Finance and Economics Discussion Series

Federal Reserve Board, Washington, D.C.
ISSN 1936-2854 (Print)
ISSN 2767-3898 (Online)

Do Banks Price Flood Risk in Mortgage Origination: Evidence from a Natural Experiment in New Orleans

David M. Arseneau and Gazi I. Kara

2025-081

Please cite this paper as:
Arseneau, David M., and Gazi I. Kara (2025). "Do Banks Price Flood Risk in Mortgage Origination: Evidence from a Natural Experiment in New Orleans," Finance and Economics Discussion Series 2025-081. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2025.081.

Do Banks Price Flood Risk in Mortgage Originations? Evidence from a Natural Experiment in New Orleans*

David M. Arseneau and Gazi I. Kara

Federal Reserve Board

September 8, 2025

Abstract

This paper uses a large-scale redrawing of flood zone maps for the City of New Orleans in 2016 to identify how banks respond to changes in perceived flood risk in residential mortgage origination. Using geo-coding, we separate loan-level data on mortgage originations into treatment versus control groups based on how individual properties were affected by the map changes. We find banks charged interest rates that were roughly 6 basis points higher for mortgages on treated properties that were removed from the special floods zones as a result of the map changes. In addition, lower loan-to-value ratios for mortgages on these properties suggest that banks also required higher downpayments. Both effects are temporary, lasting under two years. Further analysis using flood insurance claims data following a major flooding event in 2017 suggests the temporary nature of these effects may reflect learning by banks about the true extent of flood risk and insurance take-up following the map changes.

Keywords: FEMA Maps, Flood insurance, Mortgage lending
JEL Classifications: G21; Q54; R3

*The views expressed in this paper are those of the authors and do not necessarily represent those of the Federal Reserve Board of Governors or anyone in the Federal Reserve System. We thank Audrey Selley, Marcus Dockerty, and Chelsea Hunter for excellent research assistance and Antonis Kotidis, Manuel Adolino, Hyeyoon Jung, and Bhavyaa Sharma for helpful comments. We thank conference and seminar participants at the Federal Reserve Board, Wesleyan University, Koç University, Central Bank of the Republic of Türkiye, IFABS Conference in Oxford, Southern Economic Association Annual Meeting in Fort Lauderdale, IEA World Congress in Medellin and EFMA Conference in Lisbon.

1 Introduction

The U.S. residential real estate market is supported by approximately $14 trillion in mortgage debt, making it the largest non-sovereign debt market in the world. A sizeable proportion of this activity is intermediated through the banking system. Accordingly, residential mortgages make up a substantial share of assets held on bank balance sheets and are a significant source of risk exposure. This paper contributes to the literature on how banks monitor and price risk by focusing on flood risk in the context of a unique natural experiment involving the large-scale redrawing of Federal Emergency Management Agency (FEMA) flood maps for Orleans Parish, Louisiana in 2016.[1] In doing so, we provide novel insights into banks' risk management practices and their ability to adjust loan pricing in response to changes in perceived flood risk.

Prior to the 2016 redrawing, a large portion of properties in the parish were located in special flood hazard areas (SFHAs), meaning that owners of these properties were required to purchase federal flood insurance if their mortgages government-backed or from a federally regulated lender. Following the completion of flood protection infrastructure in the wake of Hurricane Katrina, which devastated New Orleans in 2005, city officials successfully lobbied FEMA to reconsider the boundaries that delineate the SFHAs in a large-scale redrawing of the flood maps. The redrawn maps were announced in early-2016 and went into effect later that year. The map changes resulted in the removal of roughly half of all properties in the parish from the special floods zones, implying that mortgage borrowers for these properties were no longer required to purchase flood insurance when financing their homes. Existing literature has shown that when flood insurance is not mandated, homeowners have a difficult time assessing their vulnerability to flood risk and often simply choose not to purchase flood insurance or they let existing policies lapse.[2]

This paper empirically examines the extent to which banks altered mortgage lending terms in response to perceived changes in the risk profile of individual properties affected by these map changes. From the lenders' point of view, the map changes could have affected the risk profile of a treated property in two potentially competing ways. On the one hand, the fact that a given property was removed from the SFHA suggests that it ought to be less flood prone and hence less risky. This is true provided the map changes implemented by FEMA accurately reflect evolving flood risk. While this is likely the case at the aggregate level, given the scale of the changes, the new flood zone boundaries might not accurately reflect evolving flood risk at the level of an individual property. Indeed, press commentary in response to the changes at the time of implementation

[1] New Orleans is a consolidated city-parish, so Orleans Parish can be used interchangeably with the City of New Orleans. Moreover, in the State of Louisiana a parish is analogous to what would be a county in other states, for New Orleans county, city, and parish all coincide.

