Do Mortgage Lenders Respond to Flood Risk?

Kristian S. Blickle*

Federal Reserve Bank of New York

Evan Perry*

Federal Reserve Bank of New York

João A. C. Santos*

Federal Reserve Bank of New York

&

Nova School of Business and Economics

First Version: March 12, 2024

This version: August 26, 2025

Abstract

Using unique property-level location, mortgage, flood risk, and FEMA flood map data for the entire US, we document that a significant number of dwellings are subject to flood risks that are not captured by official FEMA maps. Lenders are aware of this "un-mapped" risk. Large banks, in particular, originate fewer – and more expensive – loans for affected properties. In contrast, non-banks have taken the opportunity to increase their market share in these areas, while securitizing these loans more aggressively. All lenders sell large portions of their mortgages to properties in high-flood risk areas to GSEs and their conservative lending appears to have had an adverse effect in house prices in those areas.

Key Words: Floods, Flood Risk, HMDA, FEMA, Credit Constraints

*Contact: evan.perry@ny.frb.org. We thank seminar participants at the Federal Reserve Bank of New York as well as Linda Goldberg, Itay Goldstein, Stephan Luck, Matthew Plosser for valuable comments and suggestions. Any errors are our own. The views stated herein are those of the authors and are not necessarily the views of the Federal Reserve Bank of New York or the Federal Reserve System.

1 Introduction

In 2021 over 200,000 mortgages for residential homes – worth over 66 billion USD – were originated in areas covered by a FEMA 100-year flood map. Properties located in these "flood zones" are considered sufficiently at risk of water-related disasters that flood insurance is required for originated mortgages. This ensures the risk of catastrophe is not borne by the mortgage borrower or their lender alone.[1] However, flood maps are discrete with stark boundaries while flood risk itself is continuous across certain geographies. Moreover, flood maps can be outdated or inaccurate in the face of rapidly changing climate and weather patterns. Indeed, in 2021 alone 671,000 mortgages for residential properties – worth over 205 billion USD – were originated in regions that face significant flood risk but that were not covered by a FEMA flood map. Looking at the entire US, of the properties in the top 1% of CoreLogic's flood risk data, nearly two-thirds (64%) are in an official flood zone whereas the remaining third (36%) are not. Similarly, of the properties in the top 5% of the flood risk distribution, only half (52%) are in an official flood zone and the other half (48%) of these high risk properties are not. In this paper, we investigate whether mortgage lenders are aware of flood risk in these un-mapped regions and how (if at all) they manage that risk.

We build on a restricted version of the HMDA dataset that contains addresses for the years 2018-2021. This dataset enables us to match records of individual mortgage applications to property-level flood risk measures from CoreLogic as well as individually digitized flood maps. We call properties that are not covered by either a 100-year or 500-year FEMA flood map, but which face comparable flood risk to those that are covered, "un-mapped". We analyze the degree to which lenders factor this un-mapped flood risk into their lending policies by analyzing the propensity to grant mortgages (the extensive margin) and the value/prices of these mortgages (intensive margin). We also investigate the degree to which lenders pass risk to other entities – especially the GSEs – and whether lenders' management of flood risks has impacted their market shares in mortgage markets. Finally, we round out our paper by studying how lending policies affect house prices in un-mapped regions, paying attention to disentangle supply and demand effects by focusing both on house price data as well as on data that covers a lender's valuations of

[1]The insurance scheme and the consequences thereof are discussed in detail in Blickle and Santos (2021).

properties.

Our paper yields two important findings. First, lenders are aware of flood risk outside of mandated flood zones and manage their exposure to this risk. Specifically, we find that lenders are generally less likely to originate mortgages for properties that face un-mapped flood risk. For mortgages that are originated – despite the risk – lenders charge slightly higher interest rates while assigning a lower value during the inspection process. Secondly, we show that non-banks (and to a lesser extant small local banks) manage this risk through aggressive loan securitization/loan sales while absorbing market share from the more conservative larger banks. This has helped non-banks gain market share in high risk mortgage markets.

It appears that GSEs, in particular, have been absorbing mortgages with significant exposure to flood risk – but which do not appear to charge commensurate interest rates – for some time.

We document several other consequences of the reduced mortgage lending in high-flood risk areas. Specifically, the reduced availability of credit in un-mapped areas manifests as lower house prices. Given that the price reduction is smaller than the aforementioned value reduction imposed by the bank during the inspection, borrowers are forced into lower LTV loans. We also find that a lenders' reluctance to lend in un-mapped regions is attenuated in high-income census tracts. This may be because high income borrowers can more easily weather the negative impact of a flooding disaster. Similarly, we find that lenders' reluctance is lower for jumbo mortgages, also possibly because wealthier borrowers offer more innate protection in the event of flooding disasters.

Methodologically, we are able to use property-level data. As such, we can account for property and borrower characteristics in almost all regressions. Perhaps more importantly, we are able to use property-level flood risk information and place each property on a FEMA flood map. Properties with risk that are not covered by a FEMA 100-year, 500-year, or flood-way map can thus be cleanly identified. Map boundaries can move aggressively within a community, so low resolution or aggregated analyses may be insufficient for our purposes.

Empirically, we then relate loan-level outcomes to our measure of un-mapped flood risk and borrower characteristics. Given that we do not know what flood information a lender has access to, we make use of publicly available – but proprietary – risk data from CoreLogic and assume that lenders have access to similar data. We show that Corelogic risk assessments are similar to other third party assessments (such as First Street or more recent FEMA Risk indices) in an extension of the paper. Also, we cannot see how true flood risk has changed over the years. As such, we are analyzing ex-post equilibrium lending decisions for a narrow band of time during which information on flooding was **theoretically** widely available. Arguably, our results could be seen as lower bound estimates of lenders' true response to flood risk information.

Our paper is related to the recent body of research on banks' ex ante responses to climate physical risks.[2] Sastri (2021) looks at credit rationing following flood map changes in Florida. Blickle and Santos (2021) use national map changes to show that banks are less likely to lend to certain customers and in certain areas following map changes. Ortega and Petkov (2024) look at the effects of NRI 2.0 on insurance uptake. Keenan and Bradt (2020) document that local lenders transfer risk from mortgages collateralized by properties in high-risk coastal geographies in the Southeast Atlantic and Gulf Coasts through securitization, consistent with them being better informed about local risks than larger lenders with diversified portfolios. This contrasts with Keys and Mulder (2020) finding that mortgage lenders have not meaningfully changed their rates of securitization in the flood exposed areas of Florida between 2013 and 2018. Meisenzahl (2023) uses supervisory data and finds that banks retain lower credit exposures to areas subject to climate disasters/risk. Cohen et al. (2021) show the negative impact on real estate of flooding outside flood zones during hurricane Sandy in New York. Finally, Mulder (2022) analyzes the welfare implications of inaccurate flood mapping.[3]

In contrast to this body of research, we focus exclusively on flood risk outside designated flood zones to avoid the effects coming from the information associated with that designation and the mandatory flood insurance it imposes on borrowers. Meisenzahl (2023) consider climate disasters/risk throughout the US but does not distinguish the effects within vs outside flood zones. Perhaps more importantly, in contrast to them we investigate how banks manage their risk exposures through origination vs. sale and securitization decisions. Like Cohen et al. (2021) we too are interested on the impact of flood risk outside flood zones on housing values but our main focus is on the impact on property valuations banks use in mortgage applications and on their

[2]Another body of research looks at how banks change their ex post lending responses to climate disasters. A branch of this literature focuses on corporate lending (Cortes (2015), Chavaz (2016), Schüwer et al. (2018), Cortes and Strahan (2017), Rehbein and Ongena (2020), and Ivanov et al. (2022)), and another branch focuses on mortgage lending (Garbarino and Guin (2021)). Ivanov et al. (2022) and Jung et al. (2022), in turn, investigate U.S. banks' corporate lending responses to climate transition risks.

[3]The NFIP program itself as well as the implications flood risk broadly are discussed in a number of prominent papers such as Kousky (2010), Dinan et al. (2019), Kousky (2018), Kousky et al. (2020), or Gourevitch et al. (2023).

applications' decisions.

Our paper is also related to the existing literature investigating to what extent flood risks are capitalized into housing prices. Hino and Burke (2021) find that even within FEMA flood zones, housing prices fail to fully price the underlying flood risk, and that, considering only the flood risk communicated in these maps, housing in FEMA flood zones is overvalued by 32.6-55.6 billion USD. In a more recent study of housing both inside and outside of FEMA flood zones, Gourevitch et al. (2023) estimate that U.S. housing is overvalued by 121-237 billion USD due to unpriced flood risk. The authors highlight that overvalued properties are especially common in coastal areas just outside of official FEMA flood zones. Similarly, a number of studies have analyzed the pricing effects of risk revelations that are associated with the mandatory disclosure of flood maps (see for instance: Troy and Romm (2004), Pope (2008), Shr and Zipp (2019), or Gibson and Mullins (2020)). In this paper, we focus specifically on how mortgage lenders treat properties with "un-mapped" flood risks, thereby connecting this body of literature focused on flood risk in the housing market to the mortgage market and role of financial intermediaries.

Finally, our paper points to a novel explanation for the rise of non-banks in mortgage lending. Researchers, including Buchak et al. (2018), Fuster et al. (2019), Sharpe and Sherlund (2016), and Choi et al. (2022), argue that regulation together with capacity constraints and the reliance on new lending technologies explain the spectacular rise of non-banks' market share in the US mortgage market. Our evidence suggests that large banks' retreat from lending in unmapped high-flood risk areas, together with the ability to securitize mortgages, also helped non-banks increase their market share in the mortgage market.

The remaining paper is structured as follows. Section 2 introduces our data. Section 3 describes the "un-mapped" properties the analysis focuses on. Section 4 outlines our methodology. Section 5 first demonstrates that mortgage lenders are aware of flood risk beyond official flood maps – as evidenced by their origination rates, interest rates, and property valuations – and that they respond heterogeneously. Section 6 describes the implications of this lending response to flood risk for prospective homeowners through the primary mortgage market as well as for the GSEs operating in the secondary mortgage market. That section also tries to ascertain the impact of that response on house prices. Finally, section 7 concludes.

2 Data

Our primary dataset for this paper is a restricted version of the Home Mortgage Disclosure Act (HMDA) data set. We complement this dataset with proprietary, property-level measures of flood risk from CoreLogic and archived FEMA flood maps. We also use data on real estate listings from CoreLogic.

2.1 CoreLogic Flood Risk Data

We begin with a structure-level dataset of composite flood risk (including inland flooding and hurricane storm surge flooding) from CoreLogic for 2021. This dataset contains flood risk metrics for structures on over 107 million properties in the US.[4] We filter for residential properties and exclude from our data unfinished structures or secondary structures within a single parcel of land. As such, we do not conflate residences with sheds or garages that may face different risks than the residential structure itself. For each structure, the dataset includes longitude and latitude information that we use to match to flood map and mortgage data.

Our primary flood risk measure is a property's average annual loss (AAL) from flooding. This is the expected value of a structure's annual flood damages as a proportion of the estimated cost to reconstruct the entire structure. We use property-level AALs that incorporate losses due to any available type of flooding (inland, hurricane storm surge, tsunami) under current climate conditions.

A major assumption in our approach is that the risk for a property did not change between 2018 and 2021. This assumption seems reasonable because information about risks is not produced at a sufficiently rapid pace to be released annually. Nevertheless, our inability to use historically available risk data means we could be assuming a greater knowledge of risk on the part of the lender than was truly available at the start of our sample (2018). To the extent flood risk has been increasing, our estimates on a lender's response to flood risk must be viewed as lower bound estimates.

[4] CoreLogic measures flood risk by first using hydrological models that map storm events to water levels. This is then coupled with hydraulic models that map these to streamflow characteristics (e.g., the velocity of flood streams). With the mapping between storm scenarios and physical flooding outcomes, CoreLogic then uses calibrated damage functions that map these flooding conditions along with structure characteristics (e.g., number of levels, structure age, occupancy status) to estimate flood damages under a given climate scenario and timeline.

2.2 Flood Maps

To differentiate between mortgage borrowers that **are** and **are not required** to purchase flood insurance, we make use of the National Flood Hazard Layers from the Federal Emergency Management Agency (FEMA). The map layers demarcate regions at risk of experiencing a 100-year flood – that is, regions that FEMA estimates have at least a 1% probability of experiencing a severe flooding event in any given year. Importantly, borrowers for properties in these areas are required to purchase flood insurance through the National Flood Insurance Program (NFIP). We use an archived snapshot of the FEMA flood map layer from early 2022, just after our sample period (2018-2021). Given that FEMA map boundaries rarely change, it is unlikely that we classify mortgages as being outside a 100-year flood zone when they were actually within a 100-year flood zone at the time of origination.

We place each property from the CoreLogic flood risk data onto the FEMA flood maps to determine which of these properties fall into a 100-year flood zone or a 500-year flood zone and which do not fall into these zones. Borrowers in the 500-year flood zone areas are not required to purchase flood insurance but do still have a public signal of their homes' underlying flood risk. We do not want to conflate this effect with the effect of completely un-mapped flood risk on lending decisions.[5] Given that our property-level flood risk data must be anonymized after merging with the property-level mortgage lending data (more on this below), we form flood risk buckets. Just over 48% of properties have no risk (i.e. AAL = 0). Conditional on having non-zero risk, we flag properties in the upper half of this distribution (the top 26% of the full distribution) as moderate-to-high flood risk. We classify a property as "un-mapped" if it has moderate-to-high flood risk and does not fall into a FEMA-designated 100-year flood zone, 500-year flood zone, or flood way. Similarly, we classify a property as "possibly un-mapped" if it has any non-zero flood risk and is not in a FEMA-designated 100-year flood zone, 500-year flood zone, or flood way.

A potential concern with our investigation is that it relies on CoreLogic's approach to risk

[5]While considerable flood risk exists both inside and outside of official flood zones, outside of official flood zones, mortgage lenders do not benefit from the flood risk information conveyed in the maps and do not benefit from the flood insurance mandate. When controlling for flood risk, we estimate that flood insurance take-up for properties outside of the official 100-year flood zone are over 15-percentage points lower than properties inside the official 100-year flood zone. We discuss this in detail in the dedicated Appendix section A.8. Our approach gives us the opportunity to ascertain mortgage lenders' responses to flood risk that, while similar to risks present in FEMA's flood zones, is outside any official flood maps.

identification which may be unique in that it is (i) different from other flood risk data providers or (ii) overly aggressive so that areas are considered at risk that face no true flooding hazards. The second concern can be assuaged somewhat by the frequency with which areas outside of flood zones have been flooding.[6] However, we still wish to ensure that our flood data is accurate. To that end, in Appendix section A.4, we benchmark CoreLogic risks against other risk metrics including flood risk measures derived from FEMA-supported NFIP premiums data and flood risk data from First Street Foundation. We find significant overlap—even in the areas outside of FEMA flood maps that are critical to our identification strategy.

2.3 HMDA Data

Data on mortgage applications comes from a restricted version of the Home Mortgage Disclosure Act (HMDA) database available to the Federal Reserve for the years 2018 to 2021. Unlike the public HMDA dataset, our restricted version of the HMDA dataset contains specific dates of application and origination, the credit scores of mortgage applicants, and more precise values for fields such as the loan amount, loan sales, and property value. The data tracks the ultimate owner of a mortgage and does not "restart" in January.[7] Importantly, the restricted version we use contains the address of the property on the mortgage application, which allows us to geocode the mortgage records in HMDA and ultimately match these mortgage records to property-level measures of flood risk and FEMA flood map coverage. We restrict our sample to new primary home-purchase applications, excluding mortgage applications with the purpose of refinancing or home improvements.

Similar to the public version, the confidential version of HMDA contains mortgage applicant characteristics like income, gender, and race. The addition of applicant credit scores in the confidential data helps us account for borrower credit quality. HMDA data also includes details on the mortgage (e.g. the loan amount, whether it was originated, and whether it was sold to another financial institution or securitization agency) and census-tract level information on income and demographics (from the 2015 American Community Survey).

[6]Consider the recent cases in Florida that made national news:
https://www.insurancejournal.com/news/southeast/2024/08/13/788205.htm

[7]The public version of the data shows a stark decrease in securitization rates in November and December as the loan has not yet been sold and the data tracking rests in January.

Lastly, we leverage lenders' internal property valuations used in their mortgage lending decisions. This information is useful because it comes from the approval and inspection process (i.e. the inspection that goes into each mortgage origination decision). It will naturally include a bank's and an expert assessor's view of the property and inform the loan amount that a property would potentially be eligible for, no-matter the sale price. High price properties with low valuations imply the need for buyers to put in more equity (i.e. low LTV loans).

For each property, we use the address data available in the restricted version of HMDA to identify the property's longitude and latitude. We then match our HMDA data to CoreLogic and FEMA flood map data using nearest property matches. Because parcel size varies with population density, we implement match-filters that vary with population density. For properties located in census tracts that overlap with a census-designated urbanized area or cluster (i.e., "urban tracts"), we keep only matches where the geolocated HMDA property is within 250 meters of the CoreLogic property; for properties in all other census tracts ("rural tracts"), we keep only matches where the geolocated HMDA property is within 1000 meters of the CoreLogic property. Our results are unaffected by our choice of cutoff values, but these differentiated cutoffs prevent us from more systematically excluding rural properties from the sample. Of all the properties in HMDA from 2018-2021 we match over two thirds to a property in the CoreLogic flood risk data within 25 meters (i.e. the margin of error of our longitude/latitude data).

Our primary analysis makes an important set of sample restrictions. First, we consider only properties that are outside of 100-year and 500-year FEMA flood zones. This sample restriction allows us to compare lending outcomes for properties that do not have any officially mapped flood risk, but still differ in their expected flood damages. Second, we limit our primary analysis to conforming loans. Lenders may respond differently to un-mapped flood risk for jumbo loans as they do not have the option to securitize these with one of the public securitization agencies. We consider jumbo loans in Appendix A.7.

In total, we have a sample of over 13.7 million applications over the four year period between 2018-2021. Table 1 shows summary statistics for this sample.[8] Based on their AAL estimates, we classify 17% of properties in the sample as un-mapped and 47% of properties as possibly

[8]See Table A.1 in the Appendix for supplemental tables and figures for the sample that includes non-conforming (jumbo) loans and properties in 100-year or 500-year FEMA flood zones.

Table 1: *Primary Sample Descriptive Statistics*

Variable	Obs	Mean	SD	10%	Median	90%
Primary Applicant Characteristics						
Age	13,713,259	41	14	26	38	61
Credit Score	13,715,494	727	61	642	736	800
Income	13,715,494	98	75	38	78	175
Male	13,715,494	0.61	0.49	0	1	1
Hispanic	13,715,494	0.12	0.33	0	0	1
Black	13,715,494	0.09	0.28	0	0	0
White	13,715,494	0.72	0.45	0	1	1
Loan Characteristics & Outcomes						
Property Value (1000 USD)	13,585,842	300	189	116	258	532
Loan Amount (1000 USD)	13,715,494	242	137	86	220	428
Loan-to-Value Ratio	13,585,842	85	20	62	92	100
Loan Term (Months)	13,686,232	344	53	300	360	360
Interest Rate	12,281,365	3.8	1.1	2.8	3.8	5.1
Loan Approved	13,715,494	0.9	0.3	1	1	1
Loan Originated	13,715,494	0.87	0.33	0	1	1
Loan Sold or Securitized	13,715,494	0.73	0.44	0	1	1
Loan FHFA Securitized	13,715,494	0.43	0.5	0	0	1
Flood Risk						
Un-mapped	13,715,494	0.17	0.38	0	0	1
Possibly Un-mapped	13,715,494	0.47	0.5	0	0	1

Note: Descriptive statistics for our primary variables of interest and controls appear in the table above. In addition to restricting the dataset to mortgage applications for primary home purchases, the sample we use in the table keeps only mortgage applications that were either approved or denied (excluding applications that were withdrawn), are within the conforming loan limit, and lie outside of 100-year and 500-year FEMA flood zones. Applicant age, income, property value, loan amount, combined LTV, and interest rate are winsorized between the first and ninety-ninth percentiles. Our dummy for loan securitizations includes only those loans that were sold to one of the public securitization agencies within the same calendar year.

un-mapped (i.e. an additional 30% of properties, as the categories are not mutually exclusive). Conditional on a lender approving a loan application, nearly all (96%) of borrowers accept the terms and the loan is originated.[9] Some variables are reported at a slightly lower frequency, which costs us observations in specific tests.

