means of variables, with standard deviations given underneath in parentheses.

Table 5: IHP Treatment Definitions for County and ZIP Code Analyses

Treatment	Binary	Definition
treat1	✓	any claims/real property damage
treat2	✗	average real (positive) property damages
treat3	✓	real property damages above median
treat4	✓	real property damages above 75th percentile
treat5	✓	real property damages above 90th percentile
treat6	✓	at least 1% of households real property damage
treat7	✓	at least 2% of households real property damage
treat8	✓	at least 5% of households real property damage

Notes: All damages are attributed to real property. All counties with filed claims have some positive amount of real property damage. The second row indicates that treat2 is the only non-binary treatment. We define a household as 2.6 people, where population counts done at the county level using Census 2020 population estimates and American Community Survey (ACS) 2017-2021 estimates.

for FEMA assistance. There are 79,614 claims associated with positive real property damages, totaling about $1.4 billion in damages, with an average damage assessment of about $17,400. Interestingly, among those claims with personal property damage, less than 59% were classified as eligible for repairs. Finally, it is noted that among claims associated with real property damage, only around 24% had flood insurance while 61% had homeowners insurance, which typically covers wind and roof damage.

Table 5 delineates treatment definitions. To align with existing literature, we use mainly binary treatments. In the DiD estimation framework, classification of an individual home as treated hinges on its location in a county or ZIP Code satisfying that criterion. For SCM estimation, the county or ZIP Code itself serves as the focal point. There is no perfect solution to addressing the measurement error problem. One strategy is to compare several treatments, each generated by a different underlying aggregation procedure.

In Table 5, the initial definition stipulates that an area is considered treated if *any* claims in the area are present in the IHP dataset. While this allows numerous areas to be designated as treated, potentially excessively so, the IHP dataset might underestimate damages, introducing ambiguity to the net effect. Nevertheless, treat1 is viewed as the least restrictive treatment definition, consequently making it a potentially stringent control definition. Thus, setting aside the potential issue of double-counting, it remains unclear whether the claims reflected by treat1 would be associated with changes in property values. Consequently, the

Figure 2: Real Estate Market Activity at a State and County Level

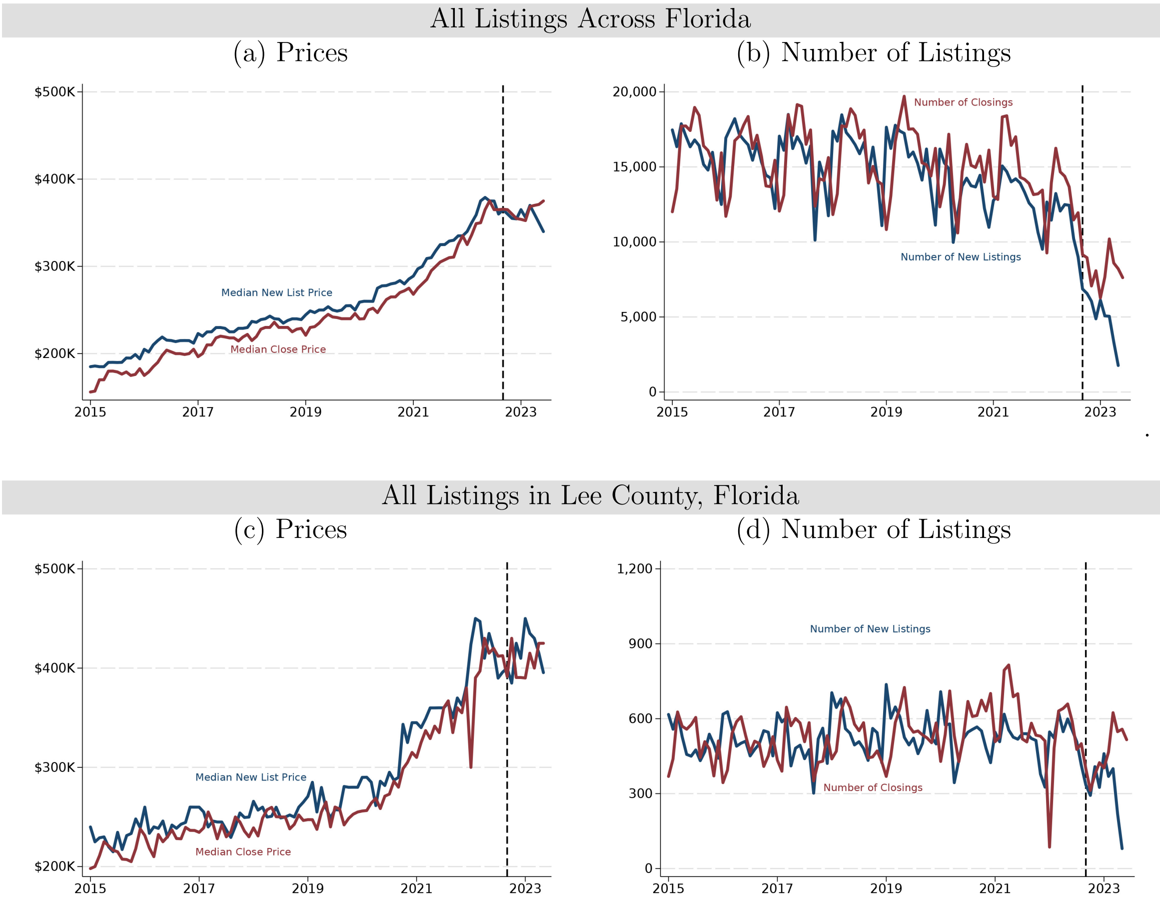

Notes: Based on author calculations using MLS data licensed from CoreLogic. The top row presents real estate activity for single-family residential homes across the entire state of Florida ($n = 1.5$ million) while the bottom row contains listings information from Lee County, Florida ($n = 73,000$) since 2015. Excluded property types are low/mid/high rise, half duplex, and villas. Panels in the first column display nominal dollar values for median new list price (navy color) and median close price (maroon color). Panels in the second column illustrate the number of new listings (navy) and closings (maroon). All series reflect a monthly time frequency without any adjustments for distributional moments, rolling windows, or seasonality. Certain periods are omitted when there are low counts (e.g., filters for fewer than 10 listings or closings prevent single listings from swaying the results).

other measures focus exclusively on claims involving real property damage. The second treatment definition (treat2) is the only non-binary measure and represents, for each area, the average dollar value of claims associated with real property damages.[18]

Among claims involving real property damage, treat3, treat4, and treat5 designate an area as treated if the property damage there exceeds (by distribution percentiles) that of other similarly affected areas. By adjusting the threshold, the objective is to observe the impact

[18]Continuous treatment measures pose econometric challenges as described by Callaway, Goodman-Bacon, and Sant'Anna (2024). Future work could utilize more recent methodology like that of Callaway, Goodman-Bacon, and Sant'Anna (2024), to handle heterogeneous treatments and damage measures.

of more stringent treatment definitions. For instance, treat5 identifies a treated area as one with real property damages at or above the 90th percentile relative to other areas, suggesting that this definition is potentially more conservative. It is important to highlight that, no adjustments have been made for variations in the populations of different areas, which is addressed by treat6, treat7, and treat8. A treatment has a value of one if a sufficient number of households in the area (according to the IHP data) experience real property damage.[19] Overall, our favored definition, treat3, seems to strike the right balance by designating an aggregated area as treated if the sum of its real property values of claims exceeds the median value of all other counties in the state of Florida with real property damage. Other candidate definitions are still included in tables and figures of results.

3.4 Preliminary Evidence

Figure 2 displays real estate market activity measured across various dimensions on a monthly basis. The vertical dashed black line in each panel marks Hurricane Ian's landfall in Lee County in September 2022. The state and county exhibit similar trends. Lee County's real estate market does not stand out as being different from other areas in the state, which gives promise for matching its trends to those of other samples.[20] The shapes of all the series remain qualitatively consistent whether displayed in raw amounts (as shown) or when using rolling averages or distributional adjustments (such as logarithmic values).

Reading from left to right across columns, the left column suggests a significant and positive impact on price appreciation from COVID-19, particularly for Lee County. However, Hurricane Ian appears to dampen price appreciation and even leads to a slight price reduction.[21]

[19] The average household is assumed to be 2.6 people to convert Census population estimates to households.

[20] Although not shown here, if we shift focus from all listings in Lee County to specific property types, the real estate market is influenced predominantly by detached single-family residences, which exhibit somewhat similar behavior to townhomes. Other property subtypes (i.e. condos, manufactured homes, or mobile homes) have limited sample size and greater volatility. The additional noise makes it difficult for statistical techniques to identify comparable samples in other areas and draw reliable conclusions for those subtypes. We present property subtype statistics and tests in the Appendix but, in the paper, combine all single-family properties together due to practical challenges.

[21] The list-to-close price ratio for each month was also constructed for both the top and bottom panel. The ratios exhibit a downward trend since 2015, indicating that close prices are progressively moving closer to initial list prices, along with rising prices in the first column. Evidence of a COVID-19 effect is also apparent, with sharp decreases in the ratio (both when computed as a mean and median), followed by increases after Ian and another decline. The decline during COVID-19 is easily noticeable with a flattened ratio perfectly equal to a unitary value, reflecting a period when homes are immediately being sold for list price or even more (as in the case of townhomes). The ratio increases just before and after Ian suggesting a return to a period of relative normalcy, reflecting pre-pandemic trends.

The right column reveals that this return to pre-pandemic behavior is short-lived, as the rapid increase in mortgage rates suppresses listings across the state and county. The stock has struggled to recover, even with buyers being in a more expensive lending environment. Figures 7 - 11 illustrate additional trends across different property types.

The various panels in these figures highlight several lessons. First, real estate market activity can be characterized differently based on whether the focus is on transacted prices or liquidity. Second, the signal can be noisy during periods of turmoil, such as COVID-19 or after a natural disaster. However, even a powerful storm like Hurricane Ian seems to have less impact than a pandemic, whether measured in immediate or long-term changes. Ian appears to slightly reduce prices, while the effect on the number of listings and transactions is not obvious visually. Third, real estate movements can vary significantly between property type categories, which is important because some categories are not as common. This note is empirically meaningful because it complicates issues like comparable sales sampling and introduces survival bias for damage estimations. Based on these series, it is unclear whether or to what extent the Lee County housing market has been affected by Hurricane Ian, despite being arguably the most impacted location by the costliest hurricane in Florida's history. In other words, if a hurricane could disrupt true real estate fundamentals, this event would be a promising candidate for study. Despite popular press reports and anecdotal observations that might suggest otherwise, the outcome is not evident with raw aggregated measures. The next two sections shed more light with sophisticated statistical techniques.

4 Price Effects with DiD Estimators

Is it possible to infer that Hurricane Ian has a property level or a marketwide price discount using DiD estimations? Over the next several subsections, we exploit an experimental framework where the unexpected natural disaster disrupts real estate prices. If Lee County resembles other parts of the state, a counterfactual control group can be designed to separate "treated" homes from "non-treated" homes using publicly-available damages data. The empirical tests consider both static treatment effects (traditional DiD) as well as the potential for a dynamic treatment effects (event study).[22] In both approaches, the DiD strategy involves attempting to find a suitable control for affected homes. Later, we try to construct

[22] The static DiD approach is implemented in STATA (v18.0) using the `didregress` command while the dynamic DiD routine is performed manually with the `reghdfe` command.

a suitable control using a SCM.[23] Throughout, both DiD and SCM approaches, we assume every treated unit becomes treated at the same time and stays treated for the remainder of the sample. This seems reasonable as Hurricane Ian has a short duration of only a week and the monthly data sample is limited to approximately one year after the disaster.[24]

4.1 Static Treatment Effects (Traditional DiD)

Our first approach lets price effects be the same, or static, at each time post-treatment period. Following the literature, we adopt the potential outcomes framework (Rubin, 1974; Angrist and Pischke, 2009; Imbens and Rubin, 2015), particularly that the stable unit treatment value assumption (SUTVA) holds. The SUTVA excludes spillovers from treated units to other units, thereby ruling out negative externalities from damaged homes to other homes.

Suppose price effects are homogenous across groups and time. If p_{igt} represents the price for home i in group g at time t, a single treatment is estimated with two-way fixed effects as:

$$ln(p_{igt}) = \alpha_g + \gamma_t + \beta D_{gt} + \theta_t x_{it} + u_{it} \tag{1}$$

where α_g denotes group fixed effects, γ_t represents time fixed effects, x_{it} includes (potentially) time-varying controls for home i, and D_{gt} corresponds to one of the treatment measures outlined in Table 5. Time is measured in months and normalized such that treatment occurs at time $t = 0$. In our analysis, time is measured at a monthly frequency and the first treatment month is September 2022.[25] Treatment $D_{gt} = treatk_g \times \mathbf{1}(t \geq 0)$, where $treatk_g$ is equal to 1 if group g satisfies the k^{th} treatment definition in Table 5. For all treatments except treat2 (a continuous measure), this is equivalent to

$$D_{gt} = \begin{cases} 1 & \text{if } t \geq 0 \text{ and group } g \text{ is treated} \\ 0 & \text{else} \end{cases} \tag{2}$$

[23]The SCM is executed in STATA using the `synth` and `synth_runner` commands, where the former gives the weights of donor units and the latter performs inference.

[24]A staggered treatment may be more appropriate for slower-moving disasters like wildfires, where the spread of the fire affects homes at different times. Future research could help determine the suitability of staggered treatments with higher-frequency data (e.g., weekly or daily).

[25]Because Ian did not occur at the initial day in a month, one is forced to misclassify some homes whether defining the first treated month as September 2022 or October 2022. We label all homes listed in September as treated, provided treat$k \neq 1$ for some k). This choice is not as problematic as might appear; it only affects homes listed in early September that are also sold before Ian struck. Fortunately, across the entire sample this subset amounted to only a few thousand homes, on the order of 0.2% of all observations.

The above specification has implicit restrictions. First, the treatment is binary, meaning there is no variation of exposure within the treated group. While exceptions exist, binary treatment is predominant in the literature. Despite having one aggregate treatment (treat2) that is non-binary, estimations adhere to the common use of binary treatments. Second, all units within the treated group undergo the same treatment effect (β), thereby precluding heterogeneity across units. In practice, accommodating both varying treatment and varying treatment effects is challenging. A recent study by Sun and Shapiro (2022) provides an example to suggest that if the level of treatment exposure and treatment effect are correlated, the estimated average treatment effect (on the treated) may fall outside the support of possible individual treatment effects in expectation. Following their approach, time-varying treatment effects are only allowed in later sections within an event-study framework. Finally, the specification requires parallel trends to hold when the dependent variable is in logs, not levels. We follow the hedonic home price literature and use log of the price as the dependent variable for the sake of comparability.

Once a group is labeled as "treated," numerous candidates are available as controls depending on sample restrictions. For instance, if a treated ZIP Code is defined by having at least 20 reports of damages (Kousky, Palim, and Pan, 2020), the control group can be sizable with all non-treated ZIP Codes across the entire state of Florida. Alternatively, it could be defined as all non-treated ZIP Codes in coastal counties or only those non-treated ZIP Codes in counties with treated ZIP Codes, depending on the imposed sample restrictions. To assess control group quality, we evaluate if parallel trends hold. This is commonly achieved through pre-trends tests, where researchers examine whether similar, parallel trends exist in the pre-treatment period. However, as previously mentioned, concerns have been raised about conventional pre-trends testing (Roth, 2022; Rambachan and Roth, 2023). Our preliminary results start a benchmark case where the control units encompass all untreated units.

Table 6 showcases the outcomes of the static treatment effect derived from estimating equation 1 with clustered standard errors at the group treatment level. A robustness check involves aggregating errors to address the limited number of groups (per Donald and Lang, 2007). Excluding the second (continuous) treatment, county level treatment effects range from 0.05 to 0.11, indicating that Hurricane Ian led to a price *increase*, contrary to the descriptive evidence. Similarly, at the ZIP Code level, treatment effects vary within a wider range. Additionally, a declining effect is observable when the definition becomes restrictive.