[2] See, for example, Atreya, Ferrerira, and Michel-Kerjan (2015), Brody, et al (2017), Cannon et al (2020), Hung (2009), Lave and Lave (1991), Li and Landery (2018), Michel-Kerjan and Kousky (2010), and Savitt (2017).

pointed to the seemingly arbitrary nature with which some properties were removed from the flood zones while others were not. For example, in some cases two properties across the street from one another ended up falling on different sides of the border delineating the new flood zone from the old one. In these cases, whether an individual property fell on one side of the border or the other likely had no material effect on its true underlying flood risk. All told, a mortgage on a treated property might have been less prone to flood risk, depending on the accuracy of the map changes at the level of an individual property. On the other hand, removing the flood insurance requirement for mortgages on treated properties may have made them potentially riskier because collateral protection for these properties now depended on voluntary insurance take-up by the borrower. These two channels—flood risk exposure and insurance take-up—work in opposite directions and which one, if any, dominates is an empirical question.

We investigate this question using confidential supervisory data on bank mortgage originations in Orleans Parish and state-of-the-art geo-coding techniques. The granular nature of our data allow us to assess how the map changes affected loan terms offered by the largest U.S. banks at the level of an individual property. Using the address field in the mortgage data, we divide individual properties into a treatment group—those that were in the SFHA prior to the map change, but were removed from the SFHA after the redrawn maps were released—and a control group—those that were in the SFHA both before and after the revision. Our analysis compares interest rates and loan-to-value ratios for treated versus control properties before and after the map changes using a dynamic differences-in-differences specification that controls for individual mortgage characteristics as well as bank-, time-, and census tract-fixed effects.

Our results show that interest rates for mortgages originated on treated properties were roughly six basis points higher relative to control properties, but this effect is temporary and only lasted for a period of about one to two years. The magnitude of the interest rate effect is in line with results found in other studies on pricing flood risk in mortgage lending (for example, Nguyen, et. al., 2022) and is economically meaningful. Compared to more traditional risk factors, the estimated effect is on par with a roughly 30 point reduction in a borrowers' FICO score for a newly originated mortgage. In addition, we find loan-to-value (LTV) ratios for treated properties temporarily declined by about 2 percentage points relative to control properties, suggesting that banks also responded by requiring larger down payments—roughly $6,600 for the average treated property. As with the interest rate effect, this response was also short-lived. These results are consistent with an interpretation that banks viewed mortgages on treated properties as riskier perhaps owing to the removal of the flood insurance requirement. A final point is that a major factor motivating city officials to lobby FEMA for the map changes in the first place was to ease the financial burden on home buyers associated with the mandatory flood insurance premiums. Based on our results, we estimate that, at least initially, roughly one-quarter of these savings were transferred back to banks in the form of higher interest rates and larger down payments to compensate for the additional risk associated with dropping the insurance requirement.

We provide additional evidence on the role of flood risk and insurance coverage in driving our results by splitting the sample into properties that were located in close proximity to the newly formed border between the new and the old flood maps and those that were farther away. Our results appear to be driven by properties located close to the border, despite the fact that there is no discernible difference in flood risk between these two sets of properties. More generally, our main findings are robust to a number of alternative fixed-effect specifications and to introducing dynamic controls to allow the effect of loan characteristics to vary over time. They are also robust to directly controlling for demographic factors, such as race, education, and median income, within a census tract as well as a proxies for flood risk as measured by ground-level or first floor elevation of the individual property.

To explain the temporary nature of the results, we examine claims data from the National Flood Insurance Program (NFIP) following a major flooding event in New Orleans in 2017—the year after the map changes were implemented. Severe flooding in August of that year caused widespread flooding in the city that took over fourteen hours to drain. This episode is informative because it offered direct insight into flood risk at the property level—it is widely held that the most common predictor of future flood risk is whether or not the property has flooded in the past (Rhodes and Besbris, 2022)—and it also revealed information about the extend of voluntary flood insurance take-up for owners of treated properties. Using claims data at the Census block level for the City of New Orleans, we find no discernible difference in either flood risk or insurance take-up for flood insurance between majority treated neighborhoods and others. A plausible explanation for the temporary nature of our results is that lenders learned from the flooding event that many owners of treated properties voluntarily retained their flood insurance.