As can be seen from Table 1 the average age of the primary applicant is 41 years. The average applicant has a credit score of 727 and total average annual income of 98,000 USD. The average value of (conforming) properties in our sample is 300,000 USD for which the applicant is seeking

[9]We consider a loan accepted only if both parties agree to the terms. The relatively high acceptance rate in our data follows from the fact that many applicants are "soft rejected" by a bank before completing the official application process.

a 242,000 USD mortgage. In our sample, 87% of loans are originated and 73% of all conforming loans are ultimately moved from the lender's balance sheet through either sale or securitization.[10]

2.4 House Price Data

To gain a more complete understanding of the impact of un-mapped flood risk throughout the mortgage lending process, we incorporate data on real estate listings. These data come from CoreLogic's Multiple Listing Services (MLS) database, which sources real estate listing and transaction records from local MLS across the country. These data allow us to observe the closing prices for individual properties for the majority of the country over the sample period. Due to confidentiality concerns we cannot match these data on house listings to the restricted HMDA data at the property level. Instead, we form a supplementary panel dataset that combines the real estate data and HMDA data after aggregating both to the census tract-quarter level. The resulting panel allows us to benchmark the pricing of flood risk in the mortgage market, as reflected by the appraisal values banks use in their lending decision, against the pricing of flood risk in the housing market. In particular, this dataset allows us to study a flood risk mitigation channel that works through the appraisal process.

Table A.2 displays summary statistics for this panel. The average number of properties purchased in each tract-quarter is 16 in both the HMDA and CoreLogic real estate data – a strong indication that we successfully capture the same properties within each tract-quarter. On average, the property values that lenders use when determining whether or not to lend is greater than the closing price of properties in our sample. It is possible that this is a purely contemporaneous feature of the sample period – a period that includes the first two years of the COVID-19 pandemic – but may just as well reflect some consistent feature of the housing and mortgage lending markets. Regardless, our identification focuses not on the difference between closing prices and appraised property values, but on the relative differences in appraised property values for un-mapped and other properties, conditional on their closing prices.

[10]Many non-bank entities first sell loans to organizations under the same umbrella before they are ultimately securitized, so we treat "sold to affiliate" and "securitized" as similar outcomes for our purposes.

3 Documenting "Un-Mapped" Properties

In this section, we demonstrate the importance of considering the flood risk borne by "un-mapped" properties. These homes which still face serious flood risk, account for a sizable market segment, and typically do not have flood insurance.

Certainly FEMA flood maps cover a large portion of the properties that are most exposed to flood risk. However, there are many properties that are not in flood zones and still bear noteworthy flood risk. Appendix Figure A.1 plots the AAL (average annual loss from CoreLogic) density distributions of properties that are un-mapped, in 100-year flood zones, and in either a 100-year or 500-year flood zone. As described in the previous section, an unmapped property is defined as a property with flood risk above the 50th pctile of all properties with non-zero AAL that is not covered by either a 100- or 500-year map. While the mean un-mapped property has a lower AAL than the mean property in the 100-year flood zones, it is also true that un-mapped properties have AALs that are comparable to a reasonable number of properties in FEMA designated flood zones (i.e. above the 25th pctile of properties in a 100 year flood zone).[11] Ultimately, it is reasonable for any risk averse lender (or borrower), operating in un-mapped regions, to concern themselves with flood risk.

Not only are un-mapped areas risky, but the volume of loans in these areas is large. Over the four-year period of our sample, 794.3 billion USD worth of mortgages were originated in areas with un-mapped risk. Appendix table A.3 displays additional details on the originated loan volumes, split by flood risk.[12] These un-mapped properties are not confined to a single region, but are spread across the country. Figure 1 displays the county-level share of properties in our sample we identify as un-mapped. We can see that affected properties can be found throughout the country along both inland waterways – such as the Mississippi – and along the coast. Un-mapped properties are also not uncommon in cities far from major waterways. As others have noted, FEMA flood maps focus on riverine and coastal flooding, but do not capture risk from pluvial flooding, created when a powerful downpour overwhelms local drainage infrastructure.[13]

[11] In our CoreLogic and FEMA matched dataset, 93.7% of properties in 100-year FEMA flood zones have an AAL that places them above the 50th pctile of non-zero AAL scores.

[12] Table A.3 displays details for just conforming mortgages. Of the 794.3 billion USD in mortgages originated for un-mapped properties from 2018-2021, 571.5 billion USD was for conforming mortgages.

[13] See, for example: https://www.washingtonpost.com/climate-environment/interactive/2022/fema-flood-risk-maps-failures/

Figure 1: *County-level % of Un-mapped Properties*

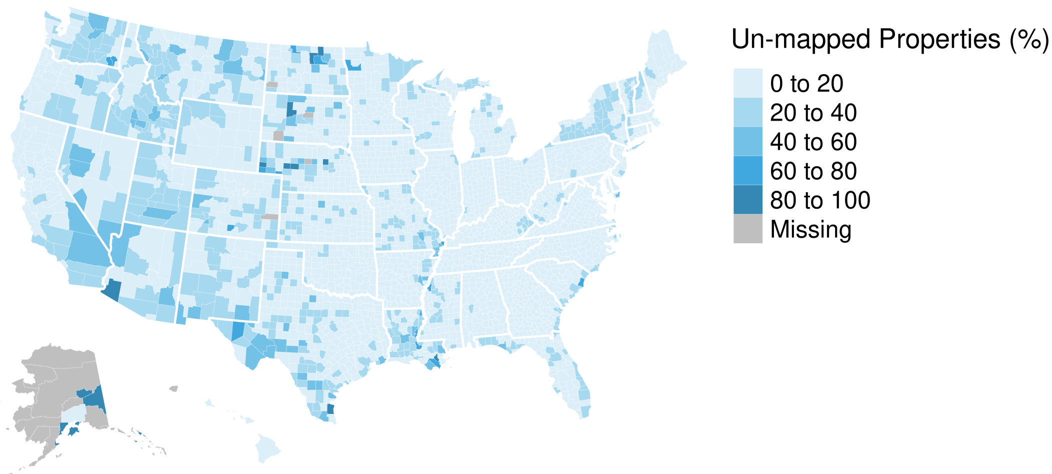

Note: Figure displays the percent of properties in the matched sample of mortgages either approved or denied that we classify as un-mapped in each county. Un-mapped properties compose the largest share of properties in counties along the coast, stretches of the Mississippi river, and across large portions of the western US.

Lastly, un-mapped properties are usually uninsured. Appendix A.8 examines differences in flood insurance take-up inside and outside of official FEMA flood zones. Overall, flood insurance take-up both inside and outside of official flood maps is low: in our analysis, we estimate that 19.89% of single-family residential properties within the 100-year flood zone have flood insurance, compared to just 1.03% outside of the 100-year flood zone. The low take-up within the 100-year flood zone, where flood insurance is required to receive a mortgage, is partially explained by the share of properties that do not have an active mortgage and non-compliance with the flood insurance mandate. Even when we control for flood risk, we estimate that insurance take-up is 15.65 percentage points lower outside of flood zones than inside flood zones. Additionally, coverage limits on the dollar insurance amount mean that even insured properties are often not "fully insured".

Given the severity of flood risk, the volume of affected mortgages, and the lack of insurance for un-mapped properties, we choose to focus on how mortgage lenders respond to this risk. Figure 2 shows an example of our empirical approach, which we outline in greater detail in the next section. The map displays Baton Rouge, Louisiana and the neighboring area alongside the Mississippi

Figure 2: *Parcels with Mapped and Un-Mapped Properties*

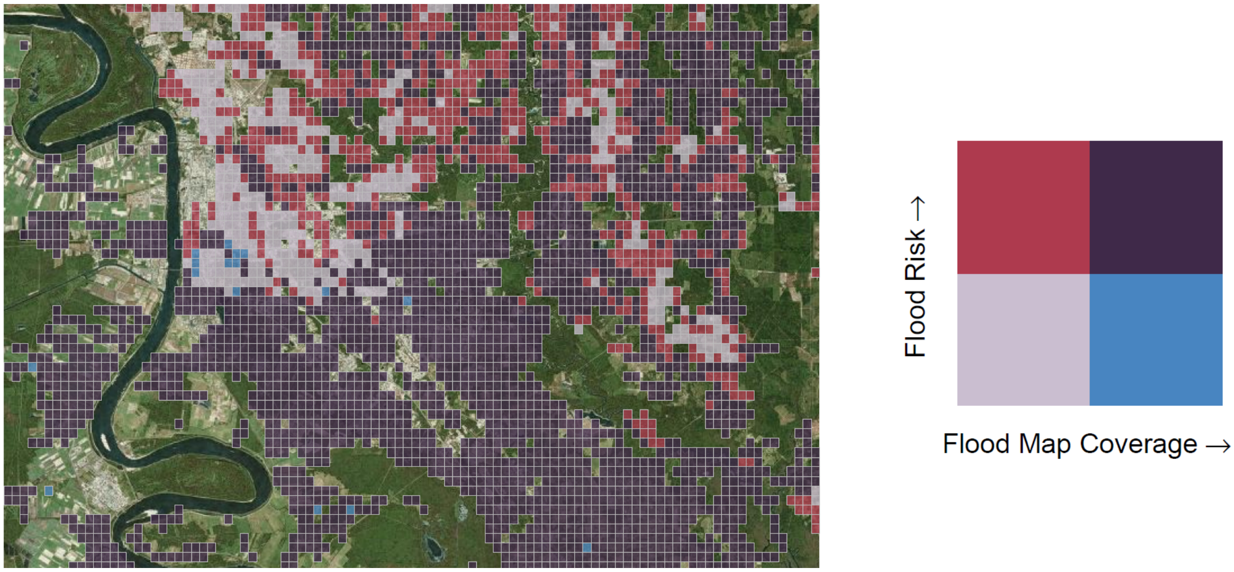

Note: This map plots flood risk measures derived from the CoreLogic data and FEMA flood map coverage along the Mississippi river by Baton Rouge, Louisiana. We overlay a 0.0025° × 0.0025° (≈ 250 × 250 meters) grid over the city. In each grid cell, we compute the (1) the proportion of properties in a 100-year FEMA flood zone, and (2) the mean composite flood risk AAL for CoreLogic properties located in the cell. For anonymity, we only display grid cells with at least five CoreLogic properties. The coloring for each grid cell is then determined by (1) whether the majority of the cell is covered by a FEMA flood map, and (2) whether the mean property in the cell has a high enough AAL that we would consider it to be "high risk." Gray cells have low flood map coverage and low flood risk; red cells have high flood risk and low flood map coverage; blue cells have low flood risk and high flood map coverage; and purple cells have high flood risk and high flood map coverage.

River. For the sake of anonymity, we do not display the property-level data we use in the analysis and instead overlay a grid. In each grid cell, we compute the proportion of properties covered by a FEMA flood map and the mean AAL. The color of the grid cell indicates (1) if most properties are in a 100-year flood zone, and (2) if the mean property has "high flood risk". Properties with non-zero AAL that are above the 50th pctile that are not covered by a flood map are considered "un-mapped" – these propererties are covered by red grids. We would consider purple and gray grid cells to be accurately mapped, as these have high risk-high flood map coverage and low risk but no flood map coverage respectively. Blues grid cells are uncommon and indicate areas that are aggressively mapped, in that the mean property does not have a high AAL but is covered by a flood map. Our empirical approach essentially compares the lending to red properties with gray properties. We ignore all properties that are blue or purple. After all, we wish to disentangle the effects of flood risk from the effects of insurance mandates.

4 Methodology

Given the granularity and detail of our data, we can afford to employ richly saturated regressions. In our main analysis, we focus on three outcomes of interest: mortgage originations, interest rates, and lenders' internal property valuations. Lenders concerned about the flood risk of a property might respond by credit rationing. Moreover, lenders may price any flood risk into the property values they use to make lending decisions. Although this is not typically an observable market outcome, lenders' internal property valuations may reveal information about their awareness of flood risk apart from how they choose to respond to it. Altogether, these outcomes help us understand lenders' risk awareness as well as their actions to mitigate this risk.

Our primary approach relates these three outcome variables to borrower and region characteristics and to our variables designating un-mapped risk. As we noted above, we explicitly focus on properties that are outside of flood zones. As such, we ignore flood mapping as this is discussed in detail in Blickle and Santos (2021).

Our specification of interest takes the following form:

$$Y_{i,p,\ell,c,t} = \alpha\,\text{Un-mapped}_p + \beta\,\text{Possibly Un-mapped}_p + \gamma X_{i,t} + \eta_c + \nu_\ell + \omega_t + \varepsilon_{i,p,\ell,c,t} \tag{1}$$

where we relate the outcome variable for a given loan to individual i, for property p, from lender ℓ in census tract c at quarter t to whether the property is inaccurately flood mapped.

As previously indicated, we consider three possible outcomes Y. First, we study loan origination decisions, where our dependent variable is a dummy variable indicating when a mortgage is originated. Second, we consider the interest rate on loans approved by the lender. Third, we consider the property value assigned by the home appraiser that the lender uses in the lending decision.

We include the two variables of flood map inaccuracy previously discussed: (i) "Un-mapped" is a binary variable equal to one if the property in question faces high flood risk (above teh 50th pctile of all non-zero AAL properties) but has no flood map coverage, and (ii) "Possibly Un-mapped" is a binary variable equal to one if the property faces some flood risk (i.e. any non-zero AAL) but has no flood map. Every property that is "Un-mapped" is also "Possibly Un-mapped," such that in all regression results reported, the estimated effect on un-mapped properties relative to properties without flood risk, will be the sum of the coefficient estimates on the "Un-mapped" and "Possibly Un-mapped" variables. X is a vector of mortgage controls that include the primary applicant's sex, ethnicity, race, credit score, income, and the loan amount.

We consider several combinations of fixed effects, including census tract fixed effects η, lender fixed effects ν, and quarter fixed effects ω. Census tract fixed effects will absorb any variation at the tract-level, which could include how desirable the location itself is. Alternatively, we present results using county $\times$ quarter fixed effects. Given that our variables of interest vary at the individual property-level, we are able to include very granular controls. In our most restrictive specification, we compare properties bought by similar borrowers, with the same loan amount, located in the same census tract, but differenced by the degree to which the property faces un-mapped flood risk.

Following our analyses on credit rationing, we investigate whether lenders use securitization to mitigate the additional risk. We test whether lenders keep loans to risky properties on their balance sheets or whether they sell/securitize these off more aggressively. We make use of data on securitization and loan sales that accurately track initial sales in HMDA data.

Finally, in extensions, we interact our variable of interest with bank and region characteristics to determine whether different lenders or different regions are more sensitive to un-mapped flood risk. Specifically, we look at non-bank entities and applicants with higher than average income for their region. Non-bank entities have differing business models than banks. Their high propensity to securitize or sell mortgages may make them less risk averse than bank.

5 Flood Risk and Mortgage Lending

In this section, we report the results of our investigation to ascertain whether lenders factor in flood risk (outside flood zones) in their mortgage lending decisions. We begin by looking at loan origination decisions, interest rates charged, and properties' valuations. After that, we look into heterogeneous effects across lenders, tracts, and borrowers.

Table 2: *Loan Origination for Conforming Loans*

	Loan Originated			
	(1)	(2)	(3)	(4)
Un-Mapped	-0.0044***	-0.0035***	-0.0024***	-0.0008***
	(0.0003)	(0.0003)	(0.0003)	(0.0003)
Possibly Un-Mapped	-0.0051***	-0.0063***	-0.0014***	-0.0012***
	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Log Loan Amount		0.0278***	0.0329***	0.0099***
		(0.0002)	(0.0002)	(0.0002)
Mortgage Controls	✓	✓	✓	✓
Quarter-Year FEs	✓	✓		✓
County-Quarter-Year FEs			✓	
Tract FEs				✓
Lender FEs				✓
Observations	13,715,494	13,715,330	13,715,330	13,715,330
R^2	0.06050	0.06350	0.08063	0.16805
Within R^2	0.05913	0.06214	0.05496	0.03396

Note: We estimate equation 1, above. Our outcome variable is binary and tracks whether a mortgage is originated (i.e. accepted by both borrower and lender) and our variables of interest are also binary and denote whether a property is un-mapped or possibly un-mapped. Each specification contains controls for an applicant's sex, ethnicity, race, credit score, and applicant income. We add additional controls and fixed effects as specified at the end of each column. We restrict our sample to conforming loans and exclude all properties that are covered by a flood zone map. Standard errors are clustered by tract and the year-quarter of origination; *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

5.1 Loan Origination

We first analyze whether mortgages are less likely to be originated for properties that face "un-mapped" flood risk. Our sample is limited specifically to those properties that are not covered by either a FEMA 100-year or 500-year flood zone. This sample restriction allows us to identify the effect of flood risk by comparing similar properties unaffected by the NFIP insurance mandate and unaffected by the public signal contained in FEMA flood maps.

From Table 2 we can see that lenders are less likely to originate loans – all else equal – for properties that have un-mapped flood risk. In column (1), a mortgage application for a property that has no flood map despite facing any amount of flood risk (i.e. a possibly un-mapped property) has a 0.5 percentage point lower chance of being originated. The effect of being un-mapped is cumulative to the baseline effect of being possibly un-mapped, such that a property that has high flood risk without a flood map is just under 1 percentage point less likely to receive a mortgage

– all else equal.[14] This is a large effect given that only 13% of applications in the sample do not result in an origination.

As we add controls for loan size in column (2) or additional fixed effects in columns (3) and (4), the magnitude of our coefficient of interest diminishes, though remains statistically significant. Ultimately, if we include census tract fixed effects and lender fixed effects, the coefficients reflects an only 0.2 percentage point additional rejection rate in areas with un-mapped risk. While still statistically significant, the effect is much less pronounced.

It is possible that the inclusion of census tract fixed effects reduces the magnitude of our estimates because banks take a region-level approach to flood risk management. Any area with un-mapped flood risk is treated somewhat similarly, with borrowers less likely to obtain a loan in these communities. Overall these results suggest that mortgage lenders are aware of some of the risks posed by possible flooding.

5.2 Interest Rate Charged

We next focus on interest rates, employing the same approach as above to determine whether banks charge higher rates for mortgages on properties with un-mapped flood risk. If loans are originated **but** at higher rates, this may still compensate the lender or loan owner for un-mapped flood risk and reveal that the lender is managing risk through pricing.

Indeed, as we can see from Table 3, mortgages which are accepted pay higher rates if they are at risk of flooding but have no flood map. Specifically, facing un-mapped flood risk is associated with an up to 2 basis points higher interest rate, depending on specification. This increase is small, representing only a minor jump over the 3% average rate. If we include more detailed controls, such as census tract fixed effects, we can see that the effect becomes smaller and (partly) insignificant. Again, this seems to be a result of banks applying risk management practices to the tract at large. It may also reflect heterogeneous risk management behavior by various lender-types (see below).

[14] The effect is cumulative and calculated as $0.0044 + 0.0051$.

Table 3: ***Interest Rates for Conforming Loans***

	Interest Rate			
	(1)	(2)	(3)	(4)
Un-Mapped	0.0197***	0.0119***	0.0014**	0.0006
	(0.0006)	(0.0006)	(0.0006)	(0.0005)
Possibly Un-Mapped	0.0004	0.0143***	-0.0017***	0.0014***
	(0.0005)	(0.0004)	(0.0004)	(0.0004)
Log Loan Amount		-0.3066***	-0.3560***	-0.2601***
		(0.0008)	(0.0010)	(0.0009)
Mortgage Controls	✓	✓	✓	✓
Quarter-Year FEs	✓	✓		✓
County-Quarter-Year FEs			✓	
Tract FEs				✓
Lender FEs				✓
Observations	12,084,309	12,084,294	12,084,294	12,084,294
R^2	0.58861	0.62054	0.63759	0.73073
Within R^2	0.05217	0.12573	0.12703	0.08723

Note: We estimate equation 1, above. The interest rate is continuous and bounded between 2 and 8. We remove properties with extreme interest rates (above the 99th percentile or below the 1st percentile). Our variables of interest are binary and denote whether a property is un-mapped or possibly un-mapped. Each specification contains controls for an applicant's sex, ethnicity, race, credit score, and income. We add additional controls and fixed effects in columns (2)–(4). We restrict our sample to conforming loans and exclude all properties that are covered by a flood zone map. Standard errors are clustered by census tract and year-quarter and shown in parentheses; *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

5.3 Property Valuation

Lenders' decisions are not limited to accepting or denying a mortgage application and the interest rates charged for originated loans. Banks also play a critical role on properties' valuations they consider in the loan application. This is important because, as we show later in the paper, property valuations are more sensitive than house prices to changes in flood risk. Under these conditions, the valuation the bank uses in the mortgage application process is an important instrument for it to manage flood risk.

