

Finance and Economics Discussion Series

Federal Reserve Board, Washington, D.C.

ISSN 1936-2854 (Print)

ISSN 2767-3898 (Online)

Flood Risk Mapping and the Distributional Impacts of Climate Information

Joakim A. Weill

2023-066

Please cite this paper as:

Weill, Joakim A. (2023). "Flood Risk Mapping and the Distributional Impacts of Climate Information," Finance and Economics Discussion Series 2023-066. Washington: Board of Governors of the Federal Reserve System, <https://doi.org/10.17016/FEDS.2023.066>.

NOTE: Staff working papers in the Finance and Economics Discussion Series (FEDS) are preliminary materials circulated to stimulate discussion and critical comment. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors. References in publications to the Finance and Economics Discussion Series (other than acknowledgement) should be cleared with the author(s) to protect the tentative character of these papers.

Flood Risk Mapping and the Distributional Impacts of Climate Information

Joakim A. Weill*

September 2023

Abstract

This paper examines the provision of official flood risk information in the United States and its distributional impacts on residential flood insurance take-up. Assembling all flood maps produced after Hurricane Katrina, I document that updated maps decreased the number of properties zoned in high-risk floodplains and incorrectly omitted five million properties, primarily in neighborhoods with more Black and Hispanic residents. Leveraging the staggered timing of map updates, I estimate they decreased flood insurance take-up and exacerbated racial disparities in insurance coverage. Correcting flood maps could increase welfare by \$20 billion annually, but past map updates distorted risk and price signals.

*Federal Reserve Board of Governors, joakim.a.weill@frb.gov. I thank the First Street Foundation, FEMA's disclosure office and Zillow for data access, as well as Marianne Bitler, Katrina Jessoe, Pierre Mérel, Ashish Shenoy and Michael Springborn for their continued support. I also thank Mark Agerton, Charlotte Ambrozek, Briana Ballis, Joshua Blonz, Celso Brunetti, Colin Cameron, Ben Collier, Miki Doan, John Driscoll, Caitlin Hesser, Caitlin Kieran, Paul Lengermann, Yanjun (Penny) Liao (discussant), Maureen Page, Emmanuel Paroissien, Francesc Ortega, Amine Ouazad, Ivan Petkov, Adam Pollack, Jeremy Porter, Jim Sanchirico, Christopher Sampson, Christopher Timmins, Takuya Ura, Katherine Wagner (discussant), Jie Yang, as well as participants at the UCLA Climate Symposium 2020, WEAI 2021, EEA 2021, the 2021 CU Environmental & Resource Economics Workshop, UEA 2022, APPAM 2022, ASSA 2023, FAERE 2023, the Wharton Risk Center, AgroParisTech-ALISS, University of Georgia, University of Alaska-Anchorage, Southern Methodist University, the Federal Reserve Board, University of Wisconsin-Madison, University of British Columbia, Harvey Mudd College, Colby College, CUNY Queens College, Freddie Mac Climate Seminar Series, FEMA Partners series, Treasury Department Office of the Comptroller of the Currency, the UCD NatuRE Policy Lab and UCD Econometrics Group for thoughtful comments. Financial support from the John Muir Institute of the Environment and the Center for Data Science and Artificial Intelligence is acknowledged. I conducted this research prior to employment at the Federal Reserve Board, while a PhD candidate at UC Davis. Data sources were obtained under purview of UC Davis licenses. The views expressed in this paper are solely my responsibility and should not be interpreted as reflecting the opinions of the Federal Reserve Board or of any other person associated with the Federal Reserve System.

1 Introduction

Climate change is increasing the frequency and impacts of extreme events, with natural disasters causing more than a trillion dollars of economic damages in the United States over the last decade. Adaptation strategies and insurance products are often available to mitigate these costs, but the lack of information about risks can inhibit climate adaptation. To address this challenge, most industrialized countries and a growing number of developing nations produce and disseminate climate risk information. Yet, whether public information provision adequately promotes climate adaptation is a priori unclear. Households and firms may be inattentive to public information; official information might not reflect the current state of scientific knowledge; and disparities in adaptive responses could raise equity concerns.

Flooding is the costliest disaster in the United States. To understand and manage flood risks, the federal government produces flood insurance rate maps, commonly called “flood maps,” that describe flood risk exposure for most properties in the country. This information also conveys regulatory conditions crucial to the functioning of the National Flood Insurance Program (NFIP), which sells around 95% of all residential flood insurance policies (Kousky, 2016; Kousky et al., 2018). Updating these maps has become one of the most expensive risk mapping efforts worldwide, but their ability to promote flood preparedness is unknown.

In this paper, I investigate how the provision of official flood risk information evolved over the past two decades and estimate its impacts on insurance take-up. I leverage the digitization of flood maps and changes to floodplain boundaries to study the role of floodplain classification and digital information access on insurance take-up. I highlight the distributional consequences of the maps, which partially explain growing neighborhood disparities in insurance coverage: between 2008 and 2018, average take-up rates went from 3.9 to 3.6 percent in neighborhoods in the top quartile of White residents, while coverage in neighborhoods with the highest share of minority residents went from 2.7 to 2.3 percent.¹

To study risk information provision and its impacts on household behavior, I consolidate the most comprehensive set of files ever assembled on flood risk by collecting all flood maps updated by the Federal Emergency Management Agency (FEMA) over a 15-year roll-out and linking them to geolocalized data on the entire residential housing stock in the contiguous US. Obtaining and compiling these data involved four separate Freedom of Information Act requests and multiple meetings with current and former government officials. Using this novel dataset, I can observe the evolution of the risk information provided to households, in particular through the rezoning of properties inside and outside of the *100-year floodplain*, or “high-risk

¹Author’s calculations, using the Decennial Census measure of “White alone” households in a neighborhood, with insurance take-up rates averaged over neighborhoods.

zone,” detailed below. I also observe discrete changes in the costs of accessing pre-existing information, through the digitization of previous floodplain boundaries. To assess the scientific accuracy of the official risk information, I compare flood maps with the First Street Foundation flood model, a recently developed, peer-reviewed, and independent alternative to official maps regarded as providing more complete risk estimates. Finally, I combine these risk estimates with a novel set of administrative records from FEMA on 60 million flood insurance policies.

I measure the effects of changes in the provision of information using the staggered roll-out of digital flood maps. While most neighborhoods are covered by a digital map in 2020, due to public funding constraints, the timing of the map updates differed between neighborhoods. This variation across space and time allows me to leverage heterogeneity-robust difference-in-differences estimators to recover the impact of map updates on the residential insurance take-up (Callaway and Sant’Anna, 2021; Sun and Abraham, 2021). To flexibly estimate the entire distribution of treatment effects and assess the robustness of the difference-in-difference estimates, I adapt recently developed synthetic control estimators (Abadie and L’Hour, 2021; Ben-Michael et al., 2021) and implement a novel approach based on the estimation of unit-specific and spatially clustered synthetic controls. I conclude by calibrating a model of insurance demand to compute back-of-the-envelope welfare estimates of the official flood risk information.

This analysis yields three novel findings. First, over the past two decades new maps rezoned 1 million properties inside the high-risk floodplain while removing more than 2.4 million properties from these areas, resulting in a net *decrease* of more than 1.4 million properties from the 100-year floodplain – primarily in neighborhoods with a higher share of Black and Hispanic residents.² Such aggregate reduction in official flood risk is incompatible with current climate science, which consistently estimates that flood risk is increasing nation-wide. This divergence between official and independent risk information is driven by two factors: improved mapping technology in areas exposed to coastal and fluvial risk, and the fact that pluvial (rain-based) risk, the fastest growing driver of flooding, continues to be overlooked in official maps.

Second, I find that map updates decreased insurance take-up nation-wide and caused disproportionate declines in neighborhoods with more Black and Hispanic residents, thus exacerbating racial disparities in flood insurance coverage. The official floodplain classification drives these impacts: despite higher premiums in the high-risk zones, for every five properties rezoned inside (outside) the high-risk zones, one household is induced to purchase (drop) flood insurance. These impacts of information can a priori occur through two channels: changing households’

²Throughout the paper, I use the following US Census Bureau racial groups: “American Indian or Alaska Native,” “Asian,” “Black or African American,” “Native Hawaiian or Other Pacific Islander,” “White,” and the “Hispanic or Latino” ethnic group. When not included in the main texts, available results for other racial and ethnic groups are presented in the appendix.

beliefs about flood risks, or expanding the enforcement of the *mandatory purchase requirement* (MPR) inside the 100-year floodplain, which nominally requires households with a federally backed mortgage to purchase insurance – although disagreement exists regarding its actual enforcement (Tobin and Calfee, 2005; Michel-Kerjan et al., 2012; National Research Council, 2015; Government Accountability Office, 2021b; Blickle and Santos, 2022). While I cannot perfectly disentangle these two mechanisms, I find evidence suggesting that both are at play: rezoning properties inside the high-risk zones causes adjacent properties outside of the high-risk zone and not subject to the MPR to purchase insurance – a local spatial spillover effect driven by beliefs about risks. Insurance take-up responses appear slightly stronger in neighborhoods with a higher share of properties with a mortgage, consistent with the MPR constraint binding for some households. In contrast, the digitization of already-available risk information does not impact insurance take-up, suggesting that the costs of accessing information did not substantially limit household’s demand for flood insurance.

Finally, I leverage the reclassifications of properties and independent estimates of flood risk in a structural model of insurance demand to get back-of-the-envelope estimates of the welfare impacts of risk information. I find that under plausible values of risk aversion parameters, the map updates conducted after Hurricane Katrina *decreased* welfare – an effect driven by the provision of incorrect risk information. These losses were spread across the income distribution. In contrast, correcting floodplain boundaries nation-wide would yield gains exceeding \$20 billion annually, primarily in wealthy and majority-white neighborhoods.

Taken together, these results suggest that publicly provided climate risk information is a key tool to promote private investments in residential flood insurance. Yet, because map updates rezoned more than one million properties outside of the high-risk zones on aggregate, the flood risk mapping program *reduced* coverage nation-wide by about 100,000 insurance policies and exacerbated racial disparities in insurance coverage. This is in stark contradiction with the objectives of the program of promoting insurance take-up through updated and digital information. While the full welfare impacts of the program are highly sensitive to modelling choices, back-of-the-envelope estimates suggest that flood information provision decreased welfare.

This work contributes to the large literature on the impacts and salience of information provision (Stigler, 1961; Cutler et al., 2004; Chetty et al., 2009; Jessoe and Rapson, 2014; Cabral and Hoxby, 2015; Davis and Metcalf, 2016; Kamenica, 2017). A small but growing literature focuses on the role of publicly provided information in the context of weather forecasts and climate risks (Rosenzweig and Udry, 2014; Shrader, 2020; Shrader et al., 2022; Molina and Rudik, 2022). I show that the coarse categorization of risk information currently used in the official maps can promote insurance uptake, whereas digitizing previously existing information

has no detectable effect. While previous work estimated the impact of flood risk information on housing values (Pope, 2008; Bin and Landry, 2013; Beltrán et al., 2018; Shr and Zipp, 2019; Gibson and Mullins, 2020; Hino and Burke, 2021; Gourevitch et al., 2023), this paper is the first to study the roll-out of flood maps and their impacts on insurance take-up. The existence of spillover effects reveals that households pay attention not only to their own floodplain classification, but also to the classifications of other properties located within the same neighborhood, implying that current hedonic estimates of the impacts of flood maps on property values might be lower bounds on the real effects.

This paper directly contributes to the climate change adaptation literature (Kahn, 2005; Annan and Schlenker, 2015; Barreca et al., 2016; Diaz and Moore, 2017; Botzen et al., 2019; Kocornik-Mina et al., 2020; Kahn, 2021; Sastry, 2021; Carleton et al., 2022; Bakkensen and Barrage, 2022; Ostriker and Russo, 2022; Fried, 2022). While post-disaster assistance is scarce (Government Accountability Office, 2020), previous work found substantial indirect public costs from extreme events (Deryugina, 2017), demonstrating that disaster impacts are not entirely borne privately. Wagner (2022) shows through a compelling analysis of the US flood insurance market how the existence of “frictions” in the demand for insurance justifies the strict implementation of an insurance mandate. Other work discusses the value of insurance to promote reconstruction or migration post-disaster (Turnham et al., 2011; Kousky, 2019; Nguyen and Noy, 2020; You and Kousky, 2023). Several papers discussed the role of information frictions in reducing insurance demand below its optimal level (Chivers and Flores, 2002; Kunreuther et al., 2013; Atreya et al., 2015; Shao et al., 2017; Bradt et al., 2021; Hu, 2022), and Mulder (2022) provides a careful study of the welfare impacts of flood risk information. While these studies are conducted in specific settings, I provide the first global assessment of the largest climate risk mapping program in the US, finding that potential gains from flood information markedly diverge from realized ones due to systematic issues in the implementation of the flood risk mapping program. Outdated flood maps are an issue known to the government (Congressional Budget Office, 2017), and several groups point to outdated maps as a main driver of climate mis-adaptation (Scata, 2017; Lehmann, 2020; Frank, 2020). I find, however, that map updates tended to further provide distorted information and price signals, with little sign of improvements over time, thus jeopardizing household-level purchase of insurance and other defensive investments.

Finally, this work contributes to the environmental inequalities and environmental justice literature (Bullard (1983), see Banzhaf et al. (2019) for a recent economic-focused review). While earlier work focused on the interaction of race, poverty, and pollution exposures, Bakkensen and Ma (2020) highlight correlations between race and flood risk. An active literature explores

the causes of environmental inequalities (Gamper-Rabindran and Timmins, 2013; Depro et al., 2015; Banzhaf and Walsh, 2016; Christensen and Timmins, 2022; Currie et al., 2022). Directly related to the present paper, Hausman and Stolper (2021) discuss the role of information as a potential driver of environmental inequalities, even when information is uniformly missing. I uncover that both information provision and omission are unequal between neighborhoods, which exacerbated an already wide gap in flood insurance coverage.

The remainder of this paper proceeds as follows. Section 2 provides additional background on the flood maps and presents the datasets used in this paper. Section 3 presents descriptive results on how the official flood risk information has evolved since Hurricane Katrina. Section 4 provides estimates of the average effects of flood map updates on insurance take-up using heterogeneity-robust difference-in-differences estimators. Section 5 introduces local synthetic controls to assess both the heterogeneity and robustness of the treatment effects. Section 6 presents a model of insurance demand and provides back-of-the-envelope welfare impacts of flood map updates. Section 7 concludes.

2 Background and data sources

2.1 A brief history of flood maps

Flood risk is hard to estimate – like other “rare” events, reliance on historical data to predict occurrence probabilities suffers from a small sample issue. The inability to predict flood events and damages led to the unraveling of the private flood insurance market following the 1927 Mississippi River flood. Since its inception in 1968, the National Flood Insurance Program (NFIP) has crucially relied on the provision and updating of Flood Insurance Rate Maps (FIRMs, colloquially called “flood maps”) with the dual goals of providing flood risk information and setting federally provided flood insurance premiums. These maps are provided at the community level, where communities can be towns, cities or counties (see Knowles and Kunreuther (2014) for a discussion of the NFIP’s history).

While flood risk is hard to assess, communicating flood risk to non-experts is an additional challenge. To convey flood risk probabilities to the public, the maps provide simplified risk estimates that are differentiated by zones. In particular, flood maps highlight the 100-year floodplain, also known as the Special Flood Hazard Area (SFHA) or high-risk zone, where the annual probability of a flood event exceeds 1%. Flood maps also highlight the 500-year floodplain (or X shaded zones) or low-risk zone, where the annual probability of flooding is estimated to be between 0.2% and 1%. Areas outside of these zones are classified as “minimal flood risk”. In practice, the 100-year floodplain plays an essential role in risk communication as

well as regulatory constraints: (i) flood insurance is mandatory for properties with a federally backed mortgage that are located within the 100-year floodplain, (ii) flood insurance premiums vary discontinuously at the 100-year floodplain borders, and (iii) urban planning often restricts development within the 100-year floodplains.³

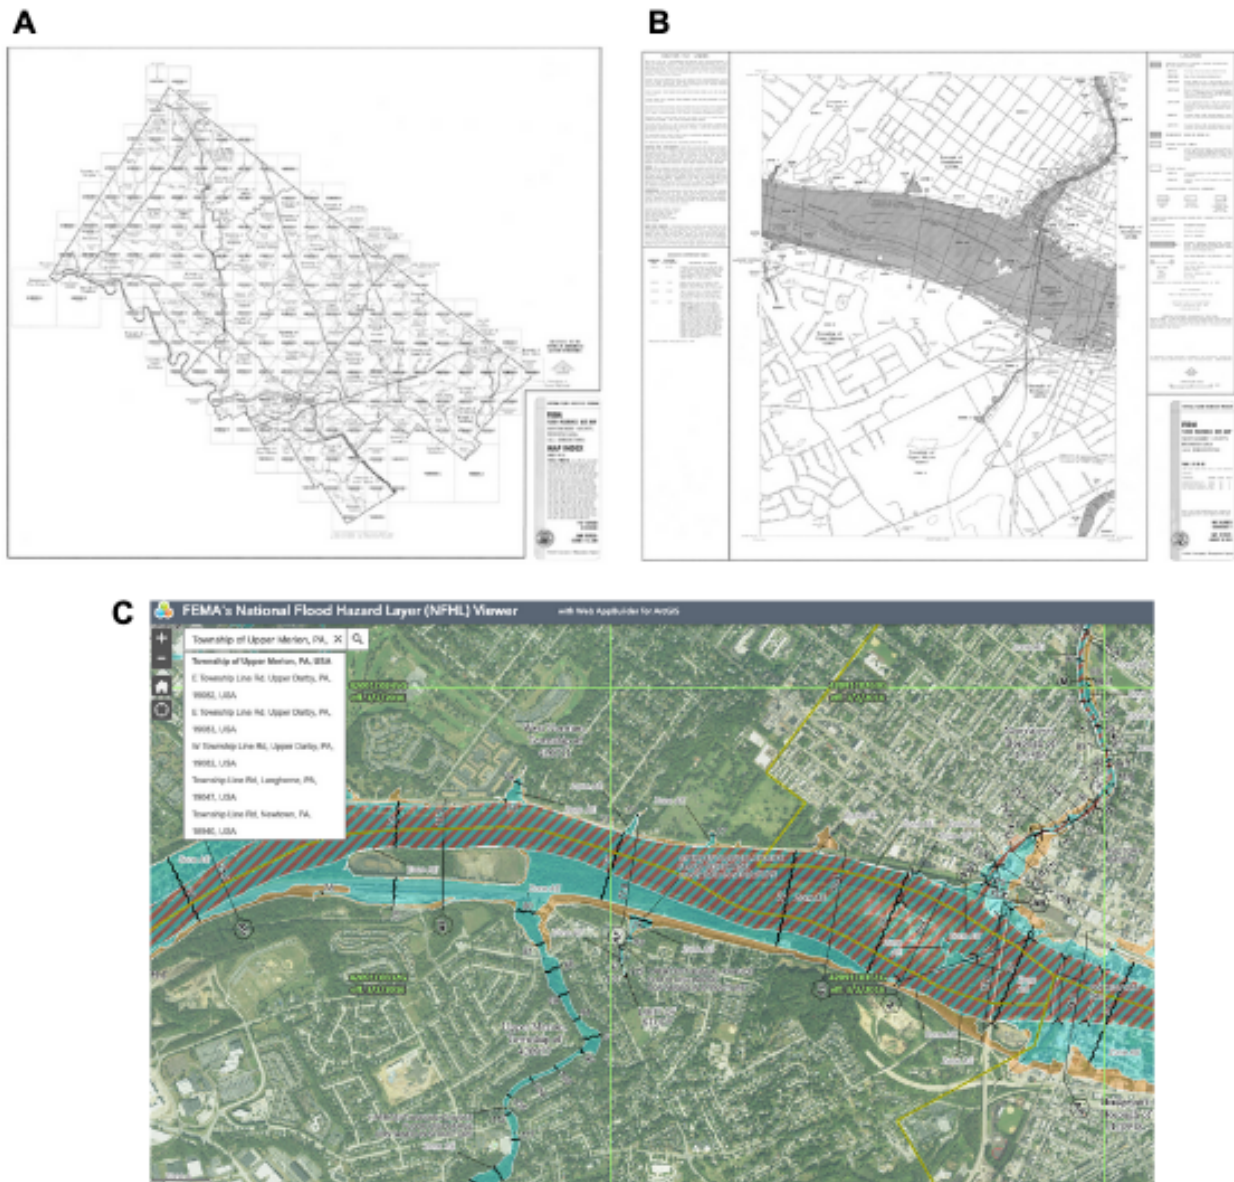
Until the early 2000s, almost all flood maps were paper-based. These maps were hard to find, hard for non-experts to read, and generally outdated. To improve flood risk identification and facilitate information access, an initiative called the Map Modernization program (MapMod) allocated more than \$1.4 billion starting in 2005 to update paper-based FIRMs to Digital Flood Insurance Rate Maps (DFIRMs, or simply “digital maps”). This made MapMod one of the most expensive public mapping programs in history (see Morrissey (2008) for detailed background on the program). In 2008, MapMod was rolled into the Risk Mapping, Assessment, and Planning Program, which essentially continued to update flood maps (FEMA, 2009), with total estimated costs of \$4 billion (Department of Homeland Security, 2012). To this day, funding is continuously allocated to update the remaining paper-based maps to digital maps.

The flood maps’ digital transition reduced the cost of accessing flood risk information and, in many cases, changed the underlying floodplain boundaries. To appreciate how digitization improved access to information, imagine that a resident living in Upper Merion Township, Pennsylvania, wanted to learn about her property’s exposure to flood risk before the Montgomery County flood map had been digitized.⁴ Because flood maps are divided into “panels,” or rectangular sections of the US, she would first have to figure out which panel number includes her property: this implies finding the Montgomery Map Index and identifying the relevant panel code (see Figure 1.A). She would then use this code to find the relevant paper-based map and attempt to locate her property relative to the 100-year floodplain. In contrast, with the updated digital flood map, she can access the National Flood Hazard Layer, a publicly available online platform, and directly locate her property by typing her address in the search bar (see Figure 1.C). Figure 1 shows a situation where floodplain boundaries did not substantially change through the map update. This is not always the case: map updates can rezone properties either outside or inside of the 100-year and 500-year floodplains. In the rest of the paper, I use pure digitizations as well as changes in floodplain boundaries to assess the impacts of both access to information and of changing risk information on the demand for flood insurance.

³In principle, flood zones could be defined for any annual probability of flooding (for instance, the flood maps could highlight the 200-year and the 1000-year floodplains within the current X zones). This has spurred criticism of the 100-year floodplain concept as both an insurance-setting and a communication tool from academics and practitioners alike (Wittenberg, 2017; Koerth, 2017; Bell and Tobin, 2007; Ludy and Kondolf, 2012; Highfield et al., 2013). A recent policy change in the NFIP, called Risk Rating 2.0, aims to better reflect flood risk in insurance premiums while still maintaining the use of 100-year floodplain boundaries to communicate risk.

⁴Hurricane Ida severely hit Montgomery County in September 2021, causing the deaths of three people.

Figure 1: Evolution of the Flood Insurance Rate Maps in Montgomery County, PA



A: Map Index from Montgomery County, PA, effective in 1999. **B:** FIRM number 42091C0351F, effective in 1999. The 100-year and 500-year floodplains are shown in dark and light grey, respectively. **C:** Screenshot of the DFIRM effective in 2016, as shown on the National Flood Hazard Layer. The 100-year floodplain is depicted in blue and blue/red stripes, while the 500-year floodplain is depicted in orange. Sources: FEMA's Map Service Center and National Flood Hazard Layer.

2.2 Digital Flood Insurance Rate Maps

I collect and compile data on all digital flood maps released between 2005 and 2019, obtained both through Freedom of Information Act requests and following meetings with former GIS analysts who performed contractual work for FEMA. These digital maps provide both the *date*

when the DFIRM became active and the *spatial polygons* of the 100-year and 500-year floodplains.⁵ Figure 2.A depicts the cumulative number of census tracts that received a DFIRM since 2000, while Figure 2.B shows the spatial distribution of the census tracts that received at least one digital flood map and the year they received it. More than 85% of census tracts are now covered by a digital map – this contributes to the credibility of the difference-in-difference estimation strategy presented in Section 5 below, as this mitigates concerns of selection into treatment (this leaves out potential concerns of selection into treatment timing, discussed below). Although flood maps are supposed to be updated at least once every five years, to date less than 20% of census tracts have received more than one DFIRM, covering only 16 million residential properties, or less than 12% of the residential housing stock. The main econometric analysis of this paper focuses on census tracts that have received their first digital flood map, as this allows me to estimate the impact of digitization.⁶

2.3 Digitized version of paper-based flood maps: the Q3 data

Most flood maps were paper-based prior to 2007. Because we are interested in how the modernization of flood maps changed the spatial extents of the floodplains, we need a GIS-ready version of these paper-based maps. Luckily, FEMA produced digitized versions of the 100-year and 500-year floodplain boundaries depicted on the paper-based maps that were effective in over 1,300 counties at the time of scanning (1996-2000). This product, known as the Q3 data, makes it possible to assess how FEMA’s floodplain designations evolved over time (FEMA, 1996).⁷ Although I can only track changes in the floodplains for which Q3 data exist, this encompasses more than 48,000 census tracts and includes the most flood-prone areas, where the majority of insurance policies are purchased.⁸

Figure 3.A shows the digital version of the 1984 flood map in New Orleans,⁹ while Figure 3.B shows the digital flood map that became effective in September 2016. In both digital products,

⁵The steps used to process these maps are provided in Appendix A.1.

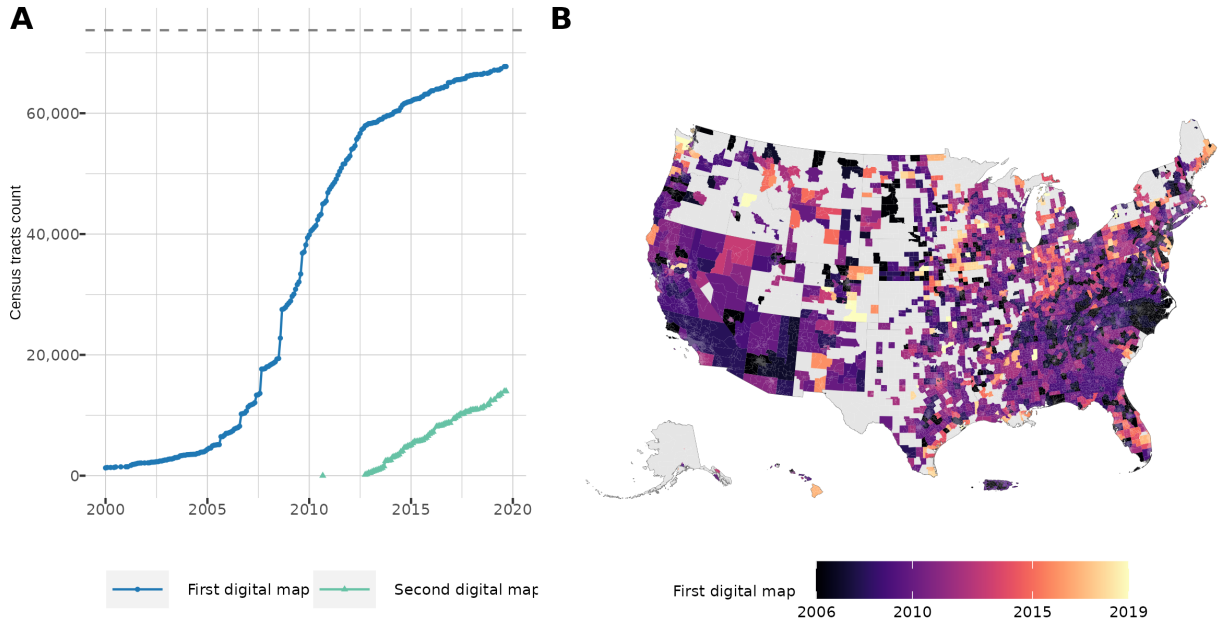
⁶This further facilitates the definition and interpretation of treatment effects: for census tracts that received more than one digital flood map, defining treatment status after the first digital map and before the second map is ambiguous. Resolving the ambiguity requires either (i) making assumptions about the dynamics of the treatment effects and leveraging an ad hoc threshold above which the first treatment effect is considered to have vanished, or (ii) defining treatment as a qualitative variable taking more than two values, which further complicates the estimation of treatment effects.

⁷The Q3 data was meant as an early digitization effort to quickly distribute the floodplain data to practitioners to facilitate disaster recovery and planning activities. It became the standard product for planning under the Disaster Mitigation Act of 2000 (Gall et al., 2007)

⁸Out of the 2,930,143 residential flood insurance policies active in January 2009, 85 % were in areas covered by a Q3 map.

⁹The original paper-based map, digitized in the Q3 product without changes in the underlying boundaries, is presented in the Appendix Figure A.3. In Figure A.4, I show a similar procedure for the south-eastern part of Broward County (FL), an area also subject to substantial flood risk.