A key takeaway from our findings is that banks appear to be quite attuned to changes in perceived risks associated with flooding and actively price them into residential mortgage originations. That said, this behavior may be unique the New Orleans real estate market given its extensive history of flooding (see Horowitz, 2020). To the extent that this behavior might not carry over to other markets, we can not rule out the possibility of broader mis-pricing of flood risk in the residential real estate market. That said, the experience of New Orleans suggests that any possible broader mis-pricing is likely to be corrected as better information becomes available allowing for more accurate assessments of flood risk at the property level.

In terms of related literature, Sastry (*forthcoming*) examines how mortgage lending terms are affected by changes in FEMA flood maps, focusing on flood map expansions in the state of Florida.

She examines county-level shocks and finds that banks manage flood risk by reducing LTV ratios when flood insurance limits bind, but they do not adjust mortgage interest rates. Our exercise differs in that we focus on a rare instance of a large-scale contraction in the flood maps and our treatment effect is at the individual property, rather than county level. Nguyen, Ongena, Si, and Sila (2022) study the effect of sea level rise on mortgage lending rates at the zip code level using mortgage data on coastal states in the continental USA. These authors find that banks charge roughly 10 basis points higher interest rates on loans in flood-prone areas. Although the quantitative magnitude of our results is quite similar, our identification strategy is very different. In a closely related study, Blickle, Perry, and Santos (2024) find that lenders are less willing to originate mortgages and charge higher rates for lower LTV loans that face un-mapped flood risk. Also relatedly, Blickle and Santos (2022) show that changes in FEMA flood maps lead to a reduction in mortgage lending by banks both at the extensive and intensive margins.

The remainder of this paper is organized as follows. The next section discusses the institutional background on FEMA flood maps and discusses how we use this episode for identification. Section 3 presents the data and discusses the econometric methodology used in the analysis. Section 4 presents the main results and discusses a series of robustness tests and Section 5 investigates the temporary nature of our results. Finally, Section 6 concludes.

2 Institutional Background and Identification Strategy

We present some institutional details on the role that FEMA flood maps play in the National Flood Insurance Program (NFIP) and briefly discuss the incentives to keep these maps up to date to accurately reflect changing flood risk. We then describe the large-scale map revisions that took place in the City of New Orleans in 2016. Finally, we explain how we use those revisions for identification in our empirical analysis.

2.1 Flood Insurance, Flood Maps, and the Map Updating Process

Flood insurance is critical for the ability of a homeowner to bounce back from a disaster and hence it matters for the credit worthiness of mortgage borrowers. Most homeowners insurance policies do not cover flood damage. Instead, homeowners that face flood risk can insure themselves through separate flood insurance policies that cover damage to buildings and property in the event of flooding. In order to facilitate the provision of flood insurance, FEMA manages the NFIP, which helps deliver flood insurance to the public through a network of private insurance companies. The NFIP makes federal flood insurance available to anyone who lives in one of the large number of communities nationwide that opt to participate in the NFIP. For many homeowners who reside in participating communities the decision to purchase flood insurance is optional and evidence suggests that uptake of flood insurance is minimal when it is not required (e.g. Kousky et al, 2020). However, homes and businesses in areas deemed to be high-risk are required to purchase flood insurance through the NFIP when obtaining a government-backed mortgage or from a federally regulated lender.

As part of its responsibilities in managing the NFIP, FEMA plays a critical role in determining who is and is not required to purchase flood insurance. FEMA does this by creating flood maps which assess an individual community's flood risk at a level of geographic detail that is sufficiently granular to identify individual properties on the maps. More generally, areas within communities are divided areas into three broad flood risk categories: (1.) Minimal Flood Hazard Areas that have a less than 0.2% change of flooding on an annual basis; (2.) Moderate Flood Hazard Areas with a 0.2% to 1% risk; and (3.) Special Flood Hazard Area (SFHA) that have a 1% or greater annual flood risk. Properties that fall into the SFHAs are the ones deemed high-risk for the purposes of the NFIP.

