

Online Appendix for “Flood Risk Mapping and the Distributional Impacts of Climate Information”

Joakim A. Weill*

A Data preparation


A.1 Map change letters to residents

A.2 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 research on flood maps relied on the “Community Status Book” (FEMA, 2020*a*) maintained by FEMA, 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., following map updates). However, because the Community Status Book records only 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


*Federal Reserve Board of Governors, joakim.a.weill@frb.gov. 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.

Figure A.1: Official FEMA template for properties mapped outside the SFHA



Source: https://www.fema.gov/sites/default/files/documents/fema_r6_guidebook_1.8_letter_to_property_owner_mapped_out_of_sfha.docx

Figure A.2: DC letter to residents for properties mapped inside the SFHA



Source: https://doee.dc.gov/sites/default/files/dc/sites/ddoe/publication/attachments/Letter_to_Property_Owners.pdf

particular, when the community is a county that 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 date.

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 Geographical Information System archives of contractors who worked with FEMA to produce these maps. I intersect the census-tract polygons with the DFIRM effective-date 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 match on only 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 DFIRM. However, this measurement error affects only 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 does not noticeably affect the model estimates. Finally, for census tracts covered by a digital 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 information provided by FEMA’s Map Service Center repository.¹ 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 flood map in 2005. The Q3 data come from three different sources: FOIA requests to FEMA (FEMA, 2024), records from the Caliper Corporation (FEMA, 2020*b*) (in particular, the list of counties with existing Q3 data), and the Princeton University Library (FEMA, 2021). As described in the main text, the Q3 maps were designed to aid floodplain

¹<https://msc.fema.gov/portal/advanceSearch>.

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.3 Note on intersecting floodplains and census-tract 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 parallelizing the intersections between the cropped polygons and census tracts.

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 postprocessing 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 run within two days.


A.4 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 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.

First, I use the entire repository of 132 million geolocalized residential properties maintained by First Street Foundation and downloaded from its API in April 2021.² This measure of population distribution presents three

²Due to the rate limitations of the API, the full download took about three weeks of uninterrupted running time.


Figure A.3: 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.

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 able to geolocalize only 95 million residential properties from ZTRAX – I discuss the ZTRAX data below). I intersect these residential properties with the Q3 flood maps, the 2012 DFIRMs, and the 2019 DFIRMs. I then aggregate the relevant changes in property 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 of the X zones); and the levee-protected flood zone.

Figure A.4: FSF vs dasymetric definition of rezoning



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.

Because the First Street Foundation residential properties data are 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, I use the 2010 dasymetric allocation of population provided by the Environmental Protection Agency's EnviroAtlas.³ This product employs land use cover and terrain slopes to reallocate the 2010 census-tract population within census tracts at the 30m-pixel level. The dasymetric layer covers the entire contiguous U.S., 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.4 plots the changes in the number of FSF residential properties zoned in the 100-year floodplain (relative to the census-tract number of prop-

³I 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 estimate of the population is greater than four times the number of properties. In practice, this concerns less than 1,500 census tracts out of 73,745.

erties) against the changes in the population in the 100-year zone using the dasymetric allocation, showing a large correlation between the two measures.

A.5 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 (Zillow, 2019), obtained through a data agreement in 2019.⁴ These data contain a large number of variables on individual structures, including their construction date. I restrict the data to geolocalized single-family residential properties that have a nonmissing construction year variable. As shown in Table A.1 in the “Unconditional” panel, the number of residential properties that satisfy these criteria is only 67 million, 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 before 2005 (third panel).

A.6 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 from 1983 through 2019 (FEMA, 2020c).⁵ After quality checks, I found that records before 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 the 2008-19 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 starting dates is not uniform within

⁴The ZTRAX program has now ended, and ZTRAX data are no longer available.

⁵In contrast, the website OpenFEMA provides publicly-available data for the 10 most recent years.

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


Floodzone	Q3	NFHL19
Unconditional:		
Inside 100-year floodplain	4,245,336	3,451,306
Outside 100-year floodplain	45,810,889	58,517,091
Not mapped	17,838,958	5,926,786
Conditional on Q3 and NFHL19:		
Inside 100-year floodplain	3,714,777	2,828,047
Outside 100-year floodplain	43,815,922	44,702,652
Conditional on Q3 and NFHL19, built prior to 2005:		
Inside 100-year floodplain	3,305,053	2,505,524
Outside 100-year floodplain	39,345,106	40,144,635

The sample is restricted to the single-family ZTRAX Assessment tables that 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 year, as policies tend to start during the summer. As a result, keeping track of only the year in which a policy was started would induce substantial measurement error. However, most policies are active for exactly 12 months, as owners typically decide whether or not to renew their policy every 365 days. As a result, creating a panel at the daily level would generate large autocorrelation 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 or 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 as 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. In practice, these data cleaning steps do not affect the sign or magnitude of the main estimates but help reduce variations that arise purely out of measurement error.

Figure A.5: Section of New Orleans' (LA) FIRM effective in 1984



FIRM section of Orleans Parish from FEMA's Map Service Center, source: <https://map1.msc.fema.gov/firm?id=2252030095E>

Figure A.6: Floodplain definitions in New Orleans (LA)




A: Q3 data product showing the flood zones effective in 1984, combined with a sample of 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 U.S. Army Corps of Engineers in flood control systems such as levees, floodgates and pump stations (FEMA, 2016).

A.7 Additional data sources

In addition to the main datasets presented above, I use data on FEMA disaster declarations (FEMA, 2020*d*), Consumer Price Index data (BLS, 2024), the YALE survey of opinions on climate change at the county-level (Howe et al., 2015), flood risk disclosure laws at the state-level (Hino and Burke, 2021), the EPA dasymetric population allocation (Environmental Protection Agency, 2016), demographic and geographic data from the US Census Bureau (United States Census Bureau, 2018, 2010*a,b*) accessed through Tidycensus (Walker and Herman, 2025)

Figure A.7: Summary of the different datasets used



B Comparing flood risk models with claims data

The first section below outlines a framework to compare the FEMA flood maps (“FEMA model”) and the First Street Foundation model (“FSF model”), highlighting what assumptions are required to perform this comparison. The second section takes this framework to the data focusing on flooding occurrence probabilities, while the third section explores flood damage estimates. Because the FEMA model does not provide damage estimates, this third section compares the FSF model with NFIP insurance claims (FEMA, 2020*e*) and SHELATUS (Spatial Hazard Events and Losses Database for the United States) hydrological perils data (Center for Emergency Management and Homeland Security, 2020).

B.1 Comparison framework

We observe a set of census tracts i with homes inside them. Each home has a probability p_F of flooding. To simplify the notation, assume that homes can be risky or safe, where $p_F = p_R$ for risky homes and $p_F = p_S$ for safe homes (the approach can be readily expanded to any number of risk classes). We have two different climate risk models M_1 and M_2 , which categorize homes as risky or safe. Conditional on climate model M_k , the number of homes within each census tract that are risky and safe are N_{kR}^i and N_{kS}^i respectively.

Let C_{kR}^{it} and C_{kS}^{it} be random variables denoting the number of insurance claims from risky and safe homes in tract i and year t under climate model M_k . Under full insurance take-up (an assumption relaxed below), each random variable follows a binomial distribution:

$$\begin{aligned} C_{kR}^{it} &\sim B(N_{kR}^i, p_R) \\ C_{kS}^{it} &\sim B(N_{kS}^i, p_S) \end{aligned} \tag{1}$$

The expected number of claims in tract i and year t if M_k is true is then

$$\mathbb{E}[C_k^{it}] = N_{kR}^i \cdot p_R + N_{kS}^i \cdot p_S \tag{2}$$

B.1.1 Testing M_k in the simple case of frequent events

In the case of auto insurance or other types of insurance with a large number of annual claims (e.g. health insurance, life insurance), equation 2 could be directly used to test different risk models. For instance, if C_k^{it} were the expected number of auto claims in tract i and year t if model k is correct, then one could directly use c^{it} (the observed yearly realizations of claims in each tract/year) and compute measures of prediction error for M_k . For instance, we can define the prediction error as $(\mathbb{E}[C_k^{it}] - c^{it})^2$, and the sum of these prediction errors could then be used to test which model is better.