Figure 2: Flood map digitization progress throughout the US



A: Cumulative number of census tracts that received a first digital flood map (in blue circles) or a second digital flood map (green triangles). The gray dotted line shows the total number of census tracts in the US (73,745 tracts, including Puerto Rico). **B:** Spatial release of the first digital flood maps, with 2006 representing all maps released in 2006 or earlier for clarity.

the 100-year floodplain is depicted in light blue, while the 500-year floodplain is shown in orange. In both panels the black dots represent residential properties, making it possible to compute changes in the number of properties located inside the 100-year floodplain (Figure 3.C). For the empirical estimation of Section 4, I aggregate these changes at the census tract level and classify census tracts depending on whether the map update rezoned more than 1% of residential properties inside or outside the 100-year floodplain on aggregate, or whether the map update did not change the 100-year classification of any property (Figure 3.D).¹⁰

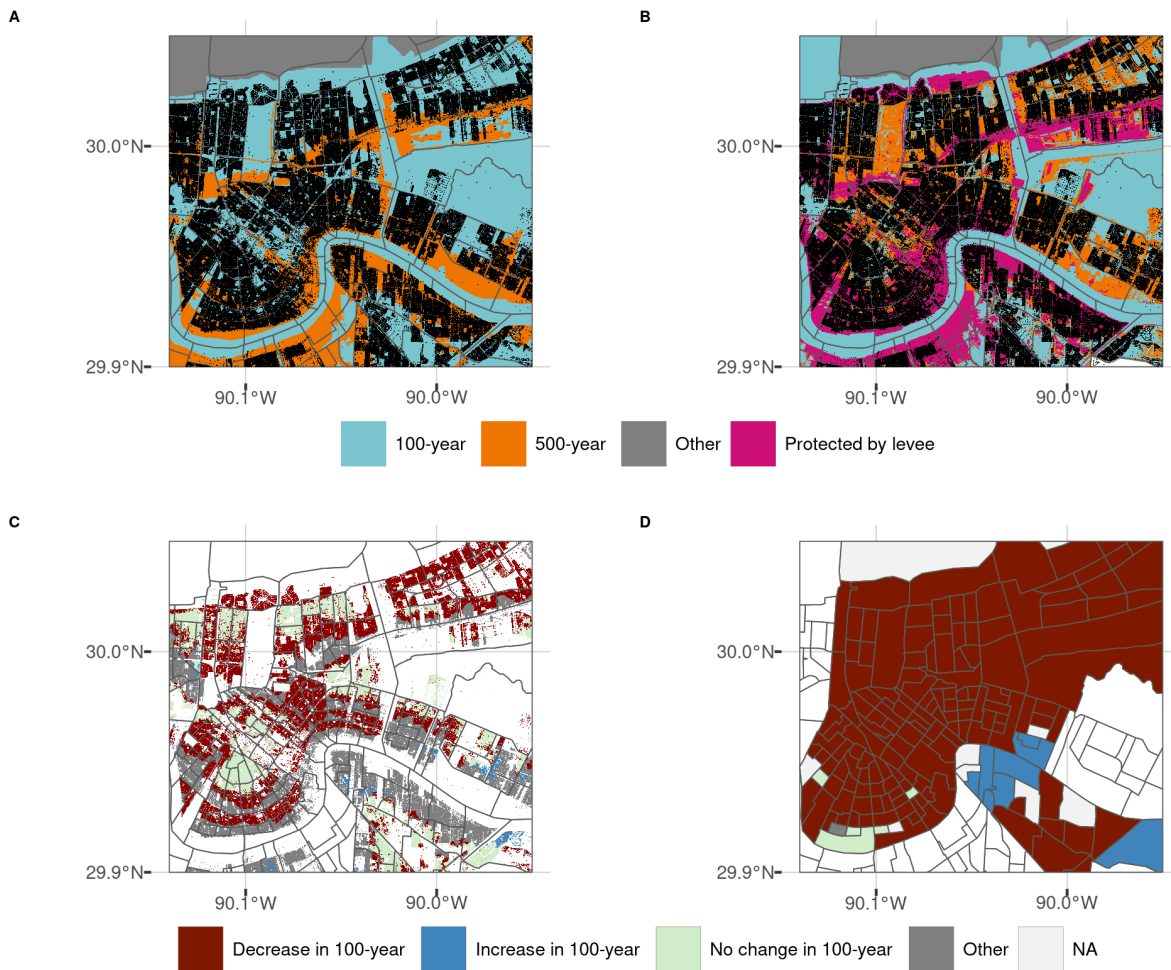
2.4 First Street Foundation data

For the past sixty years, official flood maps were the main source of flood risk information available to the public, if not the only source.¹¹ This changed in 2020 when a consortium of scientists, technologists, and communicators called the First Street Foundation (FSF) started producing and disseminating flood risk estimates for the entire contiguous US at a very fine

¹⁰The 1% cut-off is arbitrary and discussed below, before being relaxed in Section 5

¹¹Several insurance and re-insurance companies developed their own models of flood risk, but these models were not publicly available.

Figure 3: Changes in the Flood Insurance Rate Maps in New Orleans, Orleans Parish, LA



A: Q3 data product showing the flood zones effective in 1984, combined with geolocalized residential properties in Orleans Parish (black dots). **B:** Current DFIRM showing the flood zones effective as of September 2016, combined with geolocalized residential properties. **C:** Changes in the 100-year floodplain computed at the property level. **D:** Changes in the number of residential properties in the 100-year floodplain between 1984 and 2016, aggregated at the census tract level and classified into three categories (increase, decrease, or no change). In New Orleans, the map modernization program led to a substantial rezoning of properties *outside* of the 100-year floodplain. This rezoning was rationalized by substantial investments made by the US Army Corps of Engineers in flood control systems such as levees, floodgates and pump stations (FEMA, 2016).

level of spatial precision (3 meters).

The FSF flood model is built around LISFLOOD-FP, an open source hydrodynamic model developed over the last decade (Bates et al., 2010; Neal et al., 2012; De Almeida and Bates, 2013; Sampson et al., 2013, 2015).¹²

¹²The first high-resolution model of fluvial and pluvial flood hazards for the entire contiguous United States is presented in Wing et al. (2017). This modelling framework was substantially refined in Bates et al. (2020) with the inclusion of coastal risk modelling, improved data sources, and inclusion of current and future climate conditions. The cited studies underpin the FSF inundation depth estimates and have undergone extensive

In addition to differences in the hydraulic models used to simulate inundation depths, the FSF flood model and the FEMA FIRMs differ in 4 major ways: (i) FSF models pluvial risks, while such flood drivers are typically omitted in the FIRMs (House of Representatives, 2020; Government Accountability Office, 2021a); (ii) FSF considers the potential joint occurrence of pluvial, coastal and fluvial risks; (iii) FSF covers all the United states, whereas as of 2020 the FEMA FIRMs only covered about a third of the 3.5 million miles of stream (ASFPM, 2020); (iv) FSF accounts for current and modeled climate projections (using global circulation models and hurricane precipitation predictions), whereas FIRMs rely solely on historic data.¹³

Converting inundation-depth estimates into property-level damages poses an additional challenge. The FSF model employs the HAZUS-MH methodology developed by FEMA (FEMA, 2013). Significant uncertainties surround the validity of these depth-damage functions, and on-going research seeks to refine them (Pollack et al., 2022; Porter et al., 2023).

While the FSF flood model is not flawless, it is currently regarded as providing some of the best publicly available flood risk estimates in the US (First Street Foundation, 2020; Armal et al., 2020; Eby, 2023). The FSF model is now used by more than 30 federal agencies and government sponsored enterprises (First Street Foundation, 2022).

2.5 Flood insurance policies

Finally, I use data on flood insurance policies provided by FEMA covering all NFIP policies active between 2008 and 2019, obtained through Freedom of Information Act requests. Each of the 60 million observations represents an insurance policy with information on the policy active dates and the insured property (such as the construction year and the FEMA flood zone in which the property is located) as well as details about the insurance contract (coverage and cost, in particular). The dataset does not include property identifiers, and policies are not geocoded. This implies that I cannot track changes in insurance take-up at the property level. For the analyses below, I construct a panel of active policies at the census tract / month level. Details about the construction of the panel are provided in Appendix A, and the different datasets are summarized in Figure A.5.

The upper part of Table 1 presents summary statistics on the panel of insurance policies. There is a wide range in the number of active policies per census tract, varying from 0 to more

validation against historical flood event data and local hydraulic models developed by USGS where available.

¹³Providing flood hazard estimates over very large scales implies substantial computational constraints. The FSF model is tractable due to the use of LISFLOOD-FP, which solves the local inertial form of the shallow water equations in two dimensions. The main alternative to LISFLOOD-FP for large scale modelling is to use an approach based on the Height Above Nearest Drainage (Nobre et al., 2016), although this also requires strong hydrodynamic assumptions.

than 6,000. This spread will partially motivate the estimation strategy I conduct at the tract level in Section 5. In some tracts, all the policies cover properties located inside the high-risk zones, whereas in others all the policies cover only properties located outside of the 100-year floodplain. The zoning inside and outside of the 100-year floodplain explains a substantial share of the variation in average insurance premiums observed between tracts: the average policy cost is about \$596 per year, but it is \$419 outside of the 100-year floodplain and \$1,012 inside of it.

Table 1: Summary statistics, estimation panel and census tract cross section

Variable	N	mean	sd	min	max
Panel:					
Number of active policies	8200080	58.12	187.77	0	6545
Number of active policies inside 100-y	8200080	30.43	136.59	0	5399
Number of active policies outside 100-y	8200080	27.68	101.33	0	5358
Share policies inside 100-y	7297869	0.33	0.34	0	1
Average policy cost	7339805	614.93	362.93	36	10143
Average policy cost inside 100-y	4774686	1040.19	655.49	42	14048
Average policy cost outside 100-y	7043775	430.06	163.09	36	5378
Average construction of insured property	7339805	1972.5	20.26	1900	2019
Average initial insurance year	7339805	2008.57	3.74	1970	2020
Average coverage	8200080	192565.85	104842.46	0	7173600
Cost per thousand dollar insured	7339805	3.09	2.58	0.54	173.04
Cost per thousand dollar insured inside 100-y	4774652	6.65	4.31	0.1	720
Cost per thousand dollar insured outside 100-y	7043775	1.76	1.18	0.46	60
Cross section:					
Year of map modernization	65341	2009.24	3.72	1995	2019
Disaster declaration within two years of treatment	65341	0.54	0.5	0	1
Has a Q3 map	71183	0.76	0.43	0	1
Has a digital map in 2019	71183	0.92	0.27	0	1
Relative change in properties zoned inside 100-y	38773	-0.02	0.13	-1	1
Share population Black or African American	70783	0.14	0.22	0	1
Median household income, past twelve months	70440	64314.44	32176.82	2499	250001

3 Stylized facts on official flood information provision

The data consolidation of the sources presented in Section 2 makes it possible to study at scale how official flood risk information evolved over the past two decades and how it compares with independent scientific estimates. This section highlights four stylized facts: (i) new maps reduced the number of properties zoned in the 100-year floodplain over the past two decades; (ii) this reduction is primarily due to the increased complexity of flood maps, with the drawing of more floodplain boundaries; (iii) two-thirds of the rezoning that occurred is consistent with

the FSF flood model, yet a larger concern lies in the omission of rain-based flood risk during the updating process; (iv) properties rezoned outside of the 100-year floodplain as well as properties incorrectly ignored during the risk mapping process are predominantly located in more Black and Hispanic neighborhoods, which will have substantial implications on disparities in insurance take-up.

3.1 Property counts

Starting with the simplest possible exercise, Table 2 presents an accounting of the number of residential properties zoned inside the 100-year floodplain in the FEMA maps in 2005 (Q3 data) and 2019 (NFHL19) as well as the number of properties inside the 100-year floodplain based on the FSF Flood model. For all three columns and the rest of the analysis below, the properties' coordinates are taken from the 2020 First Street Foundation registry, such that changes in property counts within floodplains are entirely due to changes in the floodplain boundaries.¹⁴

The first striking finding is the *decline* in the number of properties classified inside the 100-year floodplain, despite the broader geographic coverage of the newer maps. Although the Q3 maps only covered about two-thirds of residential properties nation-wide (87 million properties compared with 116 million in NFHL19), they showed 7.2 million properties in the 100-year floodplain, or 500,000 *more* than the 6.7 million in the current digital FEMA maps. The shrinking of the 100-year floodplain is even greater if we restrict the analysis to only the tracts mapped in both the Q3 and NFHL19 products (lower panel of Table 2), which shows the Q3 number decreasing to 6.3 and the NFHL19 declining to 4.9 million.¹⁵

This reduction in the number of properties zoned in high-risk floodplains is surprising, as there is a consensus among scientists that flood risk in the US is increasing, which should lead to expansions of the 100-year floodplain in urban areas (Marsooli et al., 2019; Bates et al., 2020; Wing et al., 2022). Local reductions of the floodplains can be warranted to account for new adaptation infrastructures and improvement in mapping technologies, but the aggregate decline shown above is inconsistent with the current science.

Perhaps counter-intuitively, I find that the rezoning of properties outside the 100-year floodplain is due to the drawing of *more* floodplain boundaries. This can be summarized with the Polsby-Popper score, used in the political science literature to measure district gerrymander-

¹⁴In Appendix Table A.1, I show using data from Zillow's ZTRAX records that these patterns are not driven by differential residential construction rates inside and outside of the floodplains. The First Street Foundation registry has substantially better coverage than ZTRAX (especially in rural areas) and can directly be matched to the FSF Flood Model estimates used below, making it the preferred source of residential data.

¹⁵Figure C.6 presents the spatial distribution of increases and decreases in the number of properties located inside the 100-year floodplain.

Table 2: Residential properties in the 100-year floodplain in different risk mapping products

Flood zone	2005	NFHL 2019	First Street Model
Unconditional:			
Inside 100-year floodplain	7,209,546	6,653,511	14,776,887
Outside 100-year floodplain	79,981,116	109,614,405	117,510,574
Not mapped	45,096,799	16,019,545	0
Conditional on Q3 and NFHL19:			
Inside 100-year floodplain	6,296,693	4,917,259	9,023,220
Outside 100-year floodplain	72,746,998	74,126,432	70,020,471

The conditional counts include only those properties in census tracts that are mapped in both the Q3 and NFHL19 data products.

ing.¹⁶ I find that in 2005, the official 100-year floodplain covered an area of 255,402 km^2 , and this area slightly increased to 269,678 km^2 in 2019 (about the area of Colorado). However, the *perimeter* of the 100-year floodplain boundaries experienced a spectacular increase, from 1,289,567 to 1,766,252 km . Thus, between 2005 and 2019, the Polsby-Popper score of the 100-year floodplain changed by -44 %, revealing a surge in the complexity of the underlying polygons. As discussed below, the increased complexity is consistent with better modelling tools that allow for a finer-grain estimation of flood risk. Table B.2 in the appendix shows that such increased flood map complexity occurred in almost all US states.

3.2 Discrepancies between official and independent flood maps

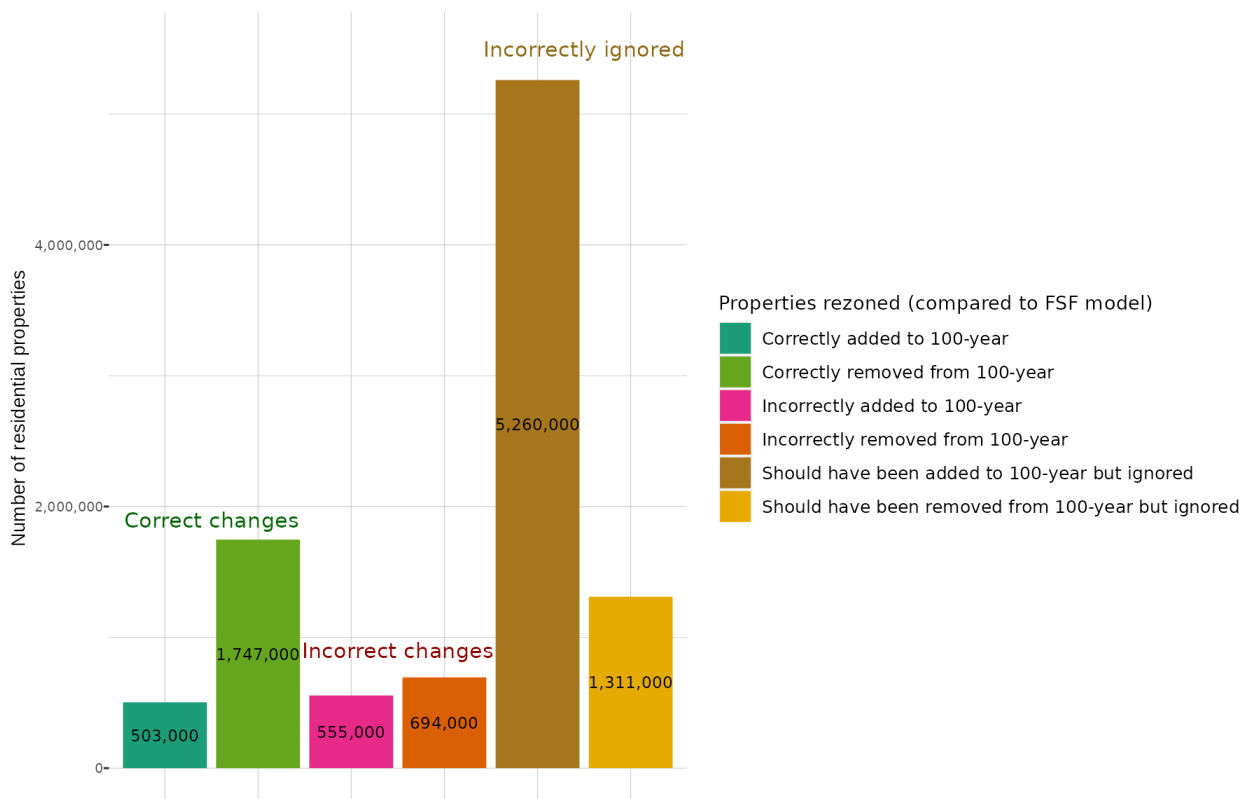
To assess the scientific accuracy of official flood information, I compare the official map updates with the First Street Foundation estimates at the property level. While the FSF model is not perfect, it provides an independent proxy of “true flood risk” based on the latest peer-reviewed flood science.

Figure 4 focuses on areas that were mapped in both the Q3 flood product (in 2005) as well as in 2019 (digital flood maps) to decompose how flood maps impacted the number of properties zoned inside the 100-year floodplain. The “correct” floodplain status is assumed to be provided by the FSF flood model. There are three important conclusions. First, comparisons with the FSF flood model indicates that out the 2.4 million properties that were rezoned outside of

¹⁶The Polsby-Popper metric of a district D is defined as $PP(D) = \frac{4\pi A(D)}{P(D)^2}$, with $A(D)$ and $P(D)$ the area and perimeter of district D , respectively. A perfect circle has a Polsby-Popper score of 1, and reduction in the “compactness” of a polygon brings the Polsby-Popper down, with a lower bound at 0.

the 100-year floodplain during the map updates, 1.7 were correctly removed. Second, half of the 1.1 million properties that were zoned inside the 100-year floodplain during the flood map update appear to have been correctly added. Lastly, over 5.2 million properties were outside the 100-year floodplain in 2005 and should have been rezoned inside the 100-year floodplain. However, they were omitted during the map update.

Figure 4: Changes in the number of residential properties in the FEMA 100-year floodplain between 2005 and 2019, compared to the FSF flood model



Counts are based on residential properties that are mapped in both the Q3 product and in the 2019 digital flood maps. The FSF flood model is assumed to provide the “correct” depiction of the 100-year floodplain. Additional counts for residential properties that remained correctly mapped inside or outside of the 100-year floodplains are omitted.

Previous work highlights the potential for biases in the official flood maps due to local political considerations (Pralle, 2019; Lea and Pralle, 2022). In particular, homeowners and city officials have incentives to lobby to remain outside of the 100-year floodplain, given that the 100-year floodplain increases flood insurance premiums and decreases property values (Hino and Burke, 2021). The results above suggest that local political motivations are unlikely to be driving the bulk of floodplain rezoning: most removals from the 100-year floodplain appear to

have been warranted, at least according to the FSF model.

These comparisons reveal that the modernization of flood maps primarily failed by *omission*: while two thirds of the removals from the official 100-year floodplain align with the FSF estimates, more than five million properties were incorrectly “ignored” during the modernization process when they should have been added to the 100-year floodplain.

Columns 5 to 8 in Table 3 show that most properties incorrectly left outside the 100-year floodplain throughout the map updates are substantially more likely to be subject to pluvial flood risk, as opposed to coastal or fluvial risk. These results complement previous work that discussed how official maps systematically omit pluvial risk (Wing et al., 2018).

Public officials recently recognized the problematic omission of pluvial risk in official flood maps, with ongoing policy discussions aiming to find ways to better account for rain-based flooding events (House of Representatives, 2020; Government Accountability Office, 2021a). The omission of pluvial risk was historically driven by the high data requirements (accurate pluvial risk assessment requires detailed topographical data, soil information, and localized rainfall patterns) and resource constraints (developing models that include pluvial risk would require significant investment in technology, training, and data collection). Climate change further complicates the assessment of rain-based events. Surprisingly, the results above reveal that the map update process does not correct the omission of pluvial risk.¹⁷

3.3 Racial disparities in floodplain exclusion and omission

Two key observations regarding the map updates over the past two decades are the removal of 1.4 million properties and the omission of 5.2 million properties from the 100-year floodplain. To assess whether the provision of official risk information could lead to unequal climate adaptation, this subsection examines where property exclusions and omissions from the 100-year floodplain occurred in the US. The main text focuses on highlighting disparities in neighborhoods with more Black and Hispanic residents, as these disparities are among the largest. I perform similar analyses in the appendix for the other main racial and ethnic groups defined by the Census Bureau.¹⁸

Columns 1 to 4 in Table 3 show properties in neighborhoods with an increasing share of Black and Hispanic residents were more likely to be removed from the 100-year floodplain, even after conditioning on FSF-estimated flood risk. These results are robust to alternative specifications (see Table C.3). Table C.4 in the appendix show that large disparities in floodplain removals

¹⁷Table C.6 in the appendix shows that the accounting of flood risk in the official flood maps does not appear to substantially improve in more recently updated flood maps.

¹⁸American Indians and Alaska Natives are omitted due to small sample sizes. Racial quantiles are defined at the neighborhood-level based on Decennial Census respondents who self-identify as single-race.

Table 3: Characteristics of properties removed and incorrectly ignored during map updates

Dependent Variables: Model:	Removed from Q3 100-year				FSF 100-year, ignored during update			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<u>Variables</u>								
Constant	0.5031 (0.0089)		0.5118 (0.0091)		0.2619 (0.0087)		0.2496 (0.0084)	
Coastal	-0.3488 (0.0097)	-0.2563 (0.0110)	-0.3425 (0.0097)	-0.2602 (0.0109)				
Inland fluvial	-0.2915 (0.0103)	-0.2590 (0.0074)	-0.2856 (0.0103)	-0.2592 (0.0074)	0.2317 (0.0102)	0.0033 (0.0137)	0.2334 (0.0100)	0.0103 (0.0137)
Inland pluvial	-0.2451 (0.0073)	-0.2101 (0.0061)	-0.2389 (0.0076)	-0.2105 (0.0060)	0.4817 (0.0083)	0.3054 (0.0125)	0.4859 (0.0083)	0.3137 (0.0125)
Black or AA 25-50%	0.0790 (0.0115)	0.0427 (0.0088)			0.0457 (0.0090)	0.0660 (0.0081)		
Black or AA 50-75%	0.1566 (0.0129)	0.0820 (0.0108)			0.0741 (0.0092)	0.1202 (0.0093)		
Black or AA 75-100%	0.1983 (0.0121)	0.0826 (0.0120)			0.0771 (0.0075)	0.1593 (0.0095)		
Hispanic 25-50%			0.0479 (0.0107)	0.0365 (0.0085)			0.0545 (0.0073)	0.0564 (0.0074)
Hispanic 50-75%			0.1343 (0.0126)	0.0579 (0.0106)			0.0721 (0.0090)	0.0999 (0.0096)
Hispanic 75-100%			0.1617 (0.0132)	0.0574 (0.0144)			0.1050 (0.0093)	0.1453 (0.0119)
<u>Fixed-effects</u>								
County FE		Yes		Yes		Yes		Yes
<u>Fit statistics</u>								
Observations	6,288,904	6,288,904	6,288,904	6,288,904	8,979,280	8,979,280	8,979,280	8,979,280
R ²	0.12240	0.30302	0.11577	0.30150	0.18612	0.29154	0.18770	0.28818
Within R ²		0.06026		0.05822		0.08598		0.08164
Clustered (Tract FE) standard-errors in parentheses								
Signif. Codes: ***: 0.01, **: 0.05, *: 0.1								

Linear probability regressions at the property level of the form $Y_i = \alpha_c + \beta X_i + \epsilon_i$, where Y_i is a dummy variable equal to 1 if the property used to be in the Q3 100-year floodplain and was removed during the map updates (columns 1 to 4), and a dummy equal to 1 when the property should be in the 100-year floodplain based on the FSF model but was left outside of the FEMA 100-year floodplain during the map update (columns 5 to 8). α_c are county fixed effects, and X_i is a vector of property and neighborhood-level characteristics. Race come from the Decennial Census. For all models, the sample is restricted to census tracts mapped in both Q3 and the NFHL19 product. In columns 1 to 4, the sample is further restricted to properties that were in the Q3 100-year floodplain, while in columns 5 to 8 the sample is restricted to properties in the FSF 100-year floodplain. “Coastal,” “Inland fluvial” and “Inland pluvial” refers to the nature of flood risk in the FSF model.

are apparent for Asian neighborhoods as well.

In addition, Columns 5 to 8 show that properties incorrectly left outside the 100-year floodplain during the mapping process are disproportionately situated in neighborhoods with higher shares of Black and Hispanic residents.¹⁹

These disparities in floodplain exclusion and omission are surprising. They are not driven by income (see Tables C.5 and C.3), but rather appear to be due to correlations between pluvial

¹⁹To ease interpretation, the racial quantiles are defined using all neighborhoods. To account for potential selection into observation, I estimated the same models but defining quantiles based on the sample of tracts for which I can observe property level changes in the floodplain classifications, and obtained similar results (see Table C.8 in the appendix).

risk, race, and pre-existing inaccuracies in the official maps. For households who purchase insurance, it is unambiguously preferable to be outside of the 100-year floodplain, as this lowers insurance premiums. However, households rezoned or left outside of the 100-year floodplain might decide to not purchase insurance, even when it might be beneficial for them to do so. The impact of map updates on consumer (and social) welfare is thus ambiguous, and will first depend on whether residents *respond* to new information. This is analyzed in Sections 4 and 5.

4 Estimating the average effects of flood map updates

4.1 Raw evidence on the impacts of flood map updates

To motivate the econometric models below, it is instructive to look at the patterns of insurance take-up already noticeable in the raw data. Figure 5 depicts the number of insurance policies covering residential properties in each month relative to January 2008, separated by the year of the map update and based on whether the map update rezoned more than 1% of properties outside (left) or inside (right) the 100-year floodplain at the tract level.²⁰

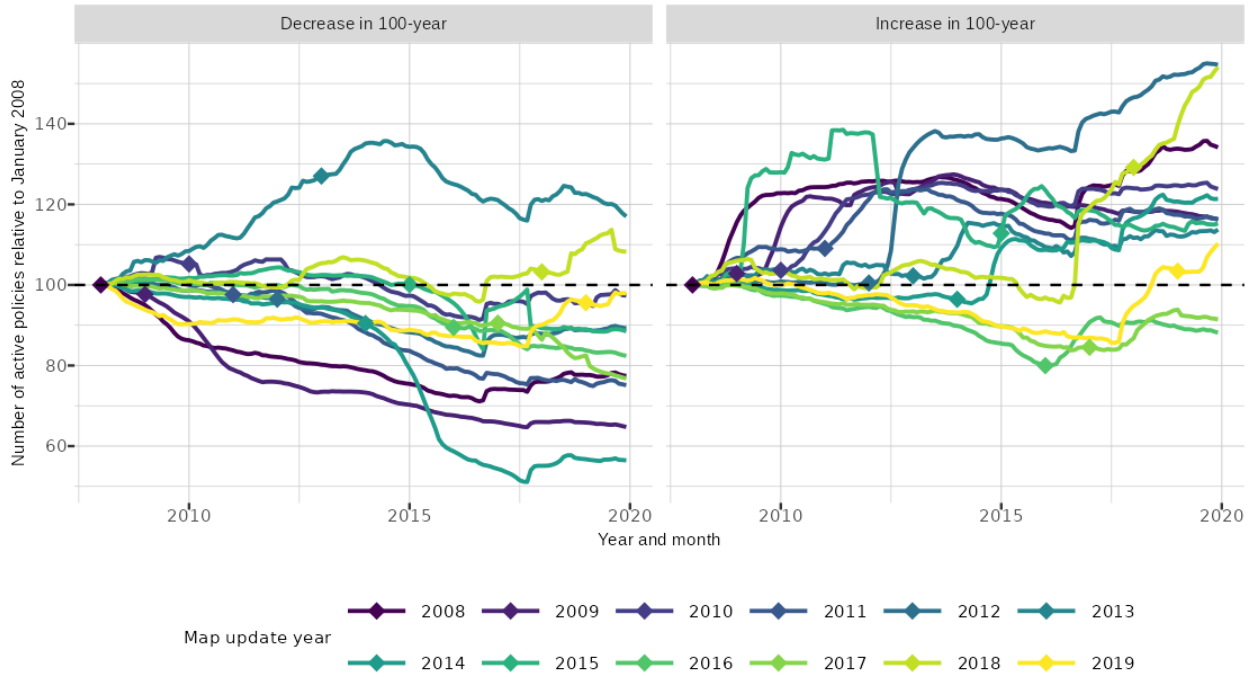
Whatever the year in which the map update occurs, post treatment flood insurance take-up appears to decrease following map updates that rezone more properties outside of the 100-year floodplain (left facet, with the exception of 2013), and it increases after the updates that rezone more properties inside the 100-year floodplain (right facet). While only suggestive, these patterns indicate that map updates may have substantial impacts on quantities of insurance policies purchased, based on the direction of rezoning of properties relative to the 100-year floodplain. In addition, Figure 5 shows that most treatment-year cohorts experience slow declines in insurance take-up prior to treatment, as well as positive shocks to insurance take-up in 2016 and 2017 (likely attributable to the hurricane seasons in these years). Unit and time fixed effects will to some extent control for these shocks, and additional methods variation, and additional methods in Section 5 will control for the potential spatial correlations in the impacts of hurricanes.

4.2 Heterogeneity-robust event studies

Our empirical setting involves a large number of units (census tracts) that get treated (receive a digital map) in different time periods (year and month). In such contexts, where the treatment is *staggered*, recent econometric work highlights important pitfalls with the use of the

²⁰Figure 5 defines treatment cohorts by year for clarity only – the empirical analyses will leverage the full variation of treatment timing by defining treatment by year and month.

Figure 5: Raw evidence on the impact of floodplain rezoning on insurance take-up



Number of active flood insurance policies covering residential properties in each month relative to January, 2008 (baseline 100). Treatment cohorts are defined at the year-level for clarity, with diamonds marking the year of the map update for each treatment-year cohort. The two facets are based on whether the map updates rezoned more than 1% of properties outside (left) or inside (right) the 100-year floodplain at the census tract level.

two-way fixed effect estimators that is still standard in the applied literature. In particular, de Chaisemartin and D’Haultfoeuille (2020) demonstrate that if the treatment effect is heterogeneous between *cohorts* (group of units treated simultaneously), then the two-way fixed effect estimator does not recover the Average Treatment Effect on the Treated (ATT).²¹ Recent applied work shows such pitfalls make two-way fixed effect estimators extremely sensitive to minor specifications of the regression model (Weill et al., 2021).

In our context, treatment effects are likely to be heterogeneous between cohorts. Section C.4 in the appendix shows that census tracts treated earlier are different from tracts treated later (potential heterogeneity due to *selection*), and tracts treated later received an updated map that tended to rezone more properties outside of the 100-year floodplain (potential heterogeneity

²¹This is true even when the identifying parallel trends assumption holds. Goodman-Bacon (2021) provides an intuitive decomposition of the two-way fixed effect estimator as a weighted average of all possible two-by-two diff-in-diff estimators. Sun and Abraham (2021) show that even the more flexible event-study specification used to estimate dynamic effects is not robust to heterogeneity between cohorts, and they provide a general formula for the bias that arises in various specifications. They further highlight that the estimated effects on leads and lags of the treatment can be contaminated by other time periods *even if* treatment effect is homogeneous.

due to varying *implementation*). In addition, the impact of a new map on insurance take-up will *mechanically* be dynamic, because insurance policies are usually purchased for an entire year.²² Purchases of policies are also not uniform within a year, with more purchases during the summer months. These aspects lead me to consider heterogeneity-robust estimators of the dynamic treatment effects.

I closely follow the approach proposed by Callaway and Sant’Anna (2021) and estimate a distinct treatment effect for each cohort and period relative to treatment. To understand how within-tract rezoning impacts the demand for insurance, I separate the models for treated tracts based on net rezoning inside or outside the 100-year floodplain at the tract level (“treatment intensity”).²³ Formally, we are interested in the Average Treatment effects on the Treated (ATTs) for each group g , time period t , and net map rezoning at the tract level x :

$$ATT(g, t, x) = \mathbb{E}(Y_t(g, x) - Y_t(0, x) | G_g = 1, X = x) \quad (1)$$

where $Y_t(g, x)$ denotes the potential outcome of units in group x at time t , *if* they were to become treated at time g . $Y_t(0, x)$ denotes the potential outcome of these same units had they not received treatment. For each unit we only observe the *realized* outcome. To identify these effects from observable data, we need to assume (i) no treatment anticipation, as well as (ii) conditional parallel trends in potential outcomes between the treated units and the group of control units.