We therefore investigate the sensitivity of lenders' internal property valuations to flood risk. These valuations are typically computed by assessors that inspect the property during the final phase of the mortgage approval process. As such, the valuation is liable to include all available information on a property, including information the assessor has on flood risk.

Table 4: ***Property Valuations***

	Log Property Value		
	(1)	(2)	(3)
Un-Mapped	-0.0233***	-0.0245***	-0.0200***
	(0.0006)	(0.0004)	(0.0003)
Possibly Un-Mapped	0.0541***	-0.0169***	-0.0119***
	(0.0005)	(0.0003)	(0.0002)
Mortgage Controls	✓	✓	✓
Quarter-Year FEs	✓		✓
County-Quarter-Year FEs		✓	
Tract FEs			✓
Lender FEs			✓
Observations	13,403,882	13,403,882	13,403,882
R^2	0.40854	0.62647	0.73750
Within R^2	0.39677	0.32120	0.20251

Note: We estimate equation 1, above. Our outcome variable is the natural logarithm of the property value the lender uses when making the lending decision. Our variables of interest are binary and denote whether a property is un-mapped or possibly un-mapped. Each specification contains controls for an applicant's sex, ethnicity, race, credit score, and income. We add additional controls and fixed effects in columns (2)–(3). We restrict our sample to conforming loans and exclude all properties that are covered by a flood zone map. Standard errors are clustered by census tract and year-quarter and shown in parentheses; *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

We find in Table 4 that properties, which face un-mapped flood risk, experience an up to 3.2% reduction in value. Importantly, unlike the effects of flood risk on interest rates discussed above, the effect on house valuation persists even if we include census tract fixed effects (column (3)). This is most likely because valuations are conducted on a case-by case basis during the assessment and reflect a property's specific flood risk – rather than risk for an entire area. The assessor is able to make use of the property's true location and risk in a way that risk-management at the bank might not.

5.4 Response Heterogeneity

Lender Types

In our analyses thus far, we have been agnostic about the type of lender making the loan. We have simply controlled for lenders themselves with lender fixed effects in our most saturated specifications. However, different types of lenders may have different approaches to flood risk management. After all, local banks may have local knowledge and large banks may be too far removed to know about local flood risks. Non-bank lenders, on the other hand, tend to originate and securitize their mortgages which may also affect the way they view flood risk. We therefore separate large banks, local banks, and non-banks from all other lenders, interacting our variables of interest with a lender-type dummy to identify whether these lenders respond differently.

For our purposes, we define the set of large banks as the bank holding companies included in the Large Institution Supervision Coordinating Committee (LISCC) program and their subsidiaries.[15] We classify lenders as "local" on a county-by-county basis. Specifically, we define a lender-county pair to be local if – at some point during the sample period – 40% or more of the lender's originated mortgage volume was in the given county. Lastly, we define non-banks as any institution not classified as a large lender or local lender, and is not a bank or credit union.[16] This group will include internet mortgage brokers and their affiliated non-bank financial institutions.

In table 5, we display results when we regress the loan origination dummy on the interactions of our lender-type and flood risk dummies. These regressions show that, broadly, all lender types engage in some amount of credit rationing in response to flood risk, but the magnitudes of these sensitivities are appreciably different. Relative to regional or state banks (the omitted lender type in the table), non-banks are slightly more willing to lend in areas with flood risk. In our most saturated specification, column (4), non-banks are just 13 basis points less likely to originate an un-mapped mortgage than an otherwise identical mortgage with no flood risk. This is also true of local banks, whose originations are similarly insensitive to un-mapped properties than non-banks.

By contrast, large banks and regional banks appear particularly averse to un-mapped risk. Again in column (4), large banks and regional banks are 58 basis points and 33 basis points less likely to originate an un-mapped mortgage than an otherwise identical mortgage with no flood risk, respectively. Although the coefficient on the interaction between large banks and the un-mapped dummy is not stable across specifications, all of the specifications in table 5 demonstrate that non-banks' and local banks' originations are less sensitive to flood risk than

[15] Currently, this list of bank holding companies includes Bank of America Corporation, The Bank of New York Mellon Corporation, Citigroup Inc., The Goldman Sachs Group, Inc., JP Morgan Chase & Co., Morgan Stanley, State Street Corporation, and Wells Fargo & Company.

[16] We identify banks and credit unions based on the entity types corresponding with their RSSD in HMDA. Specifically, we consider national banks, state member banks, cooperative banks, domestic branches of domestic banks, non-member banks, savings and loans, federal savings banks, state savings banks, uninsured branches of foreign bank offices, federal credit unions, and state credit unions.

large banks' and regional banks' originations.

Table 5: ***Loan Originated with Bank Type Interactions***

	Loan Originated			
	(1)	(2)	(3)	(4)
Un-mapped × Non-Bank	0.0078***	0.0077***	0.0072***	0.0020***
	(0.0006)	(0.0006)	(0.0006)	(0.0006)
Un-mapped × Local Bank	0.0138***	0.0125***	0.0163***	0.0027**
	(0.0013)	(0.0013)	(0.0013)	(0.0012)
Un-mapped × Large Bank	-0.0002	-0.0005	0.0035***	-0.0025*
	(0.0013)	(0.0013)	(0.0013)	(0.0013)
Un-mapped	-0.0103***	-0.0092***	-0.0078***	-0.0021***
	(0.0005)	(0.0005)	(0.0005)	(0.0005)
Possibly Un-mapped	-0.0059***	-0.0069***	-0.0013***	-0.0012***
	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Log Loan Amount		0.0253***	0.0308***	0.0099***
		(0.0002)	(0.0002)	(0.0002)
Mortgage Controls	✓	✓	✓	✓
Quarter-Year FEs	✓	✓		✓
County-Quarter-Year FEs			✓	
Tract FEs				✓
Lender FEs				✓
Observations	13,715,494	13,715,330	13,715,330	13,715,330
R^2	0.06687	0.06932	0.08636	0.16805
Within R^2	0.06551	0.06796	0.06085	0.03397

Note: We estimate equation 1, above. Our outcome variable is binary and tracks whether a mortgage is originated (i.e. accepted by both borrower and lender) and our variables of interest are also binary and denote whether a property is un-mapped or possibly un-mapped. Non-bank is a dummy variable indicating that, based on its entity type code, the lender is not a bank nor credit union. Local bank is a dummy variable indicating that at least 40% of the lender's mortgage origination occurs in the county the property is located in. Large bank is a dummy variable indicating that the lender is an entity or a subsidiary of an entity regulated under the Large Institution Supervision Coordinating Committee (LISCC) program. Each specification contains controls for an applicant's sex, ethnicity, race, credit score, and income. We add additional controls and fixed effects in columns (2)–(4). We restrict our sample to conforming loans and exclude all properties that are covered by a flood zone map. Standard errors are clustered by tract and the year-quarter or origination; *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

Apart from originations, Appendix A.6 contains additional results from similar exercises with interest rates and property values. Of the four lender categories we use, large banks' interest rates appear least sensitive to flood risk. As before though, the interest rate effects are modest and average effects for each of the lender types never differs by more than three basis points. There are more meaningful heterogeneities in the sensitivity of lenders' internal property values to flood risk. In our most saturated specification, large banks' property values are most sensitive to flood risk, and non-banks' property values are the least sensitive to flood risk. Relative to an otherwise identical mortgage with no flood risk, large banks' and non-banks' internal property value is 4.36% and 2.95% lower, respectively.

The results above raise the question of whether entities such as non-banks or even local banks are less aware of flood risk or manage their exposure to this risk differently than large banks and regional banks. We attempt to shade some light on these questions in the next section by investigating lenders' decisions to sell and securitize the mortgages their originate.

Track Characteristics and Borrower Types

Finally, we look at whether tract or borrower characteristics can impact the effect un-mapped flood risk has on a lender's decision. In Table 6 we show that mortgages for un-mapped properties are more likely to be originated if a given census tract falls into the top quartile of median family incomes across the country. We find that lenders are slightly less risk averse in high income areas, possibly due to the fact that high income borrowers can more easily weather the negative impact of a flooding disaster. These results are corroborated if we look at high credit score borrowers or borrowers that have higher incomes than the average of their county (not reported for brevity).

In Appendix A.7 we also discuss the lending responses when granting jumbo loans. We find that the lending aversion is much smaller for these larger loans. This may similarly be a side-effect of lenders being more comfortable with wealthier borrowers, whose loan to income ratios are, on average, much higher.

Summing up, the evidence above shows that mortgage lenders are aware of flood risk beyond the boundaries of FEMA's flood maps. This is apparent in their lending decisions and the interest rates they charge as well as the property valuations they consider in mortgage applications. In the next section, we investigate the implications of lenders' response to flood risk for mortgage and housing markets.

6 Implications for Mortgage and Housing Markets

In this section, we consider the implications of lenders' response to flood risk for mortgage and housing markets. Our earlier results show that while all types of lenders are aware of flood risk in un-mapped zones, local banks' and in particular non-bank lenders', have largely continued lending in un-mapped areas, raising two important questions about mortgage markets. First, do these lenders use the secondary mortgage market more aggressively to lay out their mortgage

Table 6: ***Origination with Tract Income***

	Loan Originated			
	(1)	(2)	(3)	(4)
Un-mapped × High Income Tract	-0.0011*	-0.0004	0.0040***	0.0012**
	(0.0006)	(0.0006)	(0.0006)	(0.0005)
Un-mapped × Low Income Tract	0.0056***	0.0050***	-0.0001	0.0010
	(0.0009)	(0.0009)	(0.0009)	(0.0009)
Un-mapped	-0.0046***	-0.0041***	-0.0037***	-0.0014***
	(0.0004)	(0.0004)	(0.0004)	(0.0004)
Possibly Un-mapped	-0.0049***	-0.0060***	-0.0013***	-0.0012***
	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Low Income Tract	-0.0321***	-0.0267***	-0.0134***	
	(0.0004)	(0.0004)	(0.0004)	
High Income Tract	0.0006**	-0.0062***	-0.0086***	
	(0.0002)	(0.0003)	(0.0003)	
Log Loan Amount		0.0269***	0.0328***	0.0099***
		(0.0002)	(0.0002)	(0.0002)
Mortgage Controls	✓	✓	✓	✓
Quarter-Year FEs	✓	✓		✓
County-Quarter-Year FEs			✓	
Tract FEs				✓
Lender FEs				✓
Observations	13,714,718	13,714,554	13,714,554	13,714,554
R^2	0.06144	0.06412	0.08083	0.16802
Within R^2	0.06007	0.06276	0.05516	0.03396

Note: We estimate equation 1, above. Our outcome variable is binary and tracks whether a mortgage is originated (i.e. accepted by both borrower and lender) and our variables of interest are also binary and denote whether a property is un-mapped or possibly un-mapped. Low income tract and high income tract are dummy variables that denote if a tract is in the bottom and top quartile of the tract-level median family income distribution. We add additional controls and fixed effects in columns (2)–(4). We restrict our sample to conforming loans and exclude all properties that are covered by a flood zone map. Standard errors are clustered by tract and the year-quarter or origination; *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

risks? Second, do they capitalize on their competitors', in particular large banks', conservative lending practices to increase their market shares in the primary mortgage market? More generally, do lenders' responses to un-mapped flood risk affect property values in these housing markets. We investigate these three questions next.

6.1 Lenders' Securitization and Loan Sales

The evidence reported in the previous section suggests lenders – particularly banks – attempt to mitigate un-mapped flood risk in their mortgage lending decisions. A complementary way mortgage lenders can use to manage that risk in the case of conforming mortgages, which underlie our evidence, is through the securitization and sale of mortgages in the secondary market.

HMDA data includes information on whether mortgages are securitized or sold.[17] For convenience, we have initially grouped loan sales and loan securitization together under the same umbrella. This is useful partly because some non-bank entities first "sell" a loan to a sister organization that prepares the loan for securitization and partly because we can be agnostic about the difference between sales and securitization, as either could be a risk mitigant from the perspective of the lender.

As can be seen from Table 7, we find that lenders are slightly – up to 1.2 percentage points – more likely to sell or securitize mortgages of properties with un-mapped risk. The economic magnitude of the effect is smallest when we include tract fixed effects, implying that risk management may be conducted at the level of communities (as seen above). Moreover, only the effect of being un-mapped remains, with the coefficient on "possibly un-mapped" losing significance.

We find the same patterns if we look only at securitization (see Appendix A.3). Overall, our results suggest that lenders use the secondary mortgage market to move flood risk from their balance sheet. Given the high propensity for securitizing conforming loans – i.e. 73% –, even a small increase in our measure of sold or securitized loans indicates that lenders are indeed cautious about holding such risks and are aggressive about moving these risks from their balance sheets.

Additionally, we can test for heterogeneity in secondary mortgage market behavior by lender

[17] Banks use securitization to diversify risk and often to increase their lending (Cebenoyan and Strahan (2004); Franke et al. (2022); Carbo-Valverde et al. (2015).

Table 7: *Loan Sales*

	Loan Sold or Securitized			
	(1)	(2)	(3)	(4)
Un-Mapped	0.0032***	0.0053***	0.0012***	0.0009***
	(0.0004)	(0.0003)	(0.0003)	(0.0003)
Possibly Un-Mapped	0.0087***	0.0055***	0.0013***	0.0002
	(0.0003)	(0.0003)	(0.0003)	(0.0002)
Interest Rate	-0.0941***	-0.0910***	-0.0881***	-0.0513***
	(0.0003)	(0.0002)	(0.0002)	(0.0002)
Log Loan Amount		0.0754***	0.0727***	0.0424***
		(0.0002)	(0.0003)	(0.0002)
Mortgage Controls	✓	✓	✓	✓
Quarter-Year FEs	✓	✓		✓
County-Quarter-Year FEs			✓	
Tract FEs				✓
Lender FEs				✓
Observations	11,922,096	11,922,081	11,922,081	11,922,081
R^2	0.06294	0.07817	0.11164	0.41972
Within R^2	0.04144	0.05702	0.05236	0.01798

Note: We estimate equation 1, above. The outcome variable is binary and indicates when a loan is sold, including when sold a securitization agency. Our variables of interest are binary and denote whether a property is un-mapped or possibly un-mapped. Each specification contains controls for an applicant's sex, ethnicity, race, credit score, and income. We add additional controls and fixed effects in columns (2)–(4). We restrict our sample to conforming loans and exclude all properties that are covered by a flood zone map. Standard errors are clustered by census tract and year-quarter and shown in parentheses; *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

type, as in section 5.4. This is important because as we documented earlier non-bank lenders, and to less extent local banks, appear to be less concerned with un-mapped flood risk in their mortgage origination decisions. In Table 8 we see that non-banks and local banks are indeed more aggressive in removing this risk from their books, relative to large banks and other regional banks. Non-banks in particular seem aware of property-level risks. From column (4) we can see that these lenders are aggressively securitizing/selling individual properties within a census tract that bear flood risk. Local banks, on the other hand, are very aggressive in securitizing mortgages to properties in risky regions as a whole, making fewer property-level distinctions (column (1) vs column (4)). Once we include a census tract fixed effect, we can see that the coefficient on the interaction term becomes insignificant. Ultimately, non-banks and local banks are so aggressive in their attempts to move un-mapped risk from their balance sheets, that the baseline coefficient changes sign, implying that these lenders are the drivers of the baseline-effect discussed above. Large banks are the least aggressive in securitizing loans with un-mapped risk.

Table 8: ***Loan Sales with Bank-Type Interactions***

	Loan Sold or Securitized			
	(1)	(2)	(3)	(4)
Un-mapped × Non-Bank	0.0062***	0.0059***	0.0045***	0.0031***
	(0.0008)	(0.0008)	(0.0008)	(0.0007)
Un-mapped × Local Bank	0.0218***	0.0193***	0.0248***	-0.0004
	(0.0022)	(0.0022)	(0.0021)	(0.0016)
Un-mapped × Large Bank	-0.0161***	-0.0166***	-0.0111***	-0.0020
	(0.0019)	(0.0019)	(0.0019)	(0.0016)
Un-mapped	-0.0016**	0.0003	-0.0020***	-0.0011*
	(0.0008)	(0.0008)	(0.0007)	(0.0006)
Possibly Un-mapped	0.0025***	0.0006**	0.0012***	0.0002
	(0.0003)	(0.0003)	(0.0003)	(0.0002)
Interest Rate	-0.0948***	-0.0926***	-0.0896***	-0.0513***
	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Log Loan Amount		0.0529***	0.0600***	0.0424***
		(0.0002)	(0.0002)	(0.0002)
Mortgage Controls	✓	✓	✓	✓
Quarter-Year FEs	✓	✓		✓
County-Quarter-Year FEs			✓	
Tract FEs				✓
Lender FEs				✓
Observations	11,922,096	11,922,081	11,922,081	11,922,081
R^2	0.16711	0.17448	0.20311	0.41972
Within R^2	0.14800	0.15554	0.14994	0.01799

Note: We estimate equation 1, above. The outcome variable is binary and indicates when a loan is sold, including when sales to securitization agencies. Our variables of interest are binary and denote whether a property is un-mapped or possibly un-mapped. Non-bank is a dummy variable indicating that, based on its entity type code, the lender is not a bank nor credit union. Local bank is a dummy variable indicating that at least 40% of the lender's mortgage origination occurs in the county the property is located in. Large bank is a dummy variable indicating that the lender is an entity or a subsidiary of an entity regulated under the Large Institution Supervision Coordinating Committee (LISCC) program. Each specification contains controls for an applicant's sex, ethnicity, race, credit score, and income. We add additional controls and fixed effects in columns (2)–(4). We restrict our sample to conforming loans and exclude all properties that are covered by a flood zone map. Standard errors are clustered by census tract and year-quarter and shown in parentheses; *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

Loan Destination and Risk

The evidence we presented above shows that lenders are more likely to sell or securitize mortgages for un-mapped properties than properties without flood risk, even conditional on a detailed set of mortgage, region, and time controls. Although the effects we identify are statistically significant, it is unclear whether they currently translate into appreciable changes in the distribution of flood risk across financial institutions. This is an important question given that conforming mortgages are most often securitized with a public agency.

To get some insight on who bears the flood risk in mortgages to un-mapped zones, we display in Table A.3 the distribution of loan values across flood regions and destinations for the conforming mortgages in our 2018-2021 sample. The table shows that 52% of mortgage value originated for un-mapped properties is subsequently sold to a public securitization agency. Despite the results of Table 7, this is not appreciably different than for mortgages without any flood risk: 52% of mortgage value originated for properties with low-to-no flood risk is also subsequently sold to a public securitization agency. Therefore, we do not see evidence of a large-scale redistribution of flood risk onto the GSEs. Nonetheless, GSEs already hold the majority of mortgages with notable flood risk. GSEs may even increase their exposure to flood risk if the past lending dynamics in the primary market continue into the future.

As we document in the next section, an implication of banks, in particular large banks, cutting back on their mortgage lending in un-mapped regions has been the rise of non-banks' market shares in these regions. If banks continue to cede market share to non-banks then non-banks' high propensity to securitize will mechanically increase the share of high-risk mortgages in these areas that end up with one of the GSEs. In Table A.4 we see that banks sell or securitize 63% of mortgage value for un-mapped properties (65% for low-to-no risk properties), while non-banks sell or securitize 93% of mortgage value for un-mapped properties (92% for low-to-no risk properties).