Shifting from treat3 to treat5 or from treat6 to treat8 means redefining treated homes as progressively more damaged, implying higher levels of treatment associated with less price appreciation. However, caution is warranted due to the evolving comparison group. Finally, less than half of the treatment effects are statistically significant at the 95% level, both at the county and ZIP Code levels.

The bottom of Table 6 includes robustness tests conducted commonly with static DiD models. There are two main identification conditions: the no-anticipation and the parallel trends assumptions (Roth et al., 2023). To assess the no-anticipation assumption, a Granger causality test detects any treatment effect prior to the actual treatment. Both the pre-trends and no-anticipation tests require treated units to be treated simultaneously, limiting some estimations due to specific treatment definitions. Further work could aim to balance groups to mitigate this issue, but the trade-off is a smaller sample and selection issues.

The results indicate sensitivity to the error term structure. Clustering errors at the county level, a common practice, often leads to tests failing, while aggregated errors tend to earn passing results. To test the parallel trends assumption, a (linear) pre-trends test is conducted. This involves modeling a linear time trend for each control and treated group and assessing whether their slopes differ. While the table suggests that most treatments do not reject the parallel trends test, additional testing could ascertain the accuracy of the results or assess potential low power issues (Roth, 2022). Furthermore, some sensitivity analysis is still recommended conditional on passing the pre-trends test (Rambachan and Roth, 2023)

To summarize the static results before delving into the dynamic case, the findings suggest a price premium in the range of 5% to 11%. Moreover, in some combinations of treatment definitions and estimation approaches, statistically significant treatment effects are observed, along with evidence supporting the no-anticipation assumption and passing the pre-trends test. Challenges arise if a static model is estimated when the correctly specified model should be dynamic (Sun and Abraham, 2021; De Chaisemartin and d'Haultfoeuille, 2023). Furthermore, results appear sensitive to the specific treatment and error term structure.

Table 6: Estimating a Static DiD Treatment and Verifying Valid Assumptions

	(a) Measuring ATT β at County Level								(b) Measuring ATT β at ZIP Code Level							
	treat1	treat2	treat3	treat4	treat5	treat6	treat7	treat8	treat1	treat2	treat3	treat4	treat5	treat6	treat7	treat8
β	**0.114**	–	**0.086**	**0.068**	0.050	0.057	0.055	0.046	**0.101**	–	**0.052**	**0.043**	0.009	**0.054**	0.030	0.006
(clust. errors)	**(0.039)**	–	**(0.032)**	**(0.030)**	(0.027)	(0.029)	(0.028)	(0.027)	**(0.010)**	–	**(0.009)**	**(0.011)**	(0.023)	**(0.012)**	(0.016)	(0.021)
β	**0.117**	–	**0.085**	0.064	0.048	**0.107**	0.071	0.051	**0.114**	–	**0.073**	**0.052**	0.017	**0.061**	**0.039**	0.003
(agg. errors)	**(0.022)**	–	**(0.027)**	(0.036)	(0.057)	**(0.031)**	(0.038)	(0.051)	**(0.007)**	–	**(0.008)**	**(0.010)**	(0.017)	**(0.010)**	**(0.013)**	(0.016)
	No Anticipation Test: Prob > F statistic								No Anticipation Test: Prob > F statistic							
clust. errors	–	–	0.0000	–	0.0000	–	0.0000	–	–	–	–	–	–	–	–	–
agg. errors	–	–	1.0000	–	1.0000	–	1.0000	–	–	–	–	–	–	–	–	–
	Pre-Trends Test: Prob > F statistic								Pre-Trends Test: Prob > F statistic							
clust. errors	–	–	0.0878	–	0.3201	–	0.2240	–	–	–	–	–	–	–	–	–
agg. errors	–	–	0.0103	–	0.6034	–	0.0282	–	–	–	–	–	–	–	–	–

Notes: The table is split with side-by-side panels to compare geographic differences when running estimations at the county and ZIP Code levels. The first row shows estimated value while the numbers in parentheses are standard errors clustered at that appropriate geographic level. A horizontal rule separates the second row which reports in parentheses the errors aggregated via the method of Donald and Lang (2007). The estimates correspond to the β of equation 1. A bolded result indicates statistical significance at the 95% level, i.e. p-values less than 0.05. The "No Anticipation Test" checks for Granger causality of a treatment effect prior to actual treatment. The null hypothesis is that the joint set of coefficients on all pre-treatment time periods (when assigned "treatment") is equal to zero. The "Pre-Trends Test" checks for differences between the control and treated group before treatment. Specifically, the pre-treatment trends are modeled linearly to test if slopes are different between the two groups. The null hypothesis is that the lines are parallel, i.e. have the same slope. The test is not performed for treat2 since it is a continuous treatment (coefficient magnitudes for that treatment are expressed in millions). Both tests require treatment to occur at the same time for all units. Thus, for some treatment definitions, data scarcity means there is effectively an "unbalanced" panel so that units appear to be treated later.

4.2 Dynamic Treatment Effects (Event Study)

Earlier regressions maintain that a hurricane's wrath is felt simultaneously in every location for all treated homes. Treatment effects can be altered to vary over time. More formally, our dynamic DiD specification uses an "event study" framework:

$$ln(p_{igt}) = \alpha_g + \gamma_t + \sum_{\tau=\underline{T}}^{\overline{T}} \beta_\tau D_{g,t-\tau} + \theta_t x_{it} + u_{it} \tag{3}$$

where τ is the index for relative days before and after the treatment and $\underline{T}$ and $\overline{T}$ give the range of parameters in the pre-treatment and post-treatment windows. The treatment D and variables are defined like in equation 1. The vector of coefficients β_τ replaces the singular parameter β from the static setting. The parallel trends assumption can be tested easily in the dynamic setting with more than two pre-treatment periods. Specifically, if treatment occurs at $t = 0$ and β_{-1} is normalized to 0, pre-treatment trends (of treated and control groups) are analyzed with a F-test using the hypothesis $H_0 : \beta_{\underline{T}} = \cdots = \beta_{-2} = 0$.[26]

Table 7 presents results for the dynamic treatment effect specification. The top part includes estimated treatment effects while the final three rows depict various pre-trends tests. The county level results yield several interesting findings. First, as anticipated and found earlier, treatment effects show considerable variation with treatment definitions. Second, statistical significance is achieved in many cases with positive and persistent price effects. For instance, using treat3 as the aggregation definition, there is an initial price premium of 11% with swings down to 6% and up to 15%. However, in contrast to the static DiD results, all these estimations fail common pre-trends tests, meaning more advanced pre-trends tests are not needed. These findings are robust over lengths of time windows prior to treatment. Hence, caution is warranted in making strong causal claims because the main identifying assumption merits exploration, especially about how to properly identify a suitable control.[27] While ZIP Code results are similar, some treatment definitions pass certain versions of the pre-trends test at different significance levels, unlike county versions. However, only one treatment definition (treat5) and one test (6 months prior to treatment) pass at a 5% level.

[26] Admittedly, there are known econometric challenges with pre-trends tests for two-way fixed effects (TWFE) settings, like used here and described in more detail by others (see Roth et al., 2023; Roth, 2022). Future work could assist the applied literature by offering alternative tests or improving baseline F-tests.

[27] The temporal frequency is a potential explanation. If the storm actually had a more precise and temporary impact of only a few weeks, or even days, then the monthly frequency could be attenuating the signal. Unfortunately, at the moment, the damages are not reported with enough granularity to explore this idea.

Table 7: Estimating a Dynamic Treatment at Varying Geographic Levels

	(a) Measuring ATT β at County Level								(b) Measuring ATT β at ZIP Code Level							
	treat1	treat2	treat3	treat4	treat5	treat6	treat7	treat8	treat1	treat2	treat3	treat4	treat5	treat6	treat7	treat8
Sept 2022	**0.141**	–	**0.105**	**0.107**	0.056	0.076	0.069	0.026	**0.136**	–	**0.101**	**0.111**	**0.132**	**0.206**	0.067	**0.348**
	(0.048)	–	**(0.045)**	**(0.045)**	(0.064)	(0.051)	(0.058)	(0.072)	**(0.018)**	–	**(0.017)**	**(0.021)**	**(0.018)**	**(0.110)**	(0.112)	**(0.011)**
Oct 2022	**0.126**	–	**0.085**	0.078	-0.033	0.070	0.060	-0.013	**0.126**	–	**0.081**	**0.079**	**0.132**	0.074	-0.127	**0.635**
	(0.046)	–	**(0.042)**	(0.045)	(0.040)	(0.053)	(0.060)	(0.054)	**(0.018)**	–	**(0.017)**	**(0.022)**	**(0.017)**	(0.111)	(0.132)	(0.015)
Nov 2022	**0.136**	–	**0.100**	**0.087**	0.057	**0.075**	0.063	0.058	**0.110**	–	**0.060**	**0.068**	**0.114**	-0.044	**-0.280**	**-0.157**
	(0.042)	–	**(0.037)**	**(0.034)**	(0.036)	**(0.035)**	(0.035)	(0.040)	**(0.019)**	–	**(0.020)**	**(0.024)**	**(0.019)**	(0.098)	**(0.061)**	**(0.023)**
Dec 2022	**0.155**	–	**0.122**	**0.124**	**0.111**	**0.133**	**0.116**	**0.077**	**0.121**	–	**0.087**	**0.096**	**0.130**	-0.009	-0.105	**-0.081**
	(0.043)	–	**(0.040)**	**(0.036)**	**(0.039)**	**(0.039)**	**(0.041)**	**(0.036)**	**(0.020)**	–	**(0.020)**	**(0.025)**	**(0.019)**	(0.086)	(0.088)	**(0.011)**
Jan 2023	**0.165**	–	**0.128**	**0.114**	0.064	**0.116**	**0.101**	0.049	**0.138**	–	**0.091**	**0.079**	**0.136**	-0.041	-0.050	0.044
	(0.043)	–	**(0.040)**	**(0.039)**	(0.040)	**(0.041)**	**(0.043)**	(0.037)	**(0.020)**	–	**(0.018)**	**(0.023)**	**(0.019)**	(0.117)	(0.091)	(0.029)
Feb 2023	**0.135**	–	**0.090**	0.066	0.048	0.064	0.044	0.072	**0.131**	–	**0.079**	**0.064**	**0.124**	0.095	0.021	**0.334**
	(0.037)	–	**(0.035)**	(0.035)	(0.035)	(0.037)	(0.038)	(0.042)	**(0.020)**	–	**(0.018)**	**(0.022)**	**(0.019)**	(0.096)	(0.070)	**(0.014)**
Mar 2023	**0.121**	–	**0.117**	**0.098**	0.054	**0.080**	0.074	**0.075**	**0.129**	–	**0.073**	**0.079**	**0.120**	0.126	0.088	**-0.301**
	(0.033)	–	**(0.031)**	**(0.032)**	(0.036)	**(0.035)**	(0.039)	**(0.036)**	**(0.020)**	–	**(0.023)**	**(0.020)**	**(0.020)**	(0.117)	(0.067)	**(0.015)**
Apr 2023	**0.135**	–	0.058	0.047	-0.012	0.001	-0.002	0.039	**0.118**	–	**0.076**	**0.075**	**0.121**	0.061	0.126	**0.063**
	(0.034)	–	(0.037)	(0.039)	(0.047)	(0.036)	(0.038)	(0.029)	**(0.019)**	–	**(0.018)**	**(0.022)**	**(0.018)**	(0.111)	(0.111)	**(0.016)**
May 2023	**0.155**	–	**0.101**	0.049	-0.046	0.025	-0.010	-0.037	**0.149**	–	**0.093**	**0.067**	**0.134**	0.093	0.012	**0.067**
	(0.039)	–	**(0.041)**	(0.043)	(0.041)	(0.044)	(0.045)	(0.052)	**(0.026)**	–	**(0.025)**	**(0.029)**	**(0.025)**	(0.095)	(0.079)	**(0.024)**
June 2023	**0.262**	–	0.154	**0.235**	0.106	0.150	0.191	-0.365	**0.169**	–	0.104	0.039	**0.165**	-0.282	-0.296	-
	(0.091)	–	(0.103)	**(0.104)**	(0.235)	(0.154)	(0.159)	(0.194)	**(0.076)**	–	(0.071)	(0.088)	**(0.073)**	(0.265)	(0.330)	-
	Pre-Trends Test: Prob > F statistic								Pre-Trends Test: Prob > F statistic							
All	0.0000	–	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	-	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
12 months	0.0000	–	0.0046	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	-	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
6 months	0.0097	–	0.0155	0.0443	0.0000	0.0209	0.0002	0.0000	0.0000	-	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Notes: The table is split with side-by-side panels to compare geographic differences when running estimations at the county and ZIP Code levels. Across both geographies, the top half of the table provides estimates of β_k in equation 3, while the numbers in parentheses are the associated robust standard errors, clustered at the appropriate geographic level. Bolded estimates are statistically significant at the 5% level. In the bottom half of the table, the numbers correspond to F-tests on pre-treatment coefficients of the form $H_0 : \beta_{-2} = \beta_{-3} = \cdots = 0$, where the different rows indicate how many coefficients are included in the test. All pre-treatment coefficients are tested, as well as only those coefficients in the last 6 and 12 months prior to treatment. The No Anticipation Tests were not performed in this table due to the results shown for the Pre-Trends Tests.

Sometimes results are easier to interpret when presented graphically. Figure 3 (for county level) and Figure 4 (for ZIP Code level) depict estimated coefficients across various panels for each treatment. These illustrations highlight longer historical coverage than can be easily conveyed in Table 7. Across the panels, they show significant noise and a positive uptick in the treatment effect coefficients at both group levels. The positive uptick, which crosses the zero axis, happens approximately 30 months (around March 2020) before Hurricane Ian struck, which matches the COVID-19 shock. This interesting visual finding creates a problem because it highlights the failure of the pre-trends tests: areas hit by Ian tend to be the same areas that experienced relative price gains in the pandemic. Finally, the graphs emphasize a trade-off with temporal frequency. Due to the high volatility of the indices, even after adjusting for seasonality with month fixed effects, it seems plausible that annual time frequencies may be more likely to pass pre-trends tests than monthly time frequencies, as the former have the fluctuations "smoothed out." This is an important practical note because studies often lack granular data that could perform event studies at a higher frequency. By measuring with an annual frequency, other works could be missing important signals and accidentally bypassing empirical considerations that would question causal conclusions.

4.3 Are Static and Dynamic Treatment Effects Good Enough?

Typical causal strategies may not work as well for studying the effects of natural disasters. The failure of the parallel trends assumption is concerning because there may be a lingering influence from COVID-19 or other events. Hurricane Ian could have changed market-wide price levels or price appreciation, but the identification is not convincing without a better control group for Lee County and other affected areas. Matching could occur before estimation (Abadie and Imbens, 2011; Daw and Hatfield, 2018; Ham and Miratrix, 2022). However, caution is warranted because that assumes the treatment is known with certainty. Another potential avenue is the fuzzy DiD framework proposed by De Chaisemartin and d'Haultfoeuille (2018), allowing for some units classified in the treatment group to be untreated and some in the control group to be treated.