The flood maps are important to property owners, mortgage lenders, and insurance companies alike to help understand the risk profiles of residential real estate properties and their accuracy is crucial to the NFIP. The maps have additional implications for homeowners beyond these considerations. For example, placing a property in an SFHA not only increases the cost of financing to the borrower through the insurance requirement, but it also can depress its value relative to properties outside SFHAs (see, for example, Bin, Kruse, and Landry, 2008; Ortega and Taşpınar, 2018; and Shu et al., 2022) .

As important as they are, the flood maps are understandably controversial. There is evidence that the maps are drawn using outdated methods (Wing et al., 2018; Prelle, 2019; Kousky, Palim and Pan, 2020; and Weill, 2023). Moreover, flood risk changes over time and because the maps are often not updated for long periods, they are widely regarded as out of date. Acknowledging this, FEMA started a process of map modernization in the early 2000s to bring the maps up to date on a nationwide scale. A process was put in place where the NFIP and FEMA work together with communities to identify and map flood risk on an ongoing basis. However, this process is slow and takes about three years start to finish for the average community. FEMA works with the community to draw up an initial draft of the updated maps and then releases them to the community for comment. Given their importance to homeowners, the initial drafts receive considerable scrutiny. There is a formal process in place for dealing with challenges to the maps, which happens frequently, and these challenges can be filed by individuals or by entire communities.

Overall, the process of mapping flood zones is very sticky, making it difficult to update them in a timely fashion. By 2014, roughly a decade after the modernization process began, only about one half of the process was completed.

2.2 Post-Katrina Flood Map Changes in New Orleans

The City of New Orleans is a well known example of a case in which the FEMA flood maps were redrawn at the community-level at large scale. For perspective, Figure 1 shows Orleans Parish, broken out by U.S. census tracts.

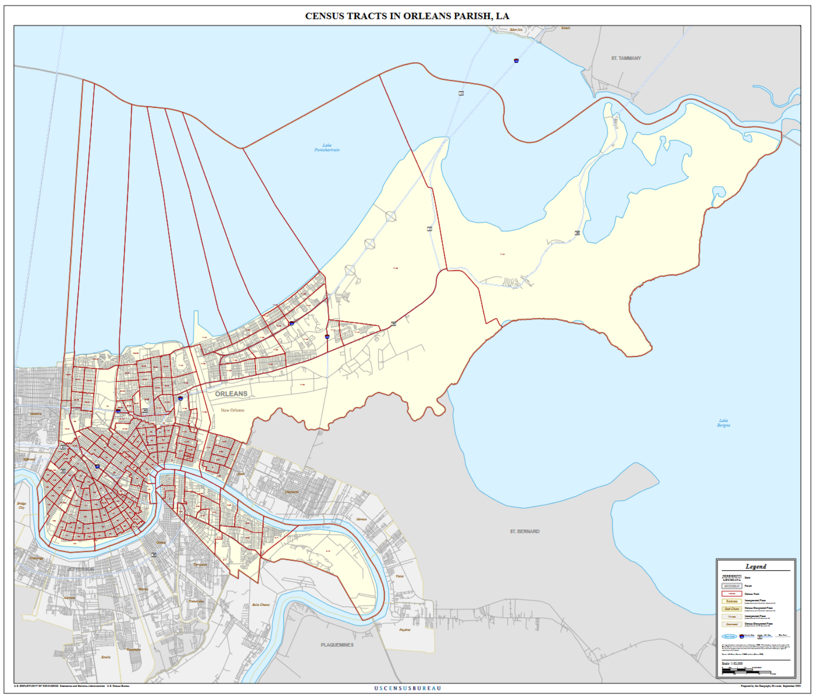

Source: Census.gov

Figure 1: The City of New Orleans, broken out by Census Tracts

Large-scale FEMA flood map revisions typically involve an expansion of flood zones, which results in the addition of new properties to the SFHA. In contrast, the New Orleans episode was one of the only instances in which the map revisions resulted in the removal of a large number of properties from the SFHA, meaning that a large number of properties which previously required flood insurance no longer did after the map revisions. This makes the New Orleans episode unique. After Hurricane Katrina, the Federal Government financed the building of a $14.5 billion flood protection system in Orleans Parish and surrounding areas, which was completed in 2013. Following the completion of the project, city officials began lobbying FEMA to reduce the size of the SFHAs in response to the new flood protection system.