B.1.2 Testing with few claims observations

Floods are “rare events”: in most years, the number of claims in a census tract is 0 (even under full insurance), and the approach above cannot be used. In order to bring equation 2 to the data, I need to adjust the approach and impose several simplifying assumptions.

Assumption 1-A. *The distributions of the number of claims under climate models M_k are stationary.*

Under assumption 1-A., I can drop the t superscript and use the number of claims averaged over time. In the analysis below, I take the average of the number of claims between 2008 and 2019 to match the insurance policy data sample.

It is important to emphasize that assumption 1-A. does not hold. The climate attribution literature highlights the role of climate change in increasing the frequency and intensity of severe flood events. While the FSF model provides different sets of estimates of flood risk in 2020, 2050, and 2100, FEMA flood maps are updated infrequently and provide only historical risk estimates (these estimates aim to be valid only for the year when they are issued). That said, local adaptation infrastructure can reduce exposure to flooding, while population movements can either increase or decrease the number of claims. Relaxing assumption 1-A. in a way that would allow us to assess forward-looking climate risk models using historical insurance claims data therefore involves a complex trade-off between a stationarity assumption and a rare observation constraint. I have not found any solution to this problem in the literature.

Assumption 1-A. makes it possible to use historical data to test models M_1 and M_2 , but it does not solve the problem of the small number of observations associated with rare events. In this case, imagine that we pool 12 years of claims data on a property exposed to a (static) 1% annual chance of flooding (and therefore in the 100-year floodplain). The probability of observing a claim on this property is only $1 - (0.99)^{12} = 11.4\%$. When the time series is relatively short (relative to the probability of occurrence), Cummins (1991) advocates using moments of the distribution rather than unit loss data. To mitigate the small number of observations problem, in the empirical sections B.2 and B.3 below I use claims data aggregated at the county level. This aggregation further allows me to ignore the correlation of claims between tracts located in the same region (at the tract level, the C_k^{it} random variables are not independent).

B.1.3 Accounting for incomplete take-up

Flood insurance take-up is extremely low in the U.S.: about 66% of census tracts have an average take-up rate over the 2008-19 period that is **less than 1%**. (Figure B.1). In these tracts, if a flood hits 99% of residential properties, none of these strikes would show up in the insurance data. Take-up is higher in the FEMA 100-year floodplain, where only 23% of census tracts have an average take-up rate over the 2008-19 period that is less than 1% (remember that although we cannot geocode insurance policies, the data report whether the policy was purchased in the FEMA 100-year floodplain). The median FEMA 100-year floodplain take-up rate is 10%, and 10% of tracts have an SFHA take-up rate of 100%.


In order to leverage insurance claims data to assess the validity of climate model M_k , we need to account for incomplete take-up. Let $\mathcal{H}^i$ be the set of households who carry flood insurance in tract i . We have

$$\mathbb{E}[C_k^i | h \notin \mathcal{H}^i] = 0 \quad (3)$$

By definition, the expected number of claims from households without insurance is 0: we can only learn about $\mathbb{E}[C_k^i | h \in \mathcal{H}^i]$.

With geolocalized flood insurance policies and claims data, we could identify the homes that carry flood insurance and construct the set $\mathcal{H}^i$ to estimate the relevant expectations. In practice, the FEMA data do not provide address-level information – we observe only the census tract, and whether or

Figure B.1: Average insurance take-up at the tract level



not the property is in the FEMA 100-year floodplain. To construct $\mathcal{H}^i$, we therefore need to make additional assumptions about **which** households are purchasing flood insurance.

Let $p_{k(1)}^i \leq p_{k(2)}^i \leq \dots \leq p_{k(H^i)}^i$ be the ordered probabilities of flooding of the H^i households in tract i under model k , and let insurance take-up in tract i be d^i .

Assumption 1-B.

$$\mathbb{E}[C_k^i | h, p_{k(h)}^i \geq p_{k(H^i - d^i + 1)}^i] \geq \mathbb{E}[C_k^i | h \in \mathcal{H}^i] \geq \mathbb{E}[C_k^i | h, p_{k(h)}^i \leq p_{k(d^i)}^i] \quad (4)$$

Under assumption 1-B. we can bound the expected number of claims in tract i by assuming that the d_i insurance policies are purchased from the riskiest properties (upper bound), or that these policies are purchased by the lowest risk properties (lower bound). Note that the inequalities in equation 4 are only correct conditional on climate model M_k being correct.

Taking this approach to the data, we can construct bounds based on observed take-up with the following procedure. Let n be the number of policies purchased in a tract and N be the number of properties. By definition we have

$$\begin{aligned} N &= N_{FEMA,100} + N_{FEMA,500} + N_{FEMA,out} = N_{FSF,100} + N_{FSF,500} + N_{FSF,out} \\ n &= n_{FEMA,100} + n_{FEMA,\neg 100} \end{aligned} \quad (5)$$

with $N_{FEMA,100}$ denoting the number of properties in the 100-year floodplain accounting to the FEMA model, and similarly for the other counts.

We are interested in bounds for $n_{FEMA,500}$, $n_{FSF,100}$ and $n_{FSF,500}$. We have:

$$\begin{aligned} n_{FEMA,500}^+ &= \min(n_{FEMA,\neg 100}, N_{FEMA,500}) \\ n_{FEMA,500}^- &= \min(n_{FEMA,\neg 100} - N_{FEMA,100} - N_{FEMA,500}, 0) \\ n_{FSF,100}^+ &= \min(n, N_{FSF,100}) \\ n_{FSF,100}^- &= \min(n - (N - N_{FSF,100} - n_{FSF,500}^-), N_{FSF,100}) \text{ if } n > (N - N_{FSF,100} - n_{FSF,500}^-), \text{ and 0 otherwise} \\ n_{FSF,500}^+ &= N_{FSF,500} \text{ if } n - n_{FSF,100}^+ > 0, \text{ and 0 otherwise} \\ n_{FSF,500}^- &= \min(n - (N - N_{FSF,100} - N_{FSF,500}), N_{FSF,500}) \text{ if } n > (N - N_{FSF,100} - N_{FSF,500}), \text{ and 0 otherwise} \end{aligned}$$

B.2 Comparison of claims occurrence predictions

Table B.1 below compares predictions of the annual number of insurance claims from single-family residential properties between FEMA flood maps and the FSF model. The rows show predictions using different assumptions on insurance take-up: the first row assumes full insurance take-up, the second row assumes that the highest-risk properties conditional on a given risk model purchase insurance (left bound of equation 4), and the third row assumes that the lowest-risk properties purchase insurance (right bound of equation 4). Rows 4 to 6 repeat the same three previous comparisons, but excluding from the sample all the properties not subject to rain-based flood risk in the FSF model.

Table B.1: Number of claims comparison

Take-up assumption	$E[C FEMA]$	$E[C FSFb]$	$E[C FSFe]$	Average number of claims
Full take-up	[75,953 ; 711,843]	[167,528 ; 1,640,766]	1,601,285	44,166
Highest risk selection	[19,576 ; 188,924]	[27,683 ; 275,105]	1,192,954	44,166
Lowest risk selection	[18,458 ; 183,334]	[6,130 ; 60,785]	32,786	44,166
Full take-up, no rain	[56,247 ; 530,225]	[76,708 ; 750,788]	784,527	44,166
Highest risk selection, no rain	[17,227 ; 166,420]	[21,554 ; 214,073]	624,286	44,166
Lowest risk selection, no rain	[16,291 ; 161,743]	[6,110 ; 60,615]	32,803	44,166

The second and third columns present intervals rather than point estimates for the number of claims. The reason is that the FEMA flood maps do not provide probability point estimates, but rather inform us about whether the probability of flooding is greater than 1% (100-year floodplain) or between 0.2% and 1% (500-year floodplain). The lower bounds in Table B.1 are obtained assuming that the 500-year and 100-year floodplains imply probabilities of flooding of exactly 0.2% and 1%, respectively, while the upper bounds are obtained assuming probabilities of 1% and 10%, respectively, in these floodplains (the correct upper bound in the 100-year floodplain is 100%, but such a bound is not informative – I use the value of 10%, as it is the average flood risk estimated in the 100-year floodplain in the FSF model). Finally, the fourth column provides estimates using the exact probability of flooding provided by the FSF model (FSF_e).

These comparisons reveal three interesting findings. First, the bounds in columns 2 and 3 are extremely wide. This finding indicates that using only the floodplains’ bounds definition of risk to predict claims provides imprecise information to predict the expected number of claims.