The parallel trend assumption requires the evolution of insurance take-up in tracts that receive a map earlier to be similar to that in “control” tracts in the absence of treatment. In a standard application of the Callaway and Sant’Anna (2021) estimator, these control units are either all those “not yet treated” at the time of the cohort considered, or all those “never treated” by the end of the sample. In our context, one might be worried that potential outcomes in tracts where the map update removes properties from the 100-year floodplain might be on a different trend than those in tracts where the update rezones properties inside the 100-year floodplain. To enhance the credibility of the identification strategy, in my preferred specifications, the ‘control’ tracts are those that have not yet been treated *but* will eventually receive a digital

²²For instance, consider a homeowner who purchased flood insurance exactly one month before the new digital flood map is released, and the new flood map induces her to stop buying flood insurance. Because canceling a flood insurance policy is costly, this homeowner may “passively decide” to remain insured for up to 11 months after the new flood map was released, and then wait until the end of the one year term to decide to not renew her policy. In this scenario, the individual decision of interest is the one to not renew the policy.

²³Few properties are zoned inside or outside the 500-year floodplain without a simultaneous change in the 100-year floodplain. For instance, most properties zoned inside the 500-year floodplain used to be zoned inside the 100-year floodplain. When estimating the effect of map updates on properties that are zoned in the 500-year floodplain without an associated change in the 100-year floodplain, estimates are small and not statistically significant.

flood map *with a similar 100-year floodplain rezoning*.²⁴ The assumptions of no anticipation and conditional parallel trends allow us to rewrite the ATT as

$$ATT(g, t, x) = \mathbb{E}(Y_t - Y_{g-1})|G_g = 1, X = x) - \mathbb{E}(Y_t - Y_{g-1})|D_t = 0, X = x) \quad (2)$$

where D_{it} is a binary variable equal to one if unit i is treated in period t and equal to zero otherwise. This expression is similar to the standard 2 periods, 2 groups difference-in-differences estimators, which evaluates the evolution of potential outcomes in group g and adjusts it with the evolution of potential outcomes in the control group. To estimate it from the data, the expectations are replaced with sample averages:

$$\widehat{ATT}(g, t, x) = (\bar{Y}_{g,t,x} - \bar{Y}_{g,g-1,x}) - (\bar{Y}_{0,t,x} - \bar{Y}_{0,g-1,x}) \quad (3)$$

where $\bar{Y}_{g,t,x} = \frac{1}{N_{gtx}} \sum_{i:G_i=g, X_i=x} Y_{i,t,x}$.²⁵

I construct the groups x to reflect treatment intensity based on how the map updates impacted rezoning into the 100-year floodplain at the tract-level. The choice of x here is arbitrary, and implies a trade-off between the degree of heterogeneity we can estimate and the quality of the treatment–control comparisons on the one hand, and the feasibility of the estimation procedure on the other hand. To see this, consider one extreme, where we seek to estimate the ATT for all tracts where exactly .75% of residential properties were removed from the 100-year floodplain. To ensure the estimation is robust to potentially heterogeneous treatment effects between cohorts, we would have to also split the estimation between units that receive treatment in different year/month, and compare each sub-sample to tracts that receive a map update *later* with exactly .75% of properties rezoned inside the 100-year floodplain. While this would allow to flexibly assess heterogeneity for varying intensity of rezoning, in many cases this approach is not feasible, because too few observations are in the estimation samples. I focus on the role of the 100-year floodplain rezoning, where x is one of {Increase in 100-year floodplain, Decrease in 100-year floodplain, No change in 100-year floodplain}. “Increase” or “decrease” in the 100-year floodplain denote tracts where more than 1% of residential properties are rezoned inside or outside of the 100-year floodplain on net respectively, while “No change in 100-year floodplain” denotes tracts with properties in the 100-year floodplain, but where the map update didn’t rezone any property relative to the 100-year floodplain. In the appendix I

²⁴For instance, consider the cohort of census tracts that receive a map in February 2012 which removes a share x of properties from the 100-year floodplain. To estimate the ATTs for this group, I use as control units all the census tracts that have not received a digital flood map as of February 2012 but that will receive one removing a share x of properties from the 100-year floodplain before the end of the sample period (2020).

²⁵Formally, when the never-treated units are part of the control group we have

$$\bar{Y}_{0,t,x} = \frac{1}{N_{0tx}} \sum_{i:G_i=g, \{X_i=x\} \cup \{X_i=\emptyset\}} Y_{i,t,x}$$

present results using a 3% rezoning cutoff (shown in Figure D.9 and Figure D.10 – as expected, treatment effects tend to increase with the magnitude of the rezoning). Section 5 is devoted to fully relaxing the arbitrary choice of x by estimating tract-specific individual treatment effects.

I again follow Callaway and Sant’Anna (2021) and aggregate cohort-specific ATTs into ATTs by length of exposure using cohort sample-size weights:

$$\theta(e, x) = \sum_g \mathbf{1}\{g + e \leq T\} \cdot P(G = g \cap X = x | G + e \leq T) \cdot ATT(g, g + e, x) \quad (4)$$

where e is the exposure time ($e = t - g$), g is an index for the cohort of flood map update, G is the time period that a unit is first treated, and T is the last time period for which ATTs are identified. The estimator is obtained by replacing ATT with its sample analogue from equation 3, \widehat{ATT} .

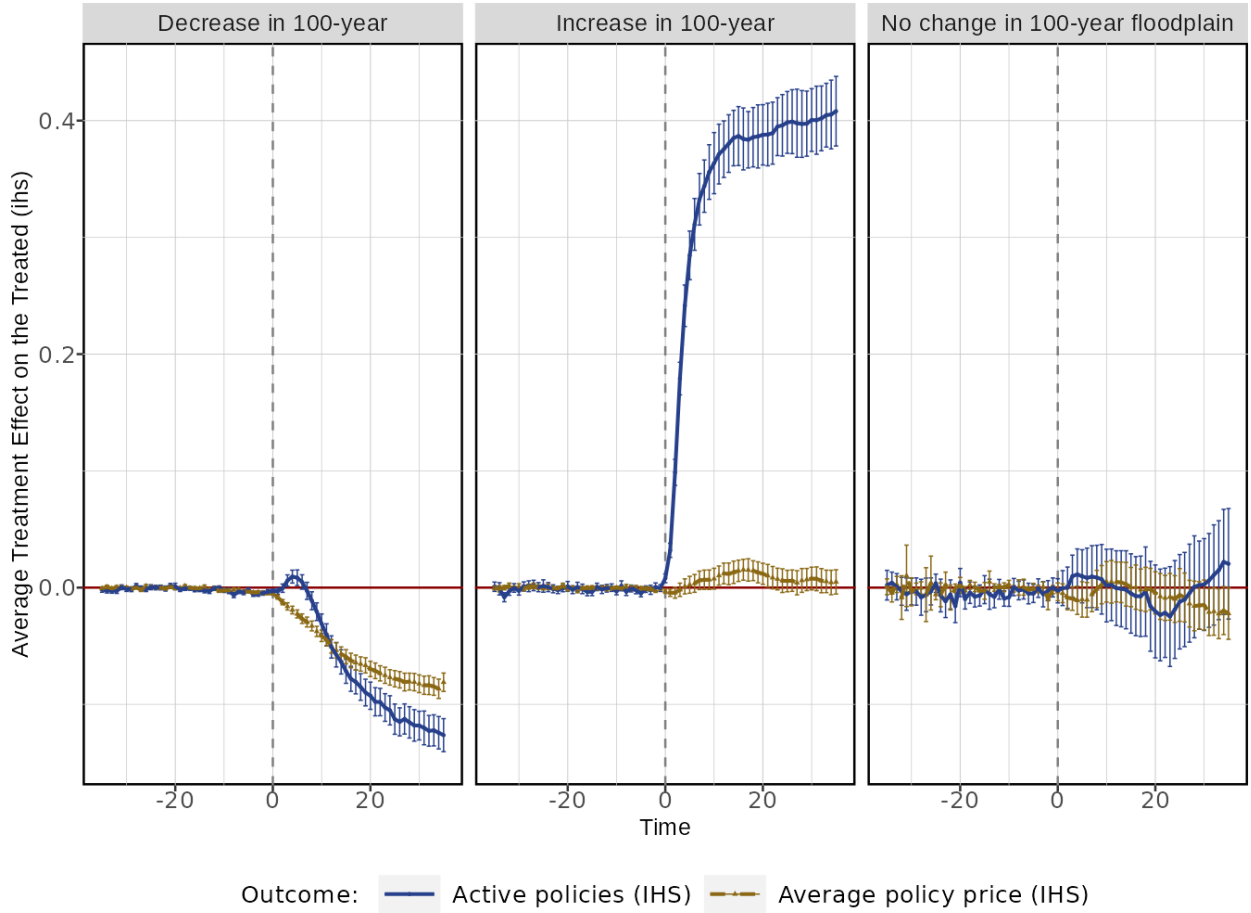
4.3 Event study estimates of the impacts of flood map updates

Figure 6 presents the ATT estimates by length of exposure, aggregated over all 117 cohorts following equation 4. The three panels estimate effects separately based on whether the map update rezoned properties outside of the 100-year floodplain (first panel), inside (second panel), and where the map update did not cause any rezoning inside or outside of the 100-year floodplain (third panel). The figure shows estimates for two different outcomes: the number of active flood insurance policies in the census tract (solid blue line) as well as the average price of the insurance policy (dashed gold line), both transformed using the inverse hyperbolic sine. For all three treated samples and both outcome variables, we note the absence of pre-trends: the estimated average effects of receiving an updated flood map on take-up and average prices is not significantly different from zero prior to the treatment date, increasing our confidence in the identifying assumption.

Post-treatment, effects are large and heterogeneous between samples, strongly suggesting that rezoning relative to the 100-year floodplain is a decisive driver of insurance take-up. Focusing first on the number of active policies, we note that among census tracts where the map modernization decreased the number of properties located inside the 100-year floodplain (first panel), the map update led to an average decline in the number of policies of about 10% after 2 years. In contrast, the map modernization caused an average increase of almost 40% in areas where the digital map rezoned properties inside the 100-year floodplain (second panel). In neighborhoods where the new maps did not rezone properties (third panel), effects are small and not statistically different from zero.

In the appendix, I show in Figure D.7 that these effects (obtained aggregating models 3 over

Figure 6: Aggregated event study estimates of the impacts of map updates



ATT estimates obtained using Callaway and Sant’Anna (2021) regressions. The outcome variables (active insurance policies in blue, average policy prices in orange) are transformed using the Inverse Hyperbolic Sine (IHS). Each facet represents average treatment effects for a different treated group, using treated census tracts where the flood map update increased, decreased, or did not change the number of properties zoned inside the 100-year floodplain (first, second and third facet respectively). The three control groups comprise not-yet-treated census tracts that later receive a flood map with a similar floodplain rezoning direction as the corresponding treated group. Error bars represent 95% confidence intervals using the multiplier bootstrap.

all treatment cohorts) also exist *within* each cohort, regardless of the treatment date. For almost all cohorts, the estimated impacts of map updates on insurance take-up are positive for tracts where the map updates rezoned properties inside the 100-year floodplain, and negative in tracts where the map rezoned properties outside of the 100-year floodplain. Such patterns confirm the role the 100-year floodplain as a major driver of heterogeneity. In the appendix I also present alternative models showing the robustness of the results presented above. Figure D.11 shows these effects are not driven by new residential constructions: the estimates are extremely similar if we only retain insurance policies covering properties that were built prior to 2008 in

the analysis.²⁶

Figure 6 shows the observed impacts on quantities are not driven by changes in prices. The impacts of flood map updates on average prices follow the same patterns as the effects on the quantities, although the effects are more modest. Due to the aggregate nature of the data, interpretation of these effects is difficult. At the census tract level, average price is an equilibrium outcome: the effects measured here include both a direct floodplain effect (everything else equal, being in the 100-year floodplain increases the price of the insurance policy), as well as a composition effect (the pool of properties purchasing insurance might change with the new flood map). I highlight these effects in the next section by further separating policies that are purchased inside and outside the 100-year floodplain.

Overall, these results rule out any large effects of the pure digitization of paper-based maps, suggesting that access to digital information was not a barrier to insurance take-up. On the other hand, the large impacts of the 100-year floodplain rezoning on insurance take-up could be due to individuals revising their beliefs about flood risk following the new map, or because of the insurance mandate requirement inside the 100-year floodplain. While I can't perfectly decompose the take-up effect between these channels, the analyses below suggest that both mechanisms are at play.

4.4 Within-census tracts spillover effects of updated flood maps

Flood maps present a discrete depiction of risk: the location of any property is either high-risk (100-year floodplain), low risk (500-year floodplain), or “minimal risk” (outside of the floodplains). Although flood maps provide some additional information within each zone, the categorical definition of flood risk has no underlying physical basis: absent specific structures, flood risk is in general continuous in space. This implies that homeowners might rationally form beliefs about their risk exposure beyond the discrete floodplain classification of their property. For instance, a homeowner could be located outside of the 100-year floodplain both before and after the map update, but might revise her beliefs about risk if the map update rezoned adjacent properties towards the 100-year floodplain.

The main challenge in testing for these “spillover effects” is that I cannot match insurance policies to geolocalized properties. However, the insurance data provide information on the floodplain in which the insured property is located *at the time of the policy purchase*. This last point bears emphasizing as it means that we can track, at any time, how many insurance

²⁶These results are also robust to changes in the transformation variable (log versus inverse hyperbolic spline), changes in the estimation sample (excluding all tracts with less than 5 active insurance policies at all times), and controls for the occurrence of presidential disaster declarations in the past 2 or 3 years (available upon request).

policies are covering properties inside and outside of the 100-year floodplain, *within* census tract. However, we cannot interpret changes in the number of policies in a particular floodplain as changes in the take-up of properties that previously purchased a policy in this floodplain, since the *definition* of the floodplain itself is changing due to the map update.

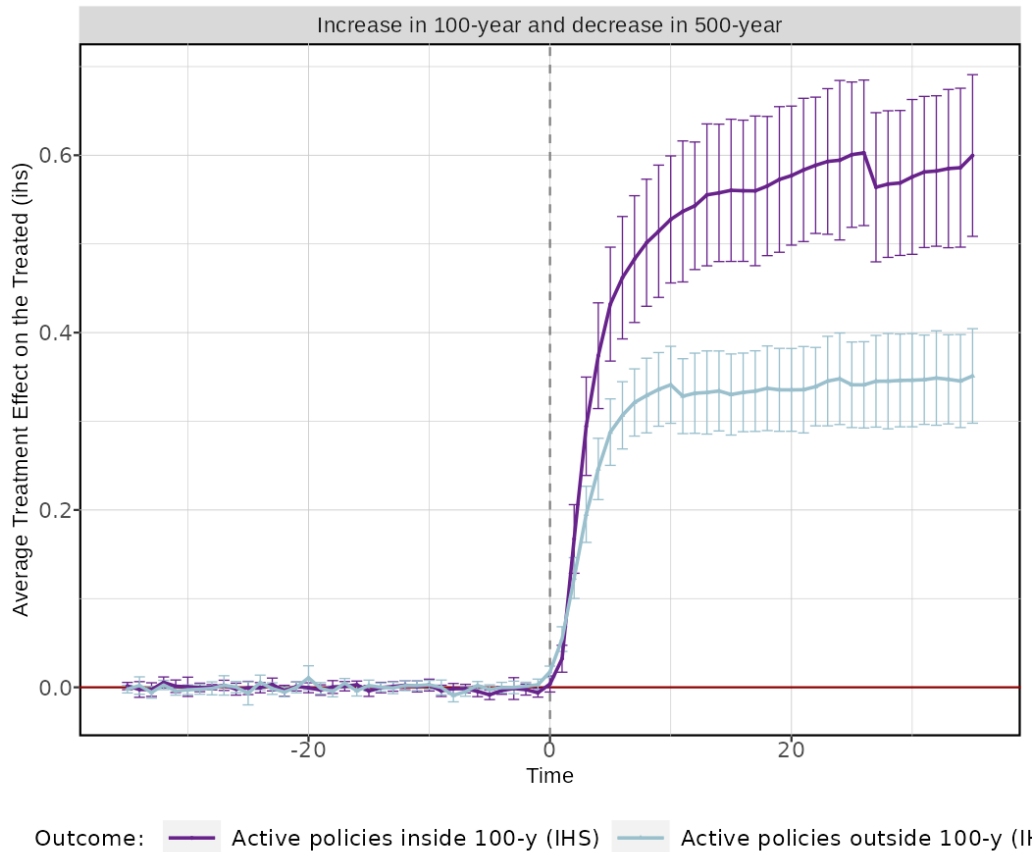
Figure 7 presents heterogeneity-robust event-study estimates for the number of active insurance policies as outcome variables, this time differentiating the number of insurance policies *by floodplains within tracts* (recall that the definition of the floodplains changes at the time of treatment, since the floodplain classification comes from the insurance data). The models are only estimated for treated census tracts where the map updates rezoned properties inside the 100-year floodplain (these treated census tracts in Figure 7 were previously part of the first panel in Figure 6), and further on tracts where properties were removed (many properties get rezoned from the 500-year floodplain to the 100-year floodplain).

I find that updates increase take-up from properties located inside the 100-year floodplain (purple line), which we would mechanically expect *even if* there is no change in tract-level take-up – this effect could be entirely due to the rezoning of properties. However, I also find these map updates also increased demand from properties located *outside* of the 100-year floodplain (light blue). This is surprising: among the treated census tracts, multiple properties were rezoned from the 500-year to the 100-year floodplain and were likely carrying flood insurance – rezoning inside the 100-year floodplain would therefore mechanically cause the number of active policies covering properties outside of the 100-year floodplain to *decrease* (recall that the number of insurance policies within each floodplain is derived from the insurance data itself and records the floodplain at the time of the purchase of the insurance policy). This means that the increase in insurance demand from properties located outside of the 100-year floodplain *more than compensates* this composition effect.²⁷

Overall, these results demonstrate that rezoning properties inside the 100-year floodplain not only has a direct impact on insurance demand (by increasing the demand for insurance from properties located within the 100-year floodplain, as previously seen in the first panel of Figure 6), but also has a large indirect impact on properties located outside of the 100-year floodplain. While I cannot completely rule out the potential role played by the insurance mandate requirement inside the 100-year floodplain, these indirect or “spatial spillover” effects outside of the 100-year floodplain strongly suggest that a substantial portion of the impacts of updated flood maps on insurance demand is driven by changing beliefs about flood risks.

²⁷Looking at tracts with a decrease in the 100-year and increase in the 500-year is complicated by coding errors in the measurement data, yet does not reveal any spillover effects (take-up decrease in the 100-year and increase outside of the floodplain).


Figure 7: Aggregated event study estimates of the impacts of flood map updates on flood insurance take-up at the census tract and flood zone-level



ATT estimates obtained using Callaway and Sant'Anna (2021), focusing on treated census tracts where the flood map update increased the number of properties in the 100-year floodplain while simultaneously rezoning properties outside of the 500-year floodplain. The outcome is measured at the *census tract / flood zone-level*, where the flood zone designations are derived from the insurance policy data. As such, they provide the number of policies that are active within each zone at the time when the policy is purchased. The control groups comprise census tracts that have not yet received a digital flood map at the time of treatment, but will receive one later with the same direction of 100-year floodplain rezoning. The outcome variables are transformed using the Inverse Hyperbolic Sine (IHS). Error bars represent 95% confidence intervals using the multiplier bootstrap.

4.5 Event-study heterogeneity

Figure 8: Event-study, heterogeneity analysis



ATT estimates obtained using Callaway and Sant’Anna (2021) regressions. The outcome variable is the number of active insurance policies in the tract, transformed using the Inverse Hyperbolic Sine (IHS). Rows represent average treatment effects for different treatment groups, using treated census tracts where the flood map update decreased (top row) or increased (bottom row) the number of residential properties in the 100-year floodplain by more than 1% relative to the total number of residential properties in the census tract. Vertical facets focus on different heterogeneity variables. The control groups comprise not-yet-treated census tracts that later receive a flood map with a similar floodplain rezoning direction as the treated groups and that are within the same heterogeneity quartile of the variable being investigated. Error bars represent 95% confidence intervals using the multiplier bootstrap.

Figure 8 continues the investigation of treatment effect heterogeneity within the event-study framework, focusing on four neighborhood-level characteristics. In the appendix, Figure D.12 present additional heterogeneity results based on race, while Figure D.13 focuses on correctness of the map change based on the FSF model, disaster experience, and disclosure laws. The control groups are further restricted to be in *same quartile* as the treatment units. For instance, estimates in the top left panel present heterogeneity based on the neighborhood-level share of Black or African American households for neighborhoods where the map update rezone more

than 1% of properties outside of the 100-year floodplain. The yellow coefficients are obtained by restricting the estimation sample to treated and control units within the top quartile of the neighborhood-African American share only.

Focusing on neighborhoods where the map updates rezoned properties outside of the 100-year floodplain (top row), Figure 8 reveals that effects are larger in neighborhoods with a higher percentage of Black residents (column 1) and Hispanic residents (column 2), poorer neighborhoods (column 3). Rezoning inside the 100-year floodplain (bottom row) caused greater increase in take-up in more Hispanic neighborhoods and areas with greater properties owned with a mortgage.

The stronger effects of map updates in areas with more properties with a mortgage provides suggestive evidence that the mandatory purchase requirement in the 100-year floodplain might be binding, at least for some households. The larger effects estimated in African American and Hispanic neighborhoods, on the other hand, could be due to greater price elasticity due to reduced budget sets. This is consistent with results in column 3 of Figure 8 and with Bradt et al. (2021), which shows that insurance uptake monotonously increases with income. This suggests that at lower level of wealth, an increase in the price of insurance (stemming from a rezoning inside the 100-year floodplain) is more likely to lead a household to drop their insurance policy. It could also be due by the disproportional share of removals from the 100-year floodplain that occurred in these areas, identified in Section 3. Tackling these questions requires to control for the *treatment intensity* at the census tract level.

5 Synthetic controls and distributional impacts

The share of properties rezoned inside and outside of the 100-year floodplain varied between neighborhoods. Yet, accounting for disparities in the type of information provided by the map updates cannot be done performed flexibly within the event-study framework, due to the small sample sizes issues discussed above. In addition, one might question the parallel trend assumption imposed above. Three features could undermine this assumption in our context. First, although there is little concern about selection into treatment (given that all tracts are scheduled to receive a map at some point), as discussed above there is some evidence of selection into treatment *timing*: units treated later appear to be exposed to more flood risk than units treated earlier. While this is partially addressed by the event studies, one might worry about diverging potential outcome trends between early and later treated units. Second, the policy environment changed during our observation window. In particular, the Biggert-Waters Act of 2012 progressively phased out insurance subsidies for a subset of properties that were

constructed prior to the community’s first flood map (called “pre-FIRM properties”) and are now located in the 100-year floodplain. Failing to account for these differential trends could attribute some impacts of the BW Act to the map updates. Third, previous work found that the occurrence of a disaster is itself a strong determinant of flood insurance take-up (Gallagher, 2014). When estimating the impacts of flood map updates, one might therefore wish to account for the occurrence of disasters.

One way to dive deeper into heterogeneous treatment effects and to construct better control groups – by accounting for the share of pre-FIRM properties, the share of properties zoned in the 100-year floodplain before the map update, and the potentially non-linear effects of disasters – would be to further split the samples of census tracts along these observable characteristics of interests. Unfortunately, this quickly becomes infeasible when using heterogeneity-robust estimators of the ATT such as the Callaway and Sant’Anna (2021) estimator, as the number of units in each $\{g, t, x\}$ cell quickly becomes too small – especially as the dimension of x (the number of characteristics along which we wish to assess heterogeneity of the treatment effects) increases.²⁸ To escape this Catch-22, I now turn to tract-specific synthetic controls.

5.1 Estimating census tract-specific synthetic controls

Synthetic controls were originally developed to estimate treatment effects when only a small number of distinct units are available, such as states or countries, and when only one unit receives treatment (Abadie and Gardeazabal, 2003; Abadie et al., 2010). Because directly comparing the outcome of the treated unit to the outcomes of the control units generally provides a poor estimation of the treatment effect, a “synthetic control” is first constructed by taking a weighted average of the different untreated units, where the weights are chosen so that the evolution of the synthetic control’s outcome pre-treatment matches as closely as possible the evolution of the outcome of the treated unit.

I estimate unit-specific synthetic controls for *each* census tract observed for at least 12 months pre-treatment and 24 months post-treatment. While most applications construct synthetic controls by only matching on the outcome variable, matching on select covariates can improve the credibility of synthetic control estimates, even if from an estimation perspective doing so can only decrease pre-treatment fit Abadie (2021).²⁹ I construct synthetic controls by

²⁸A “standard” approach to heterogeneity would be to run two-way fixed effect regressions and interact treatment with census tract-specific covariates. As discussed above, this would deliver biased and inconsistent treatment effect estimates due to both between-cohort heterogeneity in the treatment effect and treatment effect dynamics (which was the entire rationale for using the Callaway and Sant’Anna (2021) estimator in the first place).

²⁹Abadie (2021) emphasizes this aspect: “Part of the literature on synthetic controls emphasizes estimators that depend only on pre-intervention outcomes and ignore the information of other predictors [...]. This reliance

matching on pre-treatment values of the outcome variable (the number of active flood insurance policies in the census tract), as well as the share of pre-FIRM properties and the share of policies covering properties in the 100-year floodplain. Doing so will create control units that are even more likely to similarly respond to the BW Act and other idiosyncratic time shocks, matching on these two covariates generate more credible synthetic controls.

An important choice that remains concerns the definition of the “donor pool,” i.e., the group of control units that are considered when estimating the controls’ weights. In a typical synthetic control application, all untreated units are part of the donor pool. However, in our case, flooding and other spatially correlated disasters cause large increases in insurance take-up, not only in areas directly impacted by flooding, but also in neighboring areas that are located in the same television-network market (Gallagher, 2014). During our observation window, flooding occurs in multiple states, which might cause potential outcomes between treated units and their counterparts to diverge. To mitigate this issue, I estimate *spatially clustered* synthetic controls: for each treated tract, the “donor pool” contains units that do not receive treatment within our observation period and that are located within the *same* FEMA region.³⁰

Finally, to improve the quality of pre-treatment fit, I use the augmented synthetic control method of Ben-Michael et al. (2021). This approach uses an explicit outcome model to remove potential bias in the original synthetic control estimates, detailed in the appendix (Section E.1).

Moving from the estimation of average treatment effects to neighborhood-specific treatment effects is not without costs. One major drawback is that synthetic controls can perform poorly when the treatment effect is small relative to the noise in the outcome variable (Abadie, 2021). Unlike with difference-in-differences estimators, the statistical power of synthetic control estimators does not increase with sample size. In our context, this means that we are more likely to credibly estimate treatment effects where flood map updates lead to large changes in insurance take-up, and for census tracts where there is a large number of active insurance policies.³¹ I therefore estimate synthetic controls in census tracts that have at least 20 active policies at all times.³² Finally, quantifying the uncertainty around synthetic control estimates is still an active area of research. To the best of my knowledge, there is no agreement on how to account

on pre-intervention outcomes only, while adopted in many cases for technical or expositional convenience, may create the mistaken impression that other predictors play a minor role in synthetic control estimators.”

³⁰One possibility is to form the donor pool using units that are not-yet-treated, instead of units that are never-treated. This is appealing, since not-yet-treated units are more likely to resemble already-treated units. However, this makes it hard to interpret changes in the treatment effects over time, since some of these changes are driven by changes in the composition of the donor pool.

³¹Recent extensions of synthetic controls consider Average Treatment Effect estimation when several units receive treatment at different times (Xu, 2017; Robbins et al., 2017; Ben-Michael et al., 2019).

³²Due to the spatial concentration of the demand for flood insurance, excluding census tracts with less than 20 active insurance policies at all times lead to a negligible reduction in the aggregate number of active policies considered for the synthetic control estimation.

for multiple hypothesis testing when evaluating the uncertainty around the treatment effects of thousands of synthetic controls. Despite these shortcomings, I use the recently developed jackknife+ procedure from Barber et al. (2021) to construct confidence intervals.

5.2 Synthetic controls results and the distribution of treatment effects

This section presents the (ridge-augmented) synthetic control estimates for each of the 4795 census tracts that have at least 20 active policies at all times, for which I can identify changes in the 100-year floodplain, and that received a digital flood map within the period 2009-2017.

To assess pre-treatment fit of the synthetic control estimates, I compute the normalized root mean squared error (NRMSE) during the pre-treatment period: for 76 % of the estimated synthetic controls, the NRMSE is less than 0.025, meaning that synthetic controls and treated tracts deviate by less than 2.5% on the number of policies prior to treatment in most cases.³³ A histogram of the NRMSE is presented in Figure E.14 in the Appendix. In the following sections, I restrict the analysis to census tracts for which I estimated a synthetic control that provides a close match to the pre-treatment outcome, defined as an NRMSE lower than 0.05.³⁴ In the current application, the smallest and largest donor pools have 58 and 898 units, respectively. Figure E.15 shows that pre-treatment fit is not correlated with the number of units in the pools.

Figure E.18 in the appendix presents the full paths of dynamic synthetic control treatment effects for selected census tracts. While long-term treatment effects are most relevant from a public policy perspective, this requires strong assumptions and more data. For the remainder of this section I focus on the estimated treatment effects at 24 months post-treatment.³⁵

To investigate treatment effect heterogeneity and flexibly assess the role of rezoning inside and outside of the 100-year floodplain, I run the following regressions:

$$\widehat{TE}_i = \alpha + \beta Rezoning100year_i \cdot X_i + \epsilon_i \quad (5)$$

where \widehat{TE}_i is the synthetic control treatment effect estimate of the impact of the map update on the flood insurance take-up in census tract i , $Rezoning100year_i$ indicates the number of residential properties rezoned inside the 100-year floodplain with the map update (relative to the number of residential properties in the census tract), X_i is a fixed variable at the census tract level along which we wish to estimate heterogeneity. Note that because the left-hand side


³³It is smaller than .01 in 63 % of cases.

³⁴Most of the deviations between the treated units and their estimated synthetic control outcomes pre-treatment arise from the inclusion of covariates in the matching procedure: when omitting covariates, the weighting procedure generates weights that lead to NRMSE lower than 0.01 in 92% of cases.

³⁵Additional results show that effects are similar when focusing on treatment effects after 12 months or 30 months instead.

of regression 5 is estimated using synthetic controls, uncertainty surrounding the estimates need to be interpreted with caution. To the best of my knowledge, this “second stage” regression (where the first stage is the construction of the synthetic control estimates) has not been considered previously and its theoretical properties have not yet been derived. This is left for future developments.

Figure 9: Second-stage regression, marginal effects



Each dot represents a census tract-specific treatment effect estimate of the impact of the flood map update on flood insurance take-up 24 months post-treatment, using synthetic controls augmented by ridge regression. For each treated unit, the donor pool comprises never-treated census tracts within the same FEMA region. Large dots are significant treatment effects at the 10% level, using the jackknife+ procedure. Regression lines represent the marginal effects of a change in the number of properties rezoned inside the 100-year floodplain on the synthetic control treatment effect estimates, following regression 5. The marginal effects are estimated separately per quartiles.