Future work might explore the potential benefits of narrowing the treatment window to mitigate potential contamination from COVID-19. Earlier estimations partially addressed this suggestion by conducting pre-trends tests with different windows around the treatments. As demonstrated in Table 7, except for one treatment definition (treat5) at the ZIP-code level, this adjustment does not alter the qualitative outcome of the test: the pre-trends tests still indicate a failure of the parallel trends assumption. Instead of selecting certain

Figure 3: Event Study Graphs to Motivate DiD at County Level

(a) Treatment 1

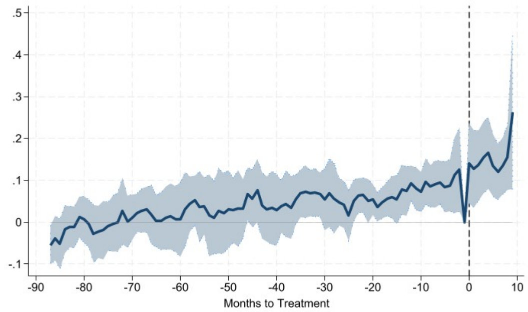

(b) Treatment 2

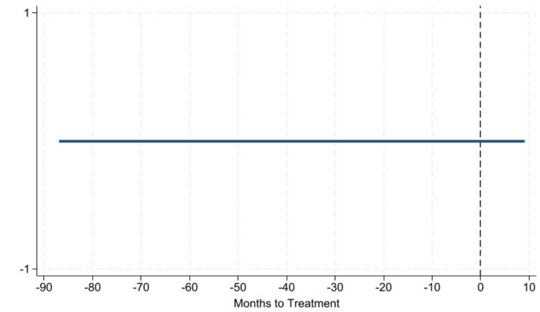

(c) Treatment 3

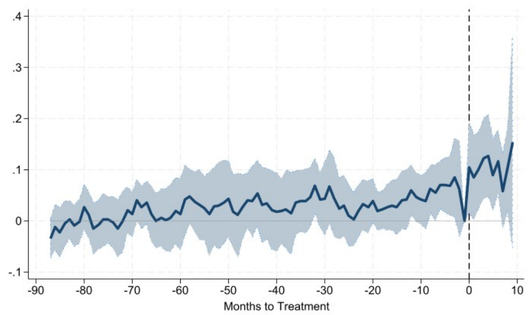

(d) Treatment 4

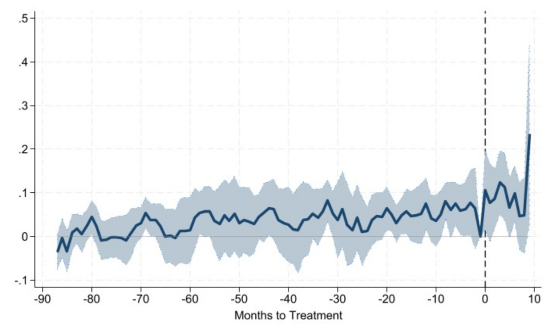

(e) Treatment 5

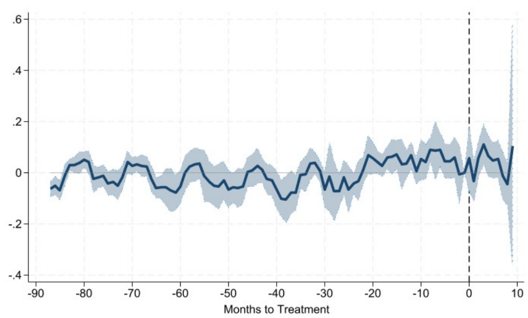

(f) Treatment 6

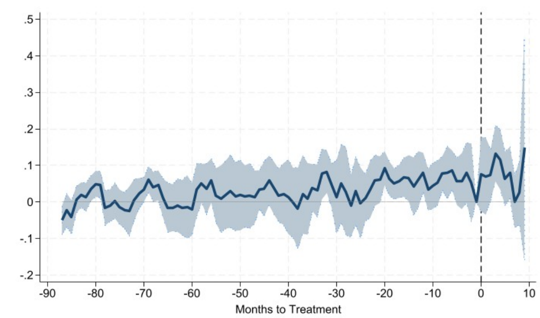

(g) Treatment 7

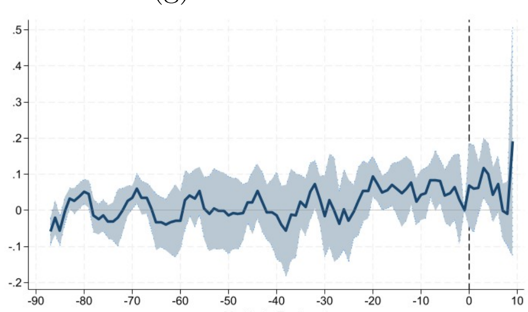

(h) Treatment 8

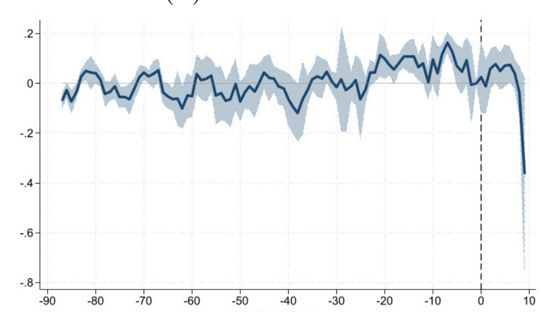

Notes: Based on author calculations. Vertical axis is the estimated treatment effect coefficient β_t from equation 3, where treatment is normalized to occur at $t = 0$ and β_{-1} is normalized to be zero. Estimated at monthly time frequency, where the geographic aggregation is at the county level.

Figure 4: Event Study Graphs to Motivate DiD at ZIP Code Level

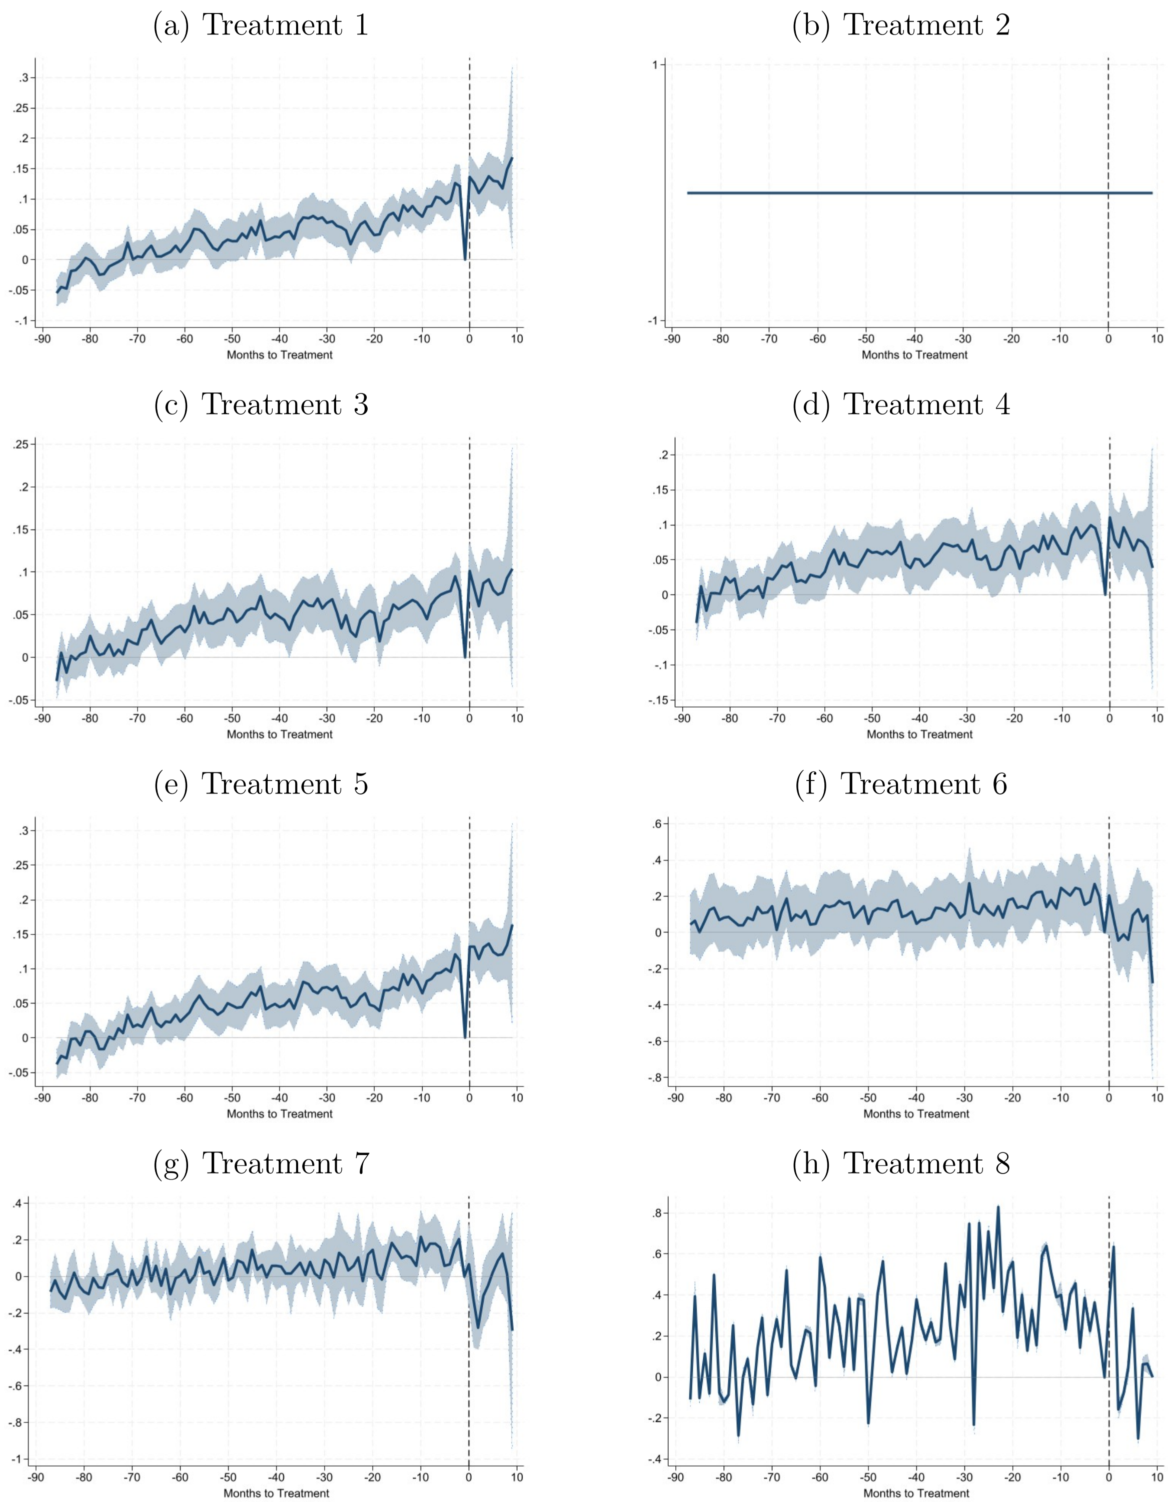

Notes: Based on author calculations. Vertical axis is the estimated treatment effect coefficient β_t from equation 3, where treatment is normalized to occur at $t = 0$ and β_{-1} is normalized to be zero. Estimated at monthly time frequency, where the geographic aggregation is at the ZIP Code level.

windows of pre-treatment coefficients, estimations could be conditioned to not even calculate coefficients prior to certain periods. Overall, these additional robustness checks may still not be enough. The subsequent section endeavors to tackle that challenge.

5 Price Effects with Synthetic Control Estimators

A variety of unexpected events have led to peculiar real estate market activity during this decade. A contribution of this paper is showing whether more advanced causal techniques might enhance results of traditional specifications. If DiD indicates a weak control, can one instead construct a suitable alternative using SCM?

For a given treatment definition, an appropriate control can be created through a weighted average of untreated "donor" regions (composed of samples of either counties or ZIP Codes). With IHP data, donors are drawn from across Florida, all coastal Gulf Coast counties, and all coastal Atlantic counties up to and including North Carolina. As a matter of practical advice, we emphasize that the pool of candidate donors changes with treatment definitions. Additionally, the empirical setup involves design-based inference rather than sampling-based inference since the sample represents the population or near population of all regions, not a random or stratified draw.[28] This design affects the reported inference later.[29]

The SCM approach has several advantages. First, it can mitigate problems with spillovers (interference).[30] Although spillover effects between aggregate units may still be present, it seems plausible that, at least for some level of aggregation, they may be minimized.[31] Resolving the issue is tangential to this present study, but future research could advance this growing econometric literature. Second, the SCM approach makes no assumptions on the observed distribution of the treated potential outcome; the only assumptions made are on the distribution of the non-treated potential outcome. Finally, synthetic controls are favored over regression approaches, as the latter might assign weights outside of $[0, 1]$ to control units, leading to an extrapolation problem.

[28] Refer to Abadie et al. (2020) for a comprehensive discussion.

[29] As emphasized by Abadie (2021), the placebo tests forming the p-values do not rely on approximations of sampling distributions of test statistics.

[30] See Imbens and Rubin (2015) for more details.

[31] Asserting that spillover effects do not exist at the state level implies that the price trends of other states would have evolved similarly had Florida not been struck by Hurricane Ian. In other words, spillover effects would exist if Floridians moved to other states after Ian, perhaps due to updated risk beliefs, thereby driving up demand and real estate prices.

As a caution, an often unstated drawback of the SCM approach is its requirement of a balanced panel.[32] This becomes particularly problematic in housing market cycles characterized by infrequent transactions. Imputation and other approximation methods seem to be the only solutions to executing the routines.

We closely follow the notation of Abadie (2021). The observed outcome Y_{it} is related to the potential outcome from treatment, Y_{it}^{I}, and the potential of not being treated, Y_{it}^{N}, through:

$$
\begin{aligned}
Y_{it} &= D_{it}Y_{it}^{I} + (1 - D_{it})Y_{it}^{N} && (4)\\
&= \tau_{it}D_{it} + Y_{it}^{N} && (5)
\end{aligned}
$$

where $D_{it} = 1$ indicates treatment for unit i at time t and 0 otherwise. The parameter of interest is denoted as $\tau_{it} = Y_{it}^{I} - Y_{it}^{N}$. The dependent variable, Y_{it}, represents a market price for area i at time t.[33] One data generating process for the potential outcome could be a linear factor model like $Y_{it}^{N} = \delta_t + \theta_t Z_i + \lambda_t \mu_i + \epsilon_{it}$ with Z_i as a vector of K observable explanatory variables and μ_i representing unobservable components, or "factor loadings."[34]

The treatment effect τ_{it} is allowed to vary over time and by unit, with the latter being particularly useful in the case of multiple treated units. Here, treatment occurs simultaneously for each treated unit and treatment time is defined as $t = 0$. For ease of exposition, we suppose that there is one treated unit and J untreated units. Without loss of generality we let the treated unit be the first unit, so that $i = 1$ is treated at time $t = 0$, while $i = 2, \ldots, J + 1$ remain untreated. In cases involving multiple treated units, such as when multiple areas are hit by a storm, synthetic controls are constructed separately for each treated unit, and the effects are subsequently aggregated to obtain an average treatment effect (on the treated).

The counterfactual non-treated outcome Y_{it}^{N} is unobserved for treated units and must be estimated from available data. A synthetic control for unit $i = 1$ is constructed as a weighted

[32] The literature does not seem to have a SCM approach that circumvents this requirement.

[33] Price is measured as both mean and median for both counties and ZIP Codes to test sensitivities. The measure underwent transformations, such as redefining as a rolling average and being expressed in logarithmic terms, but these are not reported due to qualitatively similar results but poorer fit for SCM estimations.

[34] There are multiple potential data generating processes. For instance, one might model a vector autoregressive process. Future work might explore how the options affect damage estimations from natural disasters.

average of non-treated units, resulting in the estimated treatment effect (on the treated):

$$\widehat{\tau_{1t}} = Y_{1t} - \sum_{j\geq 2} \widehat{w_j} Y_{jt} \tag{6}$$

where $\{\widehat{w_j}\}_{j\geq 2}$ are the estimated weights. The estimated treatment effects depend on both the selection of the donor pool (homes $2, \dots, J+1$) and the estimated weights. The donor pool's choice is less critical since irrelevant donors are likely to receive zero weights.

For the weights, we impose non-negativity and sum-to-one constraints: $w_j \geq 0$ for all j, and $\sum_{j\geq 2} w_j = 1$. Let $Z_0 = (Z_2, \dots, Z_{J+1)})$ represent the vector of observable characteristics for untreated units, and Z_1 is defined similarly for the treated unit. The observable characteristics $Z = (Z_0, Z_1)$ may include a market area's characteristics or lagged values of the dependent variable, but they remain time-invariant. As pointed out by Kaul et al. (2015), including other variables is redundant if all past dependent variable values are already incorporated. To optimize fit while following the approach of Ferman, Pinto, and Possebom (2020), Z includes all past dependent variable values.