Second, the hypothesis on risk selection (rows 2 and 3) has a negligible impact on the FEMA model predictions, but has a major impact on the FSF model predictions, with high and low risk selection assumptions leading to implausibly high and low predictions on the number of claims. This result is partly driven by the nature of the insurance data: policies keep track of the 100-year and 500-year FEMA floodplains, and therefore the risk selection assumptions in the FEMA model impact only properties in the 500-year floodplains. These 500-year properties have overall low impacts on the predicted number of claims. In contrast, the insurance data are not informative for risk selection inside or outside the FSF floodplains. In addition, the FSF model is more granular than the FEMA model – the estimated flood probabilities in FSF can vary substantially within a census tract, which explains the importance of knowing who is covered by insurance to inform claims predictions (see Figure B.3 below). Using exact probability estimates for the FSF model provides predictions for the number of claims that are implausibly far away from observations. Unfortunately, these exact probabilities cannot be compared with those for the FEMA model, as flood maps only provide bounds.

Finally, the observed annual number of claims (44,166) falls in the estimated ranges in columns 1 and 2 under both the highest and lowest risk selection assumptions. That is, using the same floodplain bounds criterion to compare FEMA and FSF, both models provide sensible estimates of the number of claims.

Overall, these comparisons fail to provide strong evidence against either the FSF or the FEMA model. However, they clarify both the importance of the risk selection assumption in the FSF model predictions, and the improved precision that would come with exact flood occurrence probabilities (rather than bounds) in the FEMA model.

B.3 Comparison of flood damage predictions

This section compares flood damage values in the insurance claims data, in the SHELDUS dataset, and in the FSF model (the FEMA flood maps do not provide damage estimates). Each dataset has strengths and weaknesses.

Insurance claims data provide the cleanest values of flood damages, but with two pitfalls: (i) low insurance take-up and (ii) a historical view, rather than current damages.

The SHELDUS dataset is maintained by Arizona State University’s Cen-

Table B.2: Comparison of flood damage estimates in three data sources

Sample	Insurance claims	SHELDUS	FSF	FSF no rain-based
All counties	2,490	15,048	25,599	21,751
Top 10% of county take-up	1,951	9,352	11,203	11,030
Top 5% of county take-up	1,632	7,384	9,080	9,054
Counties with high rain-based risk	481	3,504	10,084	0

All damage estimates are in million of 2020 US dollars.


ter for Emergency Management and Homeland Security. It provides the most comprehensive assessment of disaster damages in the U.S. and is available at the county level. To match the insurance claims data, I use data on the 2008-19 period focusing on flooding and hydrological damages to structures. While more comprehensive than the insurance claims data, SHELDUS damage values are not exact, but only estimates of the true damages. In particular, damages from flooding events rely explicitly on insurance claims observations.⁶ The SHELDUS estimates therefore suffer from the same pitfalls as the claims data but potentially provide greater coverage of flood events through the use of “guesstimates,” even if this possibly increased coverage comes at the cost of precision.

Finally, FSF provides property-level expected damage estimates (“annual average losses”). These estimates are based on the inundation depth estimates produced by FSF and fed through depth-damage functions developed by FEMA and the Army Corps of Engineers. FSF uses auxiliary data sources on the property value, first flood height, and presence of a basement (Armal et al., 2020).

Table B.2 presents annual average estimates of flood damages using these different sources, yielding three conclusions. First, there is a constant order-

⁶“On loss estimation: Property damage estimates should be entered as actual dollar amounts, if a reasonably accurate estimate from an insurance company or other qualified individual is available. If this estimate is not available, then the preparer has two choices: either check the “no information available” box, or make an estimate. The exception is for flood events. The Storm Data preparer must enter monetary damage amounts for flood events, even if it is a “guesstimate.” The U.S. Army Corps of Engineers requires the NWS to provide monetary damage amounts (property and/or crop) resulting from any flood event”, see Arizona State University, Center for Emergency Management and Homeland Security (2022), “Metadata,” webpage, quoted text under “Information on NCEI (formerly NCDC) Storm Data Preparation,” <https://cemhs.asu.edu/sheldus/metadata>

Figure B.2: Dominant sources of flood risk in the FSF model



These maps show the dominant source of flood risk at the tract- (left) and county- (right) level. “Pluvial tracts” (respectively counties) are defined as tracts where more than 50% of properties in the FSF 100-year floodplain are exposed to rain-based flooding, and similarly for “Fluvial” (river-based risk) and “Coastal” tracts.

ing of damages in these three datasets: FSF and insurance claims provide the largest and smallest estimates, respectively, and SHELDUS estimates always fall between the two. For data from all counties in the contiguous U.S. (first row), SHELDUS damage estimates are six times larger than insurance claims data, but only half the FSF damage estimates.

Second, focusing on counties with the largest insurance take-up rates (rows 2 and 3) substantially reduces the gap between SHELDUS and FSF estimates, without fully resolving it. This outcome is consistent with the SHELDUS data collection methodology, which uses insurance claims data to inform estimates of flood damages. Focusing on counties with the highest insurance take-up rates does not impact claim damage estimates much, reflecting the spatial concentration on insurance take-up.


Finally, the gap in estimates between data sources varies with the source of flood risk. Figure B.2 shows the prevailing source of flood risk as estimated in the FSF model. Counties exposed to high rain-based risk (fourth row of Table B.2) are responsible for most of the differences between models’ damage estimates. These counties account for more than a third of the damages

estimated by FSF, but only about 20% of insurance claims and a similar share of SHELDUS damage estimates. Column 4 provides a confirmation of this discrepancy by reporting flood damages estimates in FSF in the different set of counties (rows), but excluding all properties exposed to rain-based risk. Excluding rain-based risk from FSF provides damage estimates that are closer to the claims and SHELDUS estimates.

To summarize, despite recent work showing that intense precipitation is the fastest-growing source of flood risk (Fowler et al., 2021), these risks are not explicitly captured in the official flood maps (Government Accountability Office, 2021; FEMA, 2023). The FEMA 100-year floodplain is smaller than it should be in areas exposed to rain-based risk, which causes insurance take-up to be lower and insurance claims to underestimate flood damages.

Rain-based flood risk is not captured by the FEMA flood maps largely due to the difficulty in estimating these risks. Figure B.3 below shows that one of the challenges associated with rain-based risk estimation is spatial granularity: at a small scale, rain-based risk varies more than river-based or coastal risk. The figure shows the distributions of the share of properties estimated in the 100-year floodplain in the FEMA flood maps and in the FSF model, including all properties (top row) or excluding properties exposed to rain-based risk in the FSF model (bottom row). Compared to the FEMA model, the FSF model estimates a much greater number of tracts with an intermediary number of properties in the 100-year floodplain (top right panel). That is, while FEMA is more likely to estimate that all properties in a tract are outside or inside the 100-year floodplain, FSF is more likely to estimate that only some of these properties are in the 100-year floodplain – a direct consequence of spatial granularity. The bottom row shows that differences between distributions narrow substantially when excluding all properties exposed to rain-based risk, suggesting that spatial granularity matters most for this risk.

Figure B.3: Granularity of the 100-year floodplain in FEMA and FSF



Distributions of tract-level share of properties in the 100-year floodplain under different risk product. The y-axis is truncated at 10 for clarity, which excludes a mass of tracts at $x=0$ in the FEMA 2019 product.