Figure 9 show results for the entire estimation sample. Each dot represents a census tract-specific synthetic control estimate at 24 months post treatment. The x-axis depicts the change in the share of residential properties in the 100-year floodplain due to the map update ($Rezoning_{100year}_i$, with rezoning outside of the 100-year floodplain counted negatively), while the y-axis depicts the treatment effects TE_i expressed relative to the number of properties in

the tract. The treatment effects plotted are the same in all four vertical panels, but each panel focuses on a specific heterogeneity dimension, with different colors representing census tracts in different quartiles for this particular variable. The lines are marginal effects of increasing the share of properties rezoned in the 100-year floodplain for each of the quartile groups.

We first note a strong relationship between treatment intensity and floodplain rezoning: the more properties rezoned inside the 100-year floodplain, the larger the treatment effect. The overall effect of rezoning inside the 100-year floodplain in regression 5 is 0.17, meaning that for every 5 properties rezoned relative to the 100-year floodplain, one is induced to change insurance status. The adjusted- R^2 of regression 5 (without heterogeneity variables) is 47%, i.e., the share of properties rezoned in and out of the 100-year floodplain explain more than half of the variation in the synthetic control estimates.

As with event study models, we find that rezoning outside of the 100-year floodplain decrease insurance take-up, whereas rezoning inside the 100-year floodplain increases take-up. The synthetic control estimates further reveal that larger 100-year floodplain rezoning cause larger treatment effects, and this relationship holds over most tracts, although treatment effects are noisier when the share of properties rezoned is small.

Neighborhoods with more properties with a mortgage (second panel of Figure 9) and neighborhoods with more Hispanic residents (fourth panel) experience greater map-induced declines in insurance take-up when rezoned outside the 100-year floodplain, even when flexibly accounting for the magnitude of rezoning. In contrast, poorer neighborhoods and neighborhoods with more Black residents (first and third panels) do not appear to respond more strongly to the map update *conditional on the intensity floodplain rezoning*. Map updates did cause greater declines in insurance take-up in neighborhoods with more Black residents, but this is predominantly caused by the substantial rezoning outside of the 100-year floodplain that occurred in these areas (see Figure E.20).³⁶

On the other hand, Figure 9 reveals that income is unlikely to drive the distributional effects of the map updates. The marginal effects of rezoning properties inside the 100-year floodplain *increases* (in absolute value) with income, suggesting that higher-income households are *more* sensitive to the rezoning of properties inside and outside of the 100-year floodplain. This pattern is inconsistent with decreasing absolute risk aversion, but could be rationalized by differential access to information between poorer and wealthier neighborhoods.³⁷

Households may hold private information about their *true* risk exposure. Under this scenario, households could selectively respond to flood map updates that are “correct,” and not

³⁶In Figure E.21 appendix, I extend the analysis to other racial groups.

³⁷Outreach efforts regarding the updated flood maps may vary between neighborhoods, but despite multiple requests I was not able to obtain data on the magnitude of outreach efforts in different communities.

substantially change their demand for insurance following an incorrect update. Such differentiated response to map updates would mitigate the welfare costs of providing incorrect information. While crucial to welfare analysis, assessing the role of private information about *true risk* is challenging. True flood risk is unknown, both to the households and the econometrician. In Figure E.22, I proceed by assuming that the First Street Foundation model provides a good proxy for true flood risk. I do not find meaningful differences in the effects of map updates on insurance take-up based on whether the majority of rezoning were correct within a census tract, suggesting that official risk information is the main source of information about risk.³⁸

Overall, these results show that using local control units and matching on ex-ante measures of risk and floodplain extents generate treatment effect estimates in striking agreement with those of Section 4: rezoning inside the 100-year floodplain increase take-up, while rezoning outside reduces the demand for insurance. Due to the more numerous rezoning outside of the floodplain, map updates decreased insurance take-up on net. These declines in insurance coverage were not uniform across neighborhoods, but concentrated in neighborhoods with more Black and Hispanic residents. Individuals appear to respond to the official flood map updates whether or not the information conveyed in these maps align with the best available science. As such, their impacts on consumer welfare could be negative.

6 A model of flood insurance demand

The impact of flood information on welfare crucially depends on the *correctness* of the information. This section provides back-of-the-envelope estimates of the impacts of map updates on welfare by outlining a static model of insurance demand that isolates the correctness of the risk information, and assuming the First Street Foundation provides a measure of *true* risk. While restrictive, these assumptions allow for baseline estimates in the absence of any market distortions other than incorrect information. Other considerations of primary importance for welfare estimates are consumer's risk preferences, whether the insurance mandate requirement is binding, potential moral hazard due to post-disaster funding, adverse selection due to privately known and non-contractible costs, limited foresight, and insurance loading costs. Some of these aspects are carefully studied in recent work (Wagner, 2022; Mulder, 2022).

³⁸Figure E.22 further shows that neighborhoods hit by a disaster prior to the map update appear to respond more to the official flood map, consistent with it being the main source of information. On the other hand, map updates cause larger effects in areas where flood information disclosure is *not* required by home sellers, consistent with map updates providing more information in this context)

6.1 Model outline

A homeowner ω decides whether or not to purchase insurance, with utility $u_{\omega 1}$ if she purchases and $u_{\omega 0}$ if she does not. Let $\alpha(\omega)$ denote her (absolute) risk aversion. The homeowner will purchase insurance as long as her risk aversion is above a threshold value k_{ω} : $u_{\omega 1} > u_{\omega 0} \iff \alpha(\omega) > k_{\omega}$. The probability of ω purchasing insurance is then $p(\alpha(\omega) > k_{\omega}) \iff 1 - F_c(k_{\omega})$ for some census tract-specific cumulative distribution function F_c .

The expected social surplus obtained from homeowner ω before the map update is

$$\begin{aligned} W_{\omega,pre} &= p(\alpha(\omega) > k_{\omega,pre}) \cdot (CS_{\omega,pre} + PS_{\omega,pre}) \\ &= p(\alpha(\omega) > k_{\omega,pre}) \cdot (WTP_{\omega} - r_{\omega,pre} + r_{\omega,pre} - b_{\omega}) \\ &= p(\alpha(\omega) > k_{\omega,pre}) \cdot (WTP_{\omega} - b_{\omega}) \end{aligned} \quad (6)$$

where $CS_{\omega,pre}$ and $PS_{\omega,pre}$ are the surpluses from the insurance contract that accrue prior to the map update to the homeowner and insurer respectively,³⁹ b_{ω} is the expected (annual) cost due to flooding, $r_{\omega,pre}$ is the price of the contract prior to the map update, and WTP_{ω} is the homeowner's underlying willingness-to-pay for flood insurance. Intuitively, the welfare that arises from homeowner ω is 0 if she does not purchase insurance, while it is the difference between her willingness-to-pay and the expected costs of supplying the contract if she purchases insurance. Importantly, we consider her *accurately informed* willingness-to-pay, that is, the one that would arise *if* she had access to correct flood risk information.

Integrating over homeowners, the expected welfare at the census tract level before the map update is computed as

$$\begin{aligned} \mathbb{E}[W_{c,pre}] &= \int_{\omega} (WTP_{\omega} - b_{\omega}) \cdot \mathbb{1}(\alpha(\omega) > k_{\omega,pre}) \cdot dF_c \\ &= \int_{\alpha(\omega) > k_{\omega,pre}}^{+\infty} (WTP_{\omega} - b_{\omega}) \cdot dF_c \end{aligned} \quad (7)$$

and similarly after the map update for $\mathbb{E}[W_{c,post}]$. Finally, we sum over tracts to get the aggregate welfare change: $\Delta W = \sum_c \mathbb{E}[W_{c,post}] - \mathbb{E}[W_{c,pre}]$

In this framework, all social welfare changes attributable to the map updates arise due to changes in the *probabilities* of households purchasing flood insurance. Map updates only impact the support of the integrals by changing the cut-off values $k_{\omega,\cdot}$, while the willingness-to-pay for insurance and the expected annual costs of flooding remain constant through the map update.

³⁹These welfare definitions are net of any administrative costs of providing the contract.

6.2 Model calibration

Estimating the model requires several additional assumptions. I assume that individuals maximize expected utility, and that their preferences are captured by a Constant Absolute Risk Aversion utility function. This assumption is reasonable when changes in premiums are small relative to individuals' income, and this allows to ignore income effects associated with price changes (Einav et al., 2010). The functional form allows to numerically compute the cut-off values k_ω above which the homeowner purchases insurance and provide an expression for the willingness-to-pay of homeowner ω .

Households perceive their probability of flooding from the (official) FEMA flood maps, consistent with the results in the previous sections. When there is a map update, this potentially changes the cut-off value above which they purchase insurance, from $k_{\omega,pre}$ to $k_{\omega,post}$. I further assume that the First Street Foundation model provides a good proxy of the probability of inundation depths and the property-specific flood damages at these depths.⁴⁰ To recover the price of the insurance contract for all households, I use neighborhood-, floodplain-, and time-specific premium averages from the insurance data. I also assume that risk aversion parameters follow a Fréchet distribution within each census tract. Finally, the calibration focuses on homeowners directly impacted by map updates, and ignores the spatial spillover effects. I also compute welfare assuming that the mandatory purchase requirement constraint is not binding. Both assumptions are rejected by the analysis above, but they allow to provide transparent back-of-the-envelope estimates of the impacts of map updates nation-wide. These assumptions can be relaxed by assuming a specific shape for spatial spillover effects and the distribution of homeowners constrained by the insurance mandate within each census tract. Additional details of the estimation procedure are provided in the appendix (Section F.1).

Overall, 17.4 million unique properties have a non zero perceived flooding probability and contribute to at least one of the welfare measures.⁴¹ Computing expected welfare at the property level can be done sequentially and independently for all properties: first compute the homeowner-specific threshold values $k_{\omega,\cdot}$, then integrate out the risk aversion parameters.

⁴⁰The assumption of correctly estimated flood damages in the FSF model is strong, as it requires the depth-damage functions to be correctly specified. These damage estimates lack the validation of inundation depth estimates due to issues arising in the insurance claims data (low take-up of insurance and large adverse selection in particular).

⁴¹Welfare estimates focused on communities for which I can assign a probability of flooding pre-map update comprise 12 million unique properties.

6.3 Structural estimates

Figure 10 presents several welfare metrics under a range of assumptions. Figure 10.A presents the empirical cumulative distribution functions of relative changes in social welfare over census tracts between the paper-based and digital maps, for three different expected values for the risk aversion parameters. A total of 20,276 census tracts experience a non-zero change in social welfare. The map updates decrease welfare in almost exactly half of these census tracts, while the other half experiences gains. These relative estimates are not sensitive to the choice of risk aversion parameters. Looking at this figure, a policy maker with distributional concerns might hope that the gains are concentrated in more disadvantaged neighborhoods, or that the welfare gains are large enough to potentially compensate the welfare losses that occur in other neighborhoods.⁴² Unfortunately, the next figure reveals this is not the case.

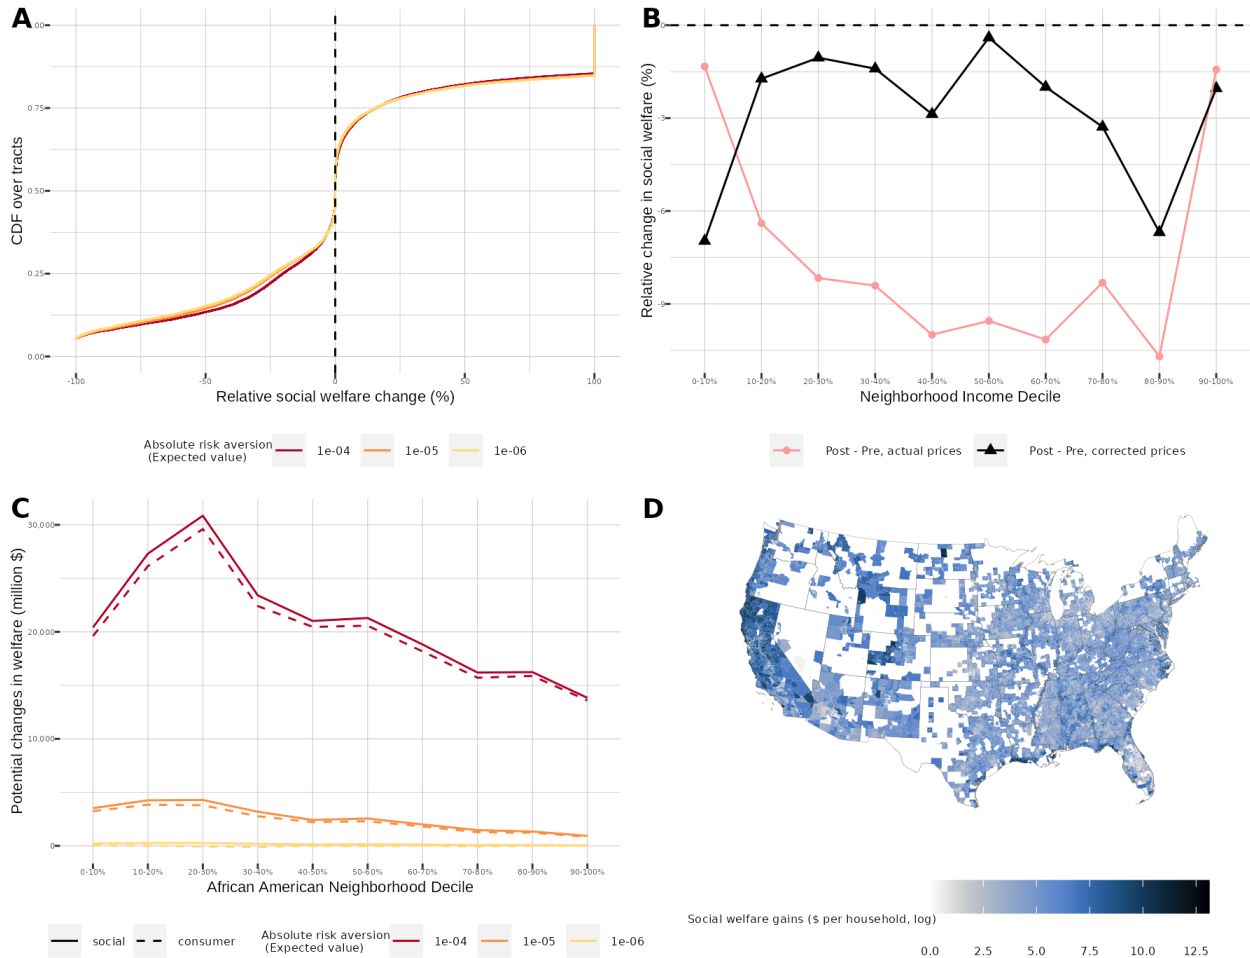
Figure 10.B depicts the relative social welfare impacts of the map update, aggregated by neighborhood income decile. The pink line presents the welfare changes that occurred, computed at the actual prices before and after the map update, while the black line presents the welfare changes *that would have occurred if* insurance premiums had been actuarially fair. First, note that for all income groups, the map updates *decreased* welfare: the relative changes vary between -2% and -10% (pink line), with no clear pattern in relation to the income group. Second, correcting insurance premiums to reflect actuarially fair prices would have mitigated some of these losses, but would not have led to substantial welfare gains. This reveals that insurance pricing is not the main driver behind the welfare losses – incorrect risk mapping is.

To assess the potential gains from improving flood risk mapping, Figure 10.C presents the welfare impacts of correcting the floodplain boundaries pre-map update using the First Street Foundation Flood Model, aggregated by African American neighborhood decile. The solid lines depict changes in social welfare while dashed lines depict changes in consumer welfare. Amounts are reported for three different expected values of the risk aversion distribution.

Figure 10.C shows that welfare gains would have been substantial if the map update had followed correct floodplain boundaries and had updated premiums to reflect actuarially fair prices. Using a plausible parameter of 10^{-5} for the expected value of the risk aversion distribution, updating the maps pre-digitization to the FSF floodplain boundaries would have yielded annual gains exceeding \$20 billion, with the largest gains concentrated in white and wealthy neighborhoods. Note that while the relative changes in welfare are not sensitive to positive values of risk aversion parameters, the absolute changes in welfare strongly depend on risk preferences and increase with risk aversion. For all three risk aversion parameters, con-

⁴²Although Sallee (2019) shows that the potential compensation criterion is problematic, as the lack of information about *who* loses from a given policy often makes compensation impractical.

Figure 10: Welfare impacts of updating flood maps



A: Cumulative distribution of social welfare changes over census tracts between the paper-based and digital map updates. Census tracts with a change of 0 are excluded, and the empirical cumulative distribution function is shown on $[-100\%, +100\%]$ for clarity. B: Social welfare impacts of updated flood maps aggregated by neighborhood income deciles, assuming an expected risk aversion value of 10^{-5} . The pink line depicts relative welfare changes using true insurance premiums (pre and post map updates), while the black line assumes actuarially fair premiums before and after the map updates. C: Potential welfare gains of updating pre-modernization flood maps following the 100-year and 500-year floodplains as estimated in the FSF model, aggregated by African American neighborhood decile and assuming actuarially fair insurance pricing in corrected floodplain boundaries. Estimates for three different values of the expected risk aversion parameters are represented, along with both social and consumer welfare measures. D: Spatial distribution of the potential social welfare gains of updating post-modernization flood maps following the 100-year and 500-year floodplains as estimated in the FSF model, assuming actuarially fair insurance pricing in corrected floodplain boundaries and an expected risk aversion value of 10^{-5} . Dollar amounts are expressed per household experiencing any change in floodplain boundaries (in log and bounded at 0 for clarity).

sumer welfare closely tracks social welfare. At very low expected value of risk aversion (10^{-4}), consumer welfare changes becomes negative, highlighting that more risk neutral homeowners would stand to lose the most from correcting floodplain boundaries and insurance pricing.

Finally, Figure 10.D shows the spatial distribution of welfare gains that would arise if we moved from the current maps and premiums to corrected floodplain boundaries and actuarially fair premiums, with the welfare gains expressed per rezoned household and assuming an expected risk aversion parameter of 10^{-5} . While flood-prone neighborhoods on the Atlantic coast are predicted to realize substantial welfare gains, households rezoned to a different floodplain in California, Oregon, and Washington State, as well as households residing inland near the Mississippi, Ohio and Tennessee rivers would also experience large gains in welfare.

7 Conclusion

This paper investigates the public provision of flood risk information and its impacts on flood insurance take-up in the United States. I compile novel data on the evolution of floodplain boundaries contained in the official flood maps, which represent the largest disaster risk mapping and information provision effort ever undertaken by a national government. I compare these flood maps with independent risk estimates provided by a state-of-the-art model, and estimate their impacts on residential flood insurance take-up.

I first document that since Hurricane Katrina, new flood maps rezoned more than one million properties outside of the high-risk floodplains. This nominal decrease in flood risk was caused by the drawing of more floodplain boundaries, which predominantly occurred in neighborhoods with greater Black and Hispanic populations. Comparing official flood maps with independent risk estimates, I find that more than five million properties were incorrectly omitted from the high-risk floodplains during the map updates. These omissions are primarily due to the continued ignorance of rain-based flood risk in the official maps, and they predominantly occurred in minority neighborhoods.

Leveraging the staggered updating of flood maps, I find that they are a crucial driver of flood insurance uptake: removing properties from the 100-year floodplain decreases insurance take-up, while rezoning properties inside the 100-year floodplain increases flood insurance take-up, both inside and outside of the 100-year floodplain. These spillover effects suggest that information plays a substantial role in the demand for flood insurance. In contrast, digitizing previously existing floodplain boundaries has insignificant impact on the demand for flood insurance, suggesting that information access costs were not limiting insurance demand. These results are robust to alternative estimation strategies based on local and clustered synthetic controls. Overall, I estimate that map updates caused 40,000 additional properties to be covered by flood insurance after two-years in areas where the 100-year floodplain was expanded, but caused more than 100,000 households to drop their flood insurance coverage in areas where

the 100-year floodplain was shrunk. Effects are larger in census tracts with a higher share of Black and Hispanic residents, which exacerbated disparities in flood risk coverage. Although removing properties from the high-risk floodplains reduced flood insurance premiums in these neighborhoods, declines in insurance coverage raise concerns regarding the vulnerability of communities already suffering from under-investments.

There are important limitations to this study. In particular, my analysis does not quantify the contributions of the different mechanisms that underlie the effect of public risk information on insurance take-up. I find that beliefs about flood risks and the mandatory purchase requirement both matter, but further understanding their relative contributions nation-wide is necessary to design policies that efficiently promote insurance coverage.

In addition, the welfare estimates computed in the paper are based on the assumption the FSF model is a decent proxy for true flood risk. This is a relatively innocuous assumption in the context of pluvial flood risk, which is systematically missing from official flood maps, but a stronger assumption in the context of fluvial and coastal flooding. The FSF does provide state-of-the-art estimates of flood risk, but the art of flood risk modelling evolves quickly. As more models become available to predict flood risk, additional work is needed to optimally combine risk assessments from various scientific sources.

A reconciliation bill approved by the House Financial Services Committee in September 2021 aims to provide three billion dollars to improve flood maps throughout the U.S (House Financial Services Committee, 2021). Such investment has the potential to benefit consumers through improved flood risk information and increased insurance take-up, but incorrect map updates could increase vulnerability to flood risk further. Given the substantial costs of risk mapping and concerns about the scientific validity of official floodplain boundaries, a promising avenue to produce official risk maps is to leverage newly available and independent models of risks as baselines for the new flood maps. Local deviations from these models could be warranted and included in official products where communities have relevant knowledge about risk or newly developed infrastructure. But as climate change further increases the importance of accurate risk information, governments can benefit from grounding official climate risk products in models that reflect the current science.

References

- Abadie, A. (2021). Using Synthetic Controls: Feasibility, Data Requirements, and Methodological Aspects. *Journal of Economic Literature* 59(2), 391–425.
- Abadie, A., A. Diamond, and J. 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, A., A. Diamond, and J. Hainmueller (2015). Comparative Politics and the Synthetic Control Method. American Journal of Political Science 59(2), 495–510.
- Abadie, A. and J. Gardeazabal (2003). The economic costs of conflict: A case study of the Basque country. American Economic Review 93(1), 113–132.
- Abadie, A. and J. L’Hour (2021). A Penalized Synthetic Control Estimator for Disaggregated Data. Journal of the American Statistical Association 116(536), 1817–1834.
- Annan, F. and W. Schlenker (2015). Federal crop insurance and the disincentive to adapt to extreme heat. American Economic Review 105(5), 262–266.
- Arnal, S., J. R. Porter, B. Lingle, Z. Chu, M. L. Marston, and O. E. Wing (2020). Assessing property level economic impacts of climate in the US, new insights and evidence from a comprehensive flood risk assessment tool. Climate 8(10), 1–20.
- ASFPM (2020). Flood Mapping for the Nation: A Cost Analysis for Completing and Maintaining the Nation’s NFIP Flood Map Inventory. Technical report, Madison, WI.
- Atreya, A., S. Ferreira, and E. Michel-Kerjan (2015). What drives households to buy flood insurance? New evidence from Georgia. Ecological Economics 117, 153–161.
- Bakkensen, L. and L. Ma (2020). Sorting over flood risk and implications for policy reform. Journal of Environmental Economics and Management 104, 22.
- Bakkensen, L. A. and L. Barrage (2022). Going Underwater? Flood Risk Belief Heterogeneity and Coastal Home Price Dynamics. Review of Financial Studies 35(8), 3666–3709.
- Banzhaf, B. H. S. and R. P. Walsh (2016). Do People Vote with Their Feet ? An Empirical Test of Tiebout’s Mechanism. American Economic Review 98(3), 843–863.
- Banzhaf, S., L. Ma, and C. Timmins (2019). Environmental Justice: The Economics of Race, Place, and Pollution. Journal of Economic Perspectives 33(1), 185–208.
- Barber, R. F., E. J. Candès, A. Ramdas, and R. J. Tibshirani (2021). Predictive inference with the jackknife+. The Annals of Statistics 49(1), 486–507.

- Barreca, A., K. Clay, O. Deschenes, M. Greenstone, and J. S. Shapiro (2016). Adapting to climate change: The remarkable decline in the US temperature-mortality relationship over the Twentieth Century. Journal of Political Economy 124(1), 105–159.
- Bates, P. D., M. S. Horritt, and T. J. Fewtrell (2010, jun). A simple inertial formulation of the shallow water equations for efficient two-dimensional flood inundation modelling. Journal of Hydrology 387(1-2), 33–45.
- Bates, P. D., N. Quinn, C. Sampson, A. Smith, O. Wing, J. Sosa, J. Savage, G. Olcese, J. Neal, G. Schumann, L. Giustarini, G. Coxon, J. R. Porter, M. F. Amodeo, Z. Chu, S. Lewis-Gruss, N. B. Freeman, T. Houser, M. Delgado, A. Hamidi, I. Bolliger, K. E. McCusker, K. Emanuel, C. M. Ferreira, A. Khalid, I. D. Haigh, A. Couasnon, R. E. Kopp, S. Hsiang, and W. F. Krajewski (2020). Combined Modeling of US Fluvial, Pluvial, and Coastal Flood Hazard Under Current and Future Climates. Water Resources Research 57(2), 1–29.
- Bell, H. M. and G. A. Tobin (2007). Efficient and effective? The 100-year flood in the communication and perception of flood risk. Environmental Hazards 7(4), 302–311.
- Beltrán, A., D. Maddison, and R. J. Elliott (2018). Is Flood Risk Capitalised Into Property Values? Ecological Economics 146, 668–685.
- Ben-Michael, E., A. Feller, and J. Rothstein (2019). Synthetic Controls and Weighted Event Studies with Staggered Adoption.
- Ben-Michael, E., A. Feller, and J. Rothstein (2021). The Augmented Synthetic Control Method. Journal of the American Statistical Association 116(536), 1789–1803.
- Bin, O. and C. 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.
- Blickle, K. and J. A. C. Santos (2022). Unintended Consequences of "Mandatory" Flood Insurance. SSRN Electronic Journal.
- Botzen, W. J., O. Deschenes, and M. Sanders (2019). The economic impacts of natural disasters: A review of models and empirical studies. Review of Environmental Economics and Policy 13(2), 167–188.
- Bradt, J. T., C. Kousky, and O. E. Wing (2021). Voluntary purchases and adverse selection in the market for flood insurance. Journal of Environmental Economics and Management 110(0), 102515.

- Bullard, R. D. (1983). Solid waste sites and the black houston community. Sociological Inquiry 53(2-3), 273–288.
- Cabral, M. and C. Hoxby (2015). The Hated Property Tax: Saliency, Tax Rates, and Tax Revolts. NBER Working Paper.
- Callaway, B. and P. H. Sant’Anna (2021). Difference-in-Differences with Multiple Time Periods. Journal of Econometrics 225(2), 200–230.
- Carleton, T., A. Jina, M. Delgado, M. Greenstone, T. Houser, S. Hsiang, A. Hultgren, R. E. Kopp, K. E. McCusker, I. Nath, J. Rising, A. Rode, H. K. Seo, A. Viaene, J. Yuan, and A. T. Zhang (2022, 04). Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits. The Quarterly Journal of Economics 137(4), 2037–2105.
- Chetty, R., A. Looney, and K. Kroft (2009). Saliency and taxation: Theory and evidence. American Economic Review 99(4), 1145–1177.
- Chivers, J. and N. E. Flores (2002). Market failure in information: The National Flood Insurance Program. Land Economics 78(4), 515–521.
- Christensen, P. and C. Timmins (2022). Sorting or Steering: The Effects of Housing Discrimination on Neighborhood Choice. Journal of Political Economy 130(8), 2110–2163.
- Congressional Budget Office (2017). Age of Flood Maps in Selected Counties That Account for Most of the Expected Claims in the National Flood Insurance Program: Supplemental Material for The National Flood Insurance Program: Financial Soundness and Affordability. Technical report.
- Currie, J., J. Voorheis, and R. Walker (2022). What Caused Racial Disparities in Particulate Exposure to Fall? New Evidence from the Clean Air Act and Satellite-Based Measures of Air Quality. American Economic Review, forthcoming.
- Cutler, D. M., R. S. Huckman, and M. B. Landrum (2004). The role of information in medical markets: An analysis of publicly reported outcomes in cardiac surgery. American Economic Review 94(2), 342–346.
- Davis, L. W. and G. E. Metcalf (2016). Does better information lead to better choices? evidence from energy-efficiency labels. Journal of the Association of Environmental and Resource Economists 3(3), 589–625.

- De Almeida, G. A. and P. Bates (2013). Applicability of the local inertial approximation of the shallow water equations to flood modeling. Water Resources Research 49(8), 4833–4844.
- de Chaisemartin, C. and X. D’Haultfoeuille (2020). Two-Way Fixed Effects Estimators With Heterogeneous Treatment Effects. American Economic Review 110(9), 2964–2996.
- Department of Homeland Security (2012). Exhibit 300 BY13 FEMA - Risk Mapping, Assessment and Planning (RISKMAP) Program.
- Depro, B., C. Timmins, and M. O’Neil (2015). White Flight and Coming to the Nuisance: Can Residential Mobility Explain Environmental Injustice? Journal of the Association of Environmental and Resource Economists 2(3), 439–468.
- Deryugina, T. (2017). The fiscal cost of hurricanes: Disaster aid versus social insurance. American Economic Journal: Economic Policy 9(3), 168–198.
- Diaz, D. and F. Moore (2017). Quantifying the economic risks of climate change. Nature Climate Change 7(11), 774–782.
- Doudchenko, N. and G. W. Imbens (2016). Balancing, Regression, Difference-in-differences and synthetic control methods: a synthesis. NBER Working Paper.
- Eby, M. (2023, March). Testimony of mr. matthew eby, chief executive officer and founder of the first street foundation. Testimony before the Senate Committee on the Budget. “Rising Seas, Rising Costs: Climate Change and the Economic Risks to Coastal Communities”.
- Einav, L., A. Finkelstein, and M. R. Cullen (2010). Estimating Welfare in Insurance Markets Using Variation in Prices. The Quarterly Journal of Economics 125(3), 877–921.
- FEMA (1996). Q3 flood data use guide. Technical Report March.
- FEMA (2009). Risk Mapping, Assessment, and Planning (Risk MAP) Multi-Year Plan: Fiscal Years 2010-2014. Technical report, Homeland Security.
- FEMA (2013). Hazus-MH Technical Manual Flood Model. Technical report.
- FEMA (2016). Flood Insurance Study - City of New Orleans and Orleans Parish, Louisiana. Technical report, FEMA.
- First Street Foundation (2020). First Street Foundation Flood Model (FSF-FM) Technical Documentation. Technical report, First Street Foundation.