In the literature, the estimated weights $\widehat{w_j}$ are determined to ensure that synthetic controls closely resemble the treated unit during the pre-treatment periods. This process involves two steps. First, the weights are estimated conditionally on a vector $v = (v_1, \dots, v_k, \dots, v_K)$ that indicates the relative importance of each covariate in the matching procedure, yielding $w^*(v)$. Subsequently, the vector v is selected to minimize the prediction error of the synthetic control against the observed value of the treated unit in the pre-treatment period. More formally, the optimal (conditional) weights $w^*(v)$ are chosen to minimize:

$$\min_{w} ||Z_1 - Z_0 w|| = \min_{(w_2,\dots,w_{J+1})} \left(\sum_{k=1}^{K} v_k (Z_{1,k} - w_2 Z_{2,k} - \cdots - w_{J+1} Z_{J+1,k})^2 \right)^{\frac{1}{2}} \tag{7}$$

where the sum is over all K characteristics in Z_j, with $Z_{j,k}$ representing the value of the kth characteristic for unit j. To determine the value of v^* and the (unconditional) weights $\{\widehat{w_j}\}_{j\geq 2}$, v is chosen to minimize the mean squared prediction error (MSPE) of the synthetic control in a subset $\mathcal{T}_0$ of the pre-treatment periods. In other words, v will minimize

$$\sum_{t\in\mathcal{T}_0} (Y_{1t} - w_2(v) Y_{2t} - \cdots - w_{J+1}(v) Y_{J+1,t})^2 \tag{8}$$

Table 8: Synthetic Weights $\widehat{w}$ for Applying Treatment to Lee County, FL

	Mean Price								Median Price							
	treat1	treat2	treat3	treat4	treat5	treat6	treat7	treat8	treat1	treat2	treat3	treat4	treat5	treat6	treat7	treat8
Bryan, GA	0.094	0.01	0.04	0.056	0.02	0.055	0.049	0.015	.205	.078	0.239	0.187	0.116	0.193	0.215	0.095
Camden, GA	–	–	–	–	–	–	–	–	–	0.023	–	–	0.01	–	–	0.019
Carteret, NC	0.122	–	–	–	–	–	–	–	0.057	–	0.017	0.068	–	0.042	0.045	–
Columbia, FL	–	–	–	–	–	–	–	–	0.049	–	0.055	0.101	0.003	0.084	0.091	0.005
Collier, FL	–	0.31	–	–	0.353	–	–	0.286	–	0.226	–	–	0.275	–	–	0.228
Fort Bend, TX	0.197	0.086	0.165	0.165	0.086	0.177	0.192	0.09	0.103	0.062	0.104	0.12	0.045	0.123	0.115	0.053
Franklin, FL	–	0.019	0.042	0.062	0.011	0.051	0.033	0.03	0.047	0.029	0.056	0.046	0.029	0.038	0.042	0.029
Gadsden, FL	0.014	0.061	–	–	0.053	–	–	0.046	–	–	–		–	–	–	–
Glades, FL	–	–	–	–	–	–	–	–	–	0.047	0.098	0.006	0.046	–	0.037	0.042
Glynn, GA	0.046	–	0.051	–	–	0.02	–	–	–	–	–	–	–	–	–	–
Gulf, FL	0.024	0.022	0.008	–	0.029	–	0.006	0.021	0.024	0.019	0.026	0.008	0.026	0.016	0.02	0.023
Hardee, FL	–	0.014	–	–	0.015	–	–	–	–	–	–	–	–	–	–	–
Hernando, FL	–	–	–	–	–	–	–	–	–	0.093	–	–	0.038	–	–	0.104
Holmes, FL	–	–	–	–	–	–	–	–	–	0.008	–	–	–	–	–	0.004
Jackson, FL	–	0.016	–	–	0.029	–	–	0.01	–	–	–	–	–	–	–	–
Jefferson, TX	–	–	–	–	–	–	–	–	–	0.015	–	–	0.008	–	–	0.004
Liberty, GA	0.025	–	0.066	0.097	–	0.066	0.012	–	–	–	–	–	–	–	–	–
Madison, FL	–	–	0.006	–	–	–	–	–	0.011	0.035	–	–	0.026	–	–	0.03
Marion, FL	–	–	–	–	–	–	–	–	0.042	0.069	–	–	0.247	–	–	0.052
Matagorda, TX	–	–	–	–	–	–	–	–	0.04	0.057	0.05	0.035	0.056	0.029	0.034	0.064
Miami-Dade, FL	0.258	–	0.066	0.03	0.029	0.038	0.021	–	–	–	0.08	–	–	–	0.041	–
Monroe, FL	–	–	0.015	–	0.003	0.007	0.001	–	–	0.002	0.024	0.034	0.01	0.039	0.03	0.004
New Hanover, NC	–	0.006	–	–	0.029	–	–	0.009	–	–	–	–	–	–	–	–
Palm Beach, FL	–	–	0.336	0.432	–	0.383	0.453	–	–	–	–	–	–	–	–	–
Pender, NC	–	0.003	–	–	0.028	–	–	–	0.083	–	0.009	–	–	–	–	–
Putnam, FL	–	0.155	–	–	0.21	–	–	0.181	–	0.011	–	–	0.006	–	–	0.029
Sarasota, FL	–	0.137	–	–	–	–	–	0.168	–	0.0156	–	–	–	–	–	0.143
St. Lucie, FL	0.066	0.079	0.119	0.052	0.033	0.094	0.109	0.055	0.257	–	0.112	0.259	–	0.275	0.209	–
Taylor, FL	0.05	–	–	–	–	0.007	0.026	–	0.001	0.022	0.03	0.04	–	0.064	.032	0.023
Terrebonne, LA	–	0.032	–	–	0.018	–	–	0.04	–	–	–	–	–	–	–	–
Waller, TX	0.038	0.024	0.041	0.05	0.034	0.043	0.042	0.023	0.021	0.025	0.031	0.026	0.036	0.027	0.022	0.026
Walton, FL	0.066	0.025	0.047	0.056	0.02	0.058	0.057	0.026	0.061	0.025	0.069	0.068	–	0.042	0.045	–

Notes: Control counties are not the same for each of the computation. The – sign denotes when a county is not used as a control.

The formula implies a division of the pre-treatment period into a training and validation period, where weight estimation utilizes data exclusively from the training pre-periods. Specifically, if there are T_{pre} periods before treatment in the sample, the training periods constitute the initial t_0 periods, and the validation periods encompass the remaining $T_{pre} - t_0$ periods before treatment. Here, the entire pre-treatment period is used as a training sample to maximize fit and ascertain the feasibility of a good fit.

5.1 Results Without Imputation

Implicit in the previously outlined methodology is the requirement for a balanced panel of observations. Instead of imputing for missing counties or ZIP Codes, unbalanced areas are omitted and the sample size reduces considerably. For instance, the county dataset has 140 groups after basic data cleaning then balancing removes 53 places (leaving 87 counties or 62%) while the ZIP Code dataset has 1,675 groups and the balancing removes 985 places (leaving 690 ZIP Codes or 41%). These balanced datasets are utilized as baselines. Results below discuss each geographic sample.

Prior to introducing results, Table 8 presents synthetic weights for donor counties matched to the treated Lee County. As anticipated, weights vary by treatment definition. For instance, under the least restrictive definition (treat1) and consequently the most stringent control group definition, the largest weight is held by Miami-Dade County. Elsewhere, it might have one of the smallest weights (treat4) or even none at all (treat2 or treat6). This variability occurs because, with a more restrictive treatment, a greater number of units become available in the donor pool, allowing for a larger choice set. When comparing treat6 through treat8, where the definition becomes progressively stricter, both Franklin County and Walton County lose their relevance when sufficiently similar counties to Lee become eligible for the donor pool. However, this monotonic pattern is not always the case, as seen with Fort Bend County, TX.

Given the highly localized nature of real estate markets, identifying places that move in parallel with an affected area and remain unaffected themselves can prove challenging. In this paper, Lee County borders Collier County, which was also affected by Ian. Therefore, Collier County could serve as a control only under a particularly stringent treatment definition that classifies Lee County as treated but not Collier. In Table 8, Collier becomes eligible to be a donor only under treat5 and treat8, but not treat3, treat4, treat6, or treat 7, which are less strict treatment definitions. Conducting a similar analysis for median prices yields comparable results, albeit with slightly different counties. This outcome highlights the functional form of the dependent variable can also be important in SCM settings.

Regarding the SCM framework, outcomes appear to yield a reasonable pre-treatment fit in Figures 5 and 6, where each panel depicts a treatment. The top (navy) line represents the treated unit (Lee County), and the bottom (maroon) line represents its synthetic control, with the vertical gap indicating the treatment effect. The figures convey a temporary and time-varying price premium that persists until the sample's end, approximately nine months. The premium appears to have peaked about three months after the storm, which may attributable to the counterfactual control experiencing a decline in prices.[35]

Table 9 presents the estimated treatment effects from equation 6, along with standardized p-values, for real estate prices across Lee County.[36] The treatment effects indicate positive and lasting price premiums due to Ian, though they are mostly statistically insignificant with respect to the standardized p-values. Specifically, none of the treatment effects for the mean are statistically significant, while definitions treat3 to treat7 for the median have some (but not all) treatment effects statistically significant at the 5% level. However, the non-standardized p-values often differ significantly from the standardized values, which raises some concerns about the results. Thus despite the apparent good fit in Figures 5 and 6, the inference statistics suggest less conclusive evidence.

A reasonable critique is the large geographic scope. A hurricane's damaging impacts might be isolated to smaller locations. Focus levels are shifted from a whole county to ZIP Codes. Synthetic control analysis is applied to two individual ZIP Codes, one in Bonita Springs (coastal-adjacent) and another in the Miromar Lakes neighborhood (more inland). Importantly, Bonita Springs experienced large aggregate damages (average real property damage claim of \$19,690.39), while Miromar Lakes had more moderate aggregate damages (\$8,665.17). Both ZIP Codes had lower damage estimates than their county average of \$28,919.60.[37]

[35]Although not depicted, if logs are used as dependent variables, the pre-treatment fits are considerably poorer and figures generally suggest positive price effects (premiums) of around 10%.

[36]For inference, p-values are derived from placebo tests, where control units are assigned treatment, and then the number of placebo effects greater than the original estimated treatment effect is calculated for each post-treatment parameter. Standardized p-values, which scale the effects by the root mean square prediction error before comparing treatment effects, are reported (Galiani and Quistorff, 2017).

[37]Graphs and tables of weights for the ZIP Code analysis are not shown, but are available upon request. We also shift the dependent variable to median price to ensure a central tendency that is less affected by extreme distributional values.

Figure 5: SCM Estimates for Lee County by Mean Closing Price

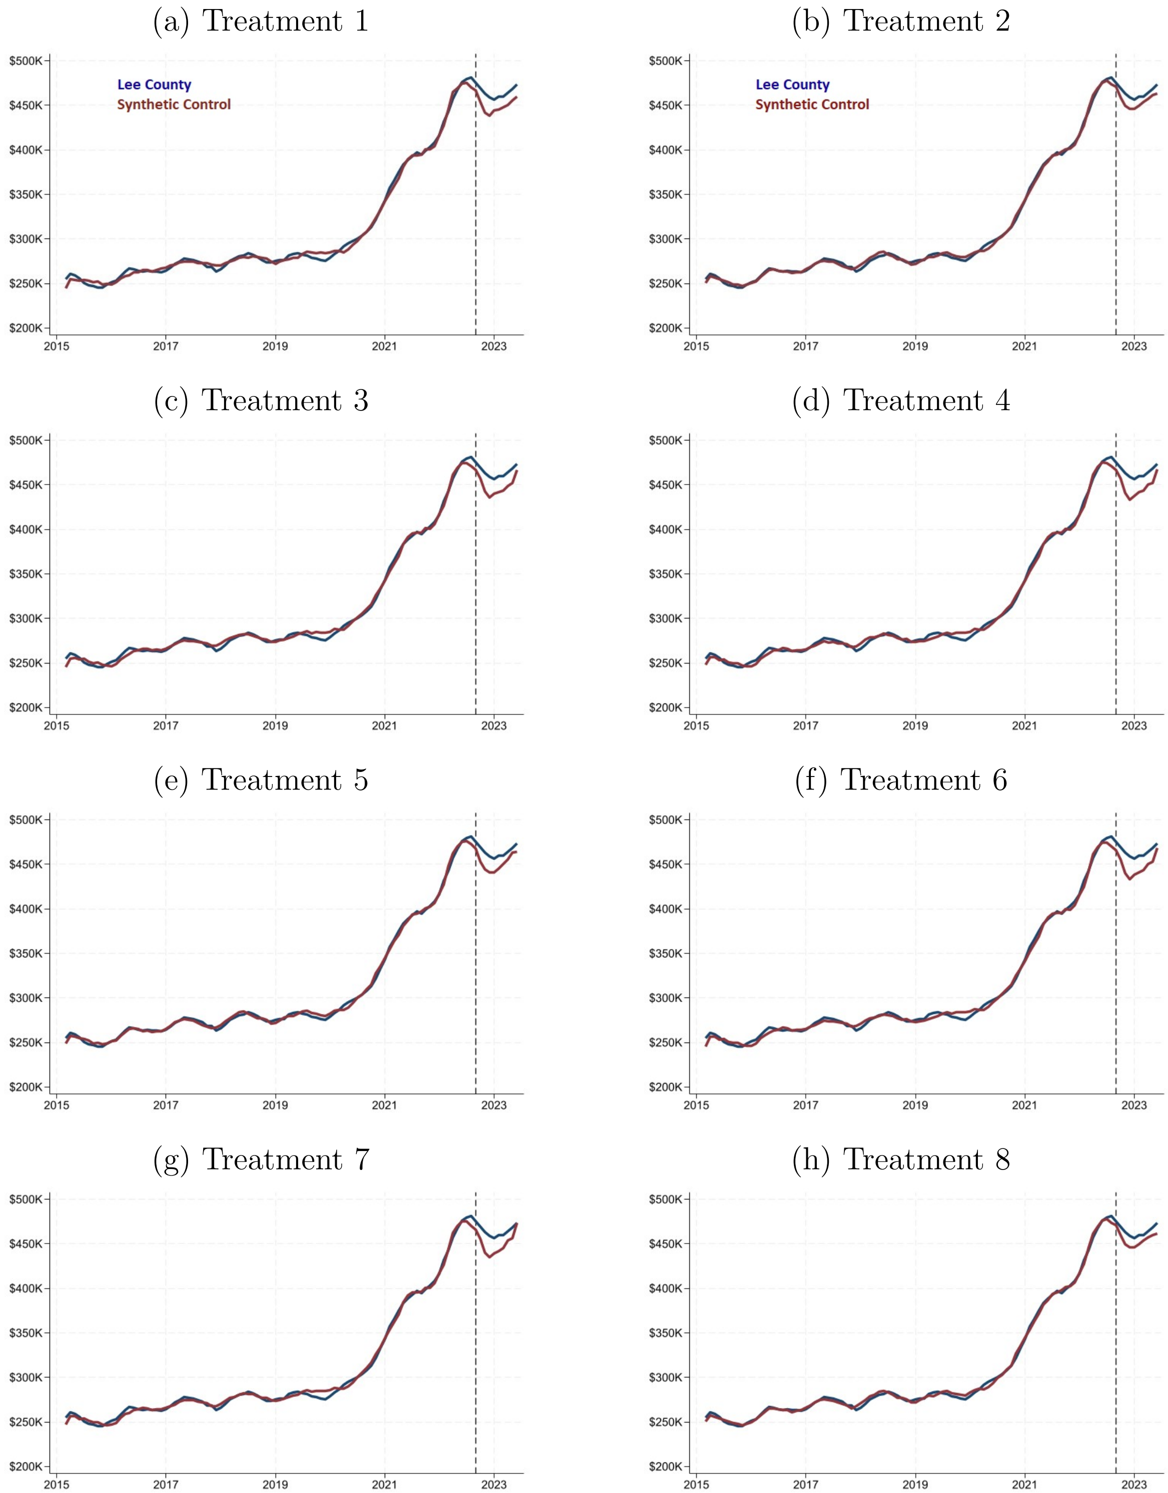

Notes: Based on author calculations. Vertical axis measures mean close price in dollars, where navy line is treated unit (Lee County) and maroon line is its synthetic control. The vertical distance between the lines after treatment corresponds to the treatment effect τ_t in equation 6.