A second dimension to consider is the interest rate on mortgages for un-mapped properties and the destination of these mortgages after origination. Table A.5 displays the mean interest rates on loans, split again by the flood risk region and the destination of the loan after origination. Across all regions, the interest rate is lowest for those loans sold to the GSEs. Loans sold to GSEs pay 0.27% less than loans retained by the issuer if they are in an un-mapped region. This of course may be related to other loan characteristics that differ systematically between those loans securitized with a GSE and those kept by the lender. We explore this more formally in Table A.8, by incorporating borrower, region, and time controls as before. Even when we condition on these other factors that may affect the pricing, we see that un-mapped loans securitized to GSEs pay around a basis point less than an otherwise identical loan the lender subsequently keeps. The effect holds when we do not include tract fixed effects, again suggesting a regional approach to risk management.

FICO scores follow a similar pattern to the rates discussed above. The highest scores are seen on borrowers whose loans are kept by the originating institution. Some of the lowest scores are again seen on loans sold to GSEs. This can be seen in Table A.6.

Overall, the results reported above add further support to our assertion that lenders are aware of the risks posed by un-mapped flood zones. They also show a potentially important implication of that awareness for GSEs. While larger banks are more likely to restrict lending (i.e. originate fewer loans) and charge higher rates to borrowers in un-mapped regions, small local banks and non-bank entities are more prone to lend in these regions and securitize the mortgages. The GSEs are the primary vector for lenders to offload this risk, meaning they have absorbed a large amount over the past few years alone.

6.2 Lenders' Market Share

Over the last decade, non-banks have increasingly gained market share in mortgage originations, going from 47.8% in 2012 to 61.6% in 2023. Researchers often point to their use of new borrower screening technologies and lower regulatory burden for this dynamic. Our evidence suggests another important driver: banks' conservative lending policies in un-mapped regions.

The results we reported in the previous section suggest that large banks and regional banks are less lenient when it comes to mortgage originations for properties exposed to un-mapped flood risk. Given this heterogeneity in lenders' response, it stands to reason that over time, the composition of the primary mortgage market in risky areas may shift towards the lenders that are least sensitive to this risk in the origination phase. Although the property-level dataset we use earlier enables a granular analysis of mortgage lending, it only covers the four year period

Figure 3: *Changes in the Primary Mortgage Market Share Against Flood Risk Rankings*

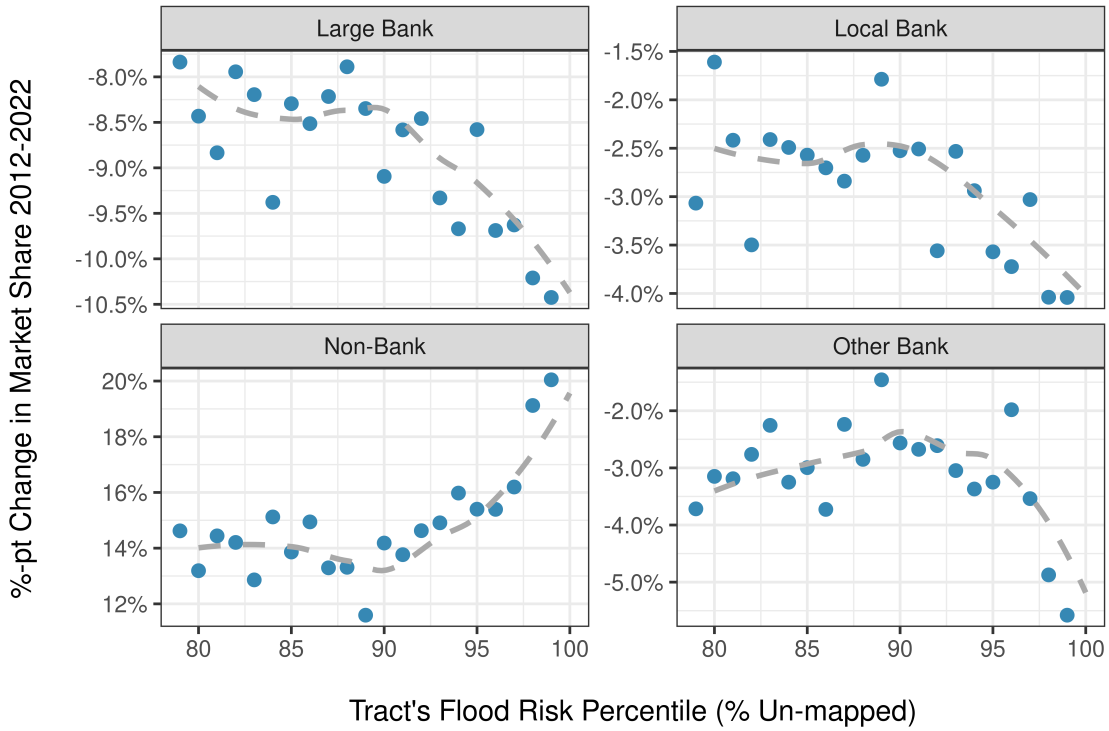

Note: This figures plots the percentage point change in each lender types' market share between 2012 and 2022 against the percentile of the tracts' flood risk. We rank tracts based on the percent of properties in the tract that are un-mapped.

from 2018-2021. To study differential changes in lender composition for areas with high flood risk, we again use HMDA data, but over a longer period from 2012-2022. We do not have this longer series of mortgage application data linked to property-level flood risk measures, so we instead aggregate these data into a census tract-year panel and match these with static tract-level measures of flood risk. In each tract-year, we compute the share of mortgage loan volume for large banks, local banks, other regional banks, and non-banks, using the same lender designations as in Table 5.[18]

For each tract, we compute the percent of properties with un-mapped flood risk, and then rank tracts by these shares such that tracts in the 99th percentile of the distribution have the highest share of un-mapped properties. For the tracts in each percentile bin, we then compute the

[18]Lenders are defined as "local" to a county if ≥40% of the lenders' originated mortgage value is located in the county. In Table 5 and elsewhere in the text, the originated mortgage volume we use to make this determination for each lender-county pair is based on originations observed in the main sample from 2018-2021. Here, we are interested in how local market shares change over time, so we opt instead to use the ex ante originated mortgage volume—from 2010 HMDA data—when designating lender-county pairs as "local".

Table 9: *Differential Shifts in Lender-Type Market Shares*

	YoY Basis Point Change in Market Share			
	Large Bank	Local Bank	Non-Bank	Other Bank
	(1)	(2)	(3)	(4)
$\mathbb{1}$(99th Un-mapped Risk Percentile)	-26.10***	-8.214	67.09***	-32.78**
	(10.01)	(7.924)	(21.71)	(16.13)
$\mathbb{1}$(95-99th Un-mapped Risk Percentile)	-7.981*	4.765	19.17***	-15.96**
	(4.551)	(3.278)	(6.773)	(7.324)
$\mathbb{1}$(90-95th Un-mapped Risk Percentile)	2.881	1.809	10.68**	-15.37***
	(2.982)	(2.133)	(5.307)	(4.708)
$\mathbb{1}$(80-90th Un-mapped Risk Percentile)	0.7655	-0.0144	1.841	-2.592
	(1.703)	(2.308)	(3.582)	(3.108)
Tract Controls	✓	✓	✓	✓
County FEs	✓	✓	✓	✓
Year FEs	✓	✓	✓	✓
Observations	742,781	742,781	742,781	742,781
R^2	0.00837	0.00968	0.02629	0.01420
Within R^2	0.00045	0.00052	0.00374	0.00236

Note: This table relates year-over-year basis point changes in the lenders' market share to tract-level flood risk rankings. The dataset is an unbalanced tract-year panel with the share of originated mortgage loan volume in the given tract-year for large banks, local banks, other regional banks, and non-banks. Year-over-year basis point changes in these market shares are computed for the years 2013-2022. In each specification we add tract-year controls, including the mean loan-to-income ratio, mean log applicant income, percent of male mortgage applicants, percent of Native American, Asian, Black, Pacific Islander, and white applicants, as well as a time-invariant count of the number of properties in the tract constructed from the CoreLogic data.

percentage point change in the market share of large banks, local banks, other regional banks, and non-banks from 2012 to 2022. Figure 3 plots this relationship between tract-level flood risk and changes in each lender type's market share in the period. Banks have lost market share to non-banks broadly, but Figure 3 shows that banks have lost even more of their market share to non-banks in the areas with the highest flood risk. Non-banks gained around 14 percentage points of market share from 2012-2022 for areas in 80-90th percentiles of the tract-level flood risk distribution but have gained closer to 20 percentage points of market share in the top percentile of the tract-level flood risk distribution.

Table 9 explores this more formally. Here, we relate year-over-year changes in lenders' tract-level market share to flood risk. In Figure 3, we see lender composition only appears to vary with flood risk at the right tail of the flood risk distribution. Motivated by this, we choose to form tract-level flood risk categories with greater granularity at the right end of the flood risk distribution. We again rank tracts by their share of un-mapped properties and use these rankings to designate tracts as being in the bottom 80 percent, the 80-89th percentile, 90-94th percentile, 95-98th percentile, or 99th percentile of the flood risk distribution. In Table 9, we regress the year-over-year changes in tract-level market share on these flood risk category dummies, incorporating tract-level controls, county fixed effects, and year fixed effects.

For tracts in the highest percentile of the flood risk distribution, large banks lost an additional 26 basis points of market share per year on average, relative to tracts in the bottom eighty percent of the flood risk distribution. Similarly, other regional banks lost on average an additional 33 basis points of market share a year in the top percentile of risk relative to tracts in the bottom eighty percent of the flood risk distribution. For both large banks and other regional banks, we see that the magnitude of this effect is monotonically increasing in the flood risk. By contrast, non-banks have gained on average an additional 67 basis points of market share a year across tracts in the top percentile of the flood risk distribution, relative to tracts in the bottom eighty percent of the flood risk distribution. This effect is also monotonically increasing in flood risk. Although Figure 3 seems to suggest that even local banks have differentially lost market share in the riskiest areas, we see here that this is not clear when we incorporate tract-level controls and county fixed effects. The coefficients in column (2) are insignificant. This corroborates earlier results in table 5 that showed a smaller reduction in lending by local banks, relative to other banks, in areas affected by flood risk.

Given the earlier evidence that non-banks sell and securitize mortgages at much higher rates than banks, the heightened growth of non-banks' market shares in predominantly un-mapped areas also suggests that sale and securitization rates in these areas may also be growing at faster rates than in areas of low flood risk. Table A.10 investigates this corollary using the same regression specification as in Table 9, but where the dependent variable is the year-over-year change in the sales rate in column (1) and the year-over-year change in the securitization rate in column (2). Consistent with Table 9, the coefficient estimates show that the average growth in the sale and securitization rates is monotonically increasing in a tract's flood risk. Tracts in the 99th percentile of the un-mapped flood risk distribution experience an additional 48 basis points of annual growth in the sales rate, relative to tracts in the bottom eighty percent of the un-mapped flood risk distribution. This evidence is consistent with the story that the differential shift in mortgage lender composition extends beyond the primary mortgage market, and has broader implications for the distribution of flood risk in even the secondary mortgage market.

Overall, our evidence shows an important change in the primary mortgage market composition in high-flood risk areas, one that is rooted in a key difference in lenders' risk management policies and which helps us understand the rise of non-banks in mortgage markets. Large banks' and other regional banks' conservative lending policies in high-flood risk areas have created a growth opportunity for non-banks in these areas. Non-banks have taken on this opportunity, building on their originate-to-distribute model and manage their exposure to flood risk by offloading the riskiest mortgages. In the process, they gained market share in high-flood risk areas and further increased their market share the market for mortgage originations. Correspondingly, the share of mortgages sold or securitized has differentially increased in areas exposed to high flood risk.

6.3 Housing Market

Our evidence on lenders' responses to un-mapped flood risk, in particular their more conservative lending in high-flood risk areas, leads to the natural question as to whether they affect property values in these housing markets. We begin our investigation of this question by studying whether housing prices reflect flood risks.

Using house listings data from CoreLogic, we relate the closing prices of transacted properties to the flood risk faced by these properties, controlling for a series of property-level and neighborhood-level characteristics. In our analysis, we are able to make use of property-level pricing data and property-level flood risk data. Since we do not need to match this with mortgage records, we are less concerned with applicant privacy and do not need to group risk into buckets. We again exclude all properties that are covered by a FEMA 100-year or 500-year flood map and focus instead on properties that have un-mapped flood risk.

The results of our investigation are reported in Table 10. We find that flood risk is negatively associated with closing house prices. In columns (1)-(3), we relate log closing prices with the average annual loss (AAL) flood risk score—a measure from CoreLogic bounded between 1 and 100. When we control for housing and tract characteristics, a one unit increase in the AAL flood

Table 10: ***Property Prices***

	Log Closing Price				
	(1)	(2)	(3)	(4)	(5)
AAL Flood Risk Score	-0.0014***	-0.0002***	-0.0006***		
	(1.82×10^{-5})	(1.46×10^{-5})	(1.15×10^{-5})		
Un-mapped				-0.0103***	-0.0213***
				(0.0008)	(0.0007)
Possibly Un-mapped				-0.0046***	-0.0099***
				(0.0005)	(0.0004)
House Controls		✓	✓	✓	✓
Tract Controls		✓		✓	
County × Quarter FEs	✓	✓	✓	✓	✓
Tract FEs			✓		✓
Observations	22,806,226	20,945,629	20,945,629	20,945,629	20,945,629
R^2	0.29054	0.46970	0.54073	0.46972	0.54071
Within R^2	0.00109	0.25863	0.17273	0.25865	0.17269

Note: We relate the natural log of closing prices on house purchases to flood risk, as measured by CoreLogic's average annual loss risk score in columns (1)-(3) and by our own "un-mapped" property designations in columns (4)-(5). We exclude properties covered by a FEMA flood map. House controls include the square feet of living area, the number of bedrooms, the number of bathrooms, the age of the house, whether the house is listed as "waterfront", and whether the house is sold as a foreclosure. Tract controls include the census tracts' median household income and whether or not the census tract is adjacent to the coast. Standard errors are clustered by census tract and year-quarter and shown in parentheses; *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

risk score corresponds with a reduction in the closing price of 0.02-0.06%. In columns (4) and (5), we make use of the same flood risk categories in the earlier analysis, designating properties as un-mapped or possibly un-mapped based on their risk exposure. Again, we add a series of property-level and neighborhood-level controls, including tract fixed-effects in column (5). Here we see that un-mapped properties close for approximately 3.1% less than otherwise comparable properties within the same tract.[19]

It should be noted that in Table 10, we cannot isolate if the sensitivity of housing prices to flood risk is a result of lenders' hesitancy to originate mortgages for these properties or is a result of home buyers' own risk sensitivity. All else equal, we would expect credit constraints and reduced lending to these properties to reduce prices to some extent. However, lenders and home owners may both be cautious, causing home buyers to bid less on homes in risky regions. To ascertain if

[19] As before, the un-mapped and possibly un-mapped designations are not mutually exclusive, such that the full coefficient on un-mapped properties is $(-0.0213) + (-0.0099)$.

mortgage lenders have any role disciplining the housing market with respect to flood risk, we look to see if housing prices and lenders' own property valuations are differentially sensitive to flood risk. If prices in the housing market are more sensitive to flood risk than mortgage lenders' property valuations, this would be suggestive that the sensitivity of housing prices to flood risk is indeed primarily related to home buyers' own risk sensitivity. On the other hand, if mortgage lenders' property valuations are more sensitive to floor risk than housing prices, this would offer additional evidence that mortgage lenders play a role in the housing markets' response to flood risk.

Table 11 explores this issue using our tract-quarter panel of mortgage originations in HMDA and house transactions from CoreLogic, introduced in section 2.4. We find that the decrease in valuation – conducted by the bank and discussed above – is larger than the decrease in prices, suggesting that banks are affecting housing prices. From Table 11 we can first see that the average value assigned to properties is strongly linked to the closing price (column (1)). However, the closing price of properties is negatively related to the share with un-mapped risk (column (2)). In fact, the higher the price of a property, the larger the deviation between value and price becomes in areas with un-mapped risk. This likely reflects a bank's aversion to being overly exposed to risky properties at high values.

As an alternative way to identify this link we look at the down payments of borrowers. If borrowers must post more equity in the property at the time of the purchase, this might also be suggestive of a lender-driven effect on housing prices. We compute a ratio of property sales prices to loan amounts at the tract level. A large deviation between price and loan amount implies that households must make use of more equity in the house purchase.[20] In column (3), we see the ratio of closing prices to loan amount grows by 1.6 percentage points in areas with low coverage.[21] This implies that – all else equal – the amount of equity households have to post for these types of buildings is higher than in comparable areas with no flood risk.

Taken together, the results in this subsection suggest banks' awareness of the flood risk in un-mapped regions has implications for the housing market. The fact that sale prices of properties

[20] As described above, we are unable to match property sales prices to mortgage data at the household level, given data concerns. Instead, we must make use of data at the tract*quarter level.

[21] Given that we have collapsed data to the tract level, "un-mapped" is no longer subsumed by "possibly un-mapped" and the effects are not cumulative.

Table 11: ***Property Value and Loan to Value***

	Log Mean Property Value		Closing Price to Loan Amount ×100
	(1)	(2)	(3)
Log Mean Closing Price	0.6301***	0.6378***	
	(0.0025)	(0.0023)	
Log Mean Closing Price × $\mathbb{1}$(Un-mapped)		-0.0703***	
		(0.0069)	
Log Mean Closing Price × $\mathbb{1}$(Possibly Un-mapped)		-0.0067**	
		(0.0031)	
Un-mapped Dummy		0.3891***	1.599***
		(0.0407)	(0.3061)
Possibly Un-mapped Dummy		0.0372**	-0.6267***
		(0.0179)	(0.2425)
County FEs	✓	✓	✓
Quarter-Year FEs	✓	✓	✓
Observations	790,681	783,942	784,367
R^2	0.84002	0.83979	0.01252
Within R^2	0.54599	0.54688	5.59×10^{-6}

Note: We construct tract-quarter dummies such that the un-mapped dummy is 1 if more than 50% of the HMDA properties in the given tract-quarter are un-mapped, and the analogous dummy variable for the possibly un-mapped group. Heteroskedasticity-robust standard-errors in parentheses; *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

in un-mapped regions are lower than comparable properties in the same tract that face no such risk could be in part the result of credit rationing in the mortgage market. The fact that the amount of equity households have to post for properties in un-mapped regions is higher than in comparable areas with no flood risk is strong evidence of that credit rationing.[22]

7 Conclusion

We make use of property-level mortgage data, property-level risk data, and country-wide FEMA flood maps to identify the effects of flood risk on mortgage lending. We focus on those properties that face flood risk but are not zoned as being in a FEMA flood zone, either the 100- or 500-year area. As such, we are abstracting from the effects of mandatory flood insurance and any other information coming from flood maps.

[22] Data for LTV calculations is at the tract-time level, meaning we cannot account for tract fixed effects in these regressions (as discussed above).

Our findings can be grouped into two broad – and novel – ideas. First, we find that lenders are aware of flood risk outside of mapped regions. Specifically, they are less likely to originate loans to un-mapped properties subject to flood risk while charging higher rates for originated loans.

Secondly, we find that there is some degree of heterogeneity between various lender types when it comes to the management of flood risks. While large banks and regional banks respond to flood risks by adopting conservative lending policies, non-banks and small local banks continue to lend in high-flood risk areas but securitizing and selling off loans with risk more aggressively. Over the 2018-2021 time period, of the roughly $5.1 trillion in mortgages originated nationwide, nearly $794 billion (16%) were originated in un-mapped regions. Small local banks and non-banks were responsible for originating 61% (3% and 58% local banks and non-banks respectively), the equivalent of about $487 billion.

We unveil some aggregate implications from our findings on banks' responses to flood risk. Larger banks' conservative response to flood risks has paved the way for non-banks to increase their market share of mortgage originations in high-flood risk areas, building on their extensive reliance on securitization. This has contributed to making the GSEs the primary recipient of loans with greater flood risk. They are especially likely to receive those loans that pay a lower rate or have lower FICO borrowers – particularly when controlling for risk. Finally, we find some suggestive evidence that banks' conservative lending policies have had a negative effect in house prices in high-flood risk areas.

References

Blickle, K. and J. Santos (2021). The unintended consequences of "mandatory" flood insurance. *NYFed Working Paper series*.