- First Street Foundation (2022, December). First street foundation’s partnerships and impact expand greatly throughout 2022.
- Frank, T. (2020, 2). Studies Sound Alarm on “Badly Out-of-Date” FEMA Flood Maps.
- Fried, S. (2022, 11). Seawalls and Stilts: A Quantitative Macro Study of Climate Adaptation. The Review of Economic Studies 89(6), 3303–3344.
- Gall, M., B. Boruff, and S. L. Cutter (2007). Assessing Flood Hazard Zones in the Absence of Digital Floodplain Maps: Comparison of Alternative Approaches. Natural Hazards Review 8(1), 1–12.
- Gallagher, J. (2014). Learning About an Infrequent Event: Evidence from Flood Insurance Take-Up in the US. American Economic Journal: Economic Policy 6(3), 206–233.
- Gamper-Rabindran, S. and C. Timmins (2013). Does cleanup of hazardous waste sites raise housing values? Evidence of spatially localized benefits. Journal of Environmental Economics and Management 65(3), 345–360.
- Gibson, M. and J. T. Mullins (2020). Climate Risk and Beliefs in New York Floodplains. Journal of the Association of Environmental and Resource Economists 7(6), 1069–1111.
- Goodman-Bacon, A. (2021). Difference-In-Differences with Variation in Treatment Timing. Journal of Econometrics 225(2), 254–277.
- Gourevitch, J. D., C. Kousky, Y. Liao, C. Nolte, A. B. Pollack, J. R. Porter, and J. A. Weill (2023). Unpriced climate risk and the potential consequences of overvaluation in us housing markets. Nature Climate Change 13, 250–257.
- Government Accountability Office (2020). Disaster Assistance: Additional Actions Needed to Strengthen FEMA’s Individuals and Households Program.
- Government Accountability Office (2021a, October). Fema flood maps: Better planning and analysis needed to address current and future flood hazards. Technical Report GAO-22-104079, U.S. Government Accountability Office.
- Government Accountability Office (2021b). National flood insurance program: Congress should consider updating the mandatory purchase requirement. Technical Report GAO-21-578, U.S. Government Accountability Office.
- Hausman, C. and S. Stolper (2021). Inequality, information failures, and air pollution. Journal of Environmental Economics and Management 110(December 2020), 102552.

- Highfield, W. E., S. A. Norman, and S. D. Brody (2013). Examining the 100-Year Floodplain as a Metric of Risk, Loss, and Household Adjustment. Risk Analysis 33(2), 186–191.
- Hino, M. and M. Burke (2021). The effect of information about climate risk on property values. Proceedings of the National Academy of Sciences 118(17).
- House Financial Services Committee (2021). Amendment in the nature of a substitute to the committee print offered by Ms. Waters of California, Title IV Subtitle A — Creating and Preserving affordable, equitable and accessible housing for the 21st Century.
- House of Representatives (2020, February). An examination of federal flood maps in a changing climate. Number 116-70. U.S. House of Representatives, Committee on Science, Space, and Technology.
- Hu, Z. (2022). Social interactions and households’ flood insurance decisions. Journal of Financial Economics 144(2), 414–432.
- Jessoe, K. and D. Rapson (2014). Knowledge is (Less) power: Experimental evidence from residential energy use. American Economic Review 104(4), 1417–1438.
- Kahn, M. (2005). The death toll from natural disasters: The role of income, geography and institutions. The Review of Economics and Statistics 87, 271–284.
- Kahn, M. E. (2021). Adapting to Climate Change: Markets and the Management of an Uncertain Future.
- Kamenica, E. (2017). Information Economics. Journal of Political Economy 125(6), 1885–1890.
- Knowles, S. G. and H. C. Kunreuther (2014). Troubled waters: The national flood insurance program in historical perspective. Journal of Policy History 26(3), 327–353.
- Kocornik-Mina, A., T. K. McDermott, G. Michaels, and F. Rauch (2020). Flooded cities. American Economic Journal: Applied Economics 12(2), 35–66.
- Koerth, M. (2017). It’s Time To Ditch The Concept Of ‘100-Year Floods’. FiveThirtyEight.
- Kousky, C. (2016). Flood Insurance: Why Don’t People Buy It? Technical report.
- Kousky, C. (2019). The Role of Natural Disaster Insurance in Recovery and Risk Reduction. Annual Review of Resource Economics 11, 399–418.
- Kousky, C., H. Kunreuther, B. Lingle, and L. Shabman (2018). The Emerging Private Residential Flood Insurance Market in the United States. (July), 1–51.

- Kunreuther, H. C., M. V. Pauly, and S. McMorrow (2013). Insurance and behavioral economics: Improving decisions in the most misunderstood industry.
- Lea, D. and S. Pralle (2022). To appeal and amend: Changes to recently updated Flood Insurance Rate Maps. Risk, Hazards and Crisis in Public Policy 13(1), 28–47.
- Lehmann, R. J. (2020). Do no harm: managing retreat by ending new subsidies. Technical report, R street.
- Ludy, J. and G. M. Kondolf (2012). Flood risk perception in lands "protected" by 100-year levees. Natural Hazards 61(2), 829–842.
- Marsooli, R., N. Lin, K. Emanuel, and K. Feng (2019). Climate change exacerbates hurricane flood hazards along US Atlantic and Gulf Coasts in spatially varying patterns. Nature Communications 10(1), 1–9.
- Mennis, J. (2017). Dasymetric Mapping. International Encyclopedia of Geography: People, the Earth, Environment and Technology, 1–10.
- Michel-Kerjan, E., S. Lemoyne de Forges, and H. Kunreuther (2012). Policy Tenure Under the U.S. National Flood Insurance Program (NFIP). Risk Analysis 32(4), 644–658.
- Molina, R. and I. Rudik (2022, Jul). The social value of hurricane forecasts.
- Morrissey, W. A. (2008). FEMA Funding for Flood Map Modernization. Technical report, Congressional Research Service.
- Mulder, P. (2022). Mismeasuring Risk: The Welfare Effects of Flood Risk Information. Working paper.
- National Research Council (2015). Affordability of National Flood Insurance Program Premiums. Technical report.
- Neal, J., G. Schumann, and P. Bates (2012). A subgrid channel model for simulating river hydraulics and floodplain inundation over large and data sparse areas. Water Resources Research 48(11), 1–16.
- Nguyen, C. N. and I. Noy (2020). Measuring the impact of insurance on urban earthquake recovery using nightlights. Journal of Economic Geography 20(3), 857–877.

- Nobre, A. D., L. A. Cuartas, M. R. Momo, D. L. Severo, A. Pinheiro, and C. A. Nobre (2016). HAND contour: A new proxy predictor of inundation extent. Hydrological Processes 30(2), 320–333.
- Ostriker, A. and A. Russo (2022). The Effects of Floodplain Regulation on Housing Markets. Working paper.
- Pollack, A. B., I. Sue Wing, N. Pinter, and C. Nolte (2022). New Empirical Models for Flood Loss Prediction and Implications for the Coterminous United States. Working Paper.
- Pope, J. C. (2008). Do seller disclosures affect property values? Buyer information and the hedonic model. Land Economics 84(4), 551–572.
- Porter, J. R., M. L. Marston, E. Shu, M. Bauer, K. Lai, B. Wilson, and M. Pope (2023). Estimating Pluvial Depth–Damage Functions for Areas within the United States Using Historical Claims Data. Natural Hazards Review 24(1), 1–10.
- Pralle, S. (2019). Drawing lines: FEMA and the politics of mapping flood zones. Climatic Change 152(2), 227–237.
- Robbins, M. W., J. Saunders, and B. Kilmer (2017). A Framework for Synthetic Control Methods With High-Dimensional, Micro-Level Data: Evaluating a Neighborhood-Specific Crime Intervention. Journal of the American Statistical Association 112(517), 109–126.
- Rosenzweig, M. R. and C. Udry (2014). Rainfall Forecasts, Weather, and Wages over the Agricultural Production Cycle. American Economic Review: Papers & Proceedings 104(5), 278–283.
- Sallee, J. M. (2019). Pigou Creates Losers: on the Implausibility of Achieving Pareto Improvements from Efficiency Enhancing Policies. WP. NBER Working Paper Series.
- Sampson, C. C., P. D. Bates, J. C. Neal, and M. S. Horritt (2013). An automated routing methodology to enable direct rainfall in high resolution shallow water models. Hydrological Processes 27(3), 467–476.
- Sampson, C. C., A. M. Smith, P. D. Bates, J. C. Neal, L. Alfieri, and J. E. Freer (2015). Water Resources Research. Water Resources Research 51, 7358–7381.
- Sastry, P. (2021). Who Bears Flood Risk? Evidence from Mortgage Markets in Florida. Working Paper.

- Scata, J. (2017). FEMA’s Outdated and Backward Looking Flood Maps.
- Shao, W., S. Xian, N. Lin, H. Kunreuther, N. Jackson, and K. Goidel (2017). Understanding the effects of past flood events and perceived and estimated flood risks on individuals’ voluntary flood insurance purchase behavior. Water Research 108, 391–400.
- Shr, Y.-H. and K. Y. Zipp (2019). The Aftermath of Flood Zone Remapping: The Asymmetric Impact of Flood Maps on Housing Prices. Land Economics 95(2), 174–192.
- Shrader, J. (2020). Improving Climate Damage Estimates by Accounting for Adaptation. SSRN Working Paper, 1–77.
- Shrader, J. G., L. Bakkensen, and D. Lemoine (2022). Fatal Errors : The Mortality Value of Accurate Weather Forecasts. Working Paper, 1–55.
- Stigler, G. J. (1961). The Economics of Information. Journal of Political Economy 69(3), 213–225.
- Sun, L. and S. Abraham (2021). Estimating Dynamic Treatment Effects in Event Studies With Heterogeneous Treatment Effects. Journal of Econometrics 225(2), 175–199.
- Tobin, R. J. and C. Calfee (2005). The National Flood Insurance Program’s Mandatory Purchase Requirement: Policies, Processes, and Stakeholders. Technical Report March, American Institutes for Research, Washington, D.C.
- Turnham, J., K. Burnett, C. Martin, T. McCall, R. Juras, and J. Spader (2011). Housing Recovery on the Gulf Coast - Phase II (Results of Property Owner Survey in Louisiana, Mississippi, and Texas). Technical report.
- Wagner, K. R. H. (2022). Adaptation and Adverse Selection in Markets for Natural Disaster Insurance. American Economic Journal: Economic Policy, forthcoming 14(3), 380–421.
- Weill, J. A., M. Stigler, O. Deschenes, and M. R. Springborn (2021). Researchers’ Degrees-of-Flexibility and the Credibility of Difference-in-Differences Estimates: Evidence From the Pandemic Policy Evaluations. NBER Working Paper, 1–72.
- Wing, O., P. D. Bates, A. M. Smith, C. C. Sampson, K. A. Johnson, J. Fargione, and P. Morefield (2018). Estimates of present and future flood risk in the conterminous United States. Environmental Research Letters 13(3), 034023.

- Wing, O. E., P. D. Bates, C. C. Sampson, A. M. Smith, K. A. Johnson, and T. A. Erickson (2017). Validation of a 30 m resolution flood hazard model of the conterminous United States. Water Resources Research 53(9), 7968–7986.
- Wing, O. E., W. Lehman, P. D. Bates, C. C. Sampson, N. Quinn, A. M. Smith, J. C. Neal, J. R. Porter, and C. Kousky (2022). Inequitable patterns of US flood risk in the Anthropocene. Nature Climate Change 12(2), 156–162.
- Wittenberg, A. (2017). The myth of the 100-year flood. EE News, GreenWire.
- Wright, J. K. (1936). A method of mapping densities of population. Geographical Review 26(1), 103–110.
- Xu, Y. (2017). Generalized synthetic control method: Causal inference with interactive fixed effects models. Political Analysis 25(1), 57–76.
- You, X. and C. Kousky (2023). Improving household and community disaster recovery: Evidence on the role of insurance. SSRN, 52.

Appendix for Online Publication

A Data preparation

A.1 Timing and content of the flood map updates

The flood insurance rate maps are provided at the community level. The community is usually a county, but can also be a city, a town, a borough or a parish. Previous work focused on flood maps relied on the “Community Status Book” maintained by FEMA,⁴³ which records the effective dates of the communities’ first FIRM (paper-based flood map) and current FIRM or DFIRM (where the D stands for *Digital*, i.e., post map update). However, because the Community Status Book only records the date of the first FIRM and current (D)FIRM, it cannot be used to identify transitions from a FIRM to a DFIRM, and from a DFIRM to a more recent DFIRM, unless the community only ever received one DFIRM – and these communities are *not* identified in the Community Status Book. In addition, although most census tracts can be easily assigned to an NFIP community (in particular when the community is a county which received a unique county-wide map), several census tracts are instead covered by a community-specific DFIRM and cannot be directly assigned a map update dates.

To circumvent these issues, I use new data on digital yearly snapshots of the National Flood Hazard Layer since 2012 to assign DFIRM effective dates at the census tract level. Finding, assembling, and cleaning these files was a substantial undertaking: most files were not maintained by FEMA and were retrieved from the GIS archives of contractors who worked with FEMA to produce these maps. I intersect the census tracts polygons with the DFIRM effective dates polygons and record the effective date indicated at the census tract centroid. I then compare the intersected DFIRM identification code with the census tract FIPS code: if the DFIRM was issued for an entire county, then the DFIRM ID and the first five-digits of the census tract FIPS code must match. For communities that are not entire counties (for instance independent Virginian cities), I only match on the State FIPS code and confirm the effective dates obtained with the Community Status Book. This procedure induces some measurement error on treatment timing for census tracts that are not entirely included in an NFIP community, since the effective date at the centroid of the census tract might differ from the date at which other parts of the census tract received their (D)FIRM. However, this measurement error only affects a minority of census tracts in sparsely populated communities, which themselves account for a minority of NFIP insurance purchases. Omitting these census tracts from the analysis do not noticeably impact the model estimates. Finally, for census tracts covered by a digital

⁴³<https://www.fema.gov/flood-insurance/work-with-nfip/community-status-book>

flood map in the 2012 National Flood Hazard Layer, I manually check that this was indeed the community’s first digital flood map by comparing the digitization date with FEMA’s Map Service Center repository⁴⁴. This platform maintains a collection of pdf files depicting historical mapping products.

To investigate the evolution of floodplain boundaries throughout the map update process, I use the Q3 data product, which depicts the floodplain boundaries of the paper-based in 2005. As described in the main text, the Q3 maps were designed to aid floodplain managers and city officials with disaster response activities, but were not widely available to the public. They however provided a digital version of floodplain boundaries as shown on the historical (paper-based) flood maps.

A.2 Note on intersecting floodplains and census tracts polygons

Intersecting the different flood zones within each census tract involves the pair-wise intersection of hundreds of thousands of polygons. This process is computationally expensive, but performance can be dramatically improved by first cropping the flood zone polygons at the county level and then run on parallel cores the intersections between the cropped polygon and the census tracts located within this county.

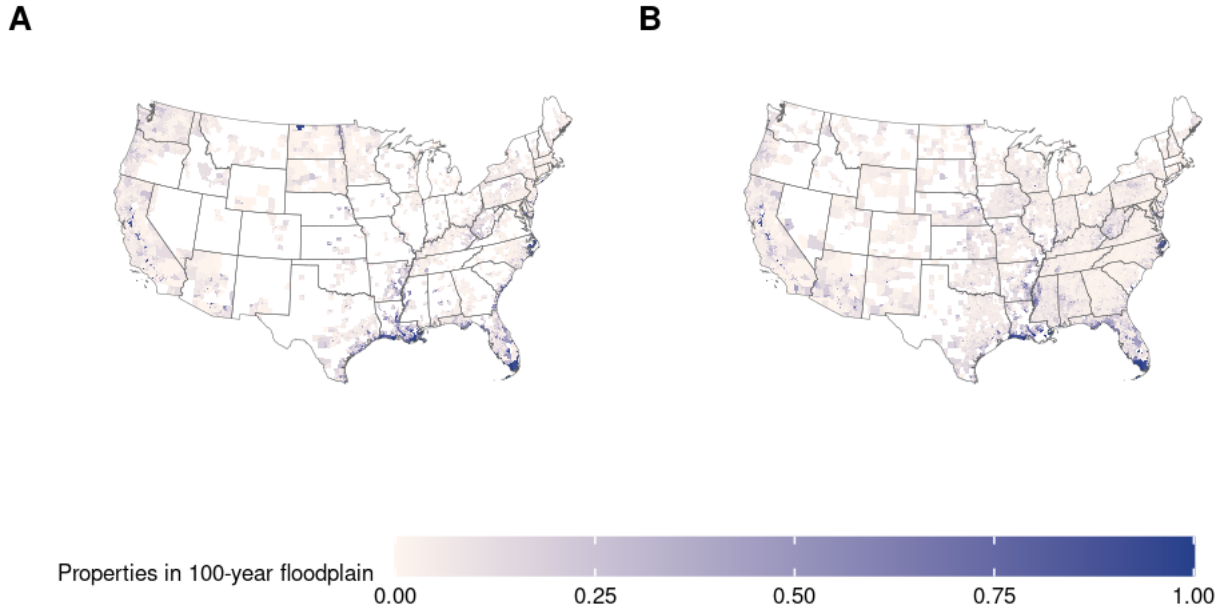
Unless a sufficient buffer is used, flood zones that barely “touch” the boundary of a census tract will induce additional intersections that are computationally costly. On the flip side, using a large buffer will induce measurement error in the computed areas. Using a minimal buffer and post-processing the intersected polygons resulted in the best performance overall. On a virtual machine with 48 vCPUS and 512 GB of RAM, all intersections can be ran within two days.

A.3 Tract-level population and property count changes within different flood zones

Each census tract covers between 1,200 and 8,000 inhabitants. In very-densely populated areas, the spatial extent of a census tract is typically small, and the (spatial) share of the census tract in different flood zones is a good approximation of the share of the census tract *population* that lives within each flood zone. This is not necessarily true in medium-sized or large census tracts, where population can often be located completely outside of any floodplains. To properly estimate population exposure to flood risk, it is therefore crucial to account for the spatial distribution of population *within* census tracts. I do so in three different ways.

⁴⁴<https://msc.fema.gov/portal/advanceSearch>

Figure A.1: Share of properties located in the 100-year floodplain



A: Census tract share of residential properties located in the 100-year floodplain based on the Q3 product. **B:** Census tract share of residential properties located in the 100-year floodplain 2019.

First, I use the entire repository of 132 millions geolocalized residential properties maintained by the First Street Foundations and downloaded from their API in April 2021.⁴⁵ This measure of population distribution presents three advantages: (i) it captures residential properties only, which makes it a suitable measure of flood risk exposure to study how flood zone changes impact the demand for *residential* flood insurance, (ii) it is consistent with the other variables that I use from First Street to estimate flood risk exposure and expected annual average losses, and (iii) its coverage is more comprehensive than the often-used ZTRAX data (for comparison, I was only able to geolocalize 95 million residential properties from ZTRAX — I discuss the ZTRAX data below). I then intersect these residential properties with the Q3 flood maps, the 2012 DFIRMs, and the 2019 DFIRMs. I then aggregate the relevant changes in properties counts at the census tract level, based on the effective date of the first DFIRM: for instance for a tract in Broward County (FL) that received its first (and unique) digital flood map in 2014, I compute the tract-level number of properties in each flood zone in the Q3 data as well as in the NFHL 2019 data, and take their difference to get the *change* in the number of properties located within this flood zone. I consider three different flood zones: the 100-year floodplain (or SFHA), which comprises all A and V flood zones; the 500-year flood zone (part


⁴⁵Due to the rate limitations of the API, the full download takes about 3 weeks of un-interrupted running time.

of the X zones), and the levee-protected flood zone.

Because the First Street Foundation residential properties data is obtained from local tax assessors offices, coverage in some specific areas can be limited. When the number of residential FSF-properties in a given census tract is too low to be representative,⁴⁶ I use the 2010 dasymetric allocation of population provided by EPA’s EviroAtlas. This product uses land use cover and terrain slopes to reallocate the 2010 census tract population within census tracts at the 30m pixel level.⁴⁷ The dasymetric layer covers the entire contiguous US, but does not differentiate well between residential properties and businesses. To transform the dasymetric-population count in a flood zone into a property count, I divide the population count by 2.5 (the average number of people living in an American household).

Figure A.2 plots the changes in the number of FSF residential properties zoned in the 100-year floodplain (relative to the census tract number of properties) against the changes in the population in the 100-year zone using the dasymetric allocation, showing a large correlation between the two measures.

Figure A.2: Comparison of the relative change in the number of properties in the 100-year floodplain using the FSF residential properties and the dasymetric population layer



⁴⁶I do not use the FSF data to compute flood risk exposure changes in census tracts where the number of FSF-properties is less than 40 and where the American Community Survey estimates of the population is greater than 4 times the number of properties. In practice, this concerns less than 1,500 census tracts out of 73,745.

⁴⁷Dasymetric mapping has a long history in geographical sciences (Wright, 1936) but remains mostly under-appreciated by economists. See (Mennis, 2017) for a review of this technique.

A.4 ZTRAX construction year robustness check

To ensure that my measures of changes in the number of properties inside and outside the 100-year floodplain are not driven by new properties being built in rezoned areas, I use Zillow’s ZTRAX Assessment data. These data contain a large number of variables on individual structures, including their construction date. I restrict the data to geolocalized residential properties that have a non-missing construction year variable. As shown in Table A.1 in the “Unconditional” panel, the number of residential properties that satisfy these criterion is only 67 million properties, almost *half* the number of geolocalized residential properties in the FSF dataset. While necessarily smaller than when using the more complete FSF properties registry, counts obtained with the ZTRAX data in Table A.1 show a similar decrease in the number of properties located inside the 100-year floodplain following the map updates. Importantly, the decline remains when restricting the counts to properties built prior to 2005 (third panel).

Table A.1: Property counts inside and outside the 100-year floodplain using Zillow’s ZTRAX

Floodzone	2005	NFHL19
Unconditional:		
Inside 100-year floodplain	4,043,702	3,294,151
Outside 100-year floodplain	45,241,341	57,872,115
Not mapped	17,841,103	5,959,880
Conditional on Q3 and NFHL19:		
Inside 100-year floodplain	3,509,070	2,672,160
Outside 100-year floodplain	43,250,294	44,087,204
Conditional on Q3 and NFHL19, built prior to 2005:		
Inside 100-year floodplain	3,121,840	2,371,970
Outside 100-year floodplain	38,854,800	39,604,670

The sample is restricted to the ZTRAX Assessment tables that (i) can be geocoded and have a non-missing “building construction year” variable. The conditional counts include only properties in census tracts that are mapped in both the Q3 and NFHL19 data products.

A.5 Panel of insurance policies at the census tract level

I use records on insurance policies obtained from FEMA through Freedom of Information Act requests for the period 1983-2019. After multiple quality checks, I found that records prior to 2008 were missing a wide number of insurance policies – given the implausibility of the missing-

at-random assumption in this context, my analysis focuses on 2008-2019 period. Insurance records in California and Minnesota contained substantial errors, and I excluded records from these states from the analysis.

Because the smallest geographical unit reported in the NFIP data is the census tract, I use this level as the panel cross-sectional unit. The choice of temporal unit, however, implies an important trade-off between granularity and effective variation. Policies can be started and terminated at any time during the year, but the distribution of the starting dates is far from being uniform within a year: in fact, policies tend to start during the summer months. As a result, keeping track of only the *year* in which a policy was started would induce substantial measurement error. On the other hand, the majority of policies are active for exactly 12 months, and most owners decide whether or not to renew their policy only once every 365 days. As a result, creating a panel at the daily level would generate a large amount of auto-correlation between observations. Finally, flood map updates occur on specific months; I therefore choose this level as the temporal unit of the panel.

The distribution of the starting dates of insurance policies is not perfectly uniform within a month: more insurance policies tend to start on exactly the 1st, the 15th/16th, or on the very last day of the month. To account for these variations, I specify the first month of activity of a policy either as the month in which it becomes effective *if* the policy became effective before the 15th, or the next month if the policy became effective after the 16th. For policies that become effective exactly on the 15th or 16th of the month, I assign their first month of activity based on the results of Bernoulli draws. I use a similar procedure for the last month of activity. In practice, these data cleaning steps do not impact the sign or magnitude of the main estimates, but help reduce variations that arise purely out of measurement error.

Figure A.3: Section of New Orleans' (LA) FIRM effective in 1984

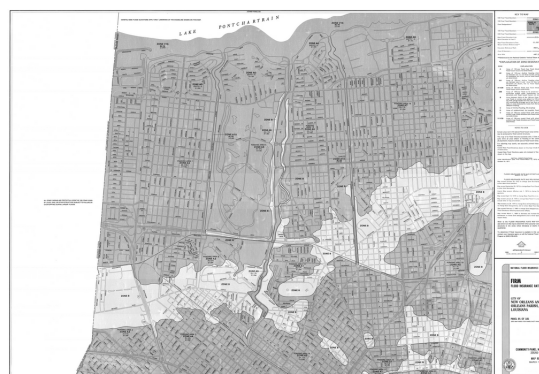
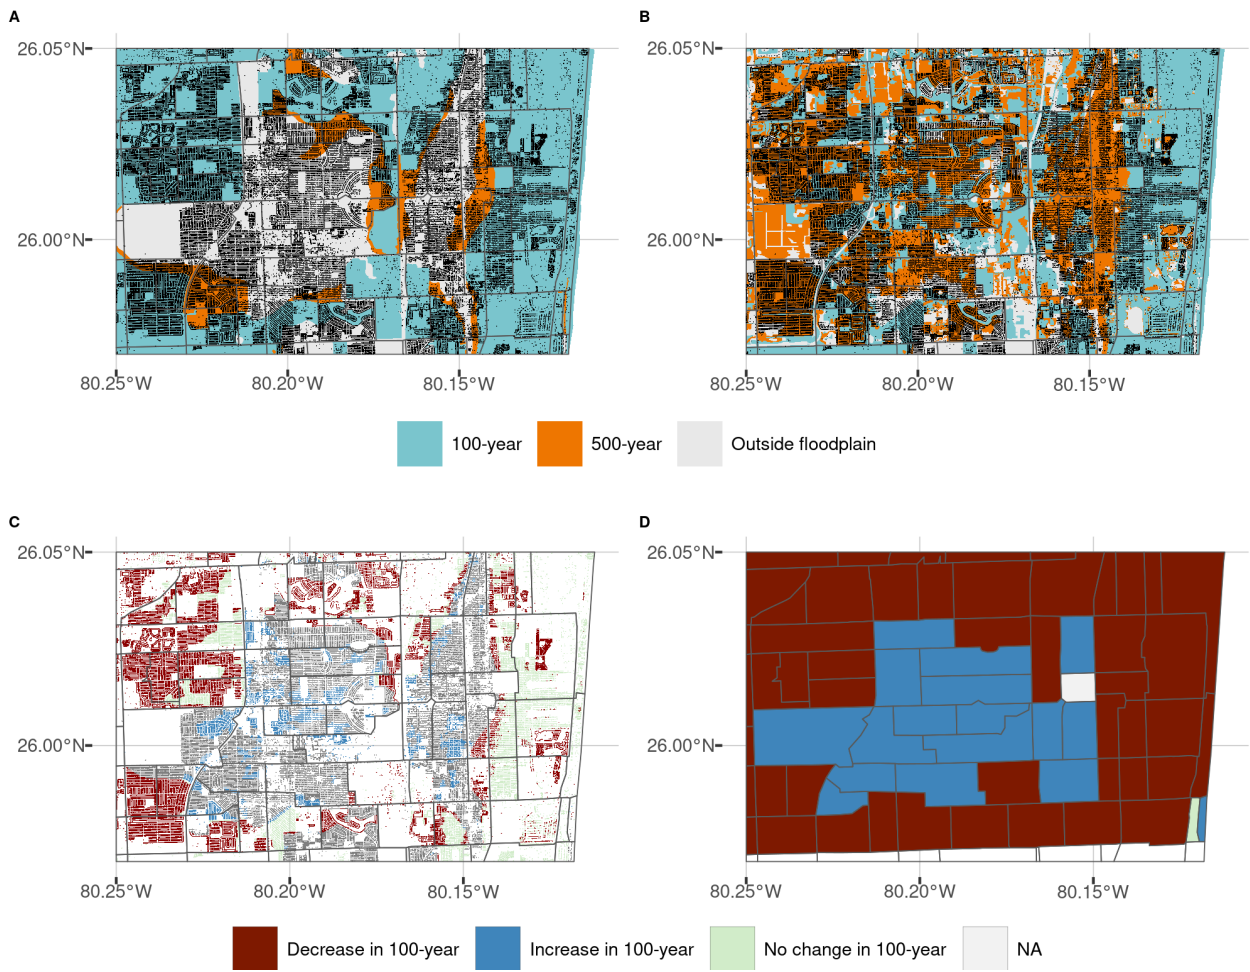
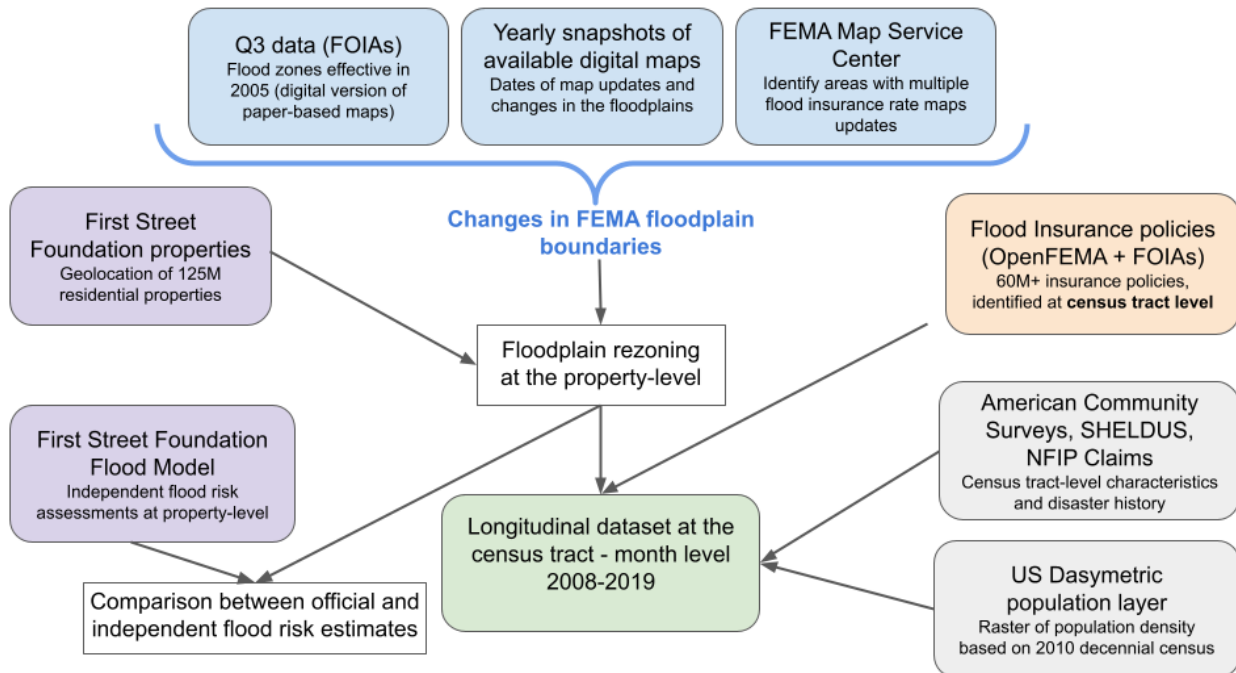


Figure A.4: Successive definition of the floodplains in South-Eastern Broward County (FL).



A: Flood map that was effective in 1992 (using the Q3 data). **B:** New flood map after modernization in August 2014. Black dots show the location of residential properties. **C:** Changes in the 100-year floodplain classification computed at the property level. **D:** Aggregated changes and classification of the change at the tract level.