Figure 6: SCM Estimates for Lee County by Median Closing Price

(a) Treatment 1

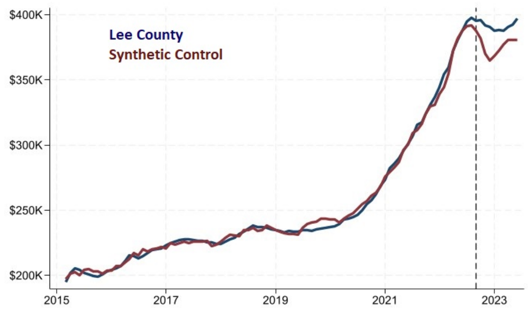

(b) Treatment 2

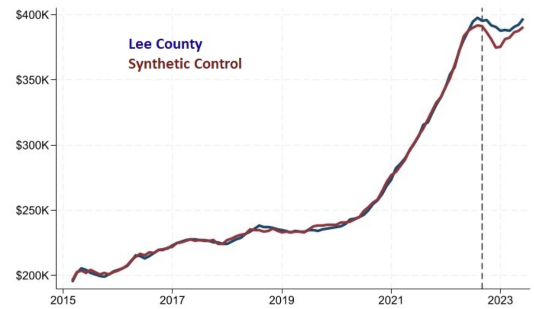

(c) Treatment 3

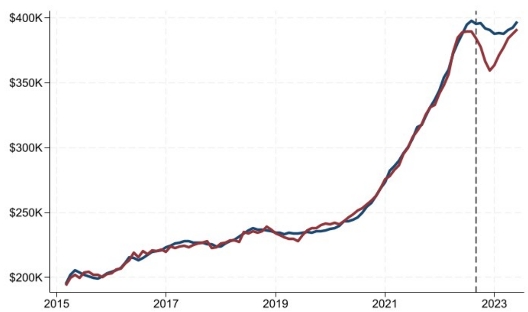

(d) Treatment 4

(e) Treatment 5

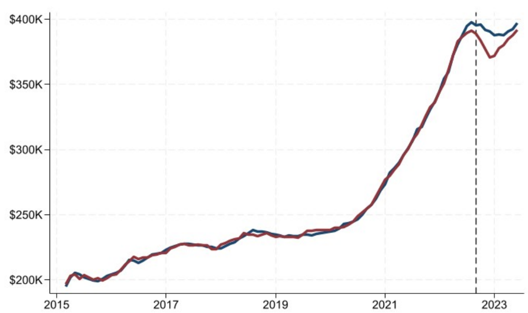

(f) Treatment 6

(g) Treatment 7

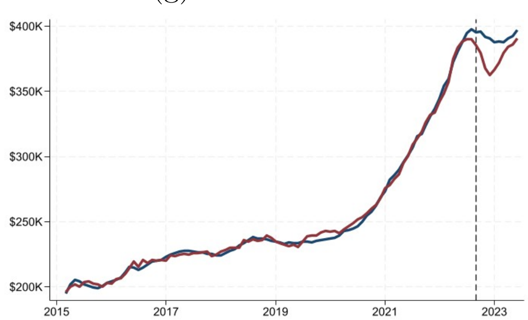

(h) Treatment 8

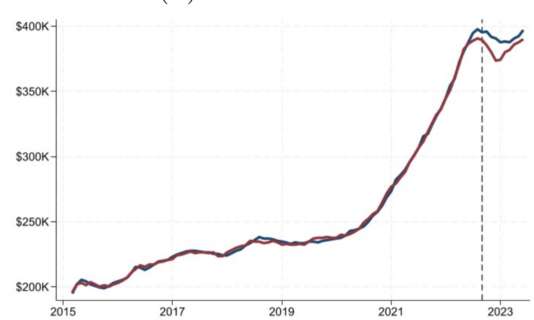

Notes: Based on author calculations. Vertical axis measures median close price in dollars, where navy line is treated unit (Lee County) and maroon line is its synthetic control. The vertical distance between the lines after treatment corresponds to the treatment effect τ_t in equation 6.

Table 9: SCM Effects for Two Price Treatments *across* the Same County

	Mean Price Effects								Median Price Effects							
	treat1	treat2	treat3	treat4	treat5	treat6	treat7	treat8	treat1	treat2	treat3	treat4	treat5	treat6	treat7	treat8
τ_0	8,582.86	4,136.80	7,988.49	7,953.71	7,742.96	8,738.88	9,018.33	4,036.98	7,318.74	4,287.57	11,376.58	10,926.90	6,460.79	11,932.29	9,953.69	5,820.12
	[0.350]	[0.523]	[0.370]	[0.413]	[0.298]	[0.364]	[0.395]	[0.536]	[0.233]	[0.349]	[0.110]	[0.113]	[0.143]	[0.091]	[0.123]	[0.167]
τ_1	15,359.79	10,515.32	12,162.64	12,267.46	16,031.85	13,420.86	14,181.28	9,204.08	13,812.25	9,027.58	18,074.46	16,788.06	12,288.31	18,102.01	16,322.15	10,541.43
	[0.183]	[0.244]	[0.329]	[0.325]	[0.107]	[0.299]	[0.284]	[0.321]	[0.067]	[0.093]	[0.041]	[0.05]	[0.083]	[0.052]	[0.074]	[0.083]
τ_2	21,402.56	13,572.41	20,746.99	22,706.36	18,854.82	23,396.78	23,049.44	13,412.27	21,307.97	10,139.98	24,869.36	25,209.96	14,316.20	26,600.18	23,914.97	11,386.11
	[0.250]	[0.198]	[0.151]	[0.113]	[0.083]	[0.143]	[0.111]	[0.226]	[0.067]	[0.163]	[0.014]	[0.038]	[0.060]	[0.039]	[0.062]	[0.107]
τ_3	20,730.75	13,076.94	23,593.76	25,981.73	18,324.46	25,773.46	24,512.07	13,403.89	25,692.47	15,465.63	30,966.04	29,413.85	19,435.04	30,022.27	28,147.51	16,832.31
	[0.133]	[0.233]	[0.178]	[0.138]	[0.131]	[0.130]	[0.148]	[0.226]	[0.067]	[0.058]	[0.000]	[0.013]	[0.024]	[0.026]	[0.037]	[0.060]
τ_4	12,409.23	10,655.77	16,929.37	19,205.90	16,225.85	18,525.51	17,142.14	10,924.17	19,002.68	12,170.85	24,074.84	22,362.61	15,684.79	21,924.86	20,992.56	13,282.17
	[0.333]	[0.279]	[0.219]	[0.163]	[0.155]	[0.208]	[0.210]	[0.298]	[0.150]	[0.163]	[0.041]	[0.088]	[0.095]	[0.130]	[0.123]	[0.155]
τ_5	14,959.49	10,161.74	18,234.29	18,509.01	14,793.44	18,941.25	17,907.67	10,075.82	15,845.83	7,216.61	17,184.54	19,120.37	10,622.75	18,689.48	16,365.56	8,199.25
	[0.267]	[0.337]	[0.219]	[0.225]	[0.190]	[0.247]	[0.259]	[0.345]	[0.183]	[0.349]	[0.123]	[0.125]	[0.226]	[0.143]	[0.185]	[0.321]
τ_6	11,826.95	5,838.84	16,686.11	16,557.29	10,015.47	16,555.09	14,460.31	6,083.57	10,638.55	5,094.70	10,385.93	10,951.15	7,616.02	9,692.92	8,267.45	5,673.22
	[0.367]	[0.593]	[0.233]	[0.275]	[0.381]	[0.312]	[0.333]	[0.536]	[0.400]	[0.547]	[0.356]	[0.338]	[0.357]	[0.377]	[0.444]	[0.476]
τ_7	13,063.51	6,383.42	15,284.80	13,372.58	8,008.76	13,718.99	10,314.54	6,927.26	9,706.07	4,267.76	6,391.74	10,956.76	5,664.50	8,527.43	6,572.22	4,580.69
	[0.317]	[0.628]	[0.260]	[0.300]	[0.500]	[0.286]	[0.531]	[0.524]	[0.450]	[0.581]	[0.548]	[0.375]	[0.500]	[0.468]	[0.556]	[0.560]
τ_8	13,512.40	6,951.47	16,366.39	16,963.12	5874.24	16,117.79	12,044.75	9,071.66	11,502.07	4,513.85	4,541.05	12,489.36	4,373.52	9,292.97	6,216.78	4,559.37
	[0.367]	[0.593]	[0.301]	[0.250]	[0.619]	[0.299]	[0.395]	[0.405]	[0.417]	[0.558]	[0.726]	[0.288]	[0.548]	[0.429]	[0.605]	[0.548]
τ_9	13,489.38	9,841.04	6,954.35	5,422.34	9,9190.29	5,022.53	-628.94	11,477.77	16,374.44	6,680.89	5,997.78	12,698.53	5,441.25	8,816.05	6,809.23	7,138.26
	[0.417]	[0.442]	[0.644]	[0.725]	[0.417]	[0.675]	[0.963]	[0.310]	[0.167]	[0.395]	[0.603]	[0.300]	[0.464]	[0.429]	[0.568]	[0.345]
#Donors	60	86	73	80	84	77	81	84	60	86	73	80	84	77	81	84

Notes: Both price treatment effects (in dollars) are measured for Lee County, FL. τ_t is the average treatment effect for the treated Lee County t months after treatment as defined in equation 6, where $t = 0$ corresponds to September 2022. Values in brackets are standardized p-values, which take into account the relative size of the treatment effect of the treated unit versus the placebos. See Galiani and Quistorff (2017) for more details. Note that p-values are given instead of more traditional standard errors due to ease of calculation from the placebo test; additionally, more care is needed due to the design-based inference (Abadie et al., 2020).

The ZIP Code SCM results are shown in Table 10 and the fit appears better than when the effects were measured across the entire county. The standardized p-values vary across time and treatment, but they are lower than their county counterparts. While not shown, the traditional p-values are all above 0.05, with most well above depending upon the treatment definition. Thus, the significant and positive estimates shown on the left side of Table 10 might have a different interpretation depending on whether a researcher places value on standardized versus traditional p-values. Regardless, the results indicate that a static treatment definition does matter, and there is mild evidence of dynamic treatment effects. However, more rigorous investigation would be useful to improve precision. Looking at the right side of the table for Miromar Lakes, the measured effects increase in similar periods as shown for Bonita Springs but the standardized p-value fit is not statistically significant.

More importantly, the ZIP Code results indicate the importance of spatial granularity. Even fixing a treatment aggregation, we see evidence of heterogeneous treatment effects for different ZIP Codes. This suggests dueling concerns when estimating price effects: SCM allows one to construct a suitable control where none exists, but at the potential cost of masking important spatial heterogeneity. Additionally, selection into the sample and aggregation bias become concerns when pooling together observations across larger geographic areas.

5.2 Are Synthetic Controls Good Enough?

The presented estimations reflect analyses that users of similar data might conduct. An immediate takeaway is that SCM results may be enhanced by altering the dependent variable measure, geographic scope, treatment definitions, imputation to improve balancing, and pre-trend matching. Each issue is addressed below.

The dependent variable measure might be improved in several ways. Besides testing with different functional forms, a more useful public measure would better describe damages. Examples could include generating a hedonic, a repeat sales, or another price index instead of relying on mean or median figures and would explicitly control for treatment at a property level. Adopting a more refined price measure seems promising for improving fit. However, a major challenge is the mandatory panel balancing. Instead of dropping so many observations, imputation or replacement of missing data are ripe for future exploration.

Table 10: Synthetic Control Effects for Two Treated ZIP Codes *within* the Same County

	Bonita Springs								Miromar Lakes							
	treat1	treat2	treat3	treat4	treat5	treat6	treat7	treat8	treat1	treat2	treat3	treat4	treat5	treat6	treat7	treat8
τ_0	18,775.62	10,834.15	13,385.48	9,394.08	8,887.29	2,791.91	8,873.20	8,887.29	1,864.29	5,106.46	1,842.05	7,249.26	6,863.92	4,822.89	6,863.92	6,230.18
	[0.263]	[0.373]	[0.567]	[0.456]	[0.436]	[0.817]	[0.463]	[0.443]	[0.791]	[0.476]	[0.841]	[0.325]	[0.335]	[0.493]	[0.336]	[0.374]
τ_1	25,616.44	15,767.03	-29,163.86	13,628.10	13,238.96	3,089.90	14,754.80	13,238.96	5,544.37	-11,121.27	10,478.85	-4,542.23	-4,562.39	-7,042.17	-4,562.39	-5,785.26
	[0.245]	[0.364]	[0.418]	[0.400]	[0.401]	[0.844]	[0.385]	[0.406]	[0.519]	[0.290]	[0.476]	[0.623]	[0.620]	[0.462]	[0.610]	[0.522]
τ_2	40,123.27	17,754.44	-39,924.32	28,051.71	22,107.69	18,951.26	29,517.29	22,107.69	14,273.15	-9,560.68	15,909.72	-6,122.94	-4,430.97	-8,711.65	-4,430.97	-4,087.06
	[0.203]	[0.401]	[0.408]	[0.228]	[0.320]	[0.390]	[0.217]	[0.333]	[0.248]	[0.435]	[0.408]	[0.596]	[0.713]	[0.461]	[0.701]	[0.726]
τ_3	24,517.14	11,832.01	-70,201.93	24,311.52	18,044.76	13,968.83	26,577.84	18,044.76	11,411.64	-12,615.13	21,364.70	-10,765.47	-9,897.41	-13,128.42	-9,897.41	-9,105.64
	[0.451]	[0.631]	[0.301]	[0.321]	[0.445]	[0.537]	[0.300]	[0.458]	[0.412]	[0.385]	[0.425]	[0.438]	[0.483]	[0.344]	[0.478]	[0.520]
τ_4	31,464.42	6,413.15	-77,706.60	17,944.25	11,987.20	4,740.06	17,058.96	11,987.20	15,201.29	-11,478.85	22,579.15	-11,848.20	-10,142.57	-14,149.56	-10,142.57	-7,460.78
	[0.388]	[0.801]	[0.334]	[0.473]	[0.655]	[0.856]	[0.506]	[0.643]	[0.349]	[0.467]	[0.446]	[0.426]	[0.520]	[0.375]	[0.498]	[0.622]
τ_5	46,372.02	25,201.89	-61,123.88	34,882.63	27,273.91	20,363.47	33,671.57	27,273.91	14,345.95	-8,843.24	34,665.35	-8,252.63	-6,299.32	-10,561.12	-6,299.32	-5,311.95
	[0.296]	[0.418]	[0.471]	[0.279]	[0.365]	[0.501]	[0.286]	[0.358]	[0.415]	[0.630]	[0.352]	[0.613]	[0.724]	[0.526]	[0.698]	[0.746]
τ_6	58,816.56	31,634.52	-51,318.79	39,960.93	34,993.65	24,150.04	37,983.09	34,993.65	10,248.28	-17,428.84	19,109.42	-15,896.23	-14,962.71	-18,234.13	-14,962.71	-12,437.32
	[0.233]	[0.350]	[0.542]	[0.260]	[0.300]	[0.438]	[0.268]	[0.292]	[0.540]	[0.372]	[0.560]	[0.387]	[0.424]	[0.327]	[0.408]	[0.480]
τ_7	27,134.87	-1,676.22	-47,376.63	5,783.72	2,230.01	-8,173.05	-1,804.93	2,230.01	885.32	-26,164.15	15,467.43	-28,943.01	-26,465.25	-31,209.56	-26,465.25	-22,661.71
	[0.525]	[0.958]	[0.585]	[0.824]	[0.945]	[0.783]	[0.952]	[0.938]	[0.958]	[0.273]	[0.643]	[0.225]	[0.250]	[0.188]	[0.253]	[0.286]
τ_8	12,977.40	-16,991.50	-23,700.96	-2,955.72	-21,072.89	-12,108.41	-13,635.79	-12,072.89	19,048.89	-823.77	29,007.99	-6,175.31	-3,753.10	-8,302.72	-3,753.10	-1,215.40
	[0.740]	[0.617]	[0.757]	[0.912]	[0.690]	[0.701]	[0.652]	[0.683]	[0.334]	[0.968]	[0.448]	[0.721]	[0.833]	[0.651]	[0.835]	[0.938]
τ_9	27,434.19	-15,958.16	19,502.02	-13,404.10	-18,705.28	-25,790.54	-26,929.49	-18,705.28	21,518.97	22,164.91	21,064.61	19,582.48	19,252.51	17,590.05	19,252.51	19,890.38
	[0.543]	[0.673]	[0.810]	[0.693]	[0.618]	[0.503]	[0.495]	[0.600]	[0.296]	[0.383]	[0.562]	[0.403]	[0.417]	[0.447]	[0.416]	[0.402]
#Donors	335	689	395	592	649	585	626	650	335	689	395	591	648	584	625	650

Notes: Both treated ZIP Codes lie within in Lee County, FL, which allows for studying what happens to synthetic control effects at increased granularity beyond what is shown in Table 9. Estimations are performed on the median price. τ_t is the average treatment effect for the treated Lee County t months after treatment as defined in equation 6, where $t = 0$ corresponds to September 2022. Values in brackets are standardized p-values, which take into account the relative size of the treatment effect of the treated unit versus the placebos. See Galiani and Quistorff (2017) for more details. Note that p-values are given instead of more traditional standard errors due to ease of calculation from the placebo test; additionally, more care is needed due to the design-based inference (Abadie et al., 2020).