C Polygon complexity table

Table C.1: Polsby-Popper metrics for the 100-year floodplain, by state


State	Perimeter Q3 (km)	Area Q3 (km2)	Perimeter NFHL 2019 (km)	Area NFHL 2019 (km2)	Change Polsby-Popper (%)
Alabama	31761.6	6129.9	46973.8	5983.5	-55.4
Arizona	42811	6121.1	61035.3	7245.3	-41.8
Arkansas	36669.7	10274.6	40736.9	10307.7	-18.7
California	91713.5	23411.7	105608	26336.9	-15.2
Colorado	10629.9	1025.4	11883.4	1067.6	-16.7
Connecticut	11231.4	815.4	11917.6	856.2	-6.7
Delaware	4845.9	991	7181.8	904.4	-58.4
District of Columbia	210.1	26.1	227	29.9	-1.7
Florida	152689.8	45751.3	290636.9	51805.2	-68.7
Georgia	33242	7486.5	50383	7622.7	-55.7
Idaho	5450.7	817.2	5601	915.7	6.1
Illinois	21640.3	2757.9	24663.6	2778.3	-22.4
Indiana	20242.5	3253.9	23767.1	3260.3	-27.3
Iowa	10556.9	1489.2	13591	1793.8	-27.3
Kansas	12660.3	2082.9	22098	2245	-64.6
Kentucky	44432.5	6092	69047.3	6843.9	-53.5
Louisiana	35298.7	25788.5	57047.1	26797.7	-60.2
Maine	14696.7	1622.5	17125.4	1580.5	-28.3
Maryland	21013.5	3559.4	26738.1	3094.3	-46.3
Massachusetts	22154.8	1930	24337.5	2107.8	-9.5
Michigan	15787.5	1524.1	18769.7	1535.9	-28.7
Minnesota	19663.3	3735.4	30587.1	4038.6	-55.3
Mississippi	32299.9	10203.8	50676.1	10534.6	-58.1
Missouri	16286.5	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	23594.1	3498.8	31624.6	3247.5	-48.3
New Mexico	5616.3	506.9	5968.7	573.7	0.2
New York	35585.6	3647.2	51098.6	3835.4	-49
North Carolina	57025.8	14108.3	77792.9	13395.3	-49
North Dakota	12033.4	1972.9	18390.6	2497	-45.8
Ohio	26025.1	2586.5	31789.8	2689.5	-30.3
Oklahoma	24779.7	3612.7	30087.3	3786.2	-28.9
Oregon	25143.3	3541.8	27669.8	3623	-15.5
Pennsylvania	67656.9	4378.8	90675	4496.5	-42.8
Rhode Island	4235.2	363.3	4504.2	375.9	-8.5
South Carolina	11134.3	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	129292.8	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	29889.4	2830	33891.9	2686.3	-26.2
Washington	22449.2	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

The 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.

D Additional descriptive statistics and robustness checks

D.1 Rezoning at the tract level

Figure D.1: 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.

D.2 Removed and ignored properties: additional results

Table D.1: Characteristics of properties removed from 100-year floodplain during map updates

Dependent Variable: Model:	(1)	(2)	(3)	Removed from Q3 100-year		(6)	(7)	(8)
				(4)	(5)			
<i>Variables</i>								
Constant	0.27 (0.01)		0.43 (0.01)		0.28 (0.01)		0.49 (0.01)	
Black or AA 25-50%	0.10 (0.01)	0.05 (0.01)	0.08 (0.01)	0.05 (0.01)				
Black or AA 50-75%	0.19 (0.01)	0.10 (0.01)	0.16 (0.01)	0.10 (0.01)				
Black or AA 75-100%	0.23 (0.01)	0.10 (0.01)	0.23 (0.01)	0.11 (0.01)				
Coastal			-0.35 (0.01)	-0.25 (0.01)			-0.34 (0.01)	-0.26 (0.01)
Inland fluvial			-0.28 (0.01)	-0.26 (0.01)			-0.28 (0.01)	-0.26 (0.01)
Inland pluvial			-0.24 (0.01)	-0.21 (0.01)			-0.24 (0.01)	-0.21 (0.01)
Income 25-50%			0.03 (0.01)	0.03 (0.01)			0.00 (0.01)	0.02 (0.01)
Income 50-75%			0.07 (0.01)	0.04 (0.01)			0.02 (0.01)	0.03 (0.01)
Income 75-100%			0.14 (0.01)	0.07 (0.01)			0.08 (0.01)	0.05 (0.01)
Hispanic 25-50%					0.04 (0.01)	0.04 (0.01)	0.04 (0.01)	0.04 (0.01)
Hispanic 50-75%					0.15 (0.01)	0.07 (0.01)	0.13 (0.01)	0.06 (0.01)
Hispanic 75-100%					0.22 (0.01)	0.08 (0.01)	0.16 (0.01)	0.07 (0.02)
<i>Fixed-effects</i>								
County FE		Yes		Yes		Yes		Yes
<i>Fit statistics</i>								
Observations	6,298,006	6,298,006	6,272,983	6,272,983	6,298,006	6,298,006	6,272,983	6,272,983
R ²	0.03	0.26	0.13	0.30	0.03	0.26	0.12	0.30
Within R ²		0.00		0.06		0.00		0.06

Clustered (Tract FE) standard-errors in parentheses

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. X_i includes the property-level source of flood risk (where applicable) estimated in the FSF model (coastal, river-based (“Inland fluvial”) or rain-based (“Inland pluvial”)), as well as tract-level racial quartile from the 2010 decennial census and Median Income from the American Community Survey 2015-2019. α_c are county fixed effects. For all models, the sample is restricted to census tracts mapped in both Q3 and the NFHL19 product and to properties that were in the Q3 100-year floodplain.

Table D.2: 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)
<i>Variables</i>						
Constant	0.38 (0.01)		0.52 (0.01)		0.25 (0.01)	
AI or AN 25-50%	0.02 (0.02)	0.00 (0.01)				
AI or AN 50-75%	0.00 (0.01)	-0.02 (0.01)				
AI or AN 75-100%	0.02 (0.02)	-0.04 (0.01)				
White 25-50%			-0.03 (0.02)	0.02 (0.01)		
White 50-75%			-0.14 (0.02)	0.00 (0.01)		
White 75-100%			-0.26 (0.01)	-0.07 (0.01)		
Asian 25-50%					0.05 (0.01)	0.04 (0.01)
Asian 50-75%					0.20 (0.01)	0.11 (0.01)
Asian 75-100%					0.37 (0.01)	0.21 (0.01)
<i>Fixed-effects</i>						
County FE		Yes		Yes		Yes
<i>Fit statistics</i>						
Observations	6,298,006	6,298,006	6,298,006	6,298,006	6,298,006	6,298,006
R ²	0.00	0.26	0.05	0.26	0.08	0.27
Within R ²		0.00		0.00		0.01

Clustered (Tract FE) standard-errors in parentheses

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. X_i includes the property-level source of flood risk (where applicable) estimated in the FSF model (coastal, river-based (“Inland fluvial”) or rain-based (“Inland pluvial”)), as well as tract-level racial quartile from the 2010 decennial census. α_c are county fixed effects. For all models, the sample is restricted to census tracts mapped in both Q3 and the NFHL19 product and to properties that were in the Q3 100-year floodplain.

Table D.3: Characteristics of FSF 100-year properties incorrectly ignored during map updates

Dependent Variable: Model:	(1)	(2)	FSF 100-year, ignored during update					
			(3)	(4)	(5)	(6)	(7)	(8)
<i>Variables</i>								
Constant	0.55 (0.01)		0.24 (0.01)		0.57 (0.01)		0.25 (0.01)	
Black or AA 25-50%	0.02 (0.01)	0.07 (0.01)	0.04 (0.01)	0.07 (0.01)				
Black or AA 50-75%	0.05 (0.01)	0.13 (0.01)	0.08 (0.01)	0.12 (0.01)				
Black or AA 75-100%	0.06 (0.01)	0.17 (0.01)	0.09 (0.01)	0.16 (0.01)				
Inland fluvial			0.23 (0.01)	0.00 (0.01)			0.23 (0.01)	0.01 (0.01)
Inland pluvial			0.48 (0.01)	0.30 (0.01)			0.48 (0.01)	0.31 (0.01)
Income 25-50%			0.01 (0.01)	0.01 (0.01)			0.00 (0.01)	0.00 (0.01)
Income 50-75%			0.00 (0.01)	-0.01 (0.01)			-0.01 (0.01)	-0.03 (0.01)
Income 75-100%			0.06 (0.01)	0.01 (0.01)			0.03 (0.01)	-0.01 (0.01)
Hispanic 25-50%					-0.01 (0.01)	0.05 (0.01)	0.05 (0.01)	0.06 (0.01)
Hispanic 50-75%					0.03 (0.01)	0.10 (0.01)	0.07 (0.01)	0.10 (0.01)
Hispanic 75-100%					0.06 (0.01)	0.14 (0.01)	0.10 (0.01)	0.14 (0.01)
<i>Fixed-effects</i>								
County FE		Yes		Yes		Yes		Yes
<i>Fit statistics</i>								
Observations	8,988,991	8,988,991	8,969,183	8,969,183	8,988,991	8,988,991	8,969,183	8,969,183
R ²	0.00	0.23	0.19	0.29	0.00	0.23	0.19	0.29
Within R ²		0.01		0.09		0.00		0.08

Clustered (Tract FE) standard-errors in parentheses

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. X_i includes the property-level source of flood risk (where applicable) estimated in the FSF model (coastal, river-based (“Inland fluvial”) or rain-based (“Inland pluvial”)), as well as tract-level racial quartile from the 2010 decennial census and Median Income from the American Community Survey 2015-2019. α_c are county fixed effects. For all models, the sample is restricted to census tracts mapped in both Q3 and the NFHL19 product and to properties that were in the Q3 100-year floodplain.