Figure A.5: Summary of the different datasets used



B Polygon complexity table

Table B.2: Gerrymandering metrics for the 100-year floodplain

State	Perimeter Q3 (km)	Area Q3 (km ²)	Perimeter NFHL 2019 (km)	Area NFHL 2019 (km ²)	Change Polsby-Popper (%)
Alabama	31757.9	6129.8	46967.8	5983.4	-55.4
Arizona	42811	6121.1	61035.3	7245.3	-41.8
Arkansas	36669.7	10274.6	40736.9	10307.7	-18.7
California	91326.8	23341.4	105136.3	26269.3	-15.1
Colorado	10625.1	1025.2	11878.5	1067.5	-16.7
Connecticut	11204.5	811.2	11896.1	851.7	-6.9
Delaware	4845.9	991	7181.8	904.4	-58.4
District of Columbia	210.1	26.1	227	29.9	-1.7
Florida	152581.8	45730.1	290274.8	51778.4	-68.7
Georgia	33234.6	7486.1	50370.3	7622.4	-55.7
Idaho	5450.7	817.2	5601	915.7	6.1
Illinois	21639.9	2757.9	24663.6	2778.3	-22.4
Indiana	20242.6	3253.9	23767.1	3260.3	-27.3
Iowa	10556.9	1489.2	13591	1793.8	-27.3
Kansas	12605.9	2075.5	21991.9	2237.4	-64.6
Kentucky	44434.3	6092	69047.3	6843.9	-53.5
Louisiana	35298.8	25788.5	57047.1	26797.7	-60.2
Maine	14696	1622.5	17125.4	1580.5	-28.3
Maryland	21009.9	3559.4	26738.1	3094.3	-46.3
Massachusetts	22154.8	1930	24337.5	2107.8	-9.5
Michigan	15673.3	1518.5	18643.2	1531.6	-28.7
Minnesota	19665.1	3735.4	30587.1	4038.6	-55.3
Mississippi	32299.9	10203.8	50676.1	10534.6	-58.1
Missouri	16287.7	2572.5	21919.8	2475	-46.9
Montana	6491.8	1236.9	7421.8	1375.7	-14.9
Nebraska	8003.2	1874.7	10111.8	1918.1	-35.9
Nevada	167	15.9	149.6	7.3	-42.6
New Hampshire	5993.2	475	6148	493.6	-1.2
New Jersey	23598.9	3498.8	31624.6	3247.5	-48.3
New Mexico	5616.3	506.9	5968.7	573.7	0.2
New York	35580.3	3647.2	51098.6	3835.4	-49
North Carolina	56975.7	14072	77695.4	13366.1	-48.9
North Dakota	12033.4	1972.9	18390.6	2497	-45.8
Ohio	26016.6	2583.9	31764.3	2687.4	-30.2
Oklahoma	24750.4	3609	30049.6	3782.7	-28.9
Oregon	25129.5	3524.9	27656.3	3606.1	-15.5
Pennsylvania	67654.7	4378.7	90668.6	4496.4	-42.8
Rhode Island	4235.2	363.3	4504.2	375.9	-8.5
South Carolina	11136.1	2947.4	17596.2	2773.7	-62.3
South Dakota	13417.4	1418.5	14585.9	1555.7	-7.2
Tennessee	25049.7	6381.2	34414.4	6370.7	-47.1
Texas	129249.6	23367.3	156263.6	24349.3	-28.7
Utah	961	256.4	941.7	248.4	0.9
Vermont	4958.6	317	5062.9	326.5	-1.2
Virginia	29161.8	2589.1	33316.1	2518.6	-25.5
Washington	22450.3	2382.1	26308.5	2351.8	-28.1
West Virginia	24394.2	1576.1	27528.2	1544.1	-23.1
Wisconsin	18669.3	2949.1	24840.8	3217.7	-38.4
Wyoming	589.1	104.2	700.7	109.5	-25.7

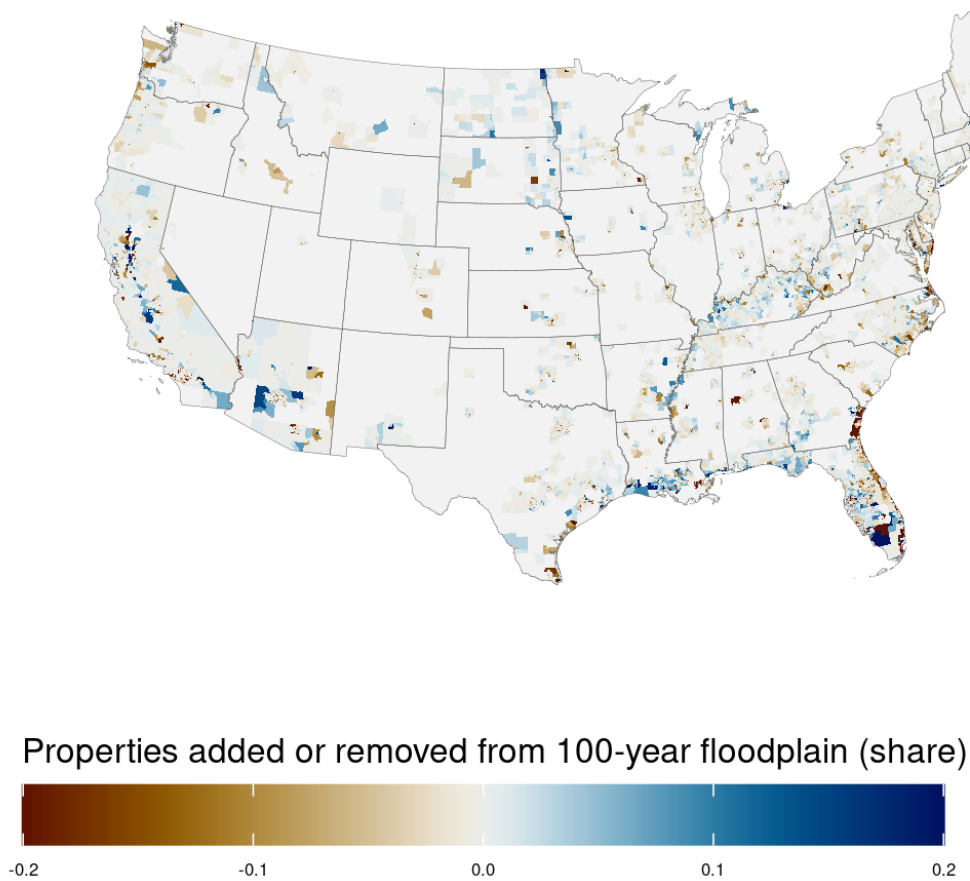
a

^aThe Polsby-Popper metric of a district D is defined as $PP(D) = \frac{4\pi A(D)}{P(D)^2}$, with $A(D)$ and $P(D)$ the area and perimeter of district D, respectively.

C Additional descriptive statistics and robustness checks

C.1 Rezoning at the tract-level

Figure C.6: Change in residential properties located in the 100-year floodplain during the map update



Changes in the number of properties zoned inside the 100-year floodplain are expressed relative to the total number of residential properties in each census tract. The figure depicts changes bounded between -0.2 and 0.2 to improve color contrasts.

C.2 Additional results on removed and ignored properties

Table C.3: Characteristics of properties removed from 100-year floodplain during map updates

Dependent Variable:	Removed from Q3 100-year							
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<u>Variables</u>								
Constant	0.2649*** (0.0079)		0.4254*** (0.0131)		0.2807*** (0.0079)		0.4885*** (0.0122)	
Black or AA 25-50%	0.0978*** (0.0134)	0.0514*** (0.0092)	0.0750*** (0.0114)	0.0475*** (0.0089)				
Black or AA 50-75%	0.1910*** (0.0149)	0.0972*** (0.0112)	0.1638*** (0.0122)	0.0955*** (0.0111)				
Black or AA 75-100%	0.2270*** (0.0137)	0.0983*** (0.0126)	0.2349*** (0.0126)	0.1121*** (0.0134)				
Coastal			-0.3468*** (0.0095)	-0.2540*** (0.0109)			-0.3411*** (0.0097)	-0.2586*** (0.0109)
Inland fluvial			-0.2771*** (0.0103)	-0.2581*** (0.0074)			-0.2791*** (0.0103)	-0.2587*** (0.0074)
Inland pluvial			-0.2407*** (0.0071)	-0.2106*** (0.0060)			-0.2370*** (0.0075)	-0.2109*** (0.0060)
Income 25-50%			0.0344*** (0.0125)	0.0300*** (0.0093)			-0.0035 (0.0125)	0.0208** (0.0092)
Income 50-75%			0.0726*** (0.0135)	0.0407*** (0.0109)			0.0248* (0.0138)	0.0275*** (0.0106)
Income 75-100%			0.1447*** (0.0137)	0.0680*** (0.0127)			0.0791*** (0.0133)	0.0462*** (0.0121)
Hispanic 25-50%					0.0418*** (0.0122)	0.0373*** (0.0090)	0.0422*** (0.0105)	0.0389*** (0.0088)
Hispanic 50-75%					0.1481*** (0.0142)	0.0664*** (0.0112)	0.1261*** (0.0122)	0.0647*** (0.0113)
Hispanic 75-100%					0.2187*** (0.0145)	0.0842*** (0.0149)	0.1603*** (0.0129)	0.0722*** (0.0155)
<u>Fixed-effects</u>								
County FE		Yes		Yes		Yes		Yes
<u>Fit statistics</u>								
Observations	6,288,904	6,288,904	6,263,881	6,263,881	6,288,904	6,288,904	6,263,881	6,263,881
R ²	0.03223	0.26155	0.13326	0.30389	0.03090	0.25955	0.11971	0.30160
Within R ²		0.00435		0.06245		0.00165		0.05937

Clustered (Tract FE) standard-errors in parentheses
Signif. Codes: ***, 0.01, **, 0.05, *, 0.1

Linear probability regressions at the property level of the form $Y_i = \alpha_c + \beta X_i + \epsilon_i$, where Y_i is a dummy variable equal to 1 if the property used to be in the Q3 100-year floodplain and was removed during the map updates. α_c are county or tract fixed effects (indicated in the first column), and X_i is a vector of property and neighborhood-level characteristics. Race come from the Decennial Census while income is taken from the American Community Survey. For all models, the model is restricted to census tracts mapped in both Q3 and the NFHL19 product and to properties that were in the Q3 100-year floodplain. “Coastal,” “Inland fluvial” and “Inland pluvial” refers to the nature of flood risk in the FSF model.

Table C.4: Characteristics of properties removed from 100-year floodplain during map updates, by race

Dependent Variable:	Removed from Q3 100-year					
Model:	(1)	(2)	(3)	(4)	(5)	(6)
<u>Variables</u>						
Constant	0.3761*** (0.0107)		0.5169*** (0.0125)		0.2534*** (0.0080)	
AI or AN 25-50%	0.0211 (0.0160)	-0.0025 (0.0097)				
AI or AN 50-75%	0.0043 (0.0147)	-0.0221** (0.0103)				
AI or AN 75-100%	0.0238 (0.0161)	-0.0394*** (0.0145)				
White 25-50%			-0.0315* (0.0166)	0.0249** (0.0117)		
White 50-75%			-0.1374*** (0.0167)	-0.0029 (0.0124)		
White 75-100%			-0.2648*** (0.0144)	-0.0653*** (0.0132)		
Asian 25-50%					0.0462*** (0.0117)	0.0388*** (0.0080)
Asian 50-75%					0.2030*** (0.0130)	0.1108*** (0.0108)
Asian 75-100%					0.3726*** (0.0145)	0.2091*** (0.0139)
<u>Fixed-effects</u>						
County FE		Yes		Yes		Yes
<u>Fit statistics</u>						
Observations	6,288,904	6,288,904	6,288,904	6,288,904	6,288,904	6,288,904
R ²	0.00044	0.25876	0.04550	0.26100	0.07649	0.26799
Within R ²		0.00059		0.00361		0.01303

Clustered (Tract FE) standard-errors in parentheses
Signif. Codes: ***, 0.01, **, 0.05, *, 0.1

Linear probability regressions at the property level of the form $Y_i = \alpha_c + \beta X_i + \epsilon_i$, where Y_i is a dummy variable equal to 1 if the property used to be in the Q3 100-year floodplain and was removed during the map updates. α_c are county or tract fixed effects (indicated in the first column), and X_i is a vector of property and neighborhood-level characteristics. Race come from the Decennial Census while income is taken from the American Community Survey. For all models, the model is restricted to census tracts mapped in both Q3 and the NFHL19 product and to properties that were in the Q3 100-year floodplain. “Coastal,” “Inland fluvial” and “Inland pluvial” refers to the nature of flood risk in the FSF model.

Table C.5: Characteristics of FSF 100-year properties incorrectly ignored during map updates

Dependent Variable:	FSF 100-year, ignored during update							
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<u>Variables</u>								
Constant	0.5524*** (0.0082)		0.2412*** (0.0113)		0.5664*** (0.0076)		0.2490*** (0.0101)	
Black or AA 25-50%	0.0231* (0.0135)	0.0668*** (0.0087)	0.0420*** (0.0090)	0.0674*** (0.0081)				
Black or AA 50-75%	0.0524*** (0.0127)	0.1296*** (0.0098)	0.0775*** (0.0091)	0.1245*** (0.0095)				
Black or AA 75-100%	0.0633*** (0.0114)	0.1723*** (0.0102)	0.0881*** (0.0079)	0.1629*** (0.0110)				
Inland fluvial			0.2351*** (0.0102)	0.0016 (0.0136)			0.2333*** (0.0100)	0.0094 (0.0136)
Inland pluvial			0.4795*** (0.0082)	0.3033*** (0.0124)			0.4830*** (0.0083)	0.3133*** (0.0124)
Income 25-50%			0.0144 (0.0093)	0.0095 (0.0087)			-0.0002 (0.0090)	-0.0050 (0.0085)
Income 50-75%			0.0050 (0.0100)	-0.0141 (0.0095)			-0.0130 (0.0097)	-0.0345*** (0.0090)
Income 75-100%			0.0576*** (0.0093)	0.0131 (0.0100)			0.0336*** (0.0085)	-0.0144 (0.0090)
Hispanic 25-50%					-0.0077 (0.0119)	0.0525*** (0.0082)	0.0510*** (0.0073)	0.0553*** (0.0076)
Hispanic 50-75%					0.0251* (0.0135)	0.1021*** (0.0105)	0.0679*** (0.0089)	0.0962*** (0.0100)
Hispanic 75-100%					0.0621*** (0.0119)	0.1392*** (0.0129)	0.1035*** (0.0093)	0.1393*** (0.0129)
<u>Fixed-effects</u>								
County FE		Yes		Yes		Yes		Yes
<u>Fit statistics</u>								
Observations	8,979,280	8,979,280	8,959,472	8,959,472	8,979,280	8,979,280	8,959,472	8,959,472
R ²	0.00252	0.23326	0.18690	0.29080	0.00279	0.22842	0.18753	0.28738
Within R ²		0.01080		0.08652		0.00454		0.08213

Clustered (Tract FE) standard-errors in parentheses
 Signif. Codes: ***: 0.01, **: 0.05, *: 0.1

Linear probability regressions at the property level of the form $Ignored_i = \alpha_c + \beta X_i + \epsilon_i$, where $Ignored_i$ is a dummy variable equal to 1 when the property should be in the 100-year floodplain based on the FSF model but was left outside of the FEMA 100-year floodplain during the map update. α_c are a county- or tract- fixed effects, and X_i is a vector of property and neighborhood-level characteristics. Race and income come from the Decennial Census and the American Community Survey, respectively. The sample is restricted to properties in the FSF 100-year floodplain in census tracts mapped in both Q3 and the NFHL19 product. “Inland fluvial” and “Inland pluvial” refers to the nature of flood risk in the FSF model, with coastal flood risk as the omitted category.

Table C.6: FSF 100-year properties incorrectly ignored during map updates, by treatment year

Dependent Variable:	FSF 100-year, ignored during update	
Model:	(1)	(2)
<u>Variables</u>		
Constant	0.7375*** (0.0217)	
Update year 2001	-0.0557 (0.1091)	-0.0830 (0.1432)
Update year 2002	-0.1238*** (0.0476)	-0.0382 (0.1193)
Update year 2003	-0.4053*** (0.0443)	-0.1656 (0.1129)
Update year 2004	-0.2257*** (0.0694)	-0.0684 (0.1252)
Update year 2005	-0.3184*** (0.0399)	-0.0391 (0.1219)
Update year 2006	-0.2152*** (0.0312)	-0.0747 (0.1231)
Update year 2007	-0.0198 (0.0251)	-0.1217 (0.1186)
Update year 2008	-0.0455* (0.0253)	-0.1380 (0.1197)
Update year 2009	-0.1641*** (0.0257)	-0.0579 (0.1183)
Update year 2010	-0.0749*** (0.0238)	-0.0444 (0.1195)
Update year 2011	-0.1018*** (0.0242)	-0.1423 (0.1202)
Update year 2012	-0.1622*** (0.0315)	-0.0944 (0.1202)
Update year 2013	-0.0941*** (0.0254)	-0.2243* (0.1191)
Update year 2014	-0.2297*** (0.0252)	-0.1867 (0.1198)
Update year 2015	-0.1818*** (0.0259)	-0.1581 (0.1183)
Update year 2016	-0.1527*** (0.0258)	-0.2271* (0.1195)
Update year 2017	-0.1993*** (0.0294)	-0.1029 (0.1202)
Update year 2018	-0.2176*** (0.0285)	-0.2042* (0.1200)
Update year 2019	-0.1865*** (0.0386)	-0.1686 (0.1211)
<u>Fixed-effects</u>		
County FE		Yes
<u>Fit statistics</u>		
Observations	8,983,650	8,983,650
R ²	0.02807	0.22695
Within R ²		0.00273

Clustered (Tract FE) standard-errors in parentheses

Signif. Codes: ***: 0.01, **: 0.05, *: 0.1

Table C.7: FSF 100-year properties incorrectly ignored during map updates, by treatment year

Dependent Variable:	FSF 100-year, ignored during update					
Model:	(1)	(2)	(3)	(4)	(5)	(6)
<u>Variables</u>						
Constant	0.5520*** (0.0076)		0.6839*** (0.0078)		0.5282*** (0.0087)	
AI or AN 25-50%	-0.0012 (0.0123)	0.0049 (0.0076)				
AI or AN 50-75%	0.0295** (0.0118)	0.0197** (0.0081)				
AI or AN 75-100%	0.1146*** (0.0142)	0.0097 (0.0113)				
White 25-50%			-0.0420*** (0.0112)	-0.0116 (0.0078)		
White 50-75%			-0.1080*** (0.0129)	-0.0625*** (0.0083)		
White 75-100%			-0.1806*** (0.0112)	-0.1642*** (0.0100)		
Asian 25-50%					-0.0086 (0.0129)	0.0416*** (0.0079)
Asian 50-75%					0.0809*** (0.0114)	0.0878*** (0.0095)
Asian 75-100%					0.2236*** (0.0113)	0.1279*** (0.0107)
<u>Fixed-effects</u>						
County FE		Yes		Yes		Yes
<u>Fit statistics</u>						
Observations	8,979,280	8,979,280	8,979,280	8,979,280	8,979,280	8,979,280
R ²	0.00831	0.22506	0.01864	0.23395	0.02882	0.22866
Within R ²		0.00021		0.01168		0.00486
<u>Clustered (Tract FE) standard-errors in parentheses</u>						
<u>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</u>						

Linear probability regressions at the property level of the form $Ignored_i = \alpha_c + \beta X_i + \epsilon_i$, where $Ignored_i$ is a dummy variable equal to 1 when the property should be in the 100-year floodplain based on the FSF model but was left outside of the FEMA 100-year floodplain during the map update. α_c are a county- or tract- fixed effects, included in some regressions only. The race and ethnicity variable from the Decennial Census 2010, focusing on single-race/ethnicity individuals (using the race / ethnicity alone variables). The sample is restricted to properties in the FSF 100-year floodplain in census tracts mapped in both Q3 and the NFHL19 product.

Table C.8: FSF 100-year properties incorrectly ignored during map updates, quantiles conditional on mapping

Dependent Variable:	FSF 100-year, ignored during update						
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<u>Variables</u>							
(Intercept)	0.3091*** (0.0081)			0.5499*** (0.0079)		0.2440*** (0.0110)	
Inland fluvial	0.2273*** (0.0104)	0.0231* (0.0138)	-0.1800*** (0.0117)			0.2337*** (0.0102)	0.0024 (0.0134)
Inland pluvial	0.4788*** (0.0085)	0.3216*** (0.0127)	0.1988*** (0.0107)			0.4780*** (0.0082)	0.3038*** (0.0122)
Black or AA 25-50% (c)				0.0434*** (0.0136)	0.0748*** (0.0086)	0.0496*** (0.0099)	0.0733*** (0.0082)
Black or AA 50-75% (c)				0.0554*** (0.0126)	0.1265*** (0.0099)	0.0729*** (0.0085)	0.1228*** (0.0097)
Black or AA 75-100% (c)				0.0658*** (0.0114)	0.1594*** (0.0100)	0.0853*** (0.0082)	0.1548*** (0.0107)
Income 25-50% (c)						0.0184** (0.0088)	0.0149* (0.0081)
Income 50-75% (c)						0.0101 (0.0099)	-0.0087 (0.0093)
Income 75-100% (c)						0.0684*** (0.0094)	0.0221** (0.0100)
<u>Fixed-effects</u>							
County FE		Yes	Yes		Yes		Yes
Tract FE			Yes				
<u>Fit statistics</u>							
Observations	8,983,650	8,983,650	8,983,650	8,979,280	8,979,280	8,959,472	8,959,472
R ²	0.18197	0.28450	0.51331	0.00299	0.23290	0.18713	0.29074
Within R ²		0.07696	0.06517		0.01033		0.08645

Clustered (Tract FE) standard-errors in parentheses

Signif. Codes: ***: 0.01, **: 0.05, *: 0.1

Linear probability regressions at the property level of the form $Ignored_i = \alpha_c + \beta X_i + \epsilon_i$, where $Ignored_i$ is a dummy variable equal to 1 when the property should be in the 100-year floodplain based on the FSF model but was left outside of the FEMA 100-year floodplain during the map update. α_c are a county- or tract- fixed effects, included in some regressions only. The race and ethnicity variable from the Decennial Census 2010, focusing on single-race/ethnicity individuals (using the race / ethnicity alone variables), and with quantiles defined conditional on the tracts observed in Q3 and NFHL19. The sample is restricted to properties in the FSF 100-year floodplain in census tracts mapped in both Q3 and the NFHL19 product.

C.3 Selection into treatment and event studies

Table C.9: Census tracts summary statistics by year of the Digital Flood Insurance Rate Map.

Treatment year	N	Policies 2010	SFHA share 2010	Policies 2010 per property	Med. income	Relative change in 100-year	Area	Density	Share Black	Has Q3 data
2005	6067	58 (181)	0.35 (0.34)	0.04 (0.11)	33939 (12511)	-0.01 (0.06)	47 (231)	1209 (1212)	0.13 (0.21)	0.89
2006	3384	48 (224)	0.36 (0.34)	0.02 (0.08)	34689 (14599)	-0.01 (0.04)	43 (161)	1265 (2370)	0.14 (0.21)	0.74
2007	6303	58 (170)	0.25 (0.31)	0.04 (0.1)	33242 (15038)	0 (0.04)	45 (227)	5105 (8570)	0.22 (0.28)	0.66
2008	5855	48 (173)	0.36 (0.34)	0.03 (0.12)	32302 (12426)	-0.01 (0.09)	64 (274)	1573 (2583)	0.15 (0.24)	0.72
2009	7547	62 (194)	0.37 (0.35)	0.04 (0.13)	32265 (12265)	-0.01 (0.14)	75 (372)	1386 (2156)	0.14 (0.21)	0.74
2010	6431	25 (93)	0.36 (0.34)	0.02 (0.09)	32182 (13876)	0 (0.04)	135 (603)	1074 (1837)	0.12 (0.21)	0.64
2011	4356	24 (82)	0.42 (0.36)	0.03 (0.11)	28920 (10417)	-0.01 (0.08)	126 (360)	718 (1101)	0.11 (0.19)	0.55
2012	4124	42 (136)	0.42 (0.35)	0.03 (0.1)	30283 (10989)	0 (0.08)	101 (337)	1049 (1512)	0.15 (0.27)	0.68
2013	1233	37 (98)	0.37 (0.34)	0.02 (0.07)	32295 (13018)	-0.01 (0.04)	99 (295)	765 (1377)	0.11 (0.19)	0.59
2014	2149	207 (396)	0.47 (0.37)	0.12 (0.21)	31496 (11985)	-0.14 (0.31)	58 (208)	1205 (1327)	0.14 (0.21)	0.8
2015	1408	75 (178)	0.48 (0.34)	0.06 (0.14)	30024 (10123)	-0.1 (0.24)	102 (488)	819 (1072)	0.15 (0.26)	0.75
2016	1146	68 (183)	0.38 (0.36)	0.05 (0.13)	32909 (11571)	-0.02 (0.08)	80 (450)	1122 (1134)	0.23 (0.3)	0.87
2017	1042	223 (373)	0.39 (0.36)	0.16 (0.25)	31864 (10253)	-0.09 (0.23)	70 (229)	1124 (1029)	0.13 (0.2)	0.81
2018	432	88 (216)	0.35 (0.34)	0.05 (0.11)	28982 (9626)	-0.01 (0.06)	103 (257)	785 (872)	0.05 (0.12)	0.74
2019	767	65 (286)	0.37 (0.35)	0.09 (0.25)	35066 (11671)	-0.01 (0.07)	77 (495)	1162 (1008)	0.08 (0.15)	0.82
	4701	76 (248)	0.43 (0.36)	0.06 (0.18)	29953 (10932)		588 (3477)	584 (956)	0.06 (0.14)	0.6

Census tracts with a flood insurance rate map prior to 2005 re-coded as 2005 for conciseness

C.4 Selection into treatment timing and implementation

While FEMA started to actively modernize flood maps in 2006, the *timing* and *implementation* of the map updates were not completely random. The Map Modernization and then Risk Mapping programs had concurrent objectives of providing digital flood maps covering the largest population possible and in areas most vulnerable to flood risk. These objectives can conceivably conflict with one another. Table C.10 offers a regression-based summary of how these conflicts were resolved. In the first two columns, the dependent variable is the year in which the census tract receives its first digital flood map, while in columns 3 and 4 the dependent variable is the net share of residential properties in the census tract that were rezoned inside the 100-year floodplain on the updated map (counting rezoning outside of the 100-year floodplain as negative). All models include fixed effects for each FEMA Region, as the roll-out of digital flood maps is primarily decentralized at this level.

Columns 1 and 2 reveal that census tracts with greater population density received a digital flood map earlier on average: a 1% increase in census tract density is associated with a treatment date that is between 0.1 and 0.2 years earlier, consistent with some targeting of populous communities in order to comply with policy mandates and deadlines.

To study whether areas more vulnerable to flooding were more likely to receive an updated flood map earlier, I use a measure of predicted average annual economic loss per property produced by the First Street Foundation (column 1). Interestingly, the point estimate is *positive*, suggesting that more flood-prone areas were treated *later*, although the effect is small and

Table C.10: Selection into treatment timing and implementation

	<i>Dependent variable:</i>			
	Treatment year		Share rezoned inside 100-year f.p.	
	(1)	(2)	(3)	(4)
Population density (IHS)	-0.204*** (0.076)	-0.272*** (0.055)	-0.012** (0.005)	-0.003* (0.002)
Median income (IHS)	-0.757*** (0.288)	-0.840*** (0.285)	-0.014* (0.009)	-0.002 (0.003)
Share African Americans (IHS)	-0.082 (0.720)	0.045 (0.695)	-0.017 (0.024)	-0.035* (0.019)
Disaster declaration prior to treatment			0.019 (0.017)	0.011 (0.011)
Average Annual Loss (IHS)	0.061 (0.049)		-0.006* (0.003)	
Insurance policies/property, 2008 (IHS)		4.624*** (1.435)		-0.538*** (0.157)
Treatment year			-0.005** (0.002)	-0.002** (0.001)
Fixed Effects	FEMA Region	FEMA Region	FEMA Region	FEMA Region
Mean outcome	2009.2	2009.2	-0.021	-0.02
Observations	51,892	52,154	28,039	28,039
R ²	0.085	0.102	0.073	0.257

Note:

*p<0.1; **p<0.05; ***p<0.01
All independent variables are transformed using the inverse hyperbolic sine
Standard errors clustered at the county level

noisy. I also estimate a model that captures flood vulnerability through the share of residential properties within a census tract that are covered by flood insurance in 2008 (column 2). This proxy for flood vulnerability is not ideal, as the main text of the paper shows that flood maps have large impacts on insurance take-up. But given that insurance take-up is a metric directly observable by policymakers, it potentially provides valuable information about where flood maps might be most used. Contrary to official policy guidelines, places that had higher flood insurance coverage in 2008 were more likely to receive their digital flood map *later*: a ten percent increase in the rate of residential property insurance coverage is associated with a treatment timing delayed by more than .4 years. Although the previous estimates are correlational, they suggest that cohorts treated in different years might respond differently to treatment.

To assess the extent of heterogeneity in policy implementation (the *intensity* of treatment), I now focus on the rezoning of properties inside or outside of the 100-year floodplain (columns 3 and 4 in table C.10). Interestingly, census tracts with higher predicted flood losses (column 3) or greater insurance coverage in 2008 (column 4) experienced *less* rezoning inside the 100-year floodplain. Census tracts that were treated earlier also saw less rezoning of properties inside the 100-year floodplain: after accounting for the other census tracts characteristics, a one-year increase in treatment year implies on average a -0.004 decline in the share of properties rezoned inside the 100-year floodplain (about 20% of the outcome's average value in the sample).