Map changes were announced by FEMA in January 2016 through an interactive risk map provided on FEMA's website where public could check the updated flood zoning of each property and see the scale of changes for the entire city. A snapshot of this website, which no longer exists, is provided in the appendix. The announced changes took effect on October 1, 2016. Figure 2 shows the old flood maps in the left panel with the SFHA zone highlighted in red, which we created by digitizing the historical flood maps provided in PDF format on FEMA's website using the ArcGIS software. The panel on the right shows the updated maps with the new SFHA highlighted in blue, which are readily available in shape file format on FEMA's website.[3] The overall size of the SFHA in the updated maps is considerably smaller, reflecting the new flood protection infrastructure, and resulted in the removal of about half of the properties in the old maps from the SFHA.[4]

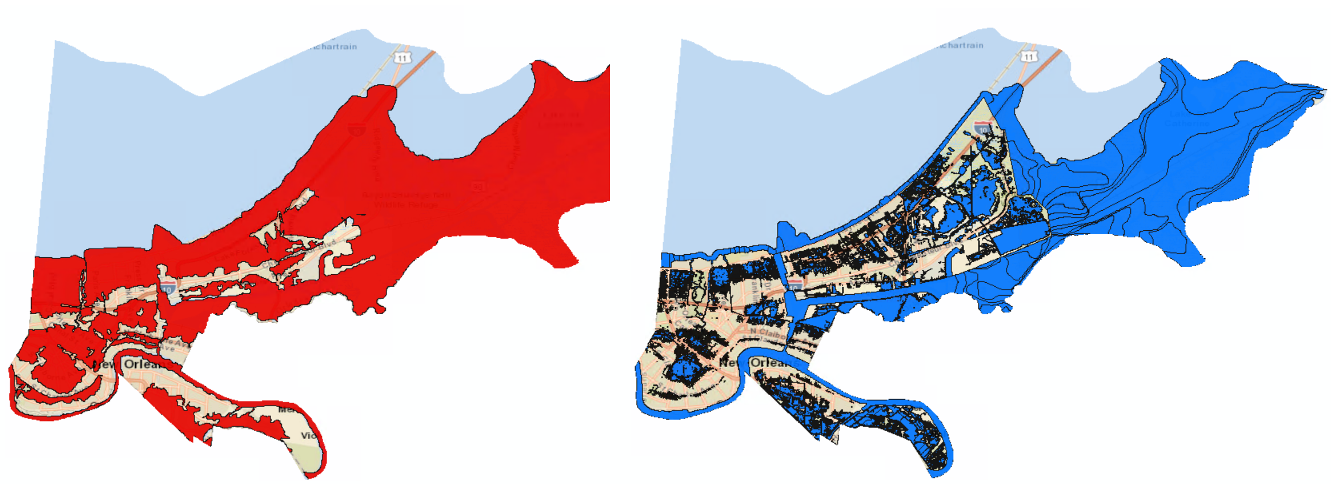

(a) Historic Flood Map

(b) Current Flood Map

Figure 2: FEMA SFHA Flood Maps for New Orleans, Old versus New Vintages

Figure 3 layers the old new map on top of the old one. Areas in which the two vintages of maps overlap (i.e., was in an SFHA in old map—red in the left panel of Figure 2—and remains in an SFHA in the new vintage—blue in the right panel of Figure 2) are shown in purple. Areas which were in the SFHA in the old maps, but were removed from the new ones remain red. Finally, because very few properties were newly introduced into the flood zone, there is no noticeable blue area in the combined map.

2.3 Identification Strategy

We use the New Orleans flood map revisions to separate individual properties in the parish into two different types of properties illustrated in Figure 4, which zooms in on a section of the overlaid maps. The dot labeled 1 in the purple portion of the Figure represents a control property that

[3]These files can be found here by searching for New Orleans: https://msc.fema.gov/portal/home.

[4]The large blue area that remains in the flood zone in the northeast of the city contains a wildlife refuge and is uninhabited.