Table D.4: FSF 100-year properties incorrectly ignored during map updates, by treatment year

Dependent Variable:	FSF 100-year, ignored during update	
Model:	(1)	(2)
<i>Variables</i>		
Constant	0.74 (0.02)	
Update year 2001	-0.06 (0.11)	-0.04 (0.15)
Update year 2002	-0.12 (0.05)	0.01 (0.12)
Update year 2003	-0.41 (0.04)	-0.17 (0.11)
Update year 2004	-0.23 (0.07)	-0.03 (0.13)
Update year 2005	-0.32 (0.04)	-0.01 (0.12)
Update year 2006	-0.22 (0.03)	-0.03 (0.13)
Update year 2007	-0.02 (0.03)	-0.08 (0.12)
Update year 2008	-0.05 (0.03)	-0.10 (0.12)
Update year 2009	-0.16 (0.03)	-0.01 (0.12)
Update year 2010	-0.07 (0.02)	0.00 (0.12)
Update year 2011	-0.10 (0.02)	-0.10 (0.12)
Update year 2012	-0.16 (0.03)	-0.05 (0.12)
Update year 2013	-0.10 (0.03)	-0.18 (0.12)
Update year 2014	-0.23 (0.03)	-0.15 (0.12)
Update year 2015	-0.18 (0.03)	-0.12 (0.12)
Update year 2016	-0.15 (0.03)	-0.18 (0.12)
Update year 2017	-0.20 (0.03)	-0.06 (0.12)
Update year 2018	-0.22 (0.03)	-0.16 (0.12)
Update year 2019	-0.19 (0.04)	-0.13 (0.12)
<i>Fixed-effects</i>		
County FE		Yes
<i>Fit statistics</i>		
Observations	8,993,361	8,993,361
R ²	0.03	0.23
Within R ²		0.00

Clustered (Tract FE) standard-errors in parentheses

Linear probability regressions at the property level. The outcome is a dummy variable equal to 1 when the property should be in the 100-year floodplain based on the FSF model but was ignored during the map update. The sample is restricted to properties in the FSF 100-year floodplain in census tracts mapped in both Q3 and the NFHL19 product. The regressors include dummies for each year of the flood map modernization, with the baseline coded as 1995-2000.

Table D.5: FSF 100-year properties incorrectly ignored during map updates, by race

Dependent Variable: Model:	FSF 100-year, ignored during update					
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Variables</i>						
Constant	0.55 (0.01)		0.68 (0.01)		0.53 (0.01)	
AI or AN 25-50%	0.00 (0.01)	0.00 (0.01)				
AI or AN 50-75%	0.03 (0.01)	0.02 (0.01)				
AI or AN 75-100%	0.11 (0.01)	0.01 (0.01)				
White 25-50%			-0.04 (0.01)	-0.01 (0.01)		
White 50-75%			-0.11 (0.01)	-0.06 (0.01)		
White 75-100%			-0.18 (0.01)	-0.16 (0.01)		
Asian 25-50%					-0.01 (0.01)	0.04 (0.01)
Asian 50-75%					0.08 (0.01)	0.09 (0.01)
Asian 75-100%					0.22 (0.01)	0.13 (0.01)
<i>Fixed-effects</i>						
County FE		Yes		Yes		Yes
<i>Fit statistics</i>						
Observations	8,988,991	8,988,991	8,988,991	8,988,991	8,988,991	8,988,991
R ²	0.01	0.22	0.02	0.23	0.03	0.23
Within R ²		0.00		0.01		0.00

Clustered (Tract FE) standard-errors in parentheses

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 D.6: FSF 100-year properties incorrectly ignored during map updates, quantiles conditional on mapping

Dependent Variable: Model:	FSF 100-year, ignored during update						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Variables</i>							
Constant	0.31 (0.01)			0.55 (0.01)		0.25 (0.01)	
Inland fluvial	0.23 (0.01)	0.02 (0.01)	-0.18 (0.01)			0.23 (0.01)	0.00 (0.01)
Inland pluvial	0.48 (0.01)	0.32 (0.01)	0.20 (0.01)			0.48 (0.01)	0.30 (0.01)
Black or AA 25-50% (c)				0.04 (0.01)	0.07 (0.01)	0.05 (0.01)	0.07 (0.01)
Black or AA 50-75% (c)				0.05 (0.01)	0.13 (0.01)	0.07 (0.01)	0.12 (0.01)
Black or AA 75-100% (c)				0.07 (0.01)	0.16 (0.01)	0.08 (0.01)	0.15 (0.01)
Income 25-50%						0.02 (0.01)	0.01 (0.01)
Income 50-75%						0.01 (0.01)	-0.01 (0.01)
Income 75-100%						0.06 (0.01)	0.01 (0.01)
<i>Fixed-effects</i>							
County FE		Yes	Yes		Yes		Yes
Tract FE			Yes				
<i>Fit statistics</i>							
Observations	8,993,361	8,993,361	8,993,361	8,988,991	8,988,991	8,969,183	8,969,183
R ²	0.18	0.28	0.51	0.00	0.23	0.19	0.29
Within R ²		0.08	0.07		0.01		0.09

Clustered (Tract FE) standard-errors in parentheses

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.

Table D.7: 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.

D.3 Selection into treatment and event studies

D.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 D.8 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 the 100-year floodplain as negative). All models include fixed effects for each FEMA region, as the rollout 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.

Table D.8: 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 (per 1,000)	−0.115 (0.028)	−0.118 (0.028)	−0.002 (0.001)	−0.001 (0.001)
Median income (per 1,000)	−0.028 (0.009)	−0.030 (0.009)	−0.0004 (0.0003)	−0.00005 (0.0001)
Share African Americans	−0.373 (0.626)	−0.322 (0.617)	−0.023 (0.023)	−0.035 (0.018)
Disaster declaration prior to treatment			0.019 (0.018)	0.010 (0.011)
Average Annual Loss	0.0001 (0.0001)		−0.00000 (0.00000)	
Insurance policies/property, 2008		3.831 (1.310)		−0.515 (0.149)
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	28,054	28,054
R ²	0.083	0.096	0.054	0.260

Note:

NA
Standard errors clustered at the county level

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 FSF (column 1). Interestingly, the point estimate is *positive*, suggesting that more flood-prone areas were treated *later*, although the effect is small and 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 10% increase in the rate of residential property insurance coverage is associated with a treatment timing delayed by more than 0.4 year. 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 the 100-year floodplain (columns 3 and 4 in Table D.8). 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-tract 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 contradicts 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 more flood-prone areas, anecdotal evidence and discussions with floodplain managers suggest that it was due to (i) the complexity of modeling flood risk in these areas and (ii) local homeowners lobbying against new flood maps.

Finally, this paper examines only flood map updates mandated by FEMA. Individuals who disagree with their floodplain designation can petition FEMA with a Letter of Map Change (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.

E Additional event-study estimates

Formally, we are interested in the ATTs for each group g , time period t , and net map rezoning at the tract level x :

$$ATT(g, t, x) = \mathbb{E}(Y_t(g, x) - Y_t(0, x) | G_g = 1, X = x) \quad (6)$$

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 observe only the realized outcome.

The assumptions of limited treatment anticipation and conditional parallel trends allow us to rewrite the ATT as

$$ATT(g, t, x, \delta) = \mathbb{E}(Y_t - Y_{g-\delta-1} | G_g = 1, X = x) - \mathbb{E}(Y_t - Y_{g-\delta-1} | D_{t+\delta} = 0, X = x) \quad (7)$$

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

$$\widehat{ATT}(g, t, x) = (\bar{Y}_{g,t,x} - \bar{Y}_{g,g-\delta-1,x}) - (\bar{Y}_{0,t,x} - \bar{Y}_{0,g-\delta-1,x}) \quad (8)$$

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


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}. Few properties are zoned inside or outside the 500-year floodplain without a simultaneous change in the 100-year floodplain (most properties zoned inside the 500-year floodplain used to be zoned inside the 100-year floodplain).