Buchak, G., G. Matvos, T. Piskorski, and A. Seru (2018). Fintech, regulatory arbitrage, and the rise of shadow banks. *Journal of financial economics 130*(3), 453–483.

Carbo-Valverde, S., H. Degryse, and F. Rodríguez-Fernández (2015). The impact of securitization on credit rationing: Empirical evidence. *Journal of Financial Stability 20*.

Cebenoyan, A. and P. Strahan (2004). Risk management, capital structure and lending at banks. *Journal of Banking & Finance 28*.

Chavaz, M. (2016). Dis-integrating credit markets: Diversification, securitization, and lending in a recovery. *Bank of England Working Paper No. 617*.

Choi, D. B., H.-s. Choi, and J.-E. Kim (2022). Clogged intermediation: Were home buyers crowded out? *Journal of Money, Credit and Banking 54*(4), 1065–1098.

Cohen, J., J. Barr, and E. Kim (2021). Storm surges, informational shocks, and the price of urban real estate: An application to the case of hurricane sandy. *Regional Science and Urban Economics 90*.

Cortes, K. (2015). Rebuilding after disaster strikes: How local lenders aid in the recovery. *FRB of Cleveland Working Paper No. 14-28*.

Cortes, K. and P. Strahan (2017). Tracing out capital flows: How financially integrated banks respond to natural disasters. *Journal of Financial Economics 125*.

Dinan, T., P. Beider, and D. Wylie (2019). The national flood insurance program: Is it financially sound? *Risk Management and Insurance Review*.

Franke, G., J. Krahnen, and P. Strahan (2022). Default risk sharing between banks and markets: The contribution of collateralized debt obligations. *NBER Working Paper No. w11741*.

Fuster, A., M. Plosser, P. Schnabl, and J. Vickery (2019). The role of technology in mortgage lending. *The Review of Financial Studies 32*(5), 1854–1899.

Garbarino, N. and B. Guin (2021). High water, no marks? biased lending after extreme weather. *Journal of Financial Stability 54*, 2–13.

Gibson, M. and J. Mullins (2020). Climate risk and beliefs in new york floodplains. *Journal of the Association of Environmental and Resource Economists 7*.

Gourevitch, J. D., C. Kousky, Y. Liao, C. Nolte, A. B. Pollack, J. R. Porter, and J. A. Weill (2023). Unpriced climate risk and the potential consequences of overvaluation in us housing markets. *Nature Climate Change 13*(3), 250–257.

Hino, M. and M. Burke (2021). The effect of information about climate risk on property values. *Proceedings of the National Academy of Sciences 118*(17), e2003374118.

Ivanov, I., M. S. Kruttli, and S. W. Watugala (2022). Banking on carbon: Corporate lending and cap-and-trade policy. *Working Paper*.

Ivanov, I., M. Macchiavelli, and J. Santos (2022). Bank lending networks and the propagation of natural disasters. *Financial Management*, 1–25.

Jung, H., J. Santos, and L. Seltzer (2022). U.s. banks' exposures to transition risks. *Federal Reserve Bank of New York Staff Report*.

Keenan, J. and J. Bradt (2020). Underwater: From theory to empiricism in regional mortgage markets in the u.s. *Climate Change 162*, 2043–2067.

Keys, B. and P. Mulder (2020). Neglected no more: Housing markets, mortgage lending, and sea level rise. *Technical report, NBER*.

Kousky, C. (2010). Learning from extreme events: Risk perceptions after the flood. *Land Economics 86*.

Kousky, C. (2018). Financing flood losses: A discussion of the national flood insurance program. *Risk Management and Insurance Review 21*.

Kousky, C., H. Kunreuther, M. LaCour-Little, and S. Wachter (2020). Flood risk and the us housing market. *Journal of Housing Research 29*.

Meisenzahl, R. (2023). How climate change shapes bank lending: Evidence from portfolio reallocation. *FRB Chicago Working Paper series*.

Mulder, P. (2022). Mismeasuring risk: The welfare effects of flood risk information. *Available at SSRN*.

Ortega, F. and I. Petkov (2024). To improve is to change? the effects of risk rating 2.0 on flood insurance demand. *Mimeo 1*.

Pope, J. (2008). Do seller disclosures affect property values? buyer information and the hedonic model. *Land Economics 84*.

Rehbein, O. and S. Ongena (2020). Flooded through the back door: The role of bank capital in local shock spillovers. *Swiss Finance Institute Research Paper No. 20-07*.

Sastri, P. (2021). Who bears flood risk? evidence from mortgage markets in florida. *Mimeo*.

Schüwer, U., C. Lambert, and F. Noth (2018). How do banks react to catastrophic events? evidence from hurricane katrina. *Review of Finance 23*(1), 75–116.

Sharpe, S. A. and S. M. Sherlund (2016). Crowding out effects of refinancing on new purchase mortgages. *Review of Industrial Organization 48*, 209–239.

Shr, Y.-H. and K. Zipp (2019). The aftermath of flood zone remapping: The asymmetric impact of flood maps on housing prices. *Land Economics*.

Troy, A. and J. Romm (2004). Learning from extreme events: Risk perceptions after the flood. *Land Economics 47*.

Appendix

A.1 Supplemental Descriptive Statistics

Table A.1: *Unrestricted Sample Descriptive Statistics*

Variable	Obs	Mean	SD	10%	Median	90%
Applicant Characteristics						
Applicant Age	18505309	42	14	26	39	62
Applicant Credit Score	16976449	729	61	644	740	800
Applicant Income	18290710	113	100	39	84	214
Applicant Male	18807654	0.61	0.49	0	1	1
Applicant Hispanic	18807654	0.12	0.33	0	0	1
Applicant Black	18807654	0.08	0.27	0	0	0
Applicant White	18807654	0.7	0.46	0	1	1
Loan Characteristics & Outcomes						
Property Value (1000 USD)	18471705	366	318	120	279	680
Loan Amount (1000 USD)	18807654	284	209	89	235	515
Loan-to-Value Ratio	18471705	84	20	60	90	100
Loan Term (Months)	18705647	340	61	276	360	360
Interest Rate	16694380	3.9	1.1	2.8	3.8	5.1
Conforming Loan	18807653	0.93	0.25	1	1	1
Loan Approved	18807654	0.9	0.31	0	1	1
Loan Originated	18807654	0.87	0.34	0	1	1
Loan Sold	18807654	0.69	0.46	0	1	1
Loan Securitized	18807654	0.4	0.49	0	0	1
Flood Risk						
100-yr Flood Zone	18187332	0.047	0.21	0	0	0
Un-mapped	18187332	0.16	0.36	0	0	1
Possibly Un-mapped	18187332	0.43	0.49	0	0	1

Note: Descriptive statistics for our primary variables of interest and controls appear in the table above. This sample keeps only mortgage applications that were either approved or denied (excluding applications that were withdrawn), but includes non-conforming (jumbo) loans and loans that lie outside of 100-year and 500-year FEMA flood zones. Applicant age, income, property value, loan amount, combined LTV, and interest rate all contain values we consider to be implausibly small and implausibly large. These variables as displayed in this table and used throughout the analysis are winsorized between the first and ninety-ninth percentiles. Our dummy for loan securitizations includes only those loans that were sold to one of the public securitization agencies within the same calendar year.

Table A.2: *Summary Statistics for Listings & HMDA Tract-Quarter Panel*

Variable	Obs	Mean	SD	10%	Median	90%
Pricing & Value						
Mean Closing Price	791113	373	308	115	284	723
Median Closing Price	791113	358	313	107	270	695
Mean Property Value	790681	394	319	134	296	764
Median Property Value	790681	374	320	125	276	718
Flood Risk						
Mean AAL $\times 100$	786197	0.11	0.13	0.0054	0.085	0.22
Mean AAL Risk Score	786197	29	25	3.1	22	69
% in 100-year FZ	784367	0.052	0.18	0	0	0.11
% in 500-year FZ	784367	0.056	0.19	0	0	0.11
% Un-mapped	784367	0.16	0.21	0	0.1	0.43
% Possibly Un-mapped	784367	0.44	0.32	0	0.38	1
Local Characteristics						
% in Foreclosure	791113	0.0071	0.052	0	0	0
Tract Median Family Income	791100	72757	34787	36500	65455	117500
Tract % Minority	791108	37	29	5.4	28	87
Tract Population	791108	4547	2097	2279	4258	7044
MSA Median Family Income	791108	69895	15178	52733	67322	92317
CL Obs. in Tract-Quarter	791113	16	17	2	12	33
HMDA Obs. in Tract-Quarter	791113	16	18	3	11	31
Tract Dummies						
Coast Within 1 mi	791113	0.045	0.21	0	0	0
Coast Within 5 mi	791113	0.085	0.28	0	0	0
Coast Within 10 mi	791113	0.13	0.34	0	0	1
100yr FZ Dummy	784367	0.038	0.19	0	0	0
Un-mapped Dummy	784367	0.061	0.24	0	0	0
Possibly Un-mapped Dummy	784367	0.34	0.47	0	0	1

Note: The table displays summary statistics for our census tract-quarter panel of housing market and lending outcomes. We restrict the sample of real estate listings in the CoreLogic MLS data to just residential properties that were sold and not rented. Similarly, we restrict the HMDA data to just those mortgages for residential properties that were successfully originated. For both closing prices and property values, we observed implausibly small and large values, so we winsorized these values between the first and ninety-ninth percentiles. Tract-quarter means and medians are taken after winsorization. The panel represents 60,621 2010 census tracts over sixteen quarters (an unbalanced panel).

Figure A.1: *AAL Risk Score Distributions by Risk Zones*

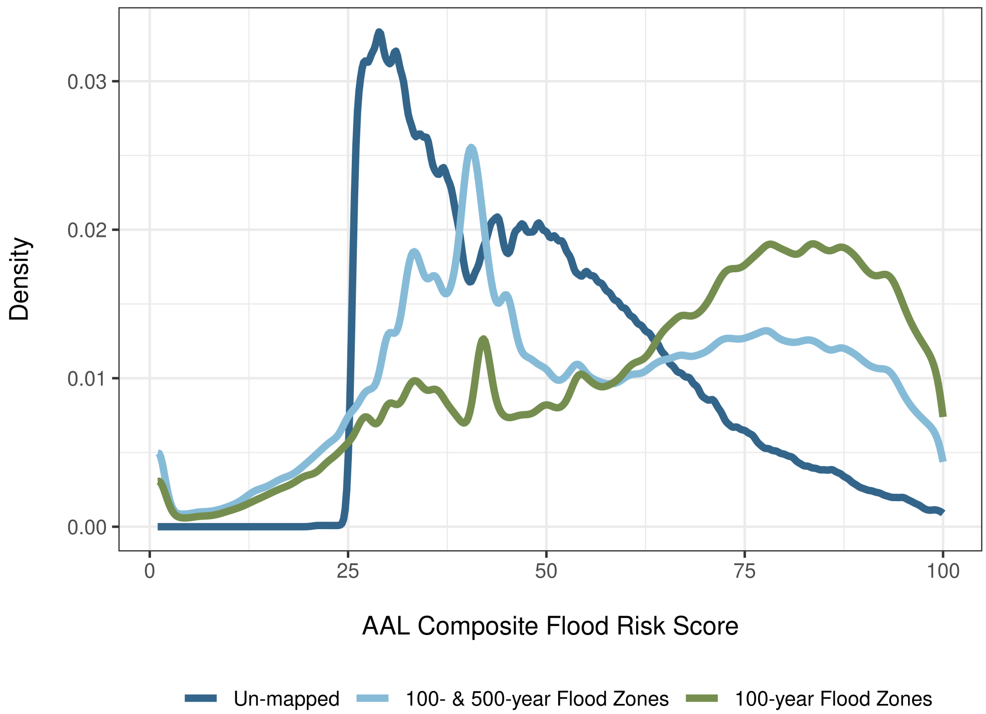

Note: The figure plots the AAL composite flood risk score empirical density distributions for properties classified as un-mapped, inside a 100- or 500-year flood zone, and inside a 100-year flood zone. All densities are conditional on the property having non-zero composite flood risk (AAL $\neq 0$). The AAL composite flood risk score describes the quantile of the AAL composite flood risk distribution of a given property such that higher AAL risk scores correspond with higher AALs and vice versa. The quantiles described by the risk scores are not simply the percentiles of the AAL distribution (i.e., a risk score of 25 corresponds to the 50th percentile of the AAL distribution for properties with non-zero risk, not the 25th percentile). The exact mapping between AALs and risk scores is proprietary.

A.2 Loan Characterizations by Destination

Table A.3: *Flood Risk-Mortgage Destination Crosstabulation: Conforming Mortgages 2018-2021*

Destination After Origination	100-yr FZ	500-yr FZ	Unmapped	Possibly Unmapped	Low Risk
Loan Kept	29.09	26.49	94.07	171.09	313.16
Share of Loans in FZ	(18%)	(13%)	(16%)	(17%)	(17%)
Share of Destination	(5%)	(4%)	(15%)	(27%)	(49%)
Sold to Bank	14.12	20.73	55.89	106.8	185.8
Share of Loans in FZ	(9%)	(10%)	(10%)	(10%)	(10%)
Share of Destination	(4%)	(5%)	(15%)	(28%)	(48%)
Sold to GSEs	86.89	108.41	299.54	532.49	954.36
Share of Loans in FZ	(53%)	(54%)	(52%)	(52%)	(52%)
Share of Destination	(4%)	(5%)	(15%)	(27%)	(48%)
Sold to Non-Bank	31.7	41.1	114.52	204.19	370.51
Share of Loans in FZ	(19%)	(20%)	(20%)	(20%)	(20%)
Share of Destination	(4%)	(5%)	(15%)	(27%)	(49%)
Sold to Private Label	2.56	4.03	7.51	12.32	18
Share of Loans in FZ	(2%)	(2%)	(1%)	(1%)	(1%)
Share of Destination	(6%)	(9%)	(17%)	(28%)	(41%)
Total	164.37	200.76	571.52	1026.9	1841.83

Note: Table displays the distribution of the total mortgage volume in billions USD across flood risk categories and destinations after origination. The share of mortgage of volume in each destination after origination is given in parentheses below the dollar value. Mortgages volumes are the sum of all mortgage loans in our sample for the years 2018-2021. Mortgage destinations are the as of the end of the calendar year.

Table A.4: *Flood Risk-Mortgage Destination Crosstabulation: Conforming Mortgages 2018-2021*

Lender	Destination After Origination	100-yr FZ	500-yr FZ	Unmapped	Possibly Unmapped	Low Risk
Banks	Loan Kept	20.22	16.75	65.35	120.67	219.86
	Share of Loans in FZ	(41%)	(36%)	(37%)	(38%)	(35%)
	Share of Destination	(5%)	(4%)	(15%)	(27%)	(50%)
	Sold to Bank	2.9	2.54	11.25	21.94	44.09
	Share of Loans in FZ	(6%)	(6%)	(6%)	(7%)	(7%)
	Share of Destination	(4%)	(3%)	(14%)	(27%)	(53%)
	Sold to GSEs	21.25	22.26	78.61	142.15	280
	Share of Loans in FZ	(43%)	(48%)	(45%)	(44%)	(45%)
	Share of Destination	(4%)	(4%)	(14%)	(26%)	(51%)
	Sold to Non-Bank	5.15	4.51	19.04	35.5	74.67
	Share of Loans in FZ	(10%)	(10%)	(11%)	(11%)	(12%)
	Share of Destination	(4%)	(3%)	(14%)	(26%)	(54%)
	Sold to Private Label	0.11	0.11	0.59	1.01	1.88
	Share of Loans in FZ	(0%)	(0%)	(0%)	(0%)	(0%)
	Share of Destination	(3%)	(3%)	(16%)	(27%)	(51%)
Non-Banks	Loan Kept	8.87	9.74	28.71	50.42	93.3
	Share of Loans in FZ	(8%)	(6%)	(7%)	(7%)	(8%)
	Share of Destination	(5%)	(5%)	(15%)	(26%)	(49%)
	Sold to Bank	11.21	18.2	44.65	84.86	141.71
	Share of Loans in FZ	(10%)	(12%)	(11%)	(12%)	(12%)
	Share of Destination	(4%)	(6%)	(15%)	(28%)	(47%)
	Sold to GSEs	65.64	86.15	220.93	390.35	674.35
	Share of Loans in FZ	(57%)	(56%)	(56%)	(55%)	(55%)
	Share of Destination	(5%)	(6%)	(15%)	(27%)	(47%)
	Sold to Non-Bank	26.55	36.59	95.48	168.69	295.84
	Share of Loans in FZ	(23%)	(24%)	(24%)	(24%)	(24%)
	Share of Destination	(4%)	(6%)	(15%)	(27%)	(47%)
	Sold to Private Label	2.46	3.92	6.92	11.31	16.12
	Share of Loans in FZ	(2%)	(3%)	(2%)	(2%)	(1%)
	Share of Destination	(6%)	(10%)	(17%)	(28%)	(40%)
All lenders	All Destinations	164.37	200.76	571.52	1026.9	1841.83

Note: Table displays the distribution of the total mortgage volume in billions USD across flood risk categories and destinations after origination. The share of mortgage of volume in each destination after origination is given in parentheses below the dollar value. Mortgages volumes are the sum of all mortgage loans in our sample for the years 2018-2021. Mortgage destinations are the as of the end of the calendar year.