Geographic scope and control location could also be improved with public data. The differences in geographic aggregations were evident as the discussion moved between county and ZIP Code results. Despite a reduced number of ZIP Codes after balancing, the quality of fit at the ZIP Code level sometimes surpassed the county level. This finding (of better fit and a more intense impact on market prices) should not seem surprising because damages are best measured with geographic precision. However, the indirect relationship (with lower geography but better fit) despite limited control groups suggests that synthetic controls are able to correctly pick weights and dynamically ignore sampled places that do not provide useful information. The two selected ZIP Codes experienced major and moderate damage—they might not be fully representative of damages across the entire county. This underscores the importance of carefully defining the sample market. Researchers might want to develop an algorithm for a more rules-based approach. Arbitrary choices can lead to match quality variations that will impact the estimated coefficients and p-values. As stressed previously, the trade-off between aggregation bias and goodness of fit should receive more attention.

Treatment definitions can lead to substantial differences even within a given dependent measure or geographic area. For instance, though subject to noise and used for illustrative purposes, in the month immediately after Hurricane Ian, altering the treatment definition in Lee County leads to varying effects ranging from 4.7% to 10.3%. Researchers need to either justify their chosen aggregate treatment definition or demonstrate the robustness of their findings across various definitions. All else equal, simpler treatment definitions tend to have price effects with clearer interpretations. We recommend using several aggregated measures of damages, along with sensitivity analysis.

Imputation is another area for improvement. Traditional statistical methods, such as multiple imputations, can be employed when some observations for the areas are available. As an example, Contat and Larson (2024) propose grouping into larger aggregates when not enough observations are observed in an area. For example, this would entail grouping counties (or any other geographical area) into "supercounties", which consist of two or more adjacent counties which individually lack enough observations but collectively have enough to perform estimations.[38] A similar approach might be used here to calculate damages and construct appropriate control samples. The SCM estimator and imputation is an under-

[38]The method relies on a substitutability and continuity assumption from urban economics, assuming that home values cannot change significantly on average when distance is changed in small increments. See Contat and Larson (2024) for more details on the algorithm.

studied area where more statistical insights would be constructive.

Pre-trend matching poses challenges for traditional treatment effects (static and dynamic) as well as synthetic control estimations.[39] While the standard recommendation is to avoid SCM if the pre-treatment fit is poor (Abadie, Diamond, and Hainmueller, 2015), recent papers have introduced methods to mitigate resulting bias. Ben-Michael, Feller, and Rothstein (2021) present the augmented synthetic control method, employing ridge regression to improve fit. However, this method may not always be preferable, as the bias reduction comes at the expense of some extrapolation. In a different approach, Ferman and Pinto (2021) propose a demeaned synthetic control estimator. They suggest that researchers should provide arguments for the absence of time-varying confounders. Additionally, they highlight that "contrasting the DID and the demeaned SCM estimators is informative about whether these conditions for unbiasedness are valid." Lastly, Arkhangelsky et al. (2021) combine DiD and SCM to reduce bias by using the synthetic control as the actual control in a DiD specification. This tactic combines the strengths of both estimators. A fuller understanding of when and why to choose these advanced methods would aid the developing literature on how to assess price effects after natural disasters.

6 Conclusion

This paper combines real estate listings with damage estimations. One contribution to the literature is showing how causal techniques can be used to estimate market disruptions. Current findings reveal positive price effects, approximately a 5% to 11% increase nine months after Hurricane Ian. Despite fluctuations over time, these effects seem relatively stable and enduring. However, given the previously mentioned limitations, caution is warranted with asserting strong causal estimates. Additionally, there are likely many means through which home prices are affected, and our results do not attempt to distinguish these different

[39] For practitioners or researchers learning these techniques, the appendix includes several robustness checks that illustrate whether the matching and results are sensitive to different study window lengths (allowing samples to fit with data from 1 year, 3 years, and 7 years prior to treatment), geographic scope (county or ZIP Code), and property type (dividing between single-family homes, condominiums, and manufactured homes). Shortening the pre-treatment windows improves matching because it avoids balancing issues with other competing treatments. Reducing geographic scope allows for better precision, but it comes at the expense of fewer sales to match. Focusing on a property subtype enhances measurement, but most sales are single-family homes and pooling helps condominiums and manufactured homes find more synthetic matches (and it is not clear that site built single-family homes are markedly different from manufactured homes).

channels.[40] While some previous literature has shown positive price effects associated with disasters, we do not have sufficient evidence to claim that Florida's costliest hurricane actually increased home prices for damaged homes. Therefore, while this research and some previous literature show positive price effects within a year following a disaster, the results should not be interpreted plainly as hurricanes increase house prices for damage homes, but rather different market mechanisms and rebuilding activities may have contributed to this result. Another contribution is pointing out the practical challenges when using various data sources. From the outset, the intent was to avoid proprietary sources and, instead, use data which are either commonly licensed by academics or publicly available.

Public data releases are a mixed blessing of equitable access at the cost of conflicting information standards. On one hand, such resources often serves as the sole source for assessing damages and pricing disaster risk, especially in the immediate aftermath of an event, making it extremely useful to many. On the other hand, challenges arise when connecting several datasets to address important issues, like understanding the impact of disasters on residential real estate prices.[41] Due to privacy concerns, address fields are expunged, which makes data aggregation necessary while creating a measurement issue. A potential solution would be combining data sources behind secure government servers and later providing the merged result in a de-identified form that could deliver a great public benefit. Currently, the absence of public property level damage data makes it difficult to assign treatments and accurately measure the extent and duration of market disruptions. We chose a hurricane for this paper's focus because there is readily available information on damages, claims, and market prices. More could be done to advance similar measures for other events.

For future research, it would be valuable to know more about proper methods for imputing and filling gaps in natural disaster or housing data. Additionally, enhancing techniques to aggregate damages or prices for the SCM approach and understanding the implications of matching before DiD could be reasonable steps forward. Furthermore, it would be a worth-

[40] As previously mentioned, one channel is that Ian decreased the supply of homes, driving up prices as homes became scarcer. Another possibility is that higher quality homes survived the storm and were able to be resold in the aftermath of Ian. We leave the identification of precise channels to future research.

[41] Sample selection complicates damage assessment across public datasets. For example, the IHP is designed for households without flood insurance or those underinsured, while NFIP's claims are for households with flood insurance. Public data could be enhanced if IHP could incorporate NFIP statistics or be completely matched to NFIP data because there are instances when claims may overlap. For instance, if a building exceeds the $250,000 limit on a NFIP policy, the owner may seek IHP assistance due to being underinsured.

while endeavor to investigate alternative options not yet considered in the applied disaster literature. Another potential avenue for research is the functional form of the dependent variable. As mentioned earlier, the pre-treatment fits of synthetic controls show variability based on whether prices are presented in dollars or log values or measured by mean or median statistics. In line with a related DiD literature, we find the parallel trends assumption does not always hold, which supports using newer SCM techniques. More work on the effects of disasters on other dependent variables would also be fruitful, including time on the market, probability of sale, and the ratio of list to sale price, among others. Finally, research could investigate differential price effects based on property types or real estate market activity. For example, unstructured textual analysis could be determine if properties are marketed differently in the wake of a disaster. In short, other fundamental questions remain.

In conclusion, this paper makes efforts to carefully combine data sources and employ cutting-edge empirical models to study whether a major hurricane may affect market prices. The analysis suggests that while individual houses may suffer considerable damages, aggregated measures are resilient to powerful natural disasters. With a rising frequency and severity of such storms, it is imperative to properly define how to monitor and measure the immediate and lasting effects on residential real estate markets. This study opens the door to further refinement of the complex interactions between natural disasters and real estate dynamics.

References

Abadie, Alberto. 2021. "Using Synthetic Controls: Feasibility, Data Requirements, and Methodological Aspects." *Journal of Economic Literature*, 59(2): 391–425.

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller. 2010. "Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program." *Journal of the American Statistical Association*, 105(490): 493–505.

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller. 2015. "Comparative Politics and the Synthetic Control Method." *American Journal of Political Science*, 59(2): 495–510.

Abadie, Alberto and Guido W. Imbens. 2011. "Bias-Corrected Matching Estimators for Average Treatment Effects." *Journal of Business & Economic Statistics*, 29(1): 1–11.

Abadie, Alberto and Javier Gardeazabal. 2003. "The Economic Costs of Conflict: A Case Study of the Basque Country." *American Economic Review*, 93(1): 113–132.

Abadie, Alberto and Jérémy L'Hour. 2021. "A Penalized Synthetic Control Estimator for Disaggregated Data." *Journal of the American Statistical Association*, 116(536): 1817–1834.

Abadie, Alberto, Susan Athey, Guido W. Imbens, and Jeffrey M. Wooldridge. 2020. "Sampling-Based Versus Design-Based Uncertainty in Regression Analysis." *Econometrica*, 88(1): 265–296.

Angrist, Joshua D. and Jörn-Steffen Pischke. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion.* Princeton University Press.

Arkhangelsky, Dmitry, Susan Athey, David A. Hirshberg, Guido W. Imbens, and Stefan Wager. 2021. "Synthetic Difference-in-Differences." *American Economic Review*, 111(12): 4088–4118.

Atreya, Ajita and Susana Ferreira. 2015. "Seeing is Believing? Evidence from Property Prices in Inundated Areas." *Risk Analysis*, 35(5): 828–848.

Ben-Michael, Eli, Avi Feller, and Jesse Rothstein. 2021. "The Augmented Synthetic Control Method." *Journal of the American Statistical Association*, 116(536): 1789–1803.

Bin, Okmyung and Craig E. Landry. 2013. "Changes in Implicit Flood Risk Premiums: Empirical Evidence from the Housing Market." *Journal of Environmental Economics and management*, 65(3): 361–376.

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro HC Sant'Anna. 2024. "Difference-in-Differences with a Continuous Treatment." National Bureau of Economic Research Working Paper Series 32117.

Contat, Justin and William D. Larson. 2024. "A Flexible Method of House Price Index Construction using Repeat-Sales Aggregates." *Real Estate Economics*, forthcoming.

Contat, Justin, Carrie Hopkins, Luis Mejia, and Matthew Suandi. 2024. "When Climate Meets Real Estate: A Survey of the Literature." *Real Estate Economics*, forthcoming.

Daw, Jamie R. and Laura A. Hatfield. 2018. "Matching and Regression to the Mean in Difference-in-Differences Analysis." *Health Services Research*, 53(6): 4138–4156.

De Chaisemartin, Clément and Xavier d'Haultfoeuille. 2023. "Two-Way Fixed Effects and Differences-In-Differences with Heterogeneous Treatment Effects: A Survey." *The Econometrics Journal*, 26(3): C1–C30.

De Chaisemartin, Clément and Xavier d'Haultfoeuille. 2018. "Fuzzy Differences-in-Differences." *The Review of Economic Studies*, 85(2): 999–1028.

Denteh, Augustine and Désiré Kédagni. 2022. "Misclassification in Difference-in-Differences Models." arXiv Working Paper 2207.11890.

Donald, Stephen G. and Kevin Lang. 2007. "Inference with Difference-in-Differences and Other Panel Data." *The Review of Economics and Statistics*, 89(2): 221–233.

Doudchenko, Nikolay and Guido W. Imbens. 2016. "Balancing, Regression, Difference-in-Differences and Synthetic Control Methods: A Synthesis." National Bureau of Economic Research Working Paper Series 22791.

Dumm, Randy, G. Stacy Sirmans, and Greg Smersh. 2012. "Building Code, Wind Contours, and House Prices." *Journal of Real Estate Research*, 34(1): 73–98.

Fang, Lu, Lingxiao Li, and Abdullah Yavas. 2023. "The Impact of Distant Hurricane on Local Housing Markets." *The Journal of Real Estate Finance and Economics*, 66: 327–372.

Ferman, Bruno and Cristine Pinto. 2021. "Synthetic Controls with Imperfect Pretreatment Fit." *Quantitative Economics*, 12(4): 1197–1221.

Ferman, Bruno, Cristine Pinto, and Vitor Possebom. 2020. "Cherry Picking with Synthetic Controls." *Journal of Policy Analysis and Management*, 39(2): 510–532.

Fisher, Jeffrey D. and Sara R. Rutledge. 2021. "The Impact of Hurricanes on the Value of Commercial Real Estate." *Business Economics*, 56(3): 129–145.

Galiani, Sebastian and Brian Quistorff. 2017. "The synth_runner Package: Utilities to Automate Synthetic Control Estimation Using Synth." *The Stata Journal*, 17(4): 834–849.

Gibson, Matthew and Jamie T. Mullins. 2020. "Climate Risk and Beliefs in New York Floodplains." *Journal of the Association of Environmental and Resource Economists*, 7(6): 1069–1111.

Gould Ellen, Ingrid and Rachel Meltzer. 2024. "Heterogeneity in the Recovery of Local Real Estate Markets after Extreme Events: The Case of Hurricane Sandy." *Real Estate Economics*, forthcoming.

Hallstrom, Daniel G. and V. Kerry Smith. 2005. "Market Responses to Hurricanes." *Journal of Environmental Economics and Management*, 50(3): 541–561.

Ham, Dae Woong and Luke Miratrix. 2022. "Benefits and Costs of Matching Prior to a Difference in Differences Analysis When Parallel Trends Does Not Hold." arXiv Working Paper 2205.08644.

Harrison, David, Greg T. Smersh, and Arthur Schwartz. 2001. "Environmental Determinants of Housing Prices: The Impact of Flood Zone Status." *Journal of Real Estate Research*, 21(1-2): 3–20.

Hennighausen, Hannah and Jordan F. Suter. 2020. "Flood Risk Perception in the Housing Market and the Impact of a Major Flood Event." *Land Economics*, 96(3): 366–383.

Hino, Miyuki and Marshall Burke. 2021. "The Effect of Information About Climate Risk on Property Values." *Proceedings of the National Academy of Sciences*, 118(17): e2003374118.

Ho, Anson T.Y., Kim P. Huynh, David T. Jacho-Chávez, and Geneviève Vallée. 2023. "We Didn't Start the Fire: Effects of a Natural Disaster on Consumers' Financial Distress." *Journal of Environmental Economics and Management*, 119: 102790.

Imbens, Guido W. and Donald B. Rubin. 2015. *Causal Inference for Statistics, Social, and Biomedical Sciences.* New York:Cambridge University Press.

Kaul, Ashok, Stefan Klößner, Gregor Pfeifer, and Manuel Schieler. 2015. "Synthetic Control Methods: Never Use All Pre-intervention Outcomes Together with Covariates." Munich Personal RePEc Archive Working Paper 83790.

Keys, Benjamin J. and Philip Mulder. 2020. "Neglected No More: Housing Markets, Mortgage Lending, and Sea Level Rise." National Bureau of Economic Research Working Paper Series 27930.