Overall, the previous results show that flood maps were first modernized in areas with slightly higher population density, consistent with the policy mandate to cover most of the nation's population with digital flood mapping products, but also in areas less vulnerable to

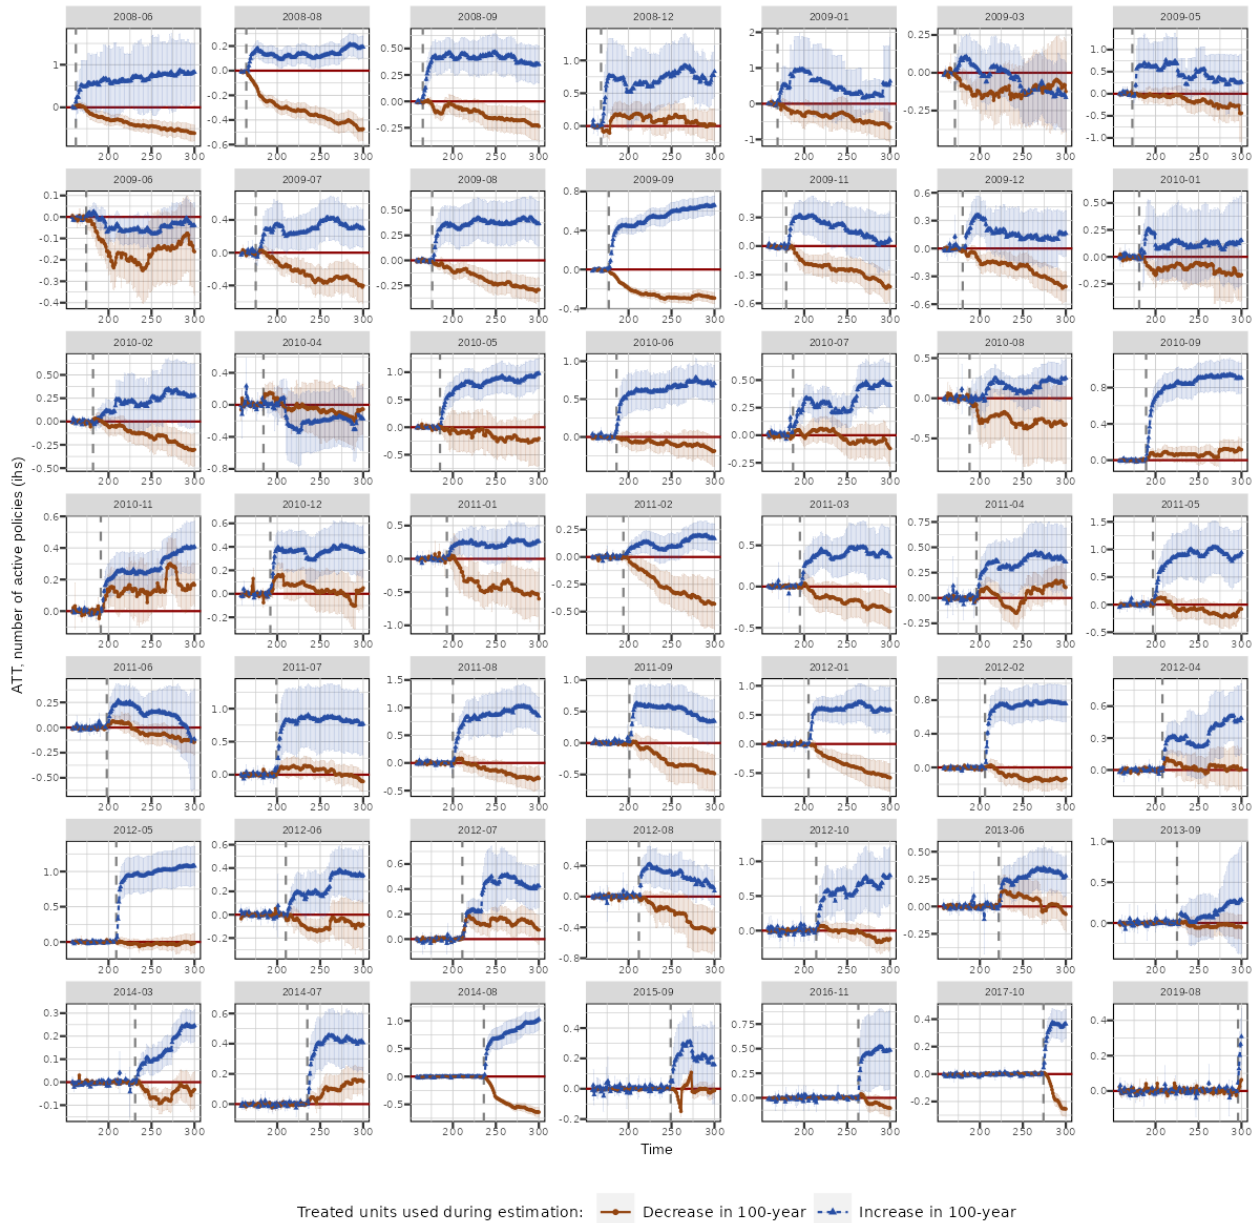
flood risk, which is in contradiction with the policy mandate to focus on flood-prone areas. Although I cannot provide credibly causal evidence on what caused the delay in flood map updates in areas more flood-prone, anecdotal evidence and discussions with floodplain managers suggest that it was due to (i) the complexity of modelling flood risk in these areas, and (ii) local homeowners lobbying against new flood maps.

Finally, this paper only examines flood map updates mandated by FEMA. Individuals who disagree with their floodplain designation can petition FEMA with a Letter of Map Changes (LOMC) to have their property removed from the SFHA, thus leading to additional removals from the 100-year floodplain. These individual-initiated map changes are outside the scope of this paper, and lead me to under-estimate the total number of properties removed from the 100-year floodplain. Such measurement error is limited by the relatively small number of LOMCs.

D Additional event-study estimates

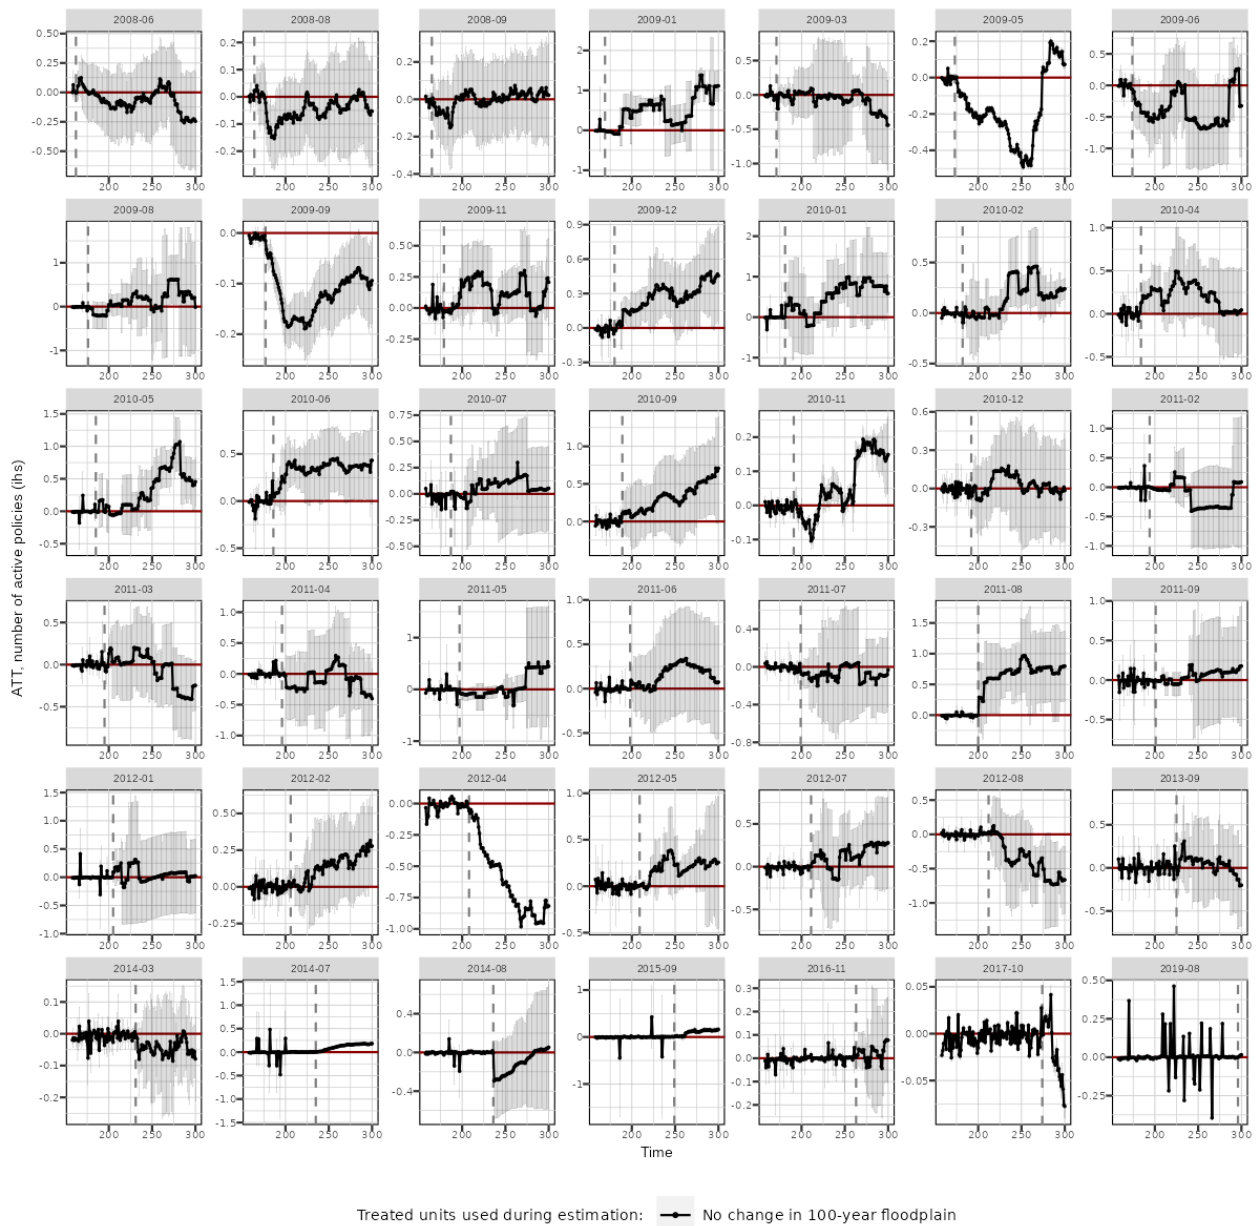
D.1 Event-study by cohort and rezoning

Figure D.7: Cohort-specific event study estimates of the impacts of flood map updates on flood insurance take-up



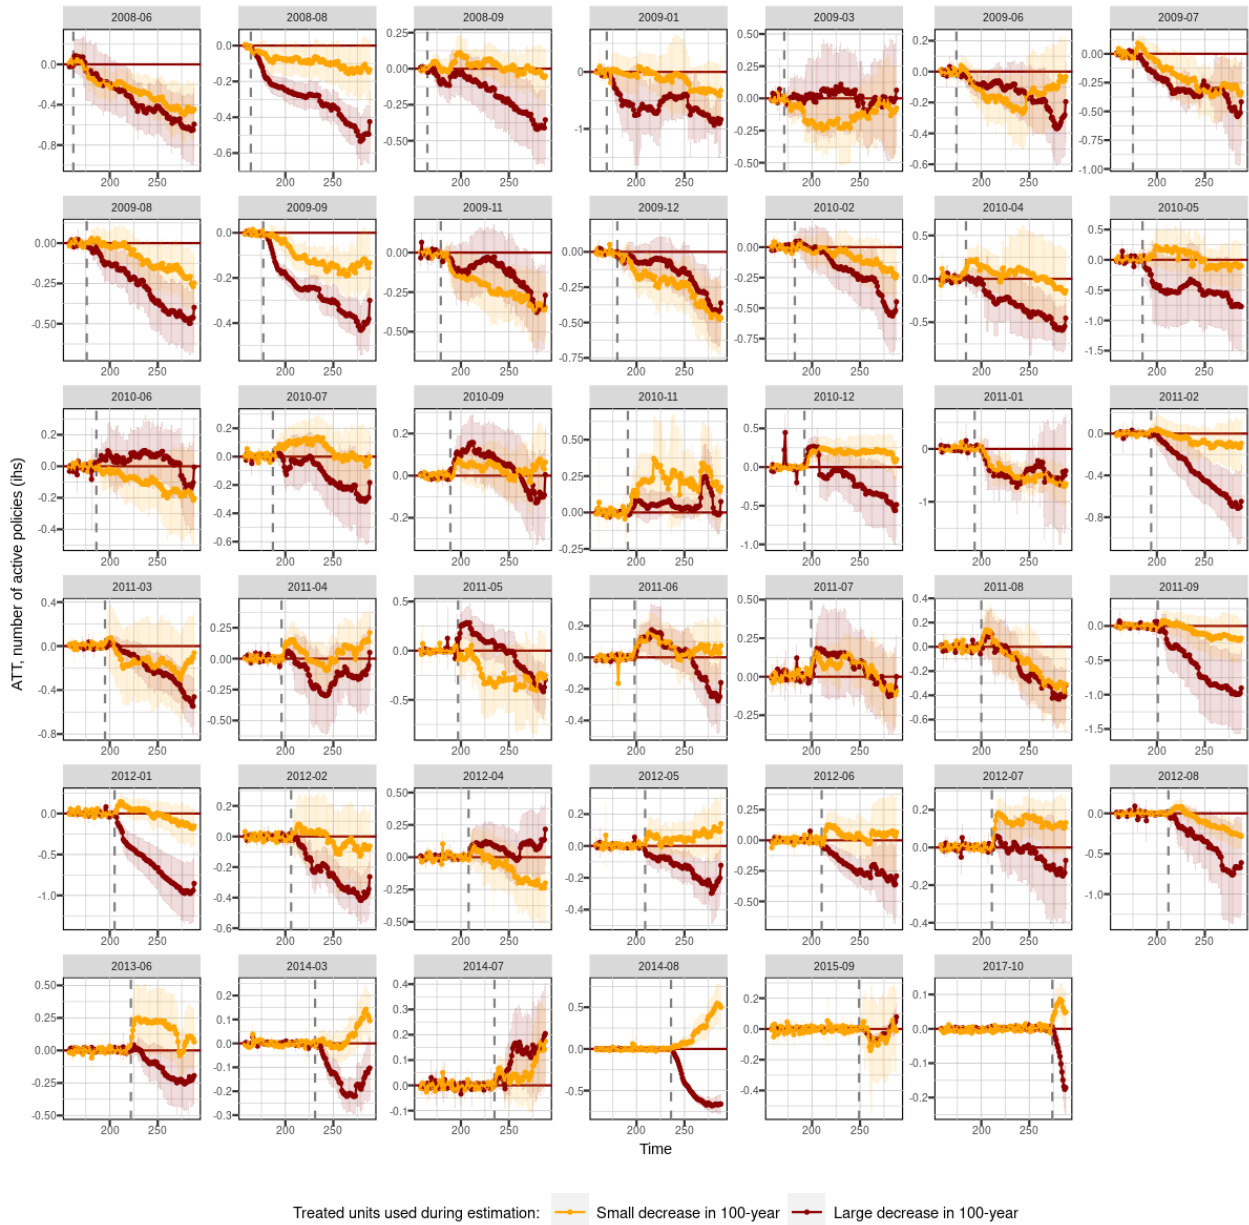
ATT estimates obtained using Callaway/Sant’Anna-type regressions. Each facet presents estimates of the average treatment effect on the treated of the impact of flood map update on insurance demand for a specific cohort (defined by the year and month of treatment). Within each cohort, the event-study are estimated separately using census tracts where the updated flood map increased (blue) or decreased (brown) the number of residential properties in the 100-year floodplain by more than 1% relative to the total number of residential properties in the census tract. The control groups comprise census tracts that have not yet received a digital flood map at the time of treatment but will receive one later with the same rezoning as the treated groups (either increase or decrease in the 100-year floodplain). Error bars represent 95% confidence intervals using the multiplier bootstrap. For clarity, only the 42 largest cohorts are represented (out of 117 cohorts).

Figure D.8: Cohort-specific event study estimates of the impacts map updates on take-up, no change in 100-year floodplain



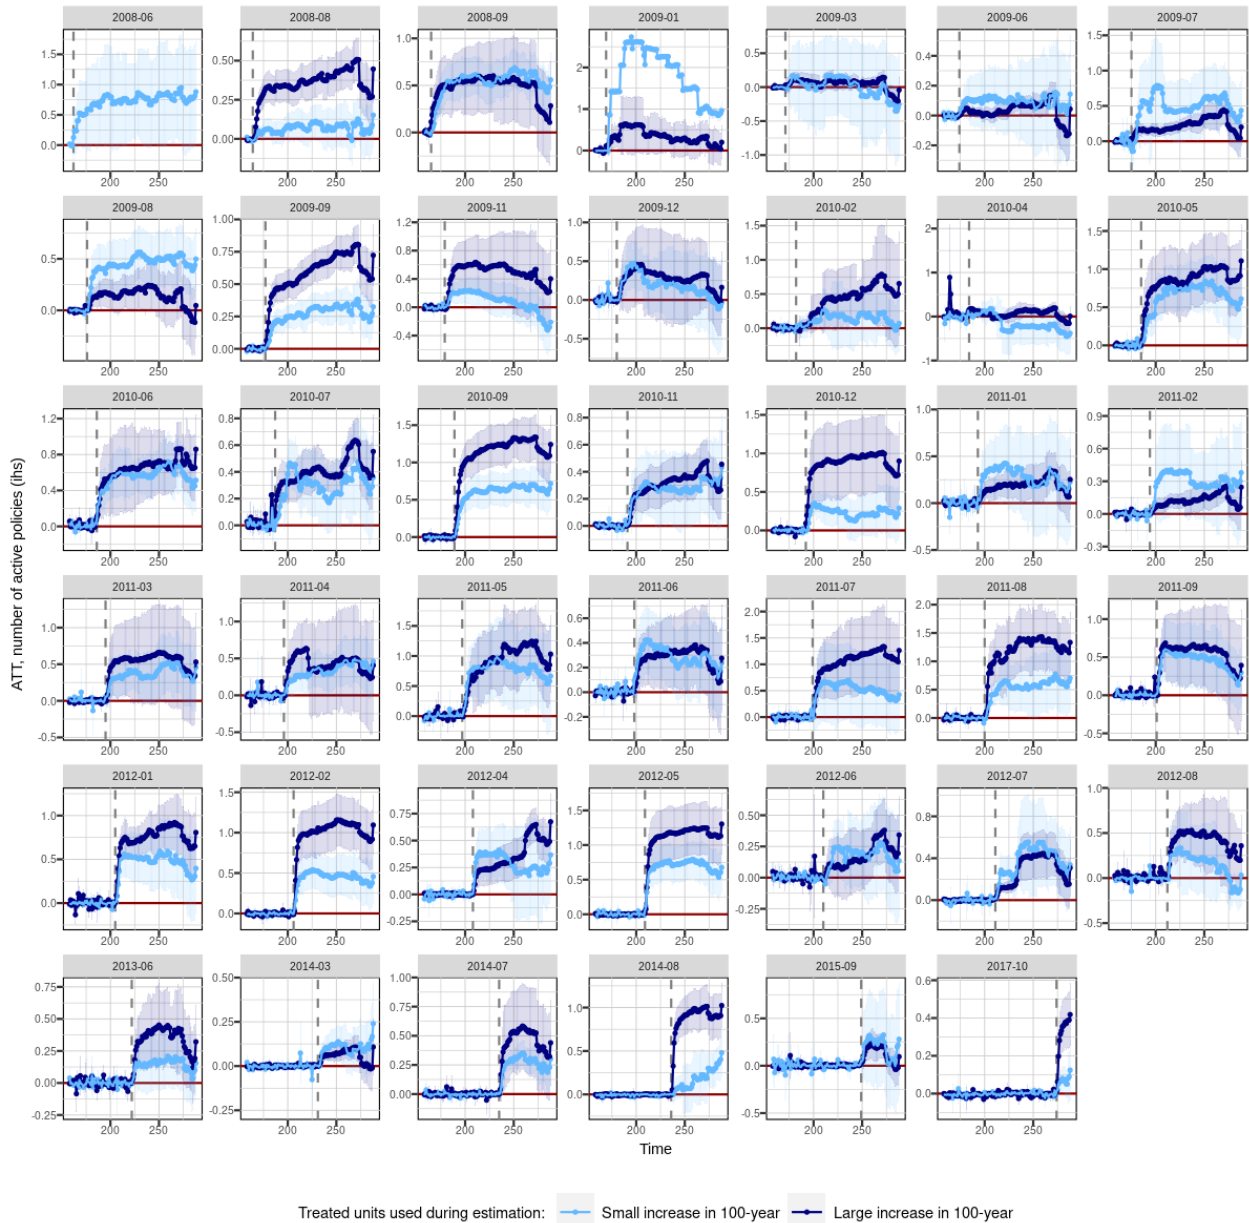
Each facet presents estimates of the average treatment effect on the treated of the impact of flood map update on insurance take-up for a specific cohort (defined by the year and month of treatment). Within each cohort, the event-study are estimated focusing on tracts where the new map did not change the number of properties in any floodplain. The control groups comprise census tracts that have not yet received a digital flood map at the time of treatment, but will receive one later with the same direction of 100-year floodplain rezoning. Error bars represent 95% confidence intervals using the multiplier bootstrap. Small within-treatment cohort can lead to missing confidence intervals.

Figure D.9: Cohort-specific event study estimates of the impacts map updates on take-up, small and large decrease



Each facet presents estimates of the average treatment effect on the treated of the impact of flood map update on insurance take-up for a specific cohort (defined by the year and month of treatment). Within each cohort, the event-study are estimated focusing on tracts where the new map removed between 1 and 3% of properties from the 100-year floodplain (“Small decrease,” in orange), or more than 3% of properties (“Large decrease,” in dark red). The control groups comprise census tracts that have not yet received a digital flood map at the time of treatment, but will receive one later with the same direction and intensity of 100-year floodplain rezoning. Error bars represent 95% confidence intervals using the multiplier bootstrap. Small within-treatment cohort can lead to missing confidence intervals.

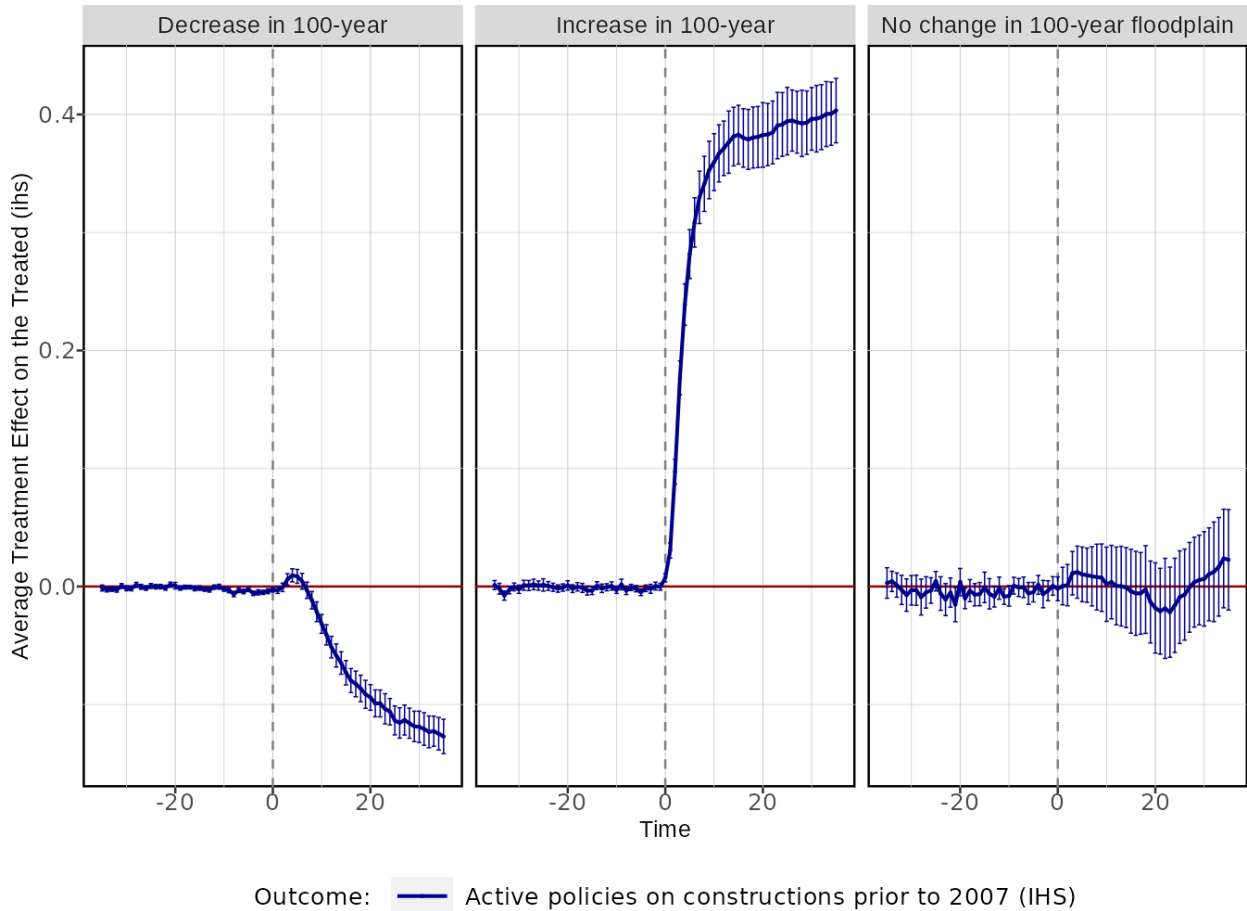
Figure D.10: Cohort-specific event study estimates of the impacts map updates on take-up, small and large increase



Each facet presents estimates of the average treatment effect on the treated of the impact of flood map update on insurance take-up for a specific cohort (defined by the year and month of treatment). Within each cohort, the event-study are estimated focusing on tracts where the new map added between 1 and 3% of properties to the 100-year floodplain (“Small increase,” in light blue), or more than 3% of properties (“Large increase,” in dark blue). The control groups comprise census tracts that have not yet received a digital flood map at the time of treatment, but will receive one later with the same direction and intensity of 100-year floodplain rezoning. Error bars represent 95% confidence intervals using the multiplier bootstrap. Small within-treatment cohort can lead to missing confidence intervals.

D.2 Event-study robustness


Figure D.11: Aggregated event study estimates of the impacts of flood map updates on flood insurance take-up, robustness tests using early constructions properties



The outcomes displayed are the number of active insurance policies covering properties that were constructed prior to 2008 (solid blue line), where the construction date is taken from the insurance data. Each facet represents average treatment effects for a different treated group, using treated census tracts where the flood map update increased, decreased, or did not change the number of properties zoned inside the 100-year floodplain (first, second and third facet respectively). The control groups comprise census tracts that have not yet received a digital flood map at the time of treatment, but will receive one later with the same direction of 100-year floodplain rezoning. The outcome variables are transformed using the Inverse Hyperbolic Sine (IHS). Error bars represent 95% confidence intervals using the multiplier bootstrap.


D.3 Event-study heterogeneity

Figure D.12: Event-study, race and ethnicity heterogeneity



ATT estimates obtained using Callaway-Sant'Anna type regressions. The outcome variable is the number of active insurance policies in the tract, transformed using the Inverse Hyperbolic Sine (IHS). Rows represent average treatment effects for different treatment groups, using treated census tracts where the flood map update decreased (top row) or increased (bottom row) the number of residential properties in the 100-year floodplain by more than 1% relative to the total number of residential properties in the census tract. Vertical facets focus on different heterogeneity variables. The control groups comprise not-yet-treated census tracts that later receive a flood map with a similar floodplain rezoning direction as the treated groups and that are within the same heterogeneity quartile of the variable being investigated. Error bars represent 95% confidence intervals using the multiplier bootstrap.

Figure D.13: Event-study, additional heterogeneity



ATT estimates obtained using Callaway-Sant'Anna type regressions. The outcome variable is the number of active insurance policies in the tract, transformed using the Inverse Hyperbolic Sine (IHS). Rows represent average treatment effects for different treatment groups, using treated census tracts where the flood map update decreased (top row) or increased (bottom row) the number of residential properties in the 100-year floodplain by more than 1% relative to the total number of residential properties in the census tract. Vertical facets focus on different heterogeneity variables. The control groups comprise not-yet-treated census tracts that later receive a flood map with a similar floodplain rezoning direction as the treated groups and that are within the same heterogeneity quartile of the variable being investigated. Error bars represent 95% confidence intervals using the multiplier bootstrap.

E Synthetic controls

E.1 Augmented synthetic controls: methodology

Formally, the synthetic control estimator proceeds in two steps. First, we solve the “standard” synthetic control program:

$$\begin{aligned} \min_{\gamma} \quad & \theta_x \|\mathbf{X}_1 - \mathbf{X}'_0 \gamma\|_2^2 + \theta_z \|\mathbf{Z}_1 - \mathbf{Z}_0 \cdot \gamma\|_2^2 + \zeta \sum_{W_i=0} f(\gamma_i) \\ \text{subject to} \quad & \sum_{W_i=0} \gamma_i = 1 \\ & \gamma_i \geq 0 \quad i : W_i = 0 \end{aligned} \tag{8}$$

where the goal is to find the vector or weights γ in the simplex

$\Delta^{N_0} = \{\gamma \in \mathbb{R}^{N_0} \mid \gamma_i \geq 0 \forall i, \sum_i \gamma_i = 1\}$ that minimizes the synthetic control objective function. This function is made of three parts: (i) the L2-norm (Euclidean distance) between the pre-treatment outcome of the treated census tract \mathbf{X}_1 and the control census tracts \mathbf{X}_0 , (ii) the L2-norm between the pre-treatment covariates of the treated census tract \mathbf{Z}_1 and the control census tracts \mathbf{Z}_0 , and (iii) a term that penalizes the dispersion of the weights assigned to control units (those with $W_i = 0$), for some function f and a positive hyperparameter ζ .⁴⁸ The weights θ_x and θ_z govern the relative importance of the deviations between lagged outcomes and covariates in the minimization program.⁴⁹

In a second step, we “augment” the synthetic control to estimate the (counterfactual) potential outcome of the treated unit:

$$\begin{aligned} \hat{Y}_{1T}^{\text{aug}}(0) &= \sum_{W_i=0} \hat{\gamma}_i^{\text{scm}} Y_{iT} + \left(\hat{m}_{1T} - \sum_{W_i=0} \hat{\gamma}_i^{\text{scm}} \hat{m}_{iT} \right) \\ &= \hat{m}_{1T} + \sum_{W_i=0} \hat{\gamma}_i^{\text{scm}} (Y_{iT} - \hat{m}_{iT}) \end{aligned} \tag{9}$$

where $\hat{\gamma}_i^{\text{scm}}$ are the solutions to the program in equation 8, Y_{iT} are the post-treatment outcomes, and \hat{m}_{iT} is an estimator of the post-treatment control potential outcome for unit i . In the standard synthetic control case, \hat{m}_{iT} is just a constant. I follow Ben-Michael et al.

⁴⁸The initial applications of the synthetic control methods did not include this penalty term – it is discussed in footnote 10 of Abadie et al. (2015) as a way to select weights when the minimization of the other parts of the objective function has multiple solutions. Different choices of penalty functions exist; see for instance Douchenko and Imbens (2016) for a discussion. In this paper I implement the Ridge-Augmented synthetic control approach, for which Ben-Michael et al. (2021) showed that the penalty term replaces the simplex constraints with the form $f(\gamma_i) = (\gamma_i - \hat{\gamma}_i^{\text{scm}})^2$ (deviations from the standard synthetic control weights are penalized).

⁴⁹Following Ben-Michael et al. (2021), in my preferred approach I use $\theta_x = \theta_z = 1$.