I then aggregate cohort-specific estimates over all cohorts to obtain the overall ATTs by length of exposure following Callaway and Sant'Anna (2021):

$$\theta(e, x) = \sum_g \mathbf{1}\{g + e \leq T\} \cdot P(G = g \cap X = x | G + e \leq T) \cdot ATT(g, g + e, x) \quad (9)$$

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, T is the last time period for which ATTs are identified, and P are cohort sample-size weights. The estimator is obtained by replacing ATT with its sample analogue from equation 8, $\widehat{ATT}$. For all the event studies, I plot coefficients in $[-36; 36]$ — or a 6 year window around the map update — as both the control and treatment groups can change abruptly and become very small further away from treatment.


Figure E.1: Event-study estimates of the impacts map updates on take-up, 3% cutoff



Callaway-Sant'Anna estimates of the average treatment effect on the treated of the impact of flood map update on insurance take-up. The x-axis depicts time in month, normalized relative to the quarter of the flood map update ($t=0$). The event-study are estimated focusing on tracts where the new map added or removed more than 3% of properties from/to the 100-year 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 and intensity of 100-year floodplain rezoning. Error bars represent 95% confidence intervals using the multiplier bootstrap.


E.1 Event-study by cohort and rezoning

Figure E.2: Cohort-specific event-study estimates, increase and decrease



ATT estimates obtained using Callaway/Sant’Anna-type regressions. The x-axes depict time in month, normalized relative to the quarter of the flood map update ($t=0$). Each panel 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 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 49 largest cohorts are represented (out of 104 cohorts).


Figure E.3: Cohort-specific event-study estimates, no change in 100-year



Each panel 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). The x-axes depict time in month, normalized relative to the quarter of the flood map update ($t=0$). 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.


E.2 Event-study robustness

Figure E.4: Event-study, robustness tests using early constructions properties



The outcomes displayed are the number of active insurance policies covering properties that were constructed prior to 2007 (solid blue line), where the construction date is taken from the insurance data. The x-axes depict time in month, normalized relative to the quarter of the flood map update ($t=0$). Each panel 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 panel 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 in levels. Error bars represent 95% confidence intervals using the multiplier bootstrap.

Figure E.5: Within-tract and between-tract spillovers



This diagram presents 3 census tracts before (left) and after (right) a flood map update. The flood map update only occurs in the tract at the bottom and rezones properties inside the 100-year floodplain.

E.3 Spillover effects

I distinguish between two different kinds of potential spillovers: within-tract spillovers and between-tract spillovers. Both kinds of spillovers could be happening in my setting. Because I use difference-in-differences estimators with tract-level observation, the SUTVA assumption (Stable Unit Treatment Value Assumption, required to causally identify the ATT) rules out the existence of between-tract spillovers. However, within-tract spillovers do not violate SUTVA. To see this point formally, note that SUTVA requires $Y_i = Y_i(0) + D_i(Y_i(1) - Y_i(0))$, with Y_i the outcome of tract i , $Y_i(0)$ and $Y_i(1)$ the untreated potential outcomes in tract i , and D_i the treatment status of tract i . Within-tract spillovers depend only on the potential outcomes of tract i and tract i treatment status. Note, however, that within-tract spillovers would violate SUTVA if i denoted a property instead of a tract. Figure E.5 presents a diagram of these spillovers.

Within-tract spillovers. The analysis presented in the main text in

Figure 6 focuses on the rezoned census tracts at the bottom of the diagram to estimate within-tract spillover effects on the red properties. Because the insurance data track whether the insurance policy is purchased inside the 100-year floodplain, I am able to do this analysis by simply changing the outcome variable (using the number of policies purchased outside the 100-year floodplain) while keeping the same research design (comparisons between rezoned tracts where the 100-year floodplain increases and not-yet-rezoned tracts where the 100-year floodplain will increase later).


Below, I directly implement the suggestion of two anonymous reviewers to test the role of distance to the 100-year floodplain. In each tract where the 100-year floodplain increased, I compute the distance of the properties that remain outside the 100-year floodplain throughout the update to the nearest floodplain boundary both before and after the map update.

Figure E.6 presents this process for Washington, D.C. The left panel shows the distance of individual properties to the 100-year floodplain before the 2010 map update. The right panel shows the new 100-year floodplain and the change in distances to the 100-year boundary at the property level.

I then compute two aggregate measures at the tract level: (1) average change in distance to the 100-year boundary (in meters, from negative to positive values), and (2) share of properties that became closer to the 100-year floodplain (in %). Figure E.7 presents estimates of the map update effects on insurance take-up outside the 100-year floodplain by quartile of change in distance to the boundary (for properties that remained outside). In the left panel, the first quartile denotes tracts where properties became closest to the border, while in the second panel the fourth quartile denotes tracts where the most properties became closer to the border. In both panels, the treatment effect is stronger in tracts where properties located outside the 100-year floodplain got closer to the 100-year floodplain boundary, suggesting that within-tract spillover effects exist and decay with distance to the floodplain boundary.


Between-tract spillovers. I estimate between-tract spillover effects using Callaway-Sant’Anna event-study regressions, comparing tracts that are adjacent to a rezoned tract where the map update increased the 100-year floodplain (N1 on Figure E.5) with their own neighbors that were not rezoned (N2), using as the treatment date the date of actual rezoning. Figure E.8 shows these census tracts in the contiguous U.S. Census tracts N1 (in orange) and N2 (in yellow) sometimes themselves receive a flood map, but they are kept in the main analysis to increase the sample size as long as the flood map

Figure E.6: Washington, D.C.: property change to border




The left panel shows the distance of individual properties to the 100-year floodplain (in blue) before the 2010 map update. The right panel shows the new 100-year floodplain (in blue), and the change in distances to the 100-year boundary at the property level. Both panels focus on tracts where the map update rezoned properties inside the 100-year floodplain only. Distances are in meters.

Figure E.7: Within-tract spillovers by distance change to border



ATT estimates obtained using Callaway-Sant'Anna type regressions. The outcome variable is the number of active insurance policies in the tract, in levels. Quartiles denote distance change to the boundary (for properties that remained outside). In the left panel the first quartile denotes tracts where properties became closest to the border, while in the second panel the fourth quartile denotes tracts where the most properties became closer to the border. The x-axes depict time in month, normalized relative to the quarter of the flood map update ($t=0$). The treated groups only include tracts where the map update rezoned properties inside the 100-year floodplain. The control groups comprise not-yet-treated census tracts that later receive a flood map with a similar floodplain rezoning as the treated group. Error bars represent 95% confidence intervals using the multiplier bootstrap.

Figure E.8: Between-tract spillovers




Spatial distribution of tracts N1 (defined as adjacent to a tract with properties rezoned inside the 100-year floodplain) and N2 (defined as adjacent to N1 tracts).

does not rezone any property either inside or outside the floodplain in N1 and N2. As shown in Figure 5, map updates that do not change the floodplain do not impact insurance take-up. In sensitivity checks, I exclude N1, N2, or both types of tracts that receive a flood map – keeping only "true" untreated tracts.

Figure E.9 presents the event-study estimates. Coefficients across specifications are small and generally not statistically significant. Keeping only "true" untreated N1 tracts that never receive a digital map in the analysis leads to wider confidence intervals (second and third panels) and, similarly, fails to detect spillover effects between tracts. These results suggest that the SUTVA assumption holds at the census-tract level.


Figure E.9: Between-tract spillovers



ATT estimates obtained using Callaway-Sant'Anna type regressions. The outcome variable is the number of active insurance policies in the tract, in levels. The top left panel includes all N1 tracts (defined as adjacent to a tract with properties rezoned inside the 100-year floodplain) and N2 tracts (defined as adjacent to N1 tracts), including those that receive a digital flood map (as long as the digital flood map does not rezone any property). The other panels only focus on N1 and/or N2 tracts that are not covered by a digital flood map, as indicated in the panel label. The x-axes depict time in month, normalized relative to the quarter of the flood map update ($t=0$). The treated groups only include tracts where the map update rezoned properties inside the 100-year floodplain. The control groups comprise not-yet-treated census tracts that later receive a flood map with a similar floodplain rezoning as the treated group. Error bars represent 95% confidence intervals using the multiplier bootstrap.