Kim, Hyejin and Jongkwan Lee. 2023. "Natural Disasters, Risk and Migration: Evidence from the 2017 Pohang Earthquake in Korea." *Journal of Economic Geography*, 23(5).

Kousky, Carolyn, Mark Palim, and Ying Pan. 2020. "Flood Damage and Mortgage Credit Risk: A Case Study of Hurricane Harvey." *Journal of Housing Research*, 29(sup1): S86–S120.

Marcus, Michelle and Pedro HC Sant'Anna. 2021. "The Role of Parallel Trends in Event Study Settings: An Application to Environmental Economics." *Journal of the Association of Environmental and Resource Economists*, 8(2): 235–275.

Miller, Douglas L. 2023. "An Introductory Guide to Event Study Models." *Journal of Economic Perspectives*, 37(2): 203–230.

Muller, Nicholas Z. and Caroline A. Hopkins. 2019. "Hurricane Katrina floods New Jersey: The Role of Information in the Market Response to Flood Risk." National Bureau of Economic Research Working Paper Series 25984.

Negi, Akanksha and Digvijay Singh Negi. 2022. "Difference-in-Differences with a Misclassified Treatment." arXiv Working Paper 2208.02412.

Rambachan, Ashesh and Jonathan Roth. 2023. "A More Credible Approach to Parallel Trends." *Review of Economic Studies*, 90(5): 2555–2591.

Roth, Jonathan. 2022. "Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends." *American Economic Review: Insights*, 4(3): 305–322.

Roth, Jonathan, Pedro HC Sant'Anna, Alyssa Bilinski, and John Poe. 2023. "What's Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature." *Journal of Econometrics*, 235(2): 2218–2244.

Rubin, Donald B. 1974. "Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies." *Journal of Educational Psychology*, 66(5): 688.

Salter, Sean and Ernest King. 2009. "Price Adjustment and Liquidity in a Residential Real Estate Market with an Accelerated Information Cascade." *Journal of Real Estate Research*, 31(4): 421–454.

Schennach, Susanne M. 2016. "Recent Advances in the Measurement Error Literature." *Annual Review of Economics*, 8: 341–377.

Strobl, Eric. 2011. "The Economic Growth Impact of Hurricanes: Evidence from US Coastal Counties." *Review of Economics and Statistics*, 93(2): 575–589.

Sun, Liyang and Jesse M. Shapiro. 2022. "A Linear Panel Model with Heterogeneous Coefficients and Variation in Exposure." *Journal of Economic Perspectives*, 36(4): 193–204.

Sun, Liyang and Sarah Abraham. 2021. "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects." *Journal of Econometrics*, 225(2): 175–199.

Turnbull, Geoffrey K., Velma Zahirovic-Herbert, and Chris Mothorpe. 2013. "Flooding and Liquidity on the Bayou: The Capitalization of Flood Risk Into House Value and Ease-Of-Sale." *Real Estate Economics*, 41(1): 103–129.

Yi, Donggyu and Hyundo Choi. 2020. "Housing Market Response to New Flood Risk Information and the Impact on Poor Tenant." *The Journal of Real Estate Finance and Economics*, 61: 55–79.

Zhang, Lei. 2016. "Flood Hazards Impact on Neighborhood House Prices: A Spatial Quantile Regression Analysis." *Regional Science and Urban Economics*, 60: 12–19.

Zivin, Joshua Graff, Yanjun Liao, and Yann Panassie. 2023. "How Hurricanes Sweep up Housing Markets: Evidence from Florida." *Journal of Environmental Economics and Management*, 118: 102770.

Appendix: Robustness Checks

Real estate activity can be conveyed with other statistics than described earlier in the paper. Figure 7 shows a common metric called the list-to-close price ratio. The market is usually viewed favorably for buyers when the ratio exceeds one. In the state and county panels, the list-to-close ratios drop to nearly one after the beginning of the COVID-19 pandemic when limited sales inventory shifts market power more to sellers. The median list-to-close price ratio is exactly equal to one and conveys that half of the listings sell precisely for their list price. Although not displayed here, townhomes even dip below one in Lee County which is consistent with anecdotal stories of escalating bidding wars of transaction prices rising above initial list prices and having a short time-on-market. The quick jump up after Hurricane Ian suggests that sellers begin to lower their initial price expectations after the storm as the market returns back to its pre-pandemic trends.

Another comparison could split real estate listings into property subtypes like detached single-family homes, townhomes, and manufactured homes. Figure 8 compares subtypes by prices and the number of listings. Price levels remain consistent across subtypes, with the greatest values in detached single-family residences then moving downward to townhomes and then manufactured homes, as noted in the top graphic of panel (a). Breaking down by subtypes indicates that Lee County's housing market is dominated by detached single-family residences with solid navy and maroon lines at the upper part of panel (b). Across years, there are fewer listings for townhomes and manufactured homes, which complicates further analysis as shown later with blank panels in Figure 12.

Several additional tests are performed to determine sensitivities to geographic scope (i.e. county or ZIP Code) as well as the starting year (i.e. 2015, 2019, or 2021) for the DiD and SCM techniques. All figures are based on median closing price and treatment 3.

Figures 9 and 10 present DiD results from event studies. The longer pre-treatment window length introduces volatility that clearly violates pre-trend assumptions. However, even when starting shortly before the storm in 2021, the disaster only has positive price effects for single-family homes while condos and manufactured homes are statistically insignificant.

Figures 11 and 12 offer SCM results to compare with DiD findings. The pre-treatment window length has a similar conclusion that shorter periods offer better fit (i.e. the colored lines are closer together when comparing 2015 to 2019 to 2021 across all property types). This is not entirely surprising because other storms and influences can muddle interpretations. The positive effect diminishes for detached single-family homes and manufactured homes have a negative impact when SCM is performed at a county level. The final figure suggests caution is needed when interpreting these results. SCM cannot be performed successfully at a ZIP Code level for most property subtypes because a lack of sales prevents weighted matchings. Both ZIP Codes can perform SCM for single-family homes. Bonita Springs shows a negative effect while Miromar Lakes has a positive effect on prices after Hurricane Ian. Recall that Bonita Springs is coastal-adjacent and experienced large aggregate damages that were, on average, twice the size as in Miromar Lakes. The more inland area of Miromar Lakes likely carries a positive price effect for both single-family homes and condos because of the lack of claims, lower damages on the claims, and the ability to sell in that area drove up demand as buyers sought locations that seemed slightly safer immediately after the storm.

Figure 7: Real Estate Market Activity Comparing List-to-Close Prices

(a) All listings across Florida

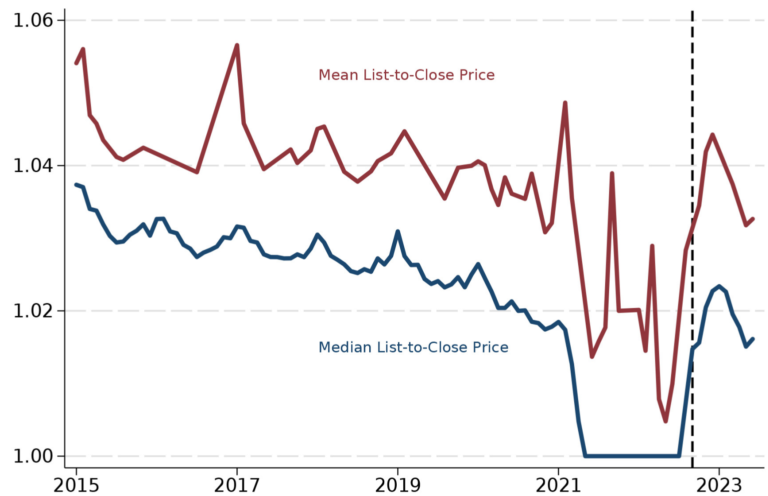

(b) All listings in Lee County, Florida

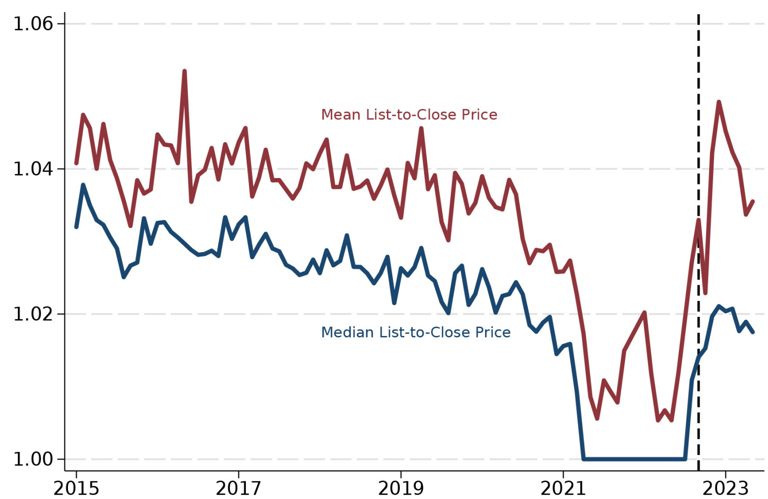

Notes: Based on author calculations using MLS data licensed from CoreLogic. The top row presents real estate activity for single-family residential homes across the entire state of Florida (n = 1.5 million) while the bottom row contains listings information from Lee County, Florida ($n = 73,000$) since 2015. Excluded property types are low/mid/high rise, half duplex, and villas. Colored lines illustrate the median ratio (navy) and mean ratio (maroon). All series reflect a monthly time frequency without any adjustments for distributional moments, rolling windows, or seasonality. Certain periods are omitted when there are low counts (e.g., filters for fewer than 10 listings or closings prevent single listings from swaying the results).

Figure 8: Real Estate Market Activity by Property Type within Lee County, Florida

(a) Prices

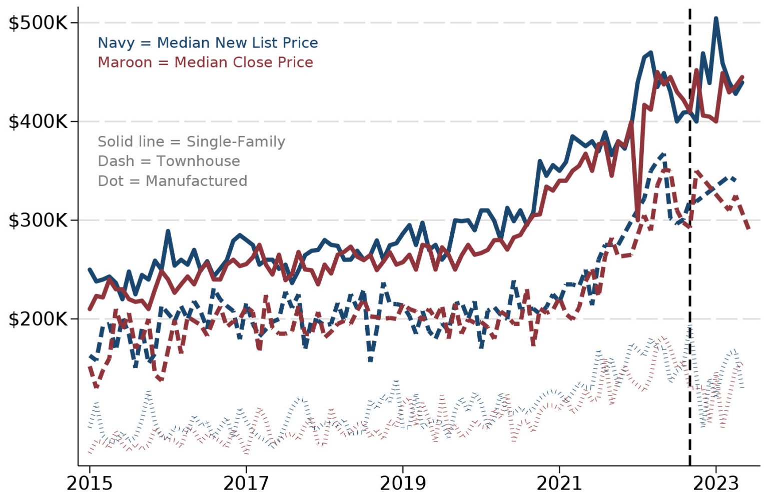

(b) Number of Listings

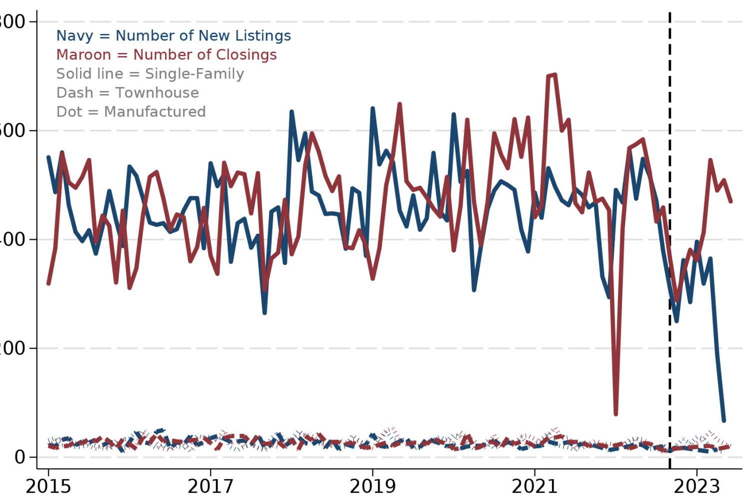

Notes: Based on author calculations using MLS data licensed from CoreLogic for Lee County, Florida ($n = 73,000$) since 2015. Excluded property types are low/mid/high rise, half duplex, and villas. The panels break the data into property subtypes where solid lines are single-family detached, dashed lines are townhomes, and dotted lines that appear slightly grayish are manufactured homes. All series reflect a monthly time frequency without any adjustments for distributional moments, rolling windows, or seasonality. Certain periods are omitted when there are low counts (filters for fewer than 10 listings or closings prevent single listings from swaying the results).

Figure 9: Event Study Graphs for Lee County by Property Type with Different Starting Years

Single-Family Homes

(a) 2015

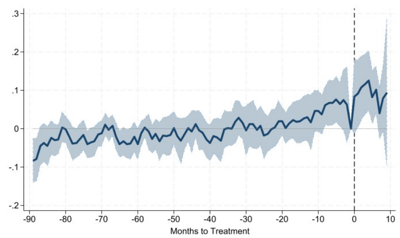

(b) 2019

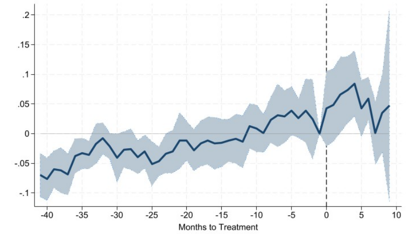

(c) 2021

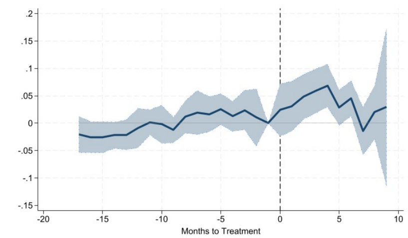

Condominiums

(d) 2015

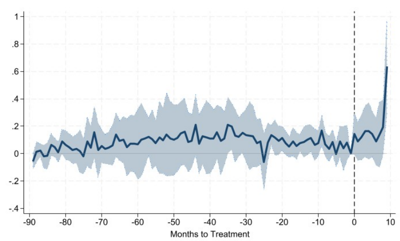

(e) 2019

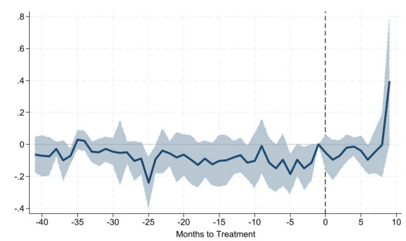

(f) 2021

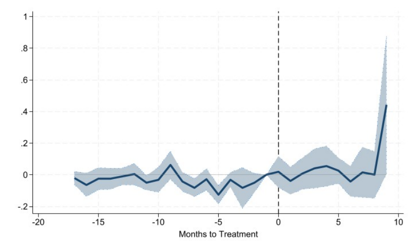

Manufactured Homes

(g) 2015

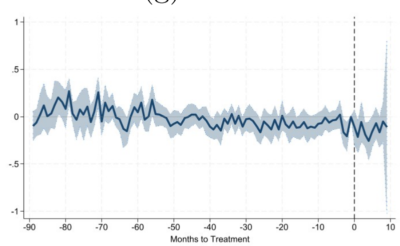

(h) 2019

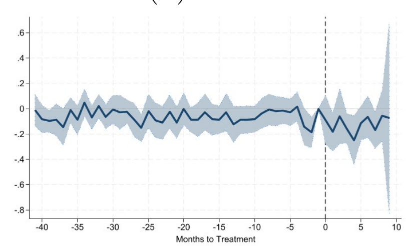

(i) 2021

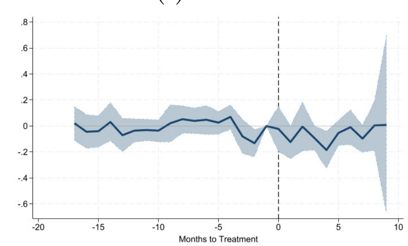

Notes: Based on author calculations. Vertical axis is the estimated treatment effect coefficient β_t from equation 3, where treatment is normalized to occur at $t = 0$ and β_{-1} is normalized to be zero. Estimated at monthly time frequency, where the geographic aggregation is at the county level. All figures are based on treatment 3.