Table A.5: *Flood Risk-Mortgage Destination Summary: Loan-Weighted Mean Interest Rate, Conforming Mortgages 2018-2021*

Destination After Origination	100-yr FZ	500-yr FZ	Unmapped	Possibly Unmapped	Low Risk
Loan Kept	3.99	4.02	3.94	3.90	3.93
Sold to Affiliate	4.08	3.97	3.90	3.87	3.77
Sold to Bank	3.93	4.03	3.93	3.89	3.95
Sold to GSEs	3.62	3.69	3.67	3.67	3.62
Sold to Non-Bank	3.92	4.00	3.98	3.98	3.82
Sold to Private Label	4.92	4.68	4.93	4.96	4.82

Table A.6: *Flood Risk-Mortgage Destination Summary: Loan-Weighted Mean FICO Score, Conforming Mortgages 2018-2021*

Destination After Origination	100-yr FZ	500-yr FZ	Unmapped	Possibly Unmapped	Low Risk
Loan Kept	748.92	746.72	747.47	747.91	744.65
Sold to Affiliate	741.94	737.77	742.19	743.22	742.03
Sold to Bank	745.92	740.40	745.09	747.38	746.16
Sold to GSEs	739.39	739.06	739.72	740.45	738.72
Sold to Non-Bank	724.23	722.07	721.71	723.10	721.15
Sold to Private Label	734.40	730.92	735.25	736.99	733.91

Table A.7: *Interest Rates with Lender Types and Mortgage Destinations*

	Interest Rate			
	Local Bank	Large Bank	Other Bank	Non-Bank
	(1)	(2)	(3)	(4)
Un-mapped × Sold to Bank/Non-Bank	0.0226***	0.0036	0.0003	0.0064**
	(0.0059)	(0.0097)	(0.0030)	(0.0032)
Un-mapped × Sold to GSEs	0.0016	0.0138***	-0.0284***	-0.0135***
	(0.0062)	(0.0036)	(0.0026)	(0.0031)
Un-mapped	-0.0115**	-0.0097***	0.0138***	0.0078**
	(0.0053)	(0.0032)	(0.0025)	(0.0031)
Possibly Un-mapped	0.0062***	0.0027*	-0.0047***	-0.0006
	(0.0023)	(0.0016)	(0.0009)	(0.0005)
Sold to Bank/Non-Bank	-0.1732***	0.0427***	-0.2857***	-0.2057***
	(0.0038)	(0.0047)	(0.0013)	(0.0014)
Sold to GSEs	-0.2047***	0.0669***	-0.3197***	-0.2729***
	(0.0045)	(0.0018)	(0.0011)	(0.0014)
Log Loan Amount	-0.2930***	-0.3222***	-0.4263***	-0.2838***
	(0.0035)	(0.0027)	(0.0013)	(0.0013)
Mortgage Controls	✓	✓	✓	✓
County-Quarter-Year FEs	✓	✓	✓	✓
Observations	389,610	513,249	3,250,507	7,930,928
R^2	0.67360	0.74584	0.59418	0.68635
Within R^2	0.13240	0.13068	0.20085	0.13119

Note: We estimate equation 1, above. The interest rate is continuous and bounded between 2 and 8. We remove properties with extreme interest rates (above the 99th percentile or below the 1st percentile). Our variables of interest are binary and denote whether a property is un-mapped or possibly un-mapped. Non-bank is a dummy variable indicating that, based on its entity type code, the lender is not a bank nor credit union. Local bank is a dummy variable indicating that at least 40% of the lender's mortgage origination occurs in the county the property is located in. Large bank is a dummy variable indicating that the lender is an entity or a subsidiary of an entity regulated under the Large Institution Supervision Coordinating Committee (LISCC) program. Each specification contains controls for an applicant's sex, ethnicity, race, credit score, and income. We add additional controls and fixed effects in columns (2)–(4). We restrict our sample to conforming loans and exclude all properties that are covered by a flood zone map. Standard errors are clustered by census tract and year-quarter and shown in parentheses; *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

Table A.8: ***Interest Rates and Mortgage Destinations***

	Interest Rate			
	(1)	(2)	(3)	(4)
Un-mapped × Private Securitizer	-0.0015	0.0041	-0.0012	-0.0044
	(0.0086)	(0.0083)	(0.0083)	(0.0069)
Un-mapped × Public Securitizer	-0.0100***	-0.0075***	-0.0106***	0.0046***
	(0.0020)	(0.0019)	(0.0018)	(0.0013)
Un-mapped × Sold to Bank/Non-Bank	0.0165***	0.0128***	0.0081***	0.0125***
	(0.0021)	(0.0020)	(0.0019)	(0.0014)
Un-mapped	0.0200***	0.0124***	0.0045**	-0.0054***
	(0.0019)	(0.0018)	(0.0018)	(0.0013)
Possibly Un-mapped	0.0019***	0.0144***	-0.0015***	0.0013***
	(0.0005)	(0.0004)	(0.0004)	(0.0004)
Private Securitizer	0.1164***	0.1934***	0.1608***	0.1907***
	(0.0040)	(0.0038)	(0.0038)	(0.0039)
Public Securitizer	-0.3044***	-0.2458***	-0.2442***	-0.1876***
	(0.0009)	(0.0008)	(0.0008)	(0.0007)
Sold to Bank/Non-Bank	-0.2392***	-0.1911***	-0.1841***	-0.1666***
	(0.0010)	(0.0009)	(0.0009)	(0.0008)
Log Loan Amount		-0.2850***	-0.3337***	-0.2518***
		(0.0008)	(0.0010)	(0.0009)
Mortgage Controls	✓	✓	✓	✓
Quarter-Year FEs	✓	✓		✓
County-Quarter-Year FEs			✓	
Tract FEs				✓
Lender FEs				✓
Observations	12,084,309	12,084,294	12,084,294	12,084,294
R^2	0.60360	0.63066	0.64712	0.73492
Within R^2	0.08671	0.14906	0.14999	0.10146

Note: We estimate equation 1, above. The interest rate is continuous and bounded between 2 and 8. Our variables of interest are binary and denote whether a property is un-mapped or possibly un-mapped. Each specification contains controls for an applicant's sex, ethnicity, race, credit score, and income. We add additional controls and fixed effects in columns (2)–(4). We restrict our sample to conforming loans and exclude all properties that are covered by a flood zone map. Standard errors are clustered by tract and the year-quarter or origination; *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

Table A.9: ***Interest Rates with Lender Types and Mortgage Destinations***

	Interest Rate		
	Loan Kept	Sold to GSEs	Sold to Bank/Non-Bank
	(1)	(2)	(3)
Un-mapped × Non-Bank	-0.0200***	-0.0016	-0.0145***
	(0.0037)	(0.0011)	(0.0021)
Un-mapped × Local Bank	-0.0348***	0.0062*	-0.0124***
	(0.0060)	(0.0033)	(0.0035)
Un-mapped × Large Bank	-0.0455***	0.0027	-0.0394***
	(0.0041)	(0.0021)	(0.0093)
Un-mapped	0.0188***	0.0014	0.0162***
	(0.0026)	(0.0010)	(0.0019)
Non-Bank	0.0173***	0.0927***	0.1026***
	(0.0017)	(0.0005)	(0.0009)
Local Bank	-0.0958***	0.0136***	-0.0215***
	(0.0031)	(0.0017)	(0.0018)
Large Bank	-0.3250***	-0.0049***	-0.1603***
	(0.0020)	(0.0009)	(0.0046)
Possibly Un-mapped	-0.0088***	0.0018***	0.0005
	(0.0015)	(0.0004)	(0.0007)
Mortgage Controls	✓	✓	✓
County-Quarter-Year FEs	✓	✓	✓
Observations	2,242,563	5,907,574	3,934,157
R^2	0.48990	0.78276	0.67915
Within R^2	0.17288	0.12037	0.12256

Note: We estimate equation 1, above. Our outcome variable is binary and tracks whether a mortgage is originated (i.e. accepted by both borrower and lender) and our variables of interest are also binary and denote whether a property is un-mapped or possibly un-mapped. Non-bank is a dummy variable indicating that, based on its entity type code, the lender is not a bank nor credit union. Local bank is a dummy variable indicating that at least 40% of the lender's mortgage origination occurs in the county the property is located in. Large bank is a dummy variable indicating that the lender is an entity or a subsidiary of an entity regulated under the Large Institution Supervision Coordinating Committee (LISCC) program. Each specification contains controls for an applicant's sex, ethnicity, race, credit score, and income. We add additional controls and fixed effects in columns (2)–(4). We restrict our sample to conforming loans and exclude all properties that are covered by a flood zone map. Standard errors are clustered by tract and the year-quarter or origination; *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

Table A.10: *Differential Shifts in Sale & Securitization Rates*

	YoY Basis Point Change	
	Sale Rate	Securitization Rate
	(1)	(2)
$\mathbb{1}$(99th Un-mapped Risk Percentile)	48.24**	25.98*
	(19.49)	(15.64)
$\mathbb{1}$(95-99th Un-mapped Risk Percentile)	28.06***	20.02***
	(9.193)	(5.709)
$\mathbb{1}$(90-95th Un-mapped Risk Percentile)	19.44***	18.29***
	(6.233)	(4.748)
$\mathbb{1}$(80-90th Un-mapped Risk Percentile)	9.202**	6.046*
	(4.091)	(3.390)
Tract Controls	✓	✓
County FEs	✓	✓
Year FEs	✓	✓
Observations	742,781	742,781
R^2	0.07361	0.06137
Within R^2	0.00661	0.00164

Note: This table relates year-over-year basis point changes in the tract-year share of mortgage loan amount that is sold or securitized to tract-level flood risk rankings. The dataset is an unbalanced tract-year panel. Year-over-year basis point changes in these market shares are computed for the years 2013-2022. In each specification we add tract-year controls, including the mean loan-to-income ratio, mean log applicant income, percent of male mortgage applicants, percent of Native American, Asian, Black, Pacific Islander, and white applicants, as well as a time-invariant count of the number of properties in the tract constructed from the CoreLogic data.

A.3 Loan Flows After Origination

Table A.11: *Loan Securitization*

	Loan Securitized			
	(1)	(2)	(3)	(4)
Un-mapped	0.0061***	0.0087***	0.0026***	0.0023***
	(0.0005)	(0.0005)	(0.0005)	(0.0004)
Possibly Un-mapped	-0.0029***	-0.0069***	0.0004	0.0003
	(0.0004)	(0.0004)	(0.0004)	(0.0003)
Interest Rate	-0.0605***	-0.0565***	-0.0587***	-0.0223***
	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Log Loan Amount		0.0944***	0.1020***	0.0721***
		(0.0003)	(0.0003)	(0.0002)
Mortgage Controls	✓	✓	✓	✓
Quarter-Year FEs	✓	✓		✓
County-Quarter-Year FEs			✓	
Tract FEs				✓
Lender FEs				✓
Observations	11,922,096	11,922,081	11,922,081	11,922,081
R^2	0.03101	0.04429	0.07488	0.43779
Within R^2	0.01397	0.02748	0.02681	0.01357

Note: We estimate equation 1, above. The outcome variable is binary and indicates when a loan is sold to a public securitization agency. Our variables of interest are binary and denote whether a property is un-mapped or possibly un-mapped. Each specification contains controls for an applicant's sex, ethnicity, race, credit score, and income. We add additional controls and fixed effects in columns (2)–(4). We restrict our sample to conforming loans and exclude all properties that are covered by a flood zone map. Standard errors are clustered by census tract and year-quarter and shown in parentheses; *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

Table A.12: ***Transaction Flows***

	Publicly Securitized (1)	Privately Securitized (2)	Sold to Bank (3)	Sold to Financial Company (4)	Sold to Life Insurance Company (5)	Sold to Affiliate (6)	Sold to Other Company (7)
Un-mapped	0.0019***	-0.0002**	-2.96×10^{-5}	-0.0008**	1.82×10^{-5}	0.0005***	-0.0014***
	(0.0004)	(8.73×10^{-5})	(0.0003)	(0.0003)	(7.06×10^{-5})	(0.0001)	(0.0003)
Possibly Un-mapped	0.0005	-3.62×10^{-6}	-0.0004*	0.0005**	-9.43×10^{-5}*	-0.0002**	-0.0003
	(0.0003)	(6.56×10^{-5})	(0.0002)	(0.0003)	(5.63×10^{-5})	(9.06×10^{-5})	(0.0002)
Interest Rate	0.0072***	0.0103***	-0.0019***	-0.0038***	-3.29×10^{-5}	0.0009***	-0.0126***
	(0.0002)	(9.92×10^{-5})	(0.0002)	(0.0002)	(4.38×10^{-5})	(7.1×10^{-5})	(0.0002)
Log Loan Amount	0.0616***	-0.0036***	0.0081***	0.0358***	0.0009***	0.0016***	-0.1044***
	(0.0003)	(5.99×10^{-5})	(0.0002)	(0.0002)	(4.61×10^{-5})	(6.27×10^{-5})	(0.0003)
Dependent Variable Mean	0.5949	0.0091	0.1183	0.1763	0.0061	0.0135	0.0819
Mortgage Controls	✓	✓	✓	✓	✓	✓	✓
Quarter-Year FEs	✓	✓	✓	✓	✓	✓	✓
Tract FEs	✓	✓	✓	✓	✓	✓	✓
Lender FEs	✓	✓	✓	✓	✓	✓	✓
Observations	9,937,282	9,937,282	9,937,282	9,937,282	9,937,282	9,937,282	9,937,282
R^2	0.54383	0.33986	0.29735	0.37682	0.32131	0.51104	0.36476
Within R^2	0.01549	0.00555	0.00596	0.01019	0.00011	0.00046	0.06730

Note: We consider seven mutually exclusive outcomes that appear at the top of the columns of the table. Variables of interest are also binary and denote whether a property is un-mapped or possibly un-mapped. We add additional controls and fixed effects in columns (2)–(4). We restrict our sample to conforming loans and exclude all properties that are covered by a flood zone map. Standard errors are clustered by tract and the year-quarter or origination; *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

A.4 Measuring Flood Risk Outside of Flood Zones

The "unmapped" categorization we rely on throughout the paper is based on differences between CoreLogic flood risk measures and official FEMA flood maps. Here, our underlying assumption is that lenders may respond to true flood risk, and that CoreLogic provides more accurate measures of true flood risk than FEMA's flood map layers. We argue that this assumption is not unreasonable. It is intuitive that flood risk varies continuously over geographies in ways that cannot be captured in FEMA's discrete flood maps. Flood maps have not been updated recently leading to stark differences in flood map coverage along what appears to be otherwise purely administrative boundaries. Figure A.2 highlights one such example just upriver from New Orleans. Still, a critical review of this might argue instead that FEMA flood maps are carefully developed with on-the-ground observation, which may be more accurate than the modeled estimates CoreLogic provides.

Here we empirically substantiate this assumption and demonstrate that the flood risk measures we rely on from CoreLogic are consistent with other measures of flood risk, including some of FEMA's own measures of flood risk. Importantly, we need to establish that this is the case for properties *outside* flood zones. To do so, we leverage policy data from the National Flood Insurance Program to derive flood risk measures from FEMA's own data that go beyond the categorical flood map designations that the flood insurance mandate varies across. We also compare CoreLogic's flood risk measures against First Street Foundation's, another source of property-level flood risk measures.

Benchmarking Against FEMA Data

In this subsection, we compare flood risk using measures from CoreLogic to flood risk measures derived from FEMA data. FEMA's flood maps demark boundaries where people are required to purchase flood insurance, but under NFIP's revised pricing scheme, the premiums for those policies can capture flood risk beyond these discrete boundaries. FEMA publishes the near universe of NFIP policies, including premiums and coverage details. For confidentiality, FEMA does not release the granular coordinates of each property associated with a flood policy, but does include information about the census blockgroup and flood zone of these properties. This means we are not able to match CoreLogic flood risk measures to flood policies at the property level, but are able to do so at a relatively granular level of geography. In the CoreLogic data, we compute the mean average annual loss (AAL) in each census tract-flood zone pair.

Figure A.3 plots the joint distribution of the mean AAL from CoreLogic and the mean flood insurance premium at the census tract-flood zone level. On first pass, we see that tracts with greater flood risk in CoreLogic indeed pay higher flood insurance premiums. This is true both within the official 100-year flood zones and outside of it. CoreLogic's AAL does a poor job of explaining the variation in the flood insurance premiums—at least at the tract-flood zone level.

There are of course other reasons why flood insurance premiums would be higher in neigh-

Figure A.2: *Mapped & Un-mapped Flood Risk Along the Mississippi River*

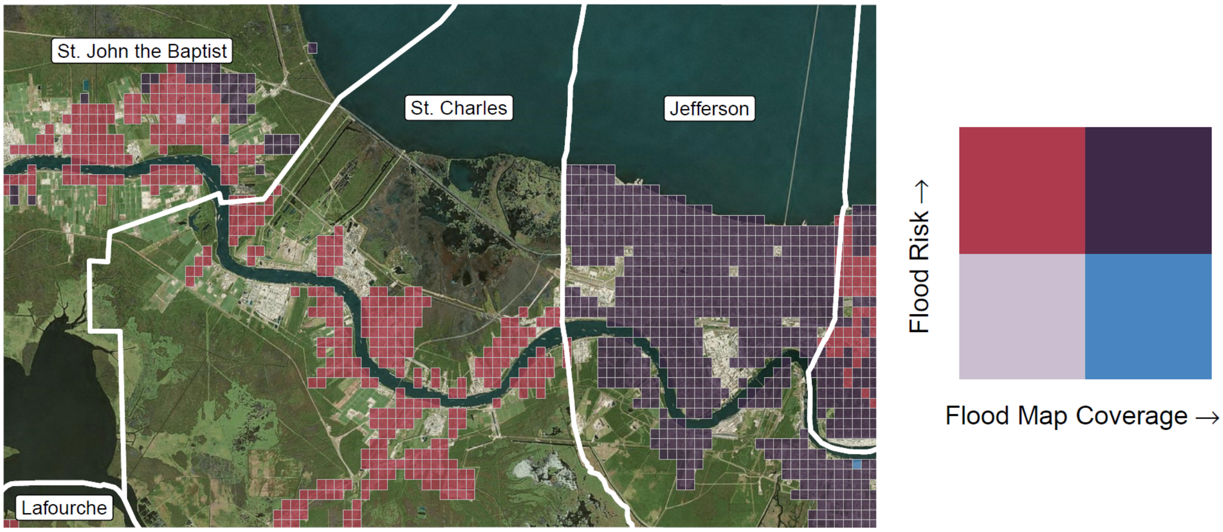

Note: This map plots CoreLogic flood risk and FEMA flood map coverage along the Mississippi river, just upriver from New Orleans. We overlay a 0.005° × 0.005° (≈ 500 × 500 meters) grid over this portion of the river, where it crosses through St. John the Bapist Parish, St. Charles Parish, and Jefferson Parish (a parish is Louisiana's county-level geography). Thick white lines denote parish boundaries. In each grid cell, we compute the (1) the area of that grid cell covered by a FEMA-designated 100-year flood zone, and (2) the mean composite flood risk AAL for CoreLogic properties located in the cell. We only display grid cells with at least five CoreLogic properties. The coloring for each grid cell is then determined by (1) whether the majority of the cell is covered by a FEMA flood map, and (2) whether the mean property in the cell has a high enough AAL that we would consider it to be "high risk." Gray cells have low flood map coverage and low flood risk; red cells have low flood map coverage and high flood risk; yellow cells have high flood map coverage and low flood risk; and blue cells have high flood map coverage and high flood risk. It follows that red cells have predominantly "un-mapped" properties. The map highlights that many properties near this portion of the Mississippi river have comparable flood risks, yet on what appears to be purely administrative boundaries, Jefferson Parish is almost entirely covered by flood maps, whereas St. John the Baptist and St. Charles Parish do not have any flood maps.

borhoods with higher flood risk (as measured by CoreLogic). Flood insurance premiums are functions of flood risk, as well as functions of coverage amounts, housing type, and structure specific characteristics. We consider two approaches to handle this. First, we regress flood insurance premiums on a host of policy-specific characteristics, as in

$$\log(\text{Premium})_{it} = \beta_1 \log(\text{Coverage Amt})_{it} + \beta_2 \text{postFIRM}_{it} + \beta_3 \text{SF}_{it} + \eta_t + \varepsilon_{it} \tag{A.1}$$

where for a policy i that is in effect at the end of year t, $\log(\text{Premium})_{it}$ is the natural logarithm of the total insurance premium, $\log(\text{Coverage Amt})_{it}$ is the natural logarithm of the total coverage amount, postFIRM_{it} is a dummy to indicate if the structure was built after the official Flood Insurance Rate Map and therefore not eligible for grandfathered policy rate, SF_{it} is a dummy to indicate if the policy is for single-family housing rather than multifamily housing, and η_t are year fixed effects. We do not control for other structure characteristics (e.g., construction date of the

Figure A.3: *Flood Premiums & CoreLogic Risk*

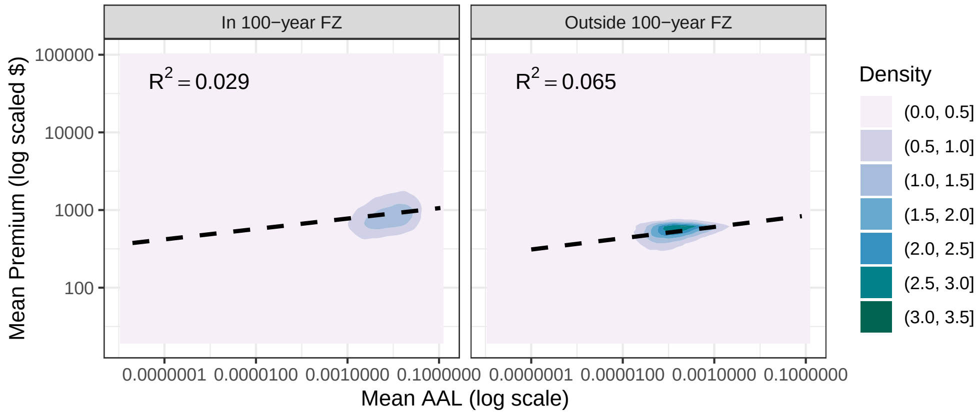

Note: The figure above plots the joint density distribution of mean flood insurance premiums and mean average annual losses (AAL) from CoreLogic. Observations are census tract-flood zone pairs. Mean insurance premiums are computed over single-family and multifamily residential properties for all policies outstanding as of December 31, 2023. Both axes are log-scaled and observations with mean premium or mean AAL equal to zero are omitted.