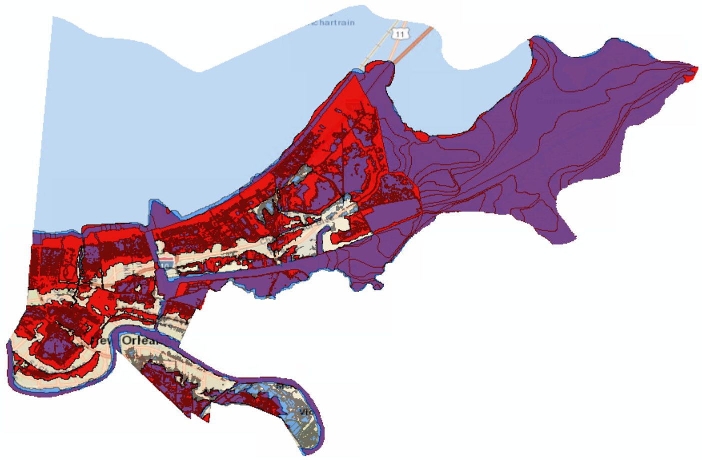

Figure 3: Historic and Current Flood Map Overlay

was in the flood zone prior to the map change and remained in the flood zone after the change. In contrast, the dot labeled 2 represents a treatment property that was previously in the SFHA, but was removed as a result of the map change.

Although the map changes were expected by the public, the exact shape of the new SFHAs took New Orleans residents and experts by surprise and started a heated debate in early-2016. As can be seen in Figure 4, even with in the same block, some properties that were previously required to purchase flood insurance were required to continue doing so (the purple areas), while their neighbors—in some cases, right across the street—were now exempt (the red areas). In other words, at the local neighborhood level many of these map changes seemed as random. An article published in the New York Times in June 2016 encapsulated the debate by questioning FEMA's methodology for generating the flood maps and objecting to the removal of many properties from SFHAs.[5]

In the analysis that follows, we use these two types of properties to tease out the impact of the map changes on residential mortgage terms of new originations. Additionally, we use the local randomness of these map changes to investigate whether distance from the new flood map borders have any effect on our results.

[5]See https://www.nytimes.com/2016/06/01/opinion/new-orleans-new-flood-maps-an-outline-for-disaster.html

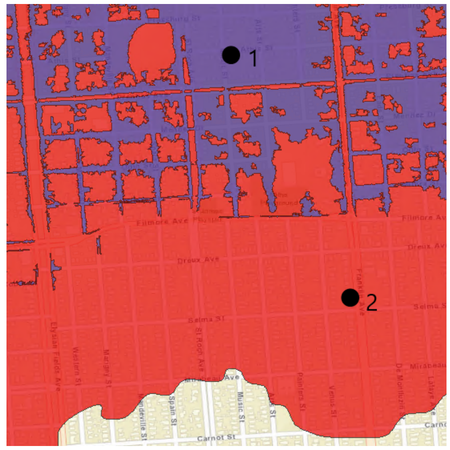

Figure 4: Illustration of treatment versus control properties in Orleans Parrish

3 Data and Methodology

Our analysis combines geographic information from the two different vintages of FEMA flood maps for New Orleans with loan-level data on residential mortgages held on the balance sheets of the largest U.S. banks.

3.1 Data

The geographic data come from the two different vintages of the FEMA flood maps drawn for the City of New Orleans discussed in the previous section. The earlier vintage reflects the flood maps in place from the beginning of our sample for the mortgage data (January 2014) until the the re-drawn maps became official on September 30th, 2016. The later vintage reflects the redrawn flood maps in place immediately following the implementation of the new maps until the end of the sample (December 2019).

We interact these flood maps with property-level address data from the Cotality (formerly CoreLogic) Real Estate Data for all residential homes in New Orleans through a process called geo-coding. Geo-coding is a method of converting the specific property address for a given residential property in the Cotality data set into latitude and longitude coordinates so that it can be placed on a map with a high degree of precision. We geo-code every residential property in New Orleans and place each on an overlay of the current vintage of the flood map over top of the earlier vintage (Figure 3 in Section 2.2). The methodology is sufficiently accurate to allow us to differentiate two different properties that are right next to each other on the same street. This degree of accuracy allows us to identify each individual property as falling into one of four mutually exclusive categories: a property that was (1.) inside the old flood zone, but outside the new flood zone (treated properties); (2.) inside the old flood zone and remains inside the new flood zone (control properties); (3.) outside the old flood zone, but inside the new flood zone; and (4.) outside the old flood zone and remains outside the new flood zone.