(2021) and use a ridge regression for the outcome model.⁵⁰

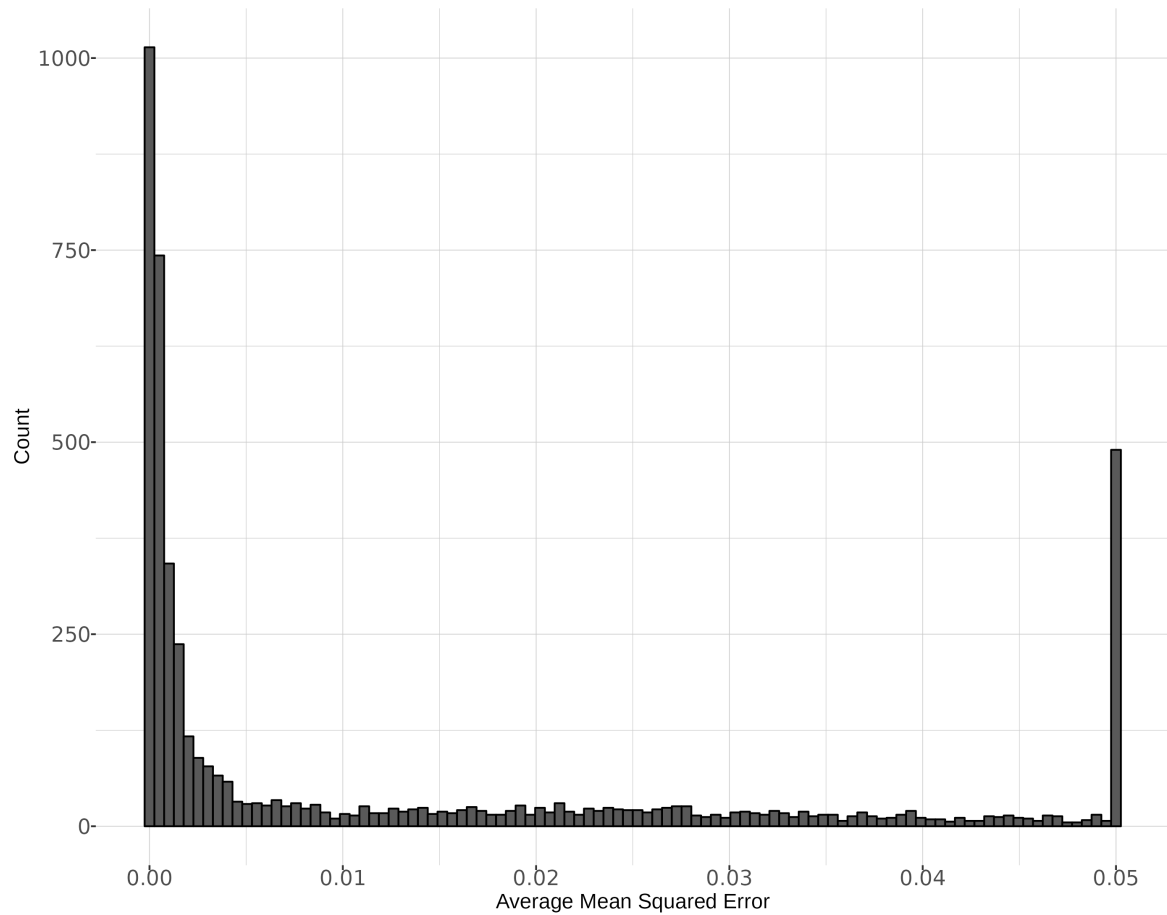
(Augmented) synthetic controls allow for the flexible estimation of treatment effect heterogeneity and can uncover whether differences in average treatment effects between groups are only driven by a small number of units within each group. They also greatly mitigate the shortcoming of the event studies presented above. First, treatment effects are estimated separately for each census tract that receives a new flood map during our observation window. Under the standard Stable Unit Treatment Value Assumption, the staggered nature of the treatment does not contaminate the estimated treatment effect.⁵¹

⁵⁰This choice of estimator has attractive properties, notably an improvement in pre-treatment fit relative to the standard synthetic control model, and a reduction in estimation error under linear and latent-factor data generating processes.

⁵¹Note that the spillover effects uncovered in Section 4.4 are not inconsistent with the Stable Unit Treatment Value Assumption, as these spillovers are found within each census tract, and not between different tracts.


E.2 Additional synthetic control results

Figure E.14: Normalized root-mean squared errors, synthetic controls



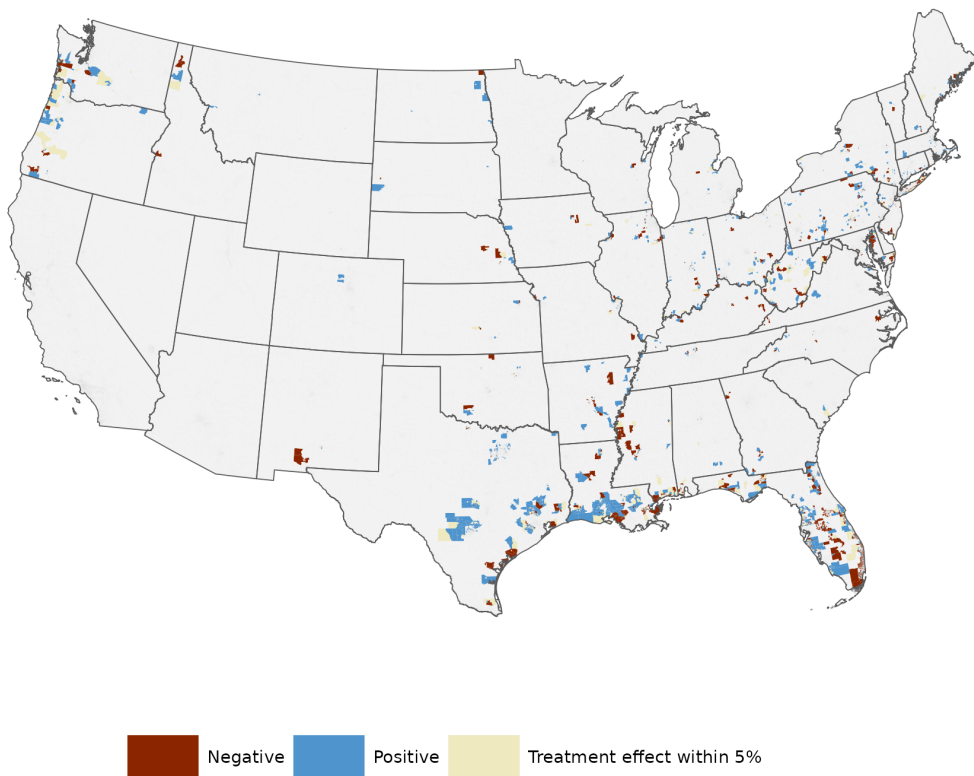
Histogram of normalized root mean squared errors between each individual tract and its (augmented) synthetic control. Synthetic controls were estimated for tracts with at least 20 insurance policies at all time, observed at least 24 months post-treatment. Matching was performed on pre-treatment outcome, share of post-FIRM policies in the tract, and share of policies in the 100-year floodplain.

Figure E.15: Normalized root-mean squared errors and number of units in synthetic control donor pool



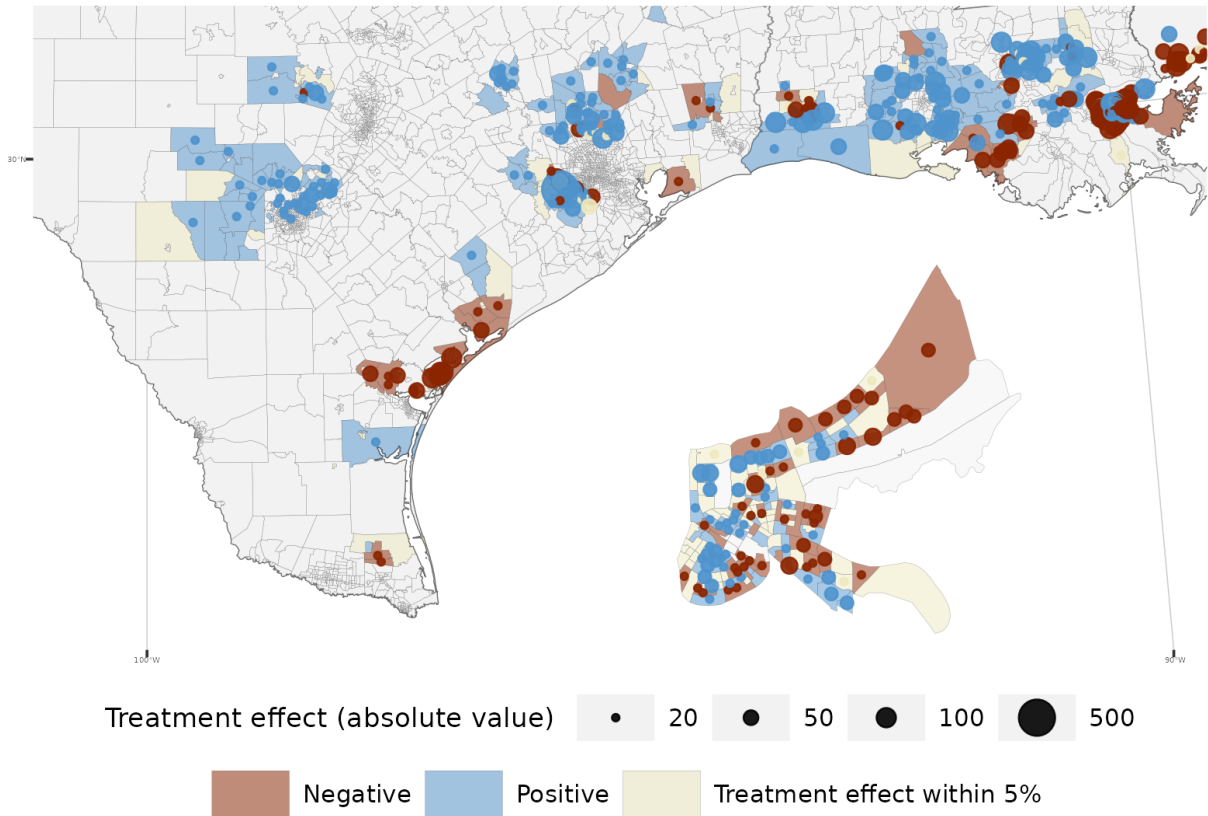
Synthetic controls were estimated for tracts with at least 20 insurance policies at all time, observed at least 24 months post-treatment. Matching was performed on pre-treatment outcome, share of post-FIRM policies in the tract, and share of policies in the 100-year floodplain.

Figure E.16: Treatment effect of the flood map update on flood insurance take-up after two years, synthetic controls



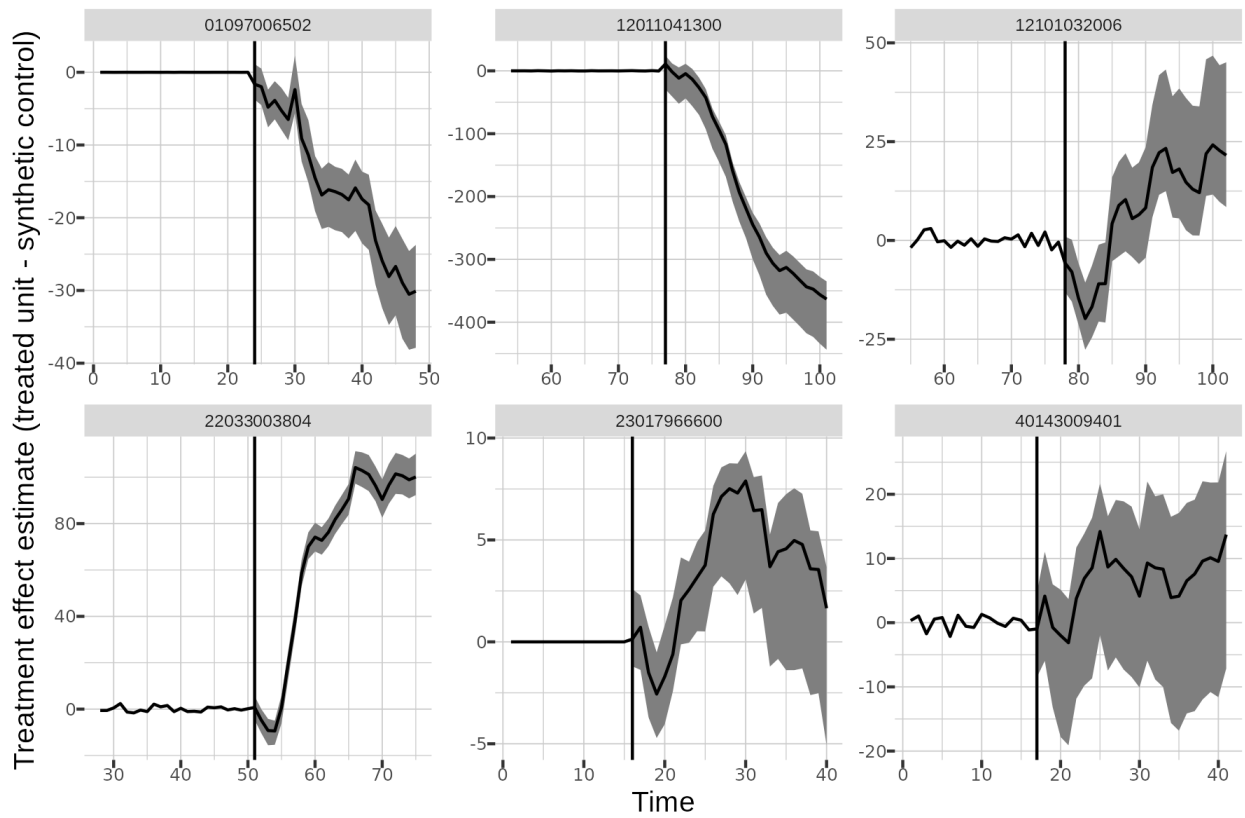
Synthetic controls were estimated for tracts with at least 20 insurance policies at all time, observed at least 24 months post-treatment. Matching was performed on pre-treatment outcome, share of post-FIRM policies in the tract, and share of policies in the 100-year floodplain.

Figure E.17: Synthetic control estimates in Texas and New Orleans




Each dot represents a census tract-specific treatment effect estimate of the impact of the flood map update on the flood insurance take-up, 24 months after the map update. Red and blue dots show negative and positive treatment effects, respectively. The size of the dot represents the absolute value of the treatment effect. An enlarged view of New Orleans is presented. Treatment effects are estimated using synthetic controls augmented by ridge regression. For each treated unit, the donor pool comprises never-treated census tracts within the same FEMA region.

Figure E.18: Synthetic control estimates for insurance take-up in six census tracts



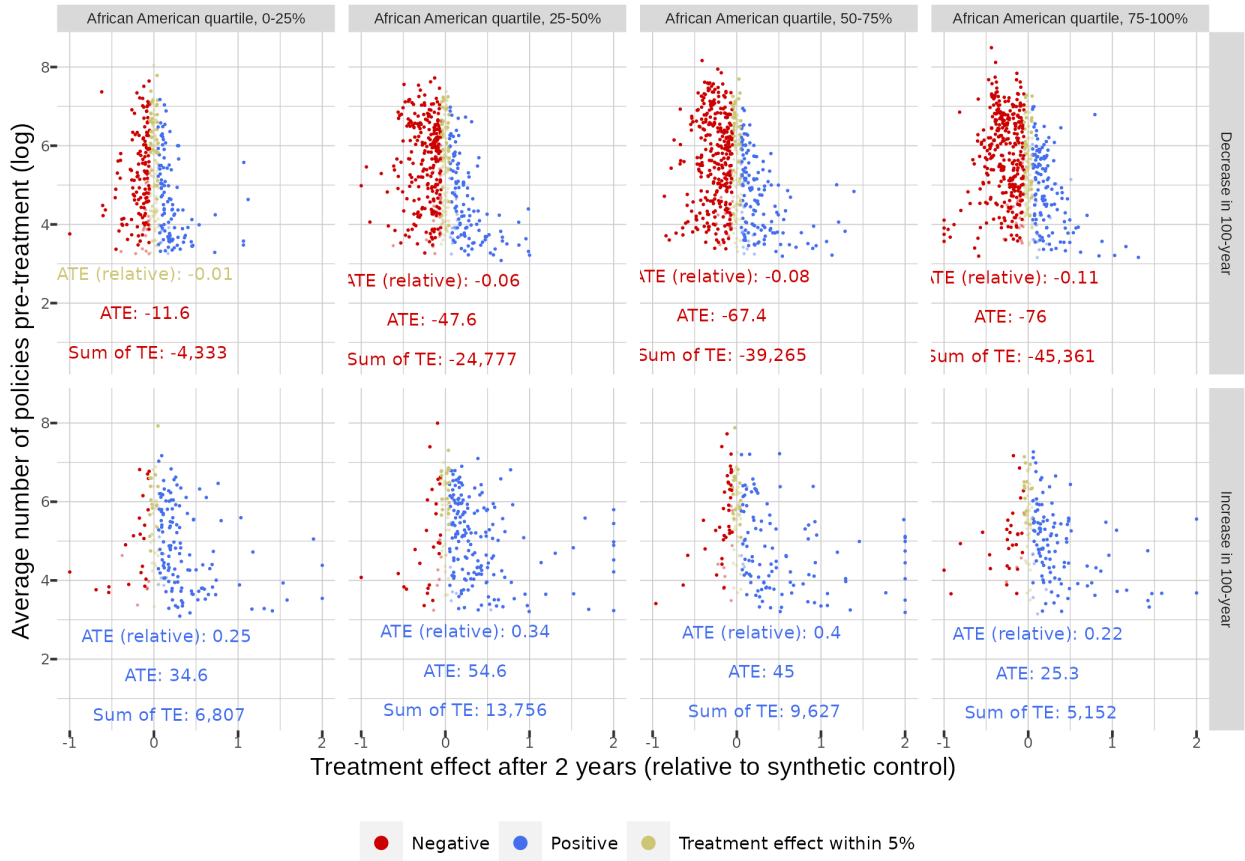
Dynamic treatment effects estimates of the impacts of flood map updates on flood insurance take-up in six census tracts, estimated by synthetic controls augmented with ridge regression. The solid black line shows the difference between the census tract outcome (number of active policies) and the synthetic control constructed for this tract. The vertical bar denotes the end of the pre-treatment optimization period, and the grey ribbon depicts 90% jackknife+ confidence intervals.

Figure E.19: Treatment effect of the flood map update on flood insurance take-up after two years, by disaster declaration history, synthetic controls




Synthetic controls were estimated for tracts with at least 20 insurance policies at all time, observed at least 24 months post-treatment. Matching was performed on pre-treatment outcome, share of post-FIRM policies in the tract, and share of policies in the 100-year floodplain.

Figure E.20: Estimated synthetic control treatment effects of the impact of the map update on demand for insurance, by quartiles of neighborhoods' share of African Americans.



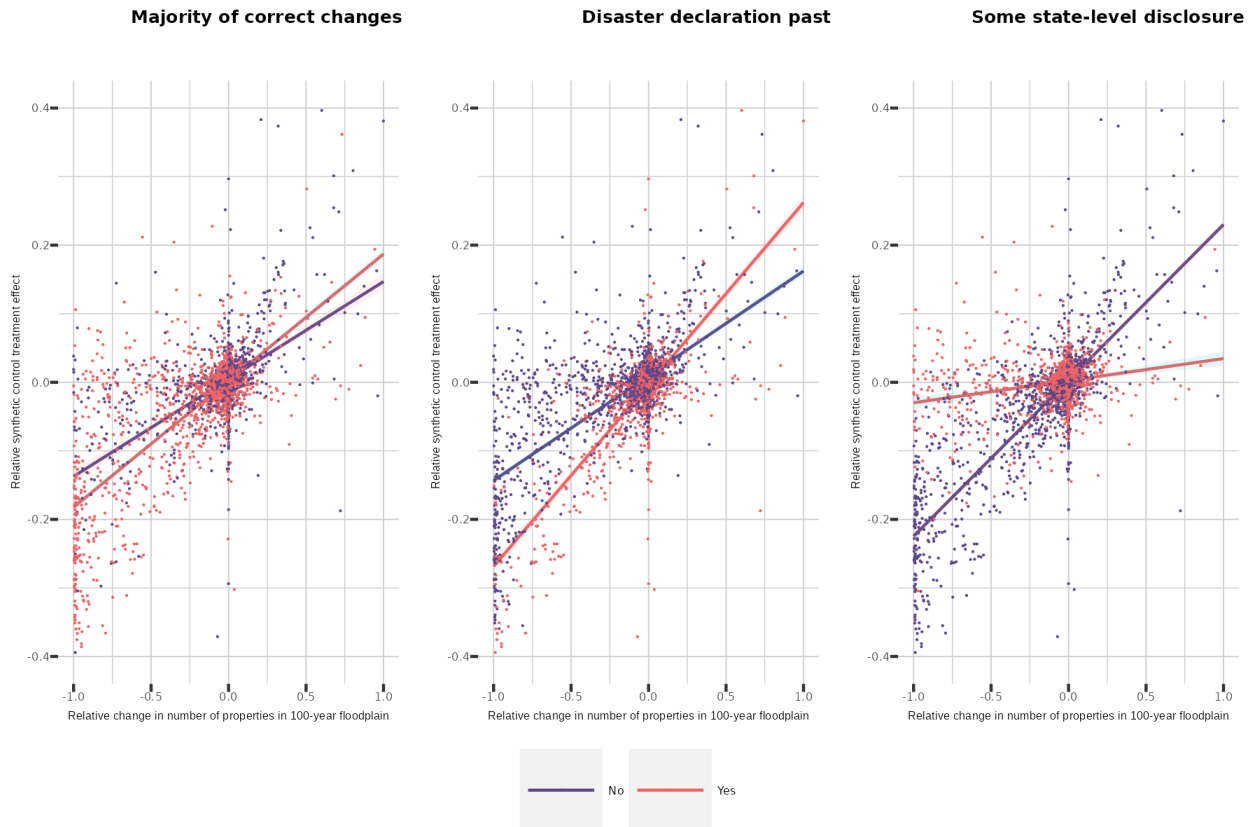
Each dot represents a census tract-specific treatment effect estimate of the impact of the flood map update on flood insurance take-up after 24 months. Red and blue dots show negative and positive treatment effects, respectively. The facets separate census tracts where the flood map update increased or decreased the number of properties zoned inside the 100-year floodplain by more than 1% (first and second rows respectively). Census tracts are further separated by quartiles of African American population in the census tract (columns). Treatment effects are estimated using synthetic controls augmented by ridge regression. For each treated unit, the donor pool comprises never-treated census tracts within the same FEMA region.

Figure E.21: Second-stage regression, race and ethnicity heterogeneity



Each dot represents a census tract-specific treatment effect estimate of the impact of the flood map update on flood insurance take-up 24 months post-treatment, using synthetic controls augmented by ridge regression. For each treated unit, the donor pool comprises never-treated census tracts within the same FEMA region. Large dots are significant treatment effects at the 10% level, using the jackknife+ procedure. Regression lines represent the marginal effects of a change in the number of properties rezoned inside the 100-year floodplain on the synthetic control treatment effect estimates, following regression 5. The marginal effects are estimated separately per quartiles.

Figure E.22: Second-stage regression, race and ethnicity heterogeneity



Each dot represents a census tract-specific treatment effect estimate of the impact of the flood map update on flood insurance take-up 24 months post-treatment, using synthetic controls augmented by ridge regression. For each treated unit, the donor pool comprises never-treated census tracts within the same FEMA region. Large dots are significant treatment effects at the 10% level, using the jackknife+ procedure. Regression lines represent the marginal effects of a change in the number of properties rezoned inside the 100-year floodplain on the synthetic control treatment effect estimates, following regression 5. The marginal effects are estimated separately for each group.

F Welfare estimates

F.1 Details on model calibration

Assuming a Constant Absolute Utility Function, and dropping the time subscripts for clarity, the threshold condition for k_ω becomes:

$$\begin{aligned}
 -\exp\left(-\alpha(\omega)(Y(\omega) - r(\omega))\right) &> -\hat{p}_\omega \exp\left(-\alpha(\omega)(Y(\omega) - L(\omega))\right) \\
 &\quad - (1 - \hat{p}_\omega) \exp\left(-\alpha(\omega)(Y(\omega))\right) \\
 \iff -\exp\left(\alpha(\omega)r(\omega)\right) &> -\hat{p}_\omega \exp\left(\alpha(\omega)L(\omega)\right) - 1 + \hat{p}_\omega \\
 \iff \alpha(\omega) &> k_\omega
 \end{aligned} \tag{10}$$

where \hat{p}_ω is the *perceived probability* of flooding, $L(\omega)$ are the *perceived* damages associated with flooding, $r(\omega)$ is the price of the insurance contract, $Y(\omega)$ is income, and $\alpha(\omega)$ is still the absolute risk aversion parameter. $L(\omega)$, $r(\omega)$ and $Y(\omega)$ are observed by households and the econometrician, whereas $\alpha(\omega)$ is known by the household only. The cutoff value k_ω does not have a closed-form solution but can be computed numerically for each property.

The willingness-to-pay is derived as the price of the insurance contract that makes the homeowner indifferent between purchasing insurance and being uninsured:

$$\begin{aligned}
 -\exp(\alpha(\omega)WTP_\omega) &= \int_0^{+\infty} -\exp\left(\alpha(\omega)D_{d,\omega}\right) dF_{d,\omega} \\
 &\quad \ln\left(\int_0^{+\infty} \exp\left(\alpha(\omega)D_{d,\omega}\right) dF_{d,\omega}\right) \\
 \iff WTP_\omega &= \frac{\ln\left(\int_0^{+\infty} \exp\left(\alpha(\omega)D_{d,\omega}\right) dF_{d,\omega}\right)}{\alpha(\omega)}
 \end{aligned} \tag{11}$$

where $D_{d,\omega}$ are the (true) expected damages due to flooding that occurs with inundation depth d for property ω , and $F_{d,\omega}$ is the probability distribution of flooding at each depth for each ω . $D_{d,\omega}$ and $F_{d,\omega}$ are both taken from the First Street Foundation Flood Model.

Assuming a map-update-invariant and household-specific willingness-to-pay for insurance allows us to view all improvements in the information provided to households as net welfare gains, *even if* the corrected risk information shows the household to be at increased risk of flooding. To clarify, imagine that Valentina is willing to pay \$10,000 for a Marc Chagall painting, while she would be willing to pay \$0 for an imitation of the same painting. If she purchases the painting for \$200 and it is later revealed that the painting is a fake, our theoretical set-up views this information revelation as (weakly) welfare increasing: Valentina experienced a net loss from the initial purchase – since all surpluses should be computed in light of her *true*

willingness-to-pay – which was zero for a fake Chagall painting. Revealing correct information does not lead to further losses.

However, assuming that the expected damages due to flooding $D_{d,\omega}$ are constant throughout the map update explicitly rules out other forms of climate adaptations, such as elevating the property or constructing flood walls, which the household might undertake in response to new risk information. Our model therefore provides a lower-bound on the welfare gains from corrected flood maps.⁵²

Consistent with the empirical evidence presented in Sections 4 and 5, I assume that households perceived probabilities come from the official FEMA flood maps: households perceive the probability of flooding to be 1% in the 100-year floodplain, 0.2% in the 500-year floodplain, and zero outside of it.⁵³ In the normative part of the analysis, I compute the welfare of correcting the official floodplain boundaries assuming that they are updated to reflect the FSF model.

I further assume that households perceive the costs of flooding based on the expected damages given by the FSF Flood Model. This is a strong assumption, which will again tend to under-estimate the impacts of new flood maps. Figure F.23 shows that results are similar if we assume instead that the perceived damages are given by the average insurance claims in the neighborhood.

To recover the price of the insurance contract for all households (including for those who do not purchase insurance), I use neighborhood-, floodplain-, and time-specific premium averages. To assess the welfare impacts of moving from the current premiums to actuarially fair premiums, I further estimate welfare changes assuming that insurance prices are provided by annual expected losses estimates in the FSF model (the expected losses to the property in a given year).

I integrate out the risk aversion parameters by assuming they follow a Fréchet distribution within each census tract. This distribution allows for a fat upper tail and is governed by two parameters for strictly positive support:⁵⁴

$$p(\alpha(\omega) \leq k_\omega) = \exp\left(-\gamma_c \left(\frac{k_\omega}{A_c}\right)^{-\theta_c}\right) \quad (12)$$

⁵²In the context of a partial equilibrium analysis only. In a general equilibrium framework that allows for household sorting and preferences defined over endogenous neighborhood amenities, the welfare effects of correcting flood risk information are a priori ambiguous.

⁵³This assumption will bias the impacts of new flood maps towards zero, as the analysis above demonstrates the existence of within-neighborhoods spatial spillover effects of flood maps. An alternative assumption would be to specify a parametric structure for these spillovers, for instance assuming that households impute their probability of flooding as linearly or exponentially decreasing based on their location relative to the nearest floodplain. While such alternative assumptions are plausible, the data do not allow me to estimate these parametric structures.

⁵⁴Consistent with the literature, I assume that all homeowners are weakly risk averse.


where θ_c and A_c are the shape and scale parameters.

Recent work by Wagner (2022) shows that the observed low level of demand for flood insurance cannot be rationalized by risk loving homeowners, but instead reveals the existence of *frictions* that limit demand below optimal levels. While the previous sections of this paper revealed that lack of information and incorrect information about flood risks are among the decisive frictions that limit demand, my results do not rule out the existence of other frictions. This poses a challenge for the structural estimation: unless *all* existing frictions are correctly specified, one cannot separately estimate frictions and risk preferences. If one were willing to assume that incorrect information is the only friction limiting demand, then data on flood insurance transactions, together with the synthetic control estimates and a functional form assumption for the distribution of risk aversion preferences are sufficient to estimate the parameters governing the distribution of risk aversion. This approach is presented in Section F.3 in the appendix.

Instead of assuming correct specifications of all the frictions in order to back out risk aversion parameters, my analysis proceeds by assuming the parameters governing the distribution of the risk aversion parameters and *then* estimate the impacts of providing correct information. Following the literature on insurance demand and Wagner (2022), in my preferred calibration I assume the distribution of risk aversion parameters has an expected value of 10^{-5} , and I assess robustness of the findings for values of 10^{-4} and 10^{-6} . The shape parameter is fixed at 2 to allow for the existence of very risk averse homeowners.

F.2 Structural estimation using perceived damages from local claims

Figure F.23: Consumer welfare impacts of updating flood maps, robustness



Consumer welfare impacts of updated flood maps aggregated by neighborhood income deciles, assuming an expected risk aversion value of 10^{-5} . The pink line with circles depicts relative welfare changes using true insurance premiums (pre and post map updates), the black line with triangles assumes actuarially fair premiums before and after the map updates, while the green line with squares assumes that households perceive flood damages from historical claims.

F.3 Backing out risk aversion parameters from the data

The approach I employ in the paper assumes that the parameters governing the distribution of risk aversion preferences are *known*. For instance, my preferred specification fixes the shape

and scale parameters of the Fréchet distribution such as to obtain an expected risk aversion value of 10^{-5} , consistent with the literature (with a shape parameter fixed at 2 to allow for a fat upper tail, although other values are plausible). Here I show that conditional on (i) knowing the functional form of this distribution and (ii) assuming that information frictions are the only distortions constraining demand, it is possible to recover the parameters governing the distribution of risk preferences. While assumption (ii) of “no omitted frictions” seems implausible in the setting of federally provided flood insurance, it might be credible in other settings.

Aggregated to the tract level, the share of homeowners who purchase a contract before the map update is

$$s_{c,pre} = \frac{\sum_{\omega} p(\alpha(\omega) > k_{\omega,pre})}{N} \quad (13)$$

and the share of individuals who purchase the contract after the map update is

$$\begin{aligned} \hat{s}_{c,post} &= \frac{\sum_{\omega} p(\alpha(\omega) > k_{\omega,post})}{N} \\ \hat{s}_{c,post} &= s_{c,pre} + \hat{\tau}_c \end{aligned} \quad (14)$$

where N is the number of residential properties in the census tract, and $\hat{\tau}_c$ is the tract-specific treatment effect of the updated flood map on insurance demand (previously estimated with synthetic controls). $s_{c,pre}$ is directly observed from the data, whereas $\hat{s}_{c,post}$ is the *predicted* share of individuals purchasing insurance after the map update, and where all changes in demand relative to $s_{c,pre}$ are caused by the map update.

Together with equations 13 and 6, this gives rise to a collection of systems of two equations (one system per census tract) which can be numerically solved to obtain $\{A_c^*, \theta_c^*\}$, the tract-specific mean risk aversion and dispersion parameters.