Figure A.4: *Flood Premium Residuals & CoreLogic Risk*

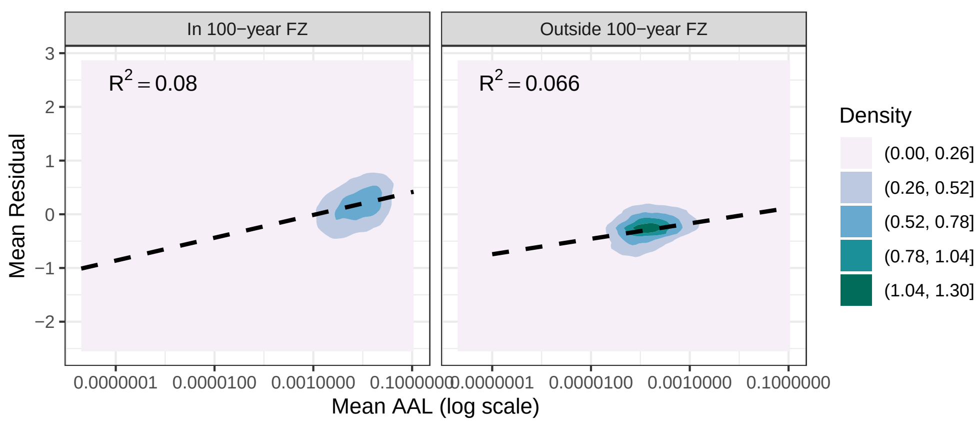

Note: The figure above plots the joint density distribution of mean residuals based on the regression in equation (A.1) and mean average annual losses (AAL) from CoreLogic. We plot only residuals from observations in 2023. Observations are census tract-flood zone pairs. Mean insurance premiums are computed over single-family and multifamily residential properties for all policies outstanding as of December 31, 2023. Both axes are log-scaled and observations with mean premium or mean AAL equal to zero are omitted.

property or structure elevation) as these are reflected in CoreLogic's AAL measures. By controlling for these policy-specific characteristics that may influence premium amounts, we assume that the unexplained variation in premiums is attributable to FEMA's internal risk measure for a property. That is, the ε_{it} should measure flood risk from FEMA, with flood risk increasing in the residual. We run this regression on the set of residential policies for the years 2013-2023. Figure A.4 plots the joint distribution of the mean residual and the mean AAL from CoreLogic in each census tract-flood zone pair. Again, we see that higher flood risk as measured by CoreLogic corresponds with higher flood risk as inferred from the FEMA's pricing of flood policies. As before, CoreLogic flood risk measures do not explain a substantial amount of the variation in the inferred flood risk from the flood policies.

As a final test, we use a change in the pricing of flood insurance policies to test if CoreLogic's flood risk measures correlate with FEMA's own flood risk measures. In 2022, FEMA began changing the pricing scheme for flood insurance policies in the NFIP. While rates had previously been determined by a schedule based of the flood zone a property was mapped to and several attributes of the policy, "Risk Rating 2.0" opted to better price insurance policies by using more granular flood risk measures (similar to what CoreLogic produces) and move towards actuarily sound rates. Starting on October 1, 2021 all new flood insurance policies were subject to Risk Rating 2.0. Existing policies were subject to Risk Rating 2.0 at the time of renewal starting on April 1, 2022 and all policies were subject to Risk Rating 2.0 by April 1, 2023. Statutory limits capped the change in premiums under Risk Rating 2.0 at 18% annually.

Prior to Risk Rating 2.0, flood insurance policies outside of the 100-year flood zone—where there is no flood insurance mandate—were given "Preferred Rate Pricing." These policies had discounted rates to encourage insurance adoption and did not differentially price flood risk, such that identical policies had identical premiums irrespective of underlying flood risk.

We leverage this policy change to infer FEMA's property-level flood risk. To do so, we begin by creating a matched sample of flood insurance policies from the NFIP policies data. We focus on just policies that existed before the first phase of Risk Rating 2.0's implementation on Ocotober 1, 2021. The raw data does not have an identifier that tracks the same policy each year it is renewed, so we instead identify the set of all policies that exist as of March 31, 2022 (pre-Risk Rating 2.0) and the set of all policies that exist as of April 1, 2023 (post-Risk Rating 2.0) and match these to each other using a set of policy and property characteristics.[1] With this set of matched policies pre- and post-implementation of Risk Rating 2.0, we compute the percent change in the total insurance premium, % Δ Premium, for each policy. We then run all policies i located in census tract-flood zone pair c and geographic region r,

$$\begin{aligned}\%\ \Delta\ \text{Premium}_{icr} = \beta_0\ \text{CoreLogic Risk}_c \times \text{Outside FZ}_i + \beta_1\ \text{CoreLogic Risk}_c + \\ \beta_2\ \text{Outside FZ}_i + \beta_3\ \text{Single Family}_i + \eta_r + \varepsilon_{icr}\end{aligned} \tag{A.2}$$

where CoreLogic Risk$_c$ is the normalized natural logarithm of the mean CoreLogic AAL within c, Outside FZ$_i$ is dummy indicating if a policy is outside the 100-year flood zone, Single Family$_i$ is a dummy indicating if a policy is for single family housing, and η_r are regional fixed effects.

Table A.13 displays the results of this regression. These indicate that higher flood risk measured

[1]Specifically, we match policies using their original construction date, original date of flood policy, property purchase date, building replacement cost, building insurance coverage, contents insurance coverage, census block group, latitude/longitude (these are rounded to 1 decimal place in the raw data), and NFIP community identifier. We exclude policies that match to multiple policies in the other period.

by CoreLogic corresponds with higher flood risk as inferred from FEMA NFIP premium changes. As expected, this is only weakly true for policies within 100-year flood zones, as prior to Risk Rating 2.0 these premiums might have already accounted for flood risk. But for policies outside of 100-year flood zones where underlying flood risk was not previously built into premiums, higher CoreLogic flood risk corresponds with substantial and precisely estimated increases in insurance premiums.

These results hold even when we layer in more granular regional fixed effects, including NFIP Community fixed effects.[2] Again, the within R^2 here indicates that census tract-flood zone level CoreLogic risk measures leave a substantial portion of the variation in premium changes unexplained. It is possible that the flood risk measures we match at the property-level in the main analysis are better at explaining this variation than the census tract-flood zone level measures we use here. Overall, CoreLogic risk measures do not seem be strong predictors of FEMA's internal risk measures, but in expectation, areas outside of official flood zones with high flood risk in CoreLogic also have higher FEMA-implied flood risk.

Table A.13: *NFIP Premium Changes and CoreLogic Risk*

	% Δ Premium			
	(1)	(2)	(3)	(4)
CoreLogic Risk × Outside FZ	2.768***	2.959***	1.948***	1.832***
	(0.7055)	(0.6212)	(0.4917)	(0.4759)
CoreLogic Risk	1.069*	0.7960	1.232***	0.8493*
	(0.6095)	(0.5611)	(0.4045)	(0.4413)
Outside FZ	4.507***	3.690***	4.503***	3.737***
	(0.8880)	(0.7972)	(0.6110)	(0.5744)
Single Family	1.170**	0.8593*	0.8179	0.6846**
	(0.5159)	(0.5109)	(0.5026)	(0.3416)
State FEs		✓		
County FEs			✓	
NFIP Community FEs				✓
Observations	2,170,381	2,170,381	2,170,381	2,170,381
R^2	0.03283	0.07053	0.13963	0.20967
Within R^2		0.02387	0.01715	0.01174

Note: Observations represent policies. Percent changes in premiums are trimmed between −50% and 50%. CoreLogic Risk is the normalized log mean composite flood risk AAL for all properties within the given policy's census tract-floodzone pair. Outside FZ is a dummy variable that indicates if a property is outside of a 100-year flood zone, and Single Family is a dummy to indicate if the policy is for single family housing rather than multifamily housing. Standard errors given in parentheses are clustered at the NFIP Community level.

[2]There are over 22,000 NFIP Communities in the U.S.. Flood map updates are made the NFIP Community level and consequently have had some role in determining the pricing of flood insurance.

Benchmarking Against First Street Foundation Flood Risk Measures

Alternatively, we can attempt to substantiate the flood risk measures we use from CoreLogic by comparing these with measures from First Street Foundation. First Street Foundation (FSF) provides conceptually similar data to CoreLogic: cross sectional, property-level flood risk measures. To compare the average annual losses (AAL) measure we use from CoreLogic with the AAL in FSF, we use our property-level CoreLogic flood risk data and aggregate it up to $0.0001^\circ \times 0.0001^\circ$ ($\approx$ 10 meters $\times$ 10 meters) grid cell-flood zone (e.g., 100-year, 500-year, or outside) pairs. These are nearly property-level matches, and in this section we refer to these grid cell-flood zone pairs as "properties," as in most all cases they include exactly one property.

Table A.14 displays summary statistics for the matched sample of grid cells. Nearly all grid cells contain exactly one property from CoreLogic and one property in the FSF data. Notably, Table A.14 shows that the FSF AAL distribution is more tail heavy, with more grid cells with zero flood risk and more cells with high flood risk.

Figure A.5 visualizes the differences in these distributions. The left panel displays the share of properties (i.e., grid cells) that each data source regards as having zero risk. Of the 54 million properties we match between the two datasets, CoreLogic regards round half as having average annual losses of zero compared to nearly nine-in-ten in FSF. The right panel displays the distribution of AALs for properties with non-zero AALs. We can see that CoreLogic has far greater weight in the bottom end of the distribution of (non-zero) flood risk.

The substantially higher share of properties with zero flood risk in the CoreLogic and FSF data raises the issue of differences in the threshold for non-zero flood risk measures. For instance, it may be the case that the distributions of AALs in CoreLogic and FSF are quite similar, but FSF has a higher threshold for measuring risk—explaining why FSF has so many more properties with AAL $= 0$ and why CoreLogic has a distribution of non-zero flood risk that is much heavier in the left tail. To explore this, Figures A.6 and A.7 replot the side-by-side distributions of AALs, but under different rules for rounding flood risk down to zero. In Figure A.6, all properties with AAL < 0.00001 are rounded down to AAL $= 0$ (a low cutoff) and in Figure A.7, all properties with AAL < 0.001 are rounded down to AAL $= 0$ (a high cutoff). The low cutoff version in Figure A.6 is still remanent of the side-by-side distributions in Figure A.5; FSF has more properties with zero flood risk and CoreLogic has more properties with non-zero but low flood risk. Still these distributions look much more comparable to each other. The high cutoff version in Figure A.7 plots two very similar distributions. Both CoreLogic and FSF have a similar share of properties with zero flood risk and the distribution of non-zero flood risk is quite close.

Finally, for the set of properties with both non-zero AALs in CoreLogic and FSF, we plot their joint distributions in Figures A.8 and A.9. If the two datasets had the same AAL for each property, their joint distribution would lie on the diagonal dashed line. We can see that the two measures are certainly correlated and that by-and-large, they are typically quite comparable. It is very rare for one measure to place high (low) flood risk on a property that does not have high (low) flood risk in the other dataset. Figure A.9 goes further to show that this is the case even outside of

Table A.14: *CoreLogic-First Street Foundation Matched Data*

Variable	% Missing	Mean	SD	Min	0.1%	1%	5%	25%	Median	75%	95%	99%	99.9%	Max
# Properties in Cell - CoreLogic	0	1.1	1.4	1	1	1	1	1	1	1	1	3	12	1495
# Properties in Cell - First Street	0	1	0.17	1	1	1	1	1	1	1	1	2	3	91
Mean AAL - CoreLogic	0	0.00042	0.003	0	0	0	0	0	0	0.000051	0.0012	0.0087	0.046	0.24
Mean AAL - First Street	0	0.00066	0.0064	0	0	0	0	0	0	0	0.00049	0.015	0.1	0.5
Observations: 54,082,840														

Note: This table provides summary statistics for matched grid cells in the CoreLogic flood risk and First Street Foundation data. The unit of observation here is a $0.0001° \times 0.0001°$ grid cell-flood zone combination.

Figure A.5: *Average Annual Loss (AAL) Distributions: CoreLogic vs. First Street Foundation*

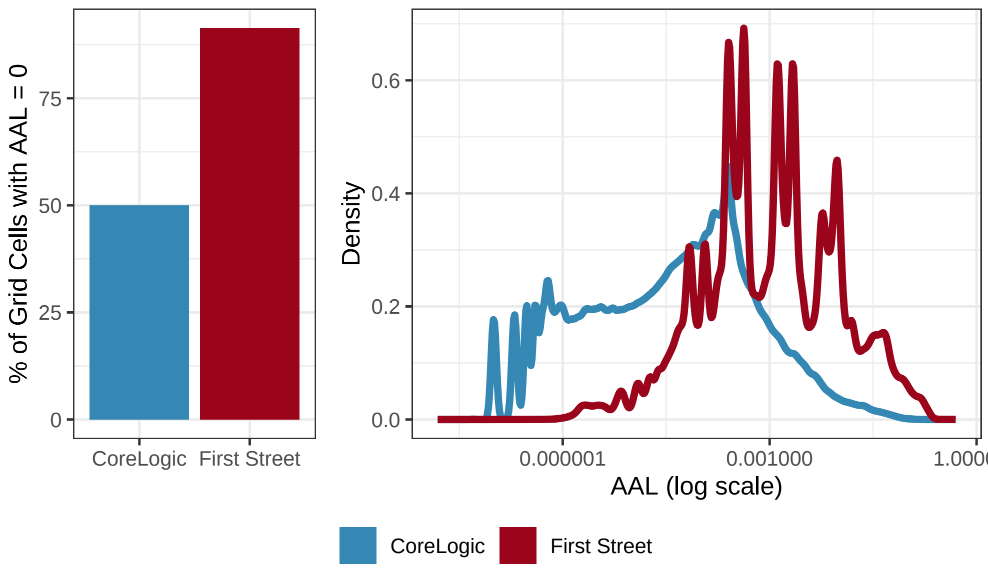

Note: The figure above compares the distribution of average annual losses in CoreLogic and First Street Foundation for the matched sample of properties. The left subfigure displays the proportion of matched properties with no flood risk (i.e., AAL = 0) from the two datasets. The right subfigure displays distribution of AALs in the two datasets, conditional on having positive flood risk.

the 100-year flood zone. The alignment of property-level CoreLogic flood risk measures with the property-level FSF flood risk measures helps substantiate our claim that CoreLogic flood risk measures capture true flood risk, even outside of the 100-year flood zone.

Figure A.6: *Average Annual Loss (AAL) Distributions: Low Cutoff*

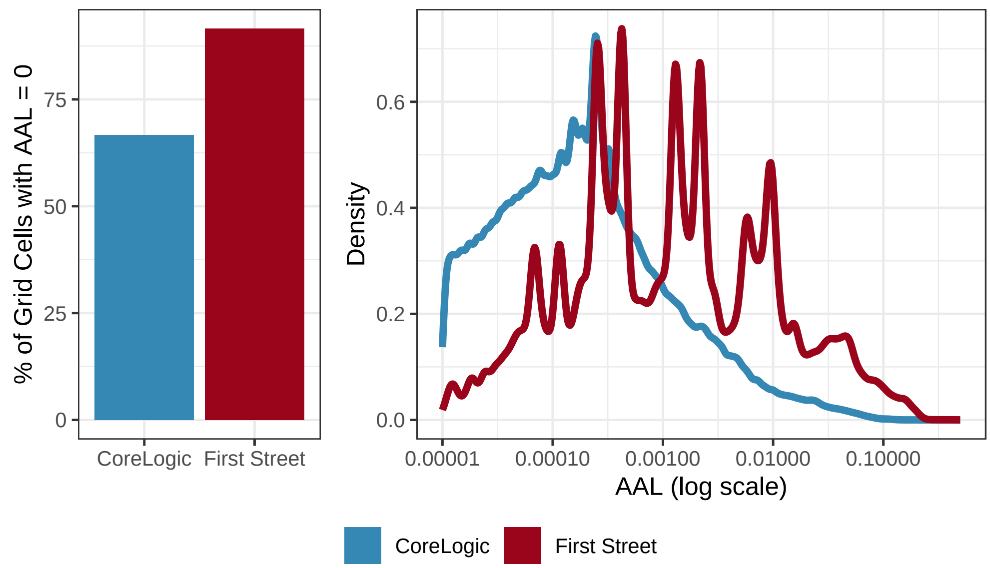

Note: The figure above compares the distribution of average annual losses in CoreLogic and First Street Foundation for the matched sample of properties. The left subfigure displays the proportion of matched properties with no flood risk (i.e., AAL = 0) from the two datasets. The right subfigure displays distribution of AALs in the two datasets, conditional on having positive flood risk.

Figure A.7: *Average Annual Loss (AAL) Distributions: High Cutoff*

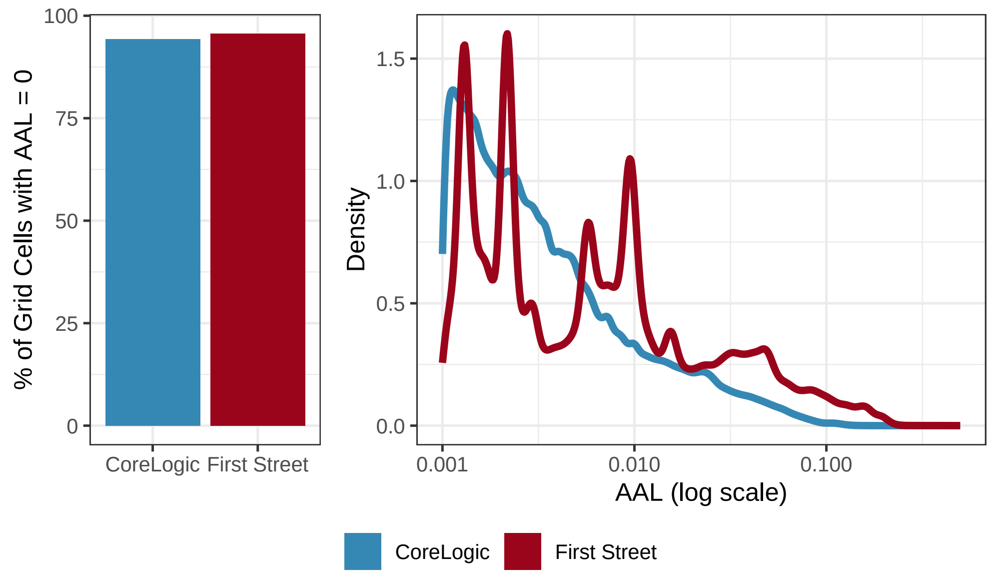

Note: The figure above compares the distribution of average annual losses in CoreLogic and First Street Foundation for the matched sample of properties. The left subfigure displays the proportion of matched properties with no flood risk (i.e., AAL = 0) from the two datasets. The right subfigure displays distribution of AALs in the two datasets, conditional on having positive flood risk.

Figure A.8: *Joint Distribution of CoreLogic and First Street AAL, Conditional on AAL > 0*

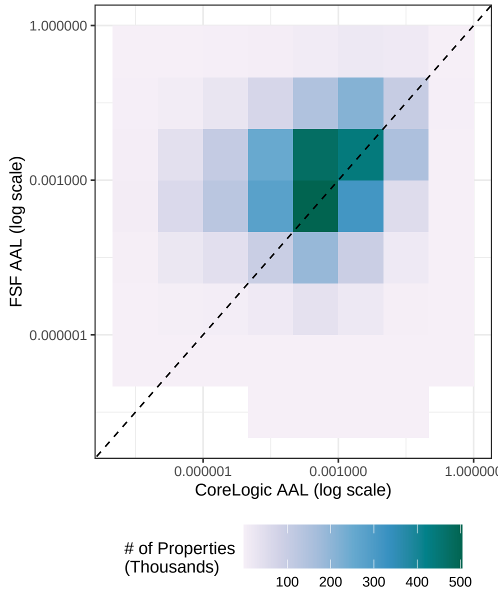

Note: The figure above plots the joint distribution of average annual losses measured using the CoreLogic data on the horizontal axis and measured using First Street Foundation on the vertical axis, conditional on both measures being positive. Both axes are log scaled. The dashed diagonal line represents where the distribution would fall along if all properties had identical AALs in the two datasets.