E.4 Event-study heterogeneity

Figure E.10: 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, in levels. The x-axes depict time in month, normalized relative to the quarter of the flood map update ($t=0$). 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 panels 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 E.11: 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, in levels. The x-axes depict time in month, normalized relative to the quarter of the flood map update ($t=0$). 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 panels 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.

F Synthetic controls and distributional impacts

The share of properties rezoned inside and outside the 100-year floodplain varied between neighborhoods. Yet, accounting for disparities in the type of information provided by the map updates cannot be done in a fully flexible way within the event-study framework because of the small sample size issue discussed above. In addition, three features could undermine the parallel assumption. 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 feature 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 my observation window. In particular, the Biggert-Waters Flood Insurance Reform Act of 2012 (BW Act) progressively phased out insurance subsidies for a subset of properties that were constructed before 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 interest. Unfortunately, this approach 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 the heterogeneity of the treatment effects) increases.⁷ To escape this Catch-22, I use tract-specific

⁷A “standard” approach to heterogeneity would be to run two-way fixed effect regressions and interact treatment with census-tract-specific covariates. As discussed above,

synthetic controls.

F.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, 2021). Because directly comparing the outcome of the treated unit with 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.

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

$$\begin{aligned} \min_{\gamma} \quad & \theta_x \|\mathbf{X}_1 - \mathbf{X}'_0 \gamma\|_2^2 + \theta_z \|\mathbf{Z}_1 - \mathbf{Z}_0 \cdot \gamma\|_2^2 + \zeta \sum_{W_i=0} f(\gamma_i) \\ \text{subject to} \quad & \sum_{W_i=0} \gamma_i = 1 \\ & \gamma_i \geq 0 \quad i : W_i = 0 \end{aligned} \tag{10}$$

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 .⁸ The weights θ_x and θ_z govern the relative im-

this approach would deliver biased and inconsistent treatment effect estimates because of 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).

⁸Different choices of penalty functions exist. In this paper, I implement the ridge-augmented synthetic control approach, for which Ben-Michael, Feller and Rothstein (2021)

portance of the deviations between lagged outcomes and covariates in the minimization program.⁹

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

$$\begin{aligned}\hat{Y}_{1T}^{\text{aug}}(0) &= \sum_{W_i=0} \hat{\gamma}_i^{\text{scm}} Y_{iT} + \left(\hat{m}_{1T} - \sum_{W_i=0} \hat{\gamma}_i^{\text{scm}} \hat{m}_{iT} \right) \\ &= \hat{m}_{1T} + \sum_{W_i=0} \hat{\gamma}_i^{\text{scm}} (Y_{iT} - \hat{m}_{iT})\end{aligned}\tag{11}$$

where $\hat{\gamma}_i^{\text{scm}}$ are the solutions to the program in equation 10, 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, Feller and Rothstein (2021) and use a ridge regression for the outcome model.¹⁰

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

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 matching on only the outcome variable, matching on selected 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).¹¹ I construct synthetic controls by matching on pre-treatment values of the outcome variable (the

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).

⁹Following Ben-Michael, Feller and Rothstein (2021), in my preferred approach I use $\theta_x = \theta_z = 1$.

¹⁰This choice of estimator has attractive properties, notably (1) an improvement in pre-treatment fit relative to the standard synthetic control model and (2) a reduction in estimation error under linear and latent-factor data-generating processes.

¹¹Abadie (2021), p. 402, emphasizes this aspect: “Part of the literature on synthetic controls emphasizes estimators that depend only on pre-intervention outcomes and ignore

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 generates 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, 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 my 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 my observation period and that are located within the *same* FEMA region.¹²

Finally, to improve the quality of the pre-treatment fit, I use the augmented synthetic control method of Ben-Michael, Feller and Rothstein (2021). This approach uses an explicit outcome model to remove potential bias in the original synthetic control estimates.

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 my context, this fact means that we are more likely to credibly estimate treatment effects where

the information of other predictors [...]. This reliance 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.”

¹²One possibility is to form the donor pool using units that are not-yet-treated, instead of units that are never-treated. This possibility is appealing, since not-yet-treated units are more likely to resemble already-treated units. However, this approach 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.

flood map updates lead to large changes in insurance take-up, and for census tracts where there are a large number of active insurance policies.¹³ I therefore estimate synthetic controls in census tracts that have at least 20 active policies at all times.¹⁴ 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 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.

F.2 Synthetic control results

This section presents the (ridge-augmented) synthetic control estimates for each of the 3095 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 2009-17 period.

To assess the pre-treatment fit of the synthetic control estimates, I compute the normalized root mean squared error (NRMSE) during the pre-treatment period: for 79% 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 before treatment in most cases.¹⁵ A histogram of the NRMSE is presented in Figure F.1. 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 0.05.¹⁶ In the current application, the smallest and largest donor pools have 58 and 898 units, respectively. Figure F.2 shows

¹³Recent extensions of synthetic controls consider Average Treatment Effect estimation when several units receive treatment at different times (Ben-Michael, Feller and Rothstein, 2022).

¹⁴Because of the spatial concentration of the demand for flood insurance, excluding census tracts with fewer than 20 active insurance policies at all times leads to a negligible reduction in the aggregate number of active policies considered for the synthetic control estimation.

¹⁵It is smaller than .01 in 67% of cases.

¹⁶Most 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 an NRMSE lower than 0.01 in 92% of cases.

that pre-treatment fit is not correlated with the number of units in the pools.

Figure F.3 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, estimating these effects requires strong assumptions and more data. For the remainder of this section, I focus on the estimated treatment effects at 24 months post-treatment.¹⁷

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


$$\widehat{TE}_i = \alpha + \beta Rezoning100year_i \cdot X_i + \epsilon_i \quad (12)$$

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), and 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 of regression 12 is estimated using synthetic controls, uncertainty surrounding the estimates needs 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.

Figure F.4 shows 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 ($Rezoning100year_i$, with rezoning outside 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.


¹⁷Additional results show that effects are similar when focusing instead on treatment effects after 12 months or 30 months.

Figure F.1: 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 F.2: 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 F.3: Synthetic control estimates for insurance take-up in four tracts



Dynamic treatment effects estimates of the impacts of flood map updates on flood insurance take-up in four 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 dashed vertical line denotes the end of the pre-treatment optimization period, while the solid vertical line denotes the effective date of the flood map. The gray ribbon depicts 90% jackknife+ confidence intervals.

Figure F.4: 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 12. The marginal effects are estimated separately per quartiles.

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 12 is 0.23 , meaning that for every five properties rezoned relative to the 100-year floodplain, one is induced to change insurance status. The adjusted R^2 of regression 12 (without heterogeneity variables) is 47%, i.e., the share of properties rezoned in and out of the 100-year floodplain half of the variation in the synthetic control estimates.

As with event-study models, we find that rezoning outside the 100-year floodplain decreases 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 causes 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 F.4) 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 of floodplain rezoning. Map updates did cause greater declines in insurance take-up in neighborhoods with more Black residents, but this result is predominantly caused by the substantial rezoning outside the 100-year floodplain that occurred in these areas (see Figure F.7).¹⁸

That said, Figure F.4 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 *increase* (in absolute value) with income, suggesting that higher-income households are *more* sensitive to the rezoning of properties inside and outside 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.¹⁹

Households may hold private information about their risk exposure. Un-

¹⁸In Figure F.8, I extend the analysis to other racial groups.


¹⁹Outreach 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.

der this scenario, households could selectively respond to flood map updates that are “correct,” and not substantially change their demand for insurance following an incorrect update. Such a 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 F.9, I use the FSF estimates as correct estimates of flood risk as an example. I do not find meaningful differences in the effects of map updates on insurance take-up based on whether the majority of rezoning was correct within a census tract, suggesting that official risk information is the main source of information about risk.²⁰

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 III: rezoning inside the 100-year floodplain increases take-up, while rezoning outside it reduces take-up. Because of the more numerous rezoning outside 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 aligns with the FSF model. As such, their impacts on consumer welfare could be negative.


²⁰Figure F.9 further shows that neighborhoods that were not hit by a disaster before the map update appear to respond more to the flood map, consistent with it being the main source of information. Map updates also appear to 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.

Figure F.5: 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 F.6: 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 F.7: 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 panels 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 F.8: 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 12. The marginal effects are estimated separately per quartiles.

Figure F.9: 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 12. The marginal effects are estimated separately for each group.

G Welfare estimates

G.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:

$$\begin{aligned}
 -\exp\left(-\alpha(\omega)(Y(\omega) - r(\omega))\right) &> -\hat{p}_\omega \exp\left(-\alpha(\omega)(Y(\omega) - L(\omega))\right) \\
 &\quad - (1 - \hat{p}_\omega) \exp\left(-\alpha(\omega)(Y(\omega))\right) \\
 \iff -\exp\left(\alpha(\omega)r(\omega)\right) &> -\hat{p}_\omega \exp\left(\alpha(\omega)L(\omega)\right) - 1 + \hat{p}_\omega \\
 \iff \alpha(\omega) &> k_\omega
 \end{aligned} \tag{13}$$

where $\hat{p}_\omega$ is the *perceived probability* of flooding, $L(\omega)$ is 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. $L(\omega), r(\omega), Y(\omega)$ 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:

$$\begin{aligned}
 -\exp(\alpha(\omega)WTP_\omega) &= \int_0^{+\infty} -\exp\left(\alpha(\omega)D_{d,\omega}\right) dF_{d,\omega} \\
 &\quad \ln\left(\int_0^{+\infty} \exp\left(\alpha(\omega)D_{d,\omega}\right) dF_{d,\omega}\right) \\
 \iff WTP_\omega &= \frac{\ln\left(\int_0^{+\infty} \exp\left(\alpha(\omega)D_{d,\omega}\right) dF_{d,\omega}\right)}{\alpha(\omega)}
 \end{aligned} \tag{14}$$

where $D_{d,\omega}$ is 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.

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. My model therefore provides a lower-bound on the welfare gains from corrected flood maps.²¹

Consistent with the empirical evidence presented in Section III:and Appendix F, 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 0% outside it.²² 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 underestimate the impacts of new flood maps. Figure G.1 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 estimates of annual expected losses in the FSF model (the expected losses to the property in a given year).

²¹My model does so in 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.

²²This assumption will bias the impacts of new flood maps towards zero, as the analysis above demonstrates the existence of within-neighborhood 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.

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:²³

$$p(\alpha(\omega) \leq k_\omega) = \exp\left(-\gamma_c\left(\frac{k_\omega}{A_c}\right)^{-\theta_c}\right) \quad (15)$$


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 fact poses a challenge for the structural estimation: unless *all* existing frictions are correctly specified, one cannot separately estimate frictions and risk preferences.

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* estimating 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 the 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.

²³Consistent with the literature, I assume that all homeowners are weakly risk averse.

Figure G.1: 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.

References

- Abadie, Alberto.** 2021. “Using Synthetic Controls: Feasibility, Data Requirements, and Methodological Aspects.” *Journal of Economic Literature*, 59(2): 391–425.
- Armal, Saman, Jeremy R. Porter, Brett Lingle, Ziyang Chu, Michael L. Marston, and Oliver E.J. 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.
- Barber, Rina Foygel, Emmanuel J. Candès, Aaditya Ramdas, and Ryan J. Tibshirani.** 2021. “Predictive inference with the jackknife+.” *The Annals of Statistics*, 49(1): 486–507.
- Ben-Michael, Eli, Avi Feller, and Jesse Rothstein.** 2021. “The Augmented Synthetic Control Method.” *Journal of the American Statistical Association*, 116(536): 1789–1803.
- Ben-Michael, Eli, Avi Feller, and Jesse Rothstein.** 2022. “Synthetic Controls with Staggered Adoption.” *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 84(2): 351–381.
- BLS.** 2024. “Bureau of Labor Statistics Consumer Price Index. Last accessed online January 2024.” <https://data.bls.gov/timeseries/CUUR0000SA0>.
- Callaway, Brantly, and Pedro H.C. Sant’Anna.** 2021. “Difference-in-Differences with Multiple Time Periods.” *Journal of Econometrics*, 225(2): 200–230.
- Center for Emergency Management and Homeland Security.** 2020. “Spatial Hazard Events and Losses Database for the United States, Version 19.0.” <https://cemhs.asu.edu/sheldus>.
- Cummins, David J.** 1991. “Statistical and Financial Models of Insurance Pricing and the Insurance Firm.” *The Journal of Risk and Insurance*, 58(2): 261–302.
- Environmental Protection Agency.** 2016. “EPA Dasymetric USA Interpolation for 2016/02/08, last accessed July 2024.” <https://edg.epa.gov/data/public/ORD/EnviroAtlas/National/ConterminousUS/>.

- FEMA.** 2016. “Flood Insurance Study - City of New Orleans and Orleans Parish, Louisiana.” FEMA.
- FEMA.** 2020*a*. “Community Status Book, accessed online January 2020.” <https://www.fema.gov/flood-insurance/work-with-nfip/community-status-book>.
- FEMA.** 2020*b*. “FEMA Q3 Flood data: version from Caliper Corporation.” <https://www.caliper.com/mapping-software-data/flood-zone-data.html>.
- FEMA.** 2020*c*. “FIMA NFIP Redacted Policies. Obtained through Freedom of Information Act Requests, received between December 2020 and June 2021.”
- FEMA.** 2020*d*. “OpenFEMA Disaster Declarations Summaries. Accessed online January 15, 2023.” <https://www.fema.gov/about/openfema/data-sets>.
- FEMA.** 2020*e*. “OpenFEMA FIMA NFIP Redacted Claims. Accessed online December 15, 2020.” <https://www.fema.gov/about/openfema/data-sets>.
- FEMA.** 2021. “FEMA Q3 Flood data: Princeton Library Repository, Accessed February 2021.” <https://maps.princeton.edu/catalog/>.
- FEMA.** 2023. “Risk Mapping, Assessment, and Planning (Risk MAP) Multi-Year Plan: 2023-2027.”
- FEMA.** 2024. “FEMA Q3 Flood data, obtained through Freedom of Information Act Requests between January 2020 and March 2024.”
- Fowler, Hayley, Geert Lenderink, Andreas F. Prein, and et al.** 2021. “Anthropogenic intensification of short-duration rainfall extremes.” *Nature Reviews Earth & Environment* volume, 2: 107–122.
- Gallagher, Justin.** 2014. “Learning About an Infrequent Event: Evidence from Flood Insurance Take-Up in the US.” *American Economic Journal: Economic Policy*, 6(3): 206–233.
- Government Accountability Office.** 2021. “FEMA FLOOD MAPS: Better Planning and Analysis Needed to Address Current and Future Flood Hazards.” U.S. Government Accountability Office GAO-22-104079.

- Hino, Miyuki, and Marshall Burke.** 2021. “Replication code for: The effect of information about climate risk on property values.” <https://purl.stanford.edu/td021fz7393>.
- Howe, Peter D, Matto Mildenerger, Jennifer R Marlon, and Anthony Leiserowitz.** 2015. “Geographic variation in opinions on climate change at state and local scales in the USA: Supplementary Information.” *Nature Climate Change*, 5(6): 596–603.
- United States Census Bureau.** 2010a. “2010 Decennial Census, Summary File 1.” <https://www.census.gov/data/developers/data-sets/decennial-census.2010.html#list-tab-517985795>.
- United States Census Bureau.** 2010b. “American Community Survey Pooled 5-year estimates, 2015.” <https://www.census.gov/acs/www/data/data-tables-and-tools/data-profiles/2015/>.
- United States Census Bureau.** 2018. “2018 TIGER/Line Shapefiles (machine readable data files).” <https://www2.census.gov/geo/tiger/TIGER2018/BG/>.
- Wagner, Katherine R. H.** 2022. “Adaptation and Adverse Selection in Markets for Natural Disaster Insurance.” *American Economic Journal: Economic Policy*, 14(3): 380–421.
- Walker, Kyle, and Matt Herman.** 2025. “tidycensus: Load US Census Boundary and Attribute Data as ‘tidyverse’ and ‘sf’-Ready Data Frames.” R package version 1.7.1.
- Zillow.** 2019. “ZTRAX: Zillow Assessor Dataset, 2019-Q4.” <http://www.zillow.com/ztrax>.