Table 1 shows that about half of the properties in New Orleans remained unaffected by the map changes, including 15,263 properties that were originally inside a SFHA prior to the map changes and remained in an SFHA after (i.e., control properties) and 44,051 properties that were originally outside before and remained outside after. The remaining roughly one-half of properties (119,573 - 15,263 - 44,051 = 60,259 out of 119,573 total) did change status as a result of the map changes. Of those, the majority, or 59,436 properties, were in an SFHA prior to map changes and were removed following the map updates (i.e., treatment properties). The remaining 823 properties were outside an SFHA before the map changes but were pulled inside an SFHA after the changes. All told, the map changes resulted in a net decline of 58,613 properties that fell into an SFHA. This decline resulted in the share of properties that fell into a FEMA-designated SFHA dropping from a little over sixty percent prior to changes (74,699 out of 119,573) to under fifteen percent after the changes (16,086 out of 119,573).[6] The relatively large decline in the fraction of properties falling into the SFHAs as a result of the map changes is not surprising given the motivation behind the city-wide redrawing of the flood maps. Local officials lobbied to allow constituents to benefit from the development of public infrastructure aimed at improving flood protection by removing them from the FEMA flood zones.

Table 1: SFHA Status for Residential Properties in New Orleans, Before and After Map Changes

		After Map Change		
		Inside SFHA	Outside SFHA	Total
Before Map Change	Inside SFHA	15,263	59,436	74,699
	Outside SFHA	823	44,051	44,874
	Total	16,086	103,487	119,573

In our analysis, we focus only on the 74,699 properties classified as in either the treatment or control group and where the borrower used a first lien residential mortgage originated by a bank that files the FR-Y14M Schedule A.[7] The mortgage data are collected monthly at the level of

[6]Treated properties represent about one-half of all residential properties in the Cotality database (59,436 out of 119,573), consistent with news reports on the scope of the map changes. See for example, https://www.npr.org/2016/09/30/495794999/new-maps-label-much-of-new-orleans-out-of-flood-hazard-area.

[7]A full description of the data can be found on the website of the Federal Reserve Board of Governors at https://www.federalreserve.gov/apps/reportforms/reportdetail.aspx?sOoYJ+5BzDYnbIw+U9pka3sMtCMopzoV

the individual loan and offer extensive information on every residential mortgage loan held on the balance sheet of or serviced by all banks that submit Y-14 data to the Federal Reserve.[8] What makes the mortgage data in the Y14 are particularly well suited for our analysis because they contain both the property address and the interest rate for the mortgage.[9] Although the FR-Y14M covers mortgages for the entire United States, our analysis restricts the sample only to newly originated fixed-rate first-lien residential mortgages within the Parish of New Orleans between January 2014 and December 2019.[10] Our final sample consists of a total of 6,063 mortgages on 5,728 individual properties (335 properties in the sample had more than one mortgage during the sample period owing to an exchange transaction or refinancing), of which we have 4,697 mortgages on treatment properties and 1,366 mortgages on control properties. The reduction in sample size relative to the Cotality data reflects that roughly ninety percent of the treatment and control properties identified in Table 1 did not have a fixed-rate first-lien mortgage originated by a Y14 bank during the time frame of our analysis.

Table 2 presents some summary statistics for our final sample as well as the means of selected mortgage or borrower characteristics broken out separately for treated and control groups. The average mortgage in our sample is a $277,851 loan at a fixed interest rate of 4.18% with a 331 month (27.7 year) maturity and has an LTV ratio of 73.[11] A typical borrower has a 742 FICO score. Just under one half of the originations are refinanced mortgages and only a small faction are for investment properties. The majority of the mortgages in our sample are securitized at some point within 12 months of origination but they are still serviced by the reporting banks. About ten percent of the loans in the sample are jumbo loans, meaning the loan amount was above the conforming loan threshold set by FHFA that makes a mortgage eligible for purchase or securitization by Government Sponsored Entities (GSEs).[12] Finally, mortgage characteristics differ importantly by loan type. In our sample, just under three-quarters of the mortgages are conventional loans without mortgage insurance. Conventional loans with mortgage insurance make up about one-fifth of the sample and the reminder are comprised of Federal Housing Administration (FHA) residential

[8]Banks with $100 billion or more in consolidated assets are required to submit these data for regulatory purposes. During our sample period, there were about 30 banks filing the Y14-M every year, but not all of these banks originated mortgages in the City of New Orleans.