Figure A.9: *Joint Distribution of CoreLogic and First Street AAL, Conditional on AAL > 0 and Split by Flood Zone*

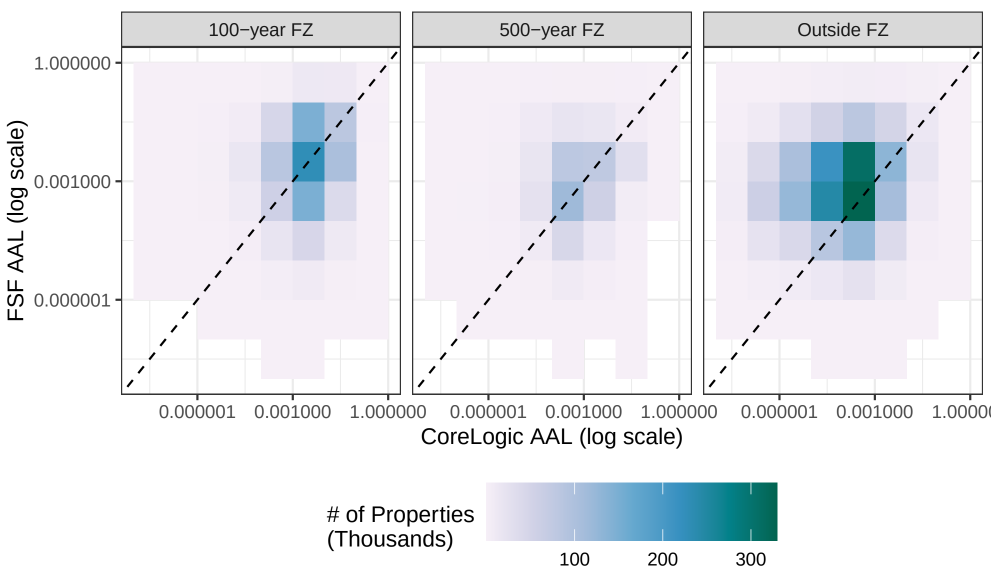

Note: The figure above plots the joint distribution of average annual losses measured using the CoreLogic data on the horizontal axis and measured using First Street Foundation on the vertical axis, conditional on both measures being positive. Both axes are log scaled. The dashed diagonal line represents where the distribution would fall along if all properties had identical AALs in the two datasets.

A.5 In & Out of Flood Zones

Table A.15: *Loan Originated: Baseline with FZ properties*

	Loan Originated			
	(1)	(2)	(3)	(4)
100-yr Flood Zone	-0.0380***	-0.0381***	-0.0261***	-0.0188***
	(0.0006)	(0.0005)	(0.0006)	(0.0006)
500-yr Flood Zone	-0.0024***	-0.0047***	-0.0052***	-0.0059***
	(0.0005)	(0.0005)	(0.0006)	(0.0007)
Un-mapped	-0.0044***	-0.0035***	-0.0026***	-0.0009***
	(0.0003)	(0.0003)	(0.0003)	(0.0003)
Possibly Un-mapped	-0.0051***	-0.0063***	-0.0015***	-0.0013***
	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Log Loan Amount		0.0279***	0.0330***	0.0097***
		(0.0002)	(0.0002)	(0.0002)
Mortgage Applications	✓	✓	✓	✓
Quarter-Year FEs	✓	✓		✓
County-Quarter-Year FEs			✓	
Tract FEs				✓
Lender FEs				✓
Observations	15,157,069	15,156,894	15,156,894	15,156,894
R^2	0.06004	0.06307	0.07948	0.16711
Within R^2	0.05871	0.06174	0.05476	0.03376

Note: We estimate equation 1, above. The dependent variable is binary, taking the value of 1 if a loan is accepted by both parties (lender and borrower). Our variables of interest are binary and denote whether a property is un-mapped or possibly un-mapped. We include properties that are in 500 year or 100 year flood zones. We add additional controls and fixed effects in columns (2)–(4). We restrict our sample to conforming loans. Standard errors are clustered by census tract and year-quarter and shown in parentheses; *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

A.6 Response Heterogeneity by Lender Type

Table A.16: ***Interest Rates with Lender Types***

	Interest Rate			
	(1)	(2)	(3)	(4)
Un-mapped × Non-Bank	-0.0124***	-0.0084***	-0.0147***	0.0003
	(0.0014)	(0.0013)	(0.0013)	(0.0010)
Un-mapped × Local Bank	-0.0039	0.0102***	-0.0098***	0.0002
	(0.0029)	(0.0028)	(0.0027)	(0.0024)
Un-mapped × Large Bank	-0.0179***	-0.0149***	-0.0295***	-0.0033*
	(0.0021)	(0.0020)	(0.0020)	(0.0019)
Un-mapped	0.0295***	0.0180***	0.0131***	0.0005
	(0.0013)	(0.0012)	(0.0012)	(0.0009)
Possibly Un-mapped	0.0019***	0.0144***	-0.0018***	0.0014***
	(0.0005)	(0.0004)	(0.0004)	(0.0004)
Non-Bank	-0.0204***	0.0386***	0.0172***	0.0109***
	(0.0006)	(0.0006)	(0.0005)	(0.0025)
Local Bank	-0.0522***	-0.0580***	-0.0336***	0.0045**
	(0.0014)	(0.0014)	(0.0013)	(0.0018)
Large Bank	-0.1844***	-0.1357***	-0.1651***	-1.124
	(0.0009)	(0.0009)	(0.0009)	(8,660.6)
Log Loan Amount		-0.3114***	-0.3579***	-0.2601***
		(0.0008)	(0.0010)	(0.0009)
Mortgage Controls	✓	✓	✓	✓
Quarter-Year FEs	✓	✓		✓
County-Quarter-Year FEs			✓	
Tract FEs				✓
Lender FEs				✓
Observations	12,084,309	12,084,294	12,084,294	12,084,294
R^2	0.59009	0.62222	0.63910	0.73073
Within R^2	0.05557	0.12960	0.13067	0.08723

Note: We estimate equation 1, above. The interest rate is continuous and bounded between 2 and 8. We remove properties with extreme interest rates (above the 99th percentile or below the 1st percentile). Our variables of interest are binary and denote whether a property is un-mapped or possibly un-mapped. Non-bank is a dummy variable indicating that, based on its entity type code, the lender is not a bank nor credit union. Local bank is a dummy variable indicating that at least 40% of the lender's mortgage origination occurs in the county the property is located in. Large bank is a dummy variable indicating that the lender is an entity or a subsidiary of an entity regulated under the Large Institution Supervision Coordinating Committee (LISCC) program. We add additional controls and fixed effects in columns (2)–(4). We restrict our sample to conforming loans and exclude all properties that are covered by a flood zone map. Standard errors are clustered by census tract and year-quarter and shown in parentheses; *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

Table A.17: ***Property Values with Lender Types***

	Log Property Value		
	(1)	(2)	(3)
Un-mapped × Non-Bank	0.0161***	0.0064***	0.0047***
	(0.0009)	(0.0007)	(0.0006)
Un-mapped × Local Bank	0.0354***	-0.0029*	-0.0010
	(0.0021)	(0.0017)	(0.0014)
Un-mapped × Large Bank	0.0109***	-0.0191***	-0.0100***
	(0.0019)	(0.0016)	(0.0013)
Un-mapped	-0.0362***	-0.0277***	-0.0227***
	(0.0008)	(0.0007)	(0.0006)
Possibly Un-mapped	0.0513***	-0.0169***	-0.0119***
	(0.0005)	(0.0003)	(0.0002)
Non-Bank	0.0954***	0.0046***	-0.0114***
	(0.0004)	(0.0003)	(0.0016)
Local Bank	-0.0261***	-0.0072***	-0.0059***
	(0.0011)	(0.0008)	(0.0011)
Large Bank	0.1402***	0.0184***	-0.4310
	(0.0008)	(0.0006)	(8,197.5)
Mortgage Controls	✓	✓	✓
Quarter-Year FEs	✓		✓
County-Quarter-Year FEs		✓	
Tract FEs			✓
Lender FEs			✓
Observations	13,403,882	13,403,882	13,403,882
R^2	0.41535	0.62653	0.73750
Within R^2	0.40372	0.32131	0.20253

Note: We estimate equation 1, above. Our outcome variable is the natural logarithm of the property value the lender uses when making the lending decision. Our variables of interest are binary and denote whether a property is un-mapped or possibly un-mapped. Non-bank is a dummy variable indicating that, based on its entity type code, the lender is not a bank nor credit union. Local bank is a dummy variable indicating that at least 40% of the lender's mortgage origination occurs in the county the property is located in. Large bank is a dummy variable indicating that the lender is an entity or a subsidiary of an entity regulated under the Large Institution Supervision Coordinating Committee (LISCC) program. We add additional controls and fixed effects in columns (2)–(3). We restrict our sample to conforming loans and exclude all properties that are covered by a flood zone map. Standard errors are clustered by census tract and year-quarter and shown in parentheses; *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

A.7 Lender Responses to Jumbo Loans

In this section, we showcase key regressions from the paper, making use of non-conforming (i.e., jumbo) loans instead of conforming loans. Jumbo loans are larger and – most importantly – not eligible for public securitization. As we can see, the effects discussed above broadly hold. Jumbo loans are less likely to be originated – because lenders are aware of the un-mapped flood risk – and charge a higher interest rate on average, as lenders manage this risk. However, the effects are significantly less pronounced than for conforming loans. In fact, if we include census tract and lender-type fixed effects, we can see that both key results are insignificant. If we interact our variables of interest with lender-type dummies, we can see that none of the various lender types respond very differently (not reported for brevity). Similarly, while we find a reduction in the value of jumbo properties, we find that this reduction in value is below 1% if we include all controls and fixed effects (results not reported for brevity). Among conforming loans on the other hand, the reduction is much larger (over 3%).

At first, this may seem somewhat surprising. After all, jumbo loans are more likely to remain on a bank's balance sheet than conforming loans. A bank has greater incentives to manage its risk exposure. However, the average income as well as the average loan-to-income ratio is significantly higher among jumbo loan borrowers than conforming borrowers. In fact, average income is 3.5 as high and LTI is over 10% higher. As such, the lender may perceive the borrowers as being more likely to weather a negative financial shock from a flood, even without insurance. Similar phenomena – with wealthier households less affected – could be observed in Blickle and Santos (2021).

Table A.18: ***Interest Rates for Jumbo Loans***

	Interest Rate			
	(1)	(2)	(3)	(4)
Un-mapped	0.0216***	0.0175***	0.0026	0.0018
	(0.0020)	(0.0019)	(0.0020)	(0.0018)
Possibly Un-mapped	-0.0130***	-0.0037**	-0.0030**	0.0019
	(0.0015)	(0.0014)	(0.0015)	(0.0014)
Log Loan Amount		-0.1972***	-0.0965***	0.0350***
		(0.0028)	(0.0035)	(0.0033)
Mortgage Controls	✓	✓	✓	✓
Quarter-Year FEs	✓	✓		✓
County-Quarter-Year FEs			✓	
Tract FEs				✓
Lender FEs				✓
Observations	797,477	797,477	797,477	797,477
R^2	0.57633	0.58001	0.60550	0.71186
Within R^2	0.07360	0.08163	0.06073	0.03949

Note: We estimate equation 1, above. The interest rate is continuous and bounded between 2 and 8. We remove properties with extreme interest rates (above the 99th percentile or below the 1st percentile). Our variables of interest are binary and denote whether a property is un-mapped or possibly un-mapped. We add additional controls and fixed effects in columns (2)–(4). We restrict our sample to non-conforming (jumbo) loans and exclude all properties that are covered by a flood zone map. Standard errors are clustered by tract and the year-quarter or origination; *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

Table A.19: ***Origination for Jumbo Loans***

	Loan Originated			
	(1)	(2)	(3)	(4)
Un-mapped	-0.0070***	-0.0098***	-0.0007	-0.0007
	(0.0011)	(0.0010)	(0.0011)	(0.0011)
Possibly Un-mapped	-0.0098***	-0.0031***	-0.0016*	-0.0017*
	(0.0008)	(0.0008)	(0.0009)	(0.0009)
Log Loan Amount		-0.1389***	-0.1914***	-0.1808***
		(0.0015)	(0.0019)	(0.0020)
Mortgage Controls	✓	✓	✓	✓
Quarter-Year FEs	✓	✓		✓
County-Quarter-Year FEs			✓	
Tract FEs				✓
Lender FEs				✓
Observations	879,145	879,145	879,145	879,145
R^2	0.02476	0.03727	0.07249	0.14183
Within R^2	0.02053	0.03309	0.03365	0.03092

Note: We estimate equation 1, above. Our outcome variable is binary and tracks whether a mortgage is originated (i.e. accepted by both borrower and lender) and our variables of interest are also binary and denote whether a property is un-mapped or possibly un-mapped. Low income tract and high income tract are dummy variables that denote . We add additional controls and fixed effects in columns (2)–(4). We restrict our sample to non-conforming (jumbo) loans and exclude all properties that are covered by a flood zone map. Standard errors are clustered by tract and the year-quarter or origination; *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

A.8 Flood Insurance In and Out of Flood Zones

A central feature of the National Flood Insurance Program (NFIP) is the mandate that mortgages for homes within the Special Flood Hazard Area (SFHA, i.e., the 100-year flood zone) must purchase flood insurance, whereas homes outside of this area are not required to purchase flood insurance. This difference in the insurance mandate motivates our analysis of what we refer to as "un-mapped" properties. It is true that even within the SFHA flood risk is not always fully insured, as the NFIP has coverage limits of $250,000. It is also not infrequent that mortgage borrowers lapse in their coverage when both the borrower fails to renew their flood insurance policy and the mortgage servicer fails to force-put flood insurance coverage. Still, as we demonstrate in this section, there is a considerable take-up gap between areas in the SFHA and outside the SFHA in the NFIP—even when we condition on underlying flood risk.

In this section, we explore this difference in insurance take-up and discuss the risk implications for areas outside of the 100-year flood zone. To study this, we use FEMA's NFIP redacted policies data, which contains the near universe of NFIP policies and combine this with property-level data from CoreLogic matched to FEMA flood zones. This allows use to determine both the number of properties and active policies within a given flood zone and level of geography. Nationally, we find that 19.89% of single-family residential properties within the 100-year flood zone have flood insurance, whereas just 1.03% of single-family residential properties outside the 100-year flood zone have flood insurance through the NFIP. Our own estimates are comparable to others in the literature. In a similar exercise with property-level data from First Street Foundation, Sastri (2021) finds NFIP take-up rates inside and outside of the 100-year flood zone of 25.8% and 6.65% respectively in Florida. These correspond well with estimates in our own data of 33.58% and 4.55% for Florida.

Figure A.10 displays county-level flood insurance take-up rates inside and outside 100-year flood zones for the contiguous U.S.. Overall, we see that even in areas where flood insurance is required for mortgaged properties, insurance take-up is often low. Take-up in the mandatory flood zone is higher in coastal and riverine areas—suggestive evidence that homeowners are either more willing to maintain flood insurance in areas where flood risk is more salient or are sensitive to within-flood zone variation in flood risk. As low as flood insurance take-up is inside the 100-year flood zone, it appears to be almost universally lower outside the 100-year flood zone. There are some areas, particularly in Florida and around the mouth of the Mississippi River, where this voluntary insurance take-up is non-trivial. Apart from these select areas—where flood risk is certainly salient—flood insurance take-up outside the 100-year flood zone is rare.

A clear caveat is that properties inside the 100-year flood zone are higher-risk by design than properties outside the 100-year flood zone, so differences in flood insurance take-up rates cannot speak to un-mapped areas being more or less insured than mapped areas. To get at this question of uninsurance, we incorporate CoreLogic flood risk measures. These allow us to condition policy take-up on the underlying flood risk. We take the crosssection of all CoreLogic properties and

Figure A.10: *County-Level Flood Insurance Take-Up Rates*

(a) *Take-Up Inside Flood Zones* **(b)** *Take-Up Outside Flood Zones*

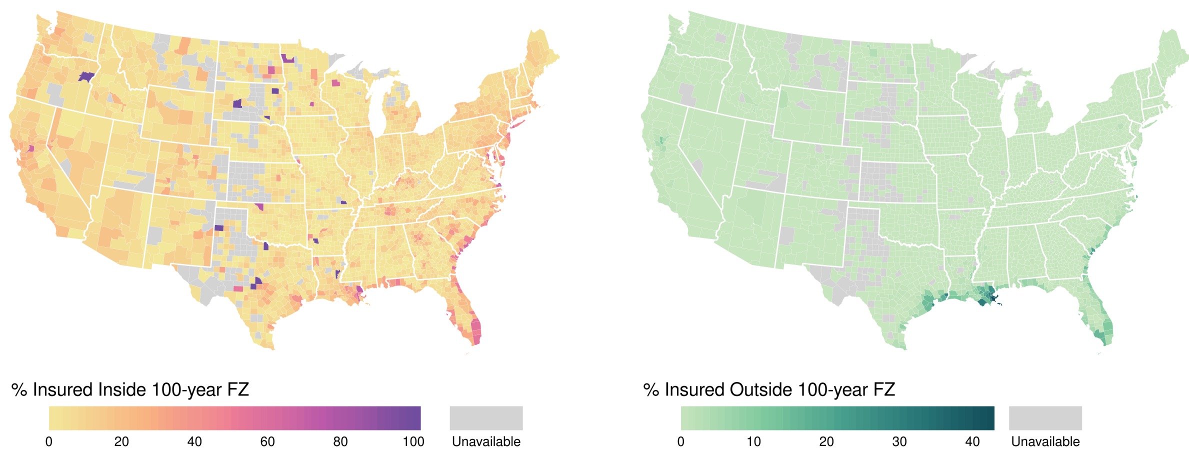

Note: The figure displays the rate of flood insurance take-up rate inside the official SFHA (i.e., 100 year flood zones) and outside the SFHA. Take-up rates are computed as the number of flood insurance policies divided by the total number of residential properties in each county-flood zone pair. We derive the number of flood insurance policies in each county-flood zone pair from the NFIP redacted policies data by counting all flood insurance policies for single-family residential properties in the county-flood zone pair in effect as of December 31, 2023. We derive the number of residential properties in each county-flood zone pair by counting all residential properties in the CoreLogic data by county-flood zone pair. Take-up rates are unavailable for some counties where there are no CoreLogic properties in our dataset within the given county-flood zone. Small samples in counties with low populations result in extreme values; we winsorize ratios above 100 in the figure.

compute the mean AAL for each census tract-flood zone pair. We cannot cleanly match CoreLogic properties with properties in the NFIP policy data, but because we use tract-flood zone-level flood risk measures, it is not material for us to match properties-to-properties. Instead, we assign NFIP policies to properties in CoreLogic such that the number of properties in our CoreLogic property data with polices in each tract-flood zone matches what we observe in the NFIP policies data. Then, for property i located in census tract c and flood zone f, we run

$$\mathbb{1}(\text{Policy})_{icf} = \alpha + \beta\, \mathbb{1}(\text{Outside FZ})_{icf} + \gamma \text{ Flood Risk Score}_{cf} + \varepsilon_{icf} \tag{A.3}$$

where $\mathbb{1}(\text{Policy})_{icf}$ is a dummy indicating the property has a policy, $\mathbb{1}(\text{Outside FZ})_{icf}$ is a dummy indicating the property is outside the 100-year flood zone, and Flood Risk Score$_{cf}$ is the mean CoreLogic AAL Composite Flood Risk Score over the census tract-flood zone pair.

Table A.20 displays results for this regression. We see that even when conditioning on census tract-floodzone risk, flood insurance take-up is 15.65 percentage points lower outside of the flood zone than inside the flood zone. We still have the caveat that the flood risk measures we use here are not at the property level, so this estimate might overstate the true size of the insurance gap between mapped and un-mapped properties. For this reason, we do not wish to emphasize the point estimate itself. We do believe though that this is suggestive that the insurance mandate is binding. That is, because it is not required to get a mortgage, flood insurance take-up outside the flood zone will be weakly less than inside the flood zone—even conditional on underlying flood risk.

Table A.20: *Insurance Gap*

	$\mathbb{1}$(Policy)	
	(1)	(2)
$\mathbb{1}$(Outside FZ)	-0.2462***	-0.1565***
	(0.0037)	(0.0040)
Flood Risk Score		0.0017***
		(4.26×10^{-5})
Constant	✓	✓
Observations	107,370,854	107,370,854
R^2	0.11554	0.13134
Adjusted R^2	0.11554	0.13134

Note: We estimate equation A.3, above on a cross-section of properties from CoreLogic. The dependent variable is a dummy indicating whether or not a property has an active flood insurance policy from the National Flood Insurance Program (NFIP). Standard errors (in parentheses) are clustered by census tract; *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.