Figure 10: Event Study Graphs for ZIP Codes by Property Type with Different Starting Years

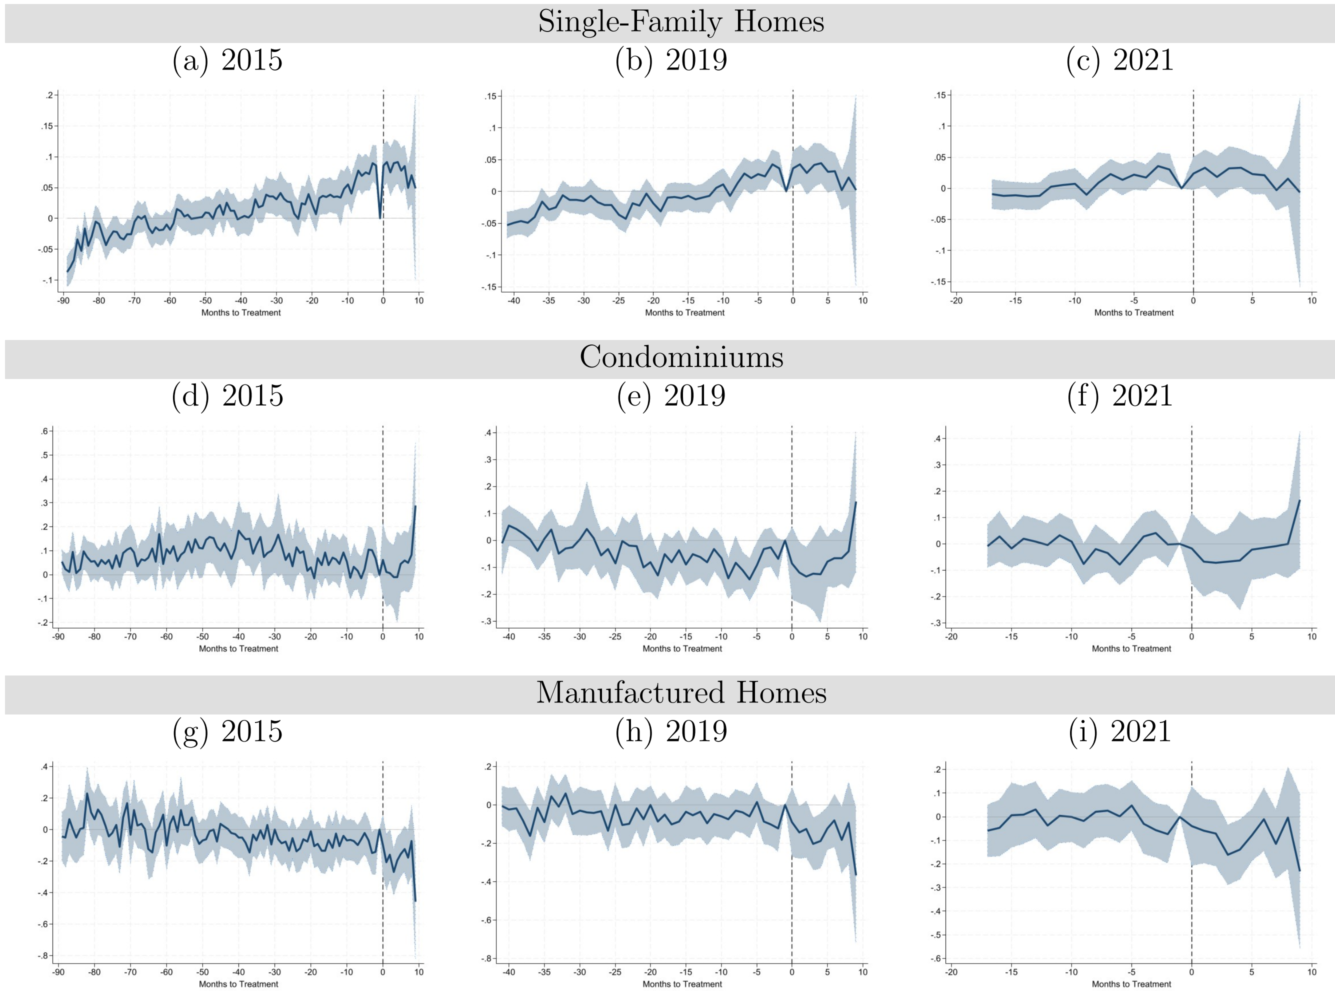

Notes: Based on author calculations. Vertical axis is the estimated treatment effect coefficient β_t from equation 3, where treatment is normalized to occur at $t = 0$ and β_{-1} is normalized to be zero. Estimated at monthly time frequency, where the geographic aggregation is at the ZIP Code level. All figures are based on treatment 3.

Figure 11: SCM Estimates for Lee County by Property Type with Different Starting Years

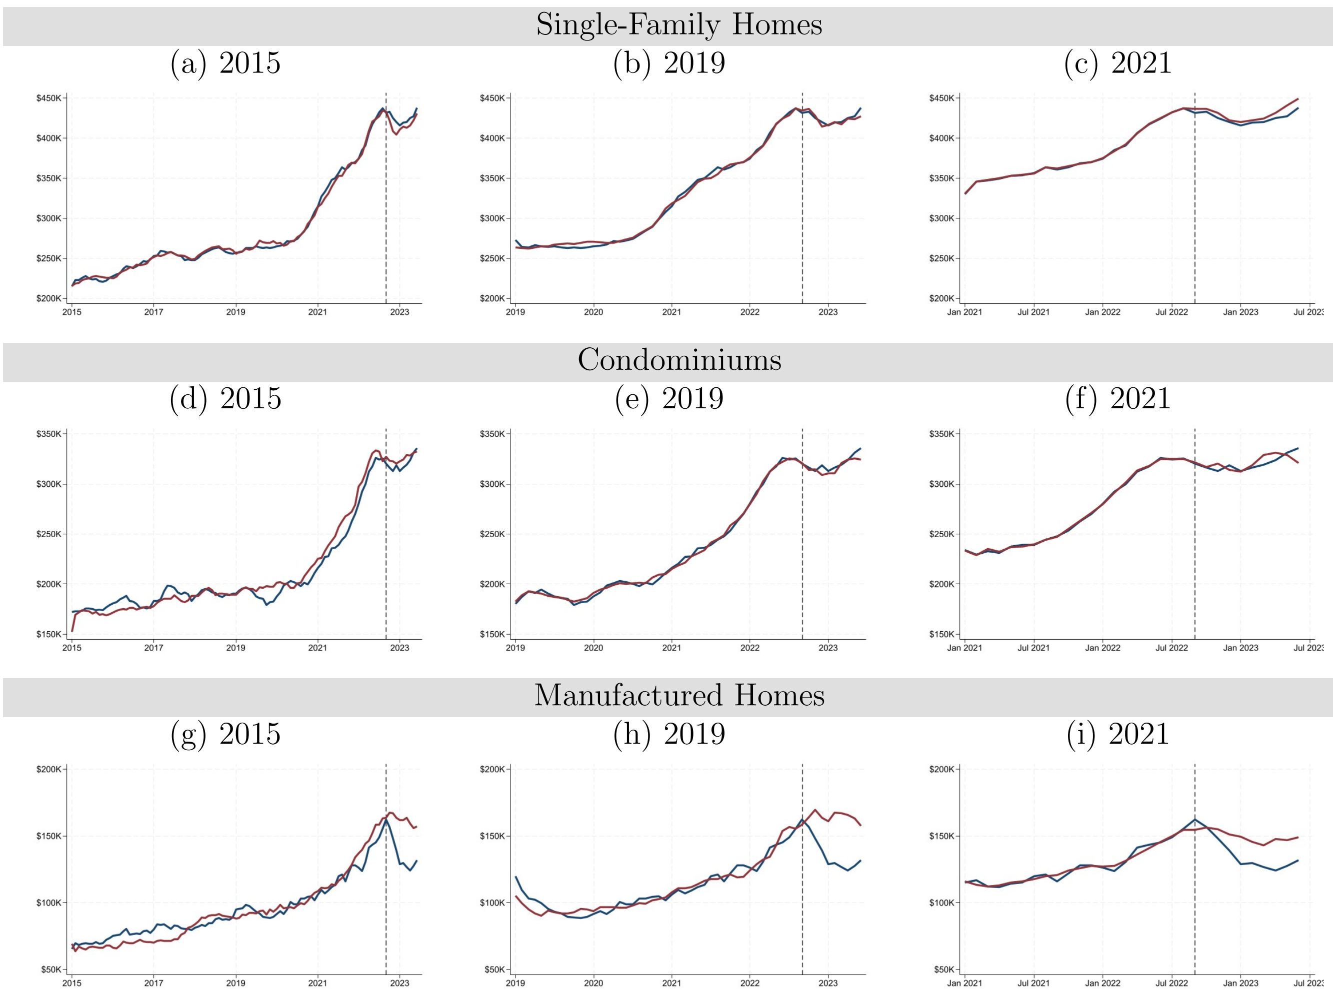

Notes: Based on author calculations. Vertical axis measures median close price in dollars, where navy line is treated unit (Lee County) and maroon line is its synthetic control. The vertical distance between the lines after treatment corresponds to the treatment effect τ_t in equation 6.

Figure 12: SCM Estimates for ZIP Codes by Property Type Starting in 2019

Single-Family Homes

(a) Bonita Springs

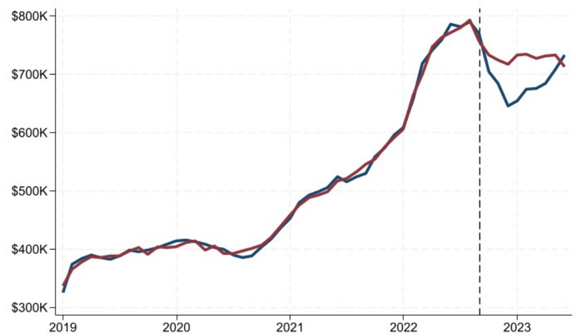

(b) Miromar Lakes

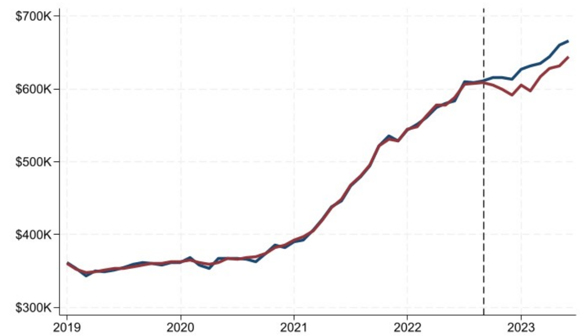

Condominiums

(c) Bonita Springs

Insufficient data.
Unable to display graph.

(d) Miromar Lakes

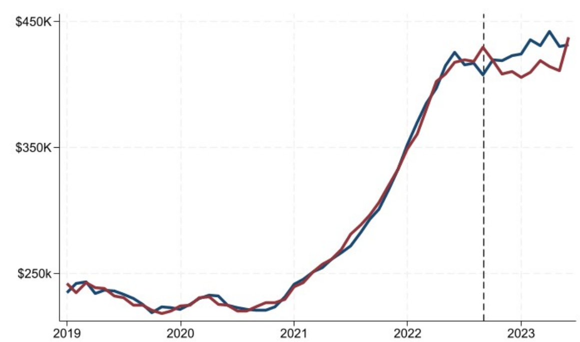

Manufactured Homes

(e) Bonita Springs

Insufficient data.
Unable to display graph.

(f) Miromar Lakes

Insufficient data.
Unable to display graph.

Notes: Based on author calculations. Vertical axis measures median close price in dollars, where navy line is treated unit (Lee County) and maroon line is its synthetic control. The vertical distance between the lines after treatment corresponds to the treatment effect τ_t in equation 6.

Table 11: Comparing Housing Characteristic between Control and Treated Groups

	treat1		treat2		treat3		treat4	
	Control	Treatment	Control	Treatment	Control	Treatment	Control	Treatment
beds	3.1	3.0	-	-	3.0	3.1	3.1	3.1
	(0.9)	(1.0)	-	-	(1.0)	(0.9)	(1.0)	(0.9)
baths	2.3	2.4	-	-	2.3	2.4	2.3	2.4
	(1.2)	(2.3)	-	-	(1.5)	(2.3)	(1.9)	(0.9)
home size (sqft)	1,848	1,907	-	-	1,886	1,848	1,885	1,829
	(1,902)	(39,371)	-	-	(31,842)	(981)	(29,579)	(1,032)
home age (years)	27	26	-	-	27	25	27	25
	(23)	(21)	-	-	(22)	(21)	(22)	(21)
cooling (yes/no)	99.5%	99.7%	-	-	99.6%	99.7%	99.6%	99.7%
	(6.9%)	(5.6%)	-	-	(6.5%)	(5.8%)	(6.4%)	(5.8%)
time-on-market (days)	118	119	-	-	120	114	119	118
	(155)	(160)	-	-	(158)	(157)	(157)	(159)
price	12.35	12.42	-	-	12.36	12.45	12.37	12.46
	(0.83)	(0.74)	-	-	(0.83)	(0.65)	(0.81)	(0.68)
N	997,466	858,057	-	-	1,350,102	505,421	1,547,469	308,054

	treat5		treat6		treat7		treat8	
	Control	Treatment	Control	Treatment	Control	Treatment	Control	Treatment
beds	3.1	2.9	3.1	2.9	3.1	2.9	3.1	2.9
	(1.0)	(0.8)	(1.0)	(0.8)	(1.0)	(0.8)	(1.0)	(0.7)
baths	2.3	2.3	2.3	2.3	2.3	2.3	2.3	2.3
	(1.8)	(0.7)	(1.9)	(0.8)	(1.8)	(0.8)	(1.8)	(0.7)
house size (sqft)	1,881	1,790	1,885	1,801	1,884	1,793	1,878	1,801
	(27,888)	(1,067)	(28,772)	(971)	(28,406)	(996)	(27,561)	(755)
home age (years)	27	24	27	26	27	25	27	21
	(22)	(20)	(22)	(20)	(22)	(20)	(22)	(19)
cooling (yes/no)	99.6%	99.9%	99.6%	99.7%	99.6%	99.7%	99.6%	99.8%
	(6.4%)	(3.7%)	(6.4%)	(5.3%)	(6.4%)	(5.2%)	(6.4%)	(4.2%)
time-on-market (days)	118	127	118	122	118	128	118	130
	(157)	(163)	(157)	(161)	(157)	(163)	(157)	(161)
sale price (logged)	12.37	12.55	12.37	12.48	12.37	12.53	12.38	12.46
	(0.79)	(0.68)	(0.80)	(0.71)	(0.79)	(0.71)	(0.79)	(0.68)
N	1,739,643	115,880	1,629,790	225,733	1,669,642	185,881	1,778,117	77,406

Notes: Based on author calculations using MLS data licensed from CoreLogic. This table compares summary statistic values for treatment and control groups, across all treatments except treat2 (which is not a binary measure). Variables are selected based on being used later as controls for difference-in-difference estimations. Each variable presents mean values with standard deviations below in parentheses. Beds and baths are the number of units reported in the home. Home size is the structure's total livable space in square feet (sqft). The control standard deviations are large, but acceptable because values at the 99th percentiles do not exceed 5,000 sqft. Home age is defined as the difference between the closing date and year built. Cooling indicates (yes/no) if a cooling system is present. Time-on-market is the number of days from list date to contract date. Sale price (logged) is the natural log of the close price (nominal dollars) of the home.

Figure 13: Median Sale Price Time Series Plot: Treatment vs. Control

(a) Treatment 1

(b) Treatment 2 (Not Applicable)[a]

Treatment is non-binary.
Unable to display graph.

(c) Treatment 3

(d) Treatment 4

(e) Treatment 5

(f) Treatment 6

(g) Treatment 7

(h) Treatment 8

Notes: Based on author calculations from MLS data licensed from CoreLogic. Vertical axis measures median close price in dollars, where the navy line is the median for the corresponded units in treated counties and the maroon line corresponds to the control group. The dashed line corresponds to September 2022, which is the month that Hurricane Ian made landfall in Florida. Treatment 2's graph is omitted because it is not a binary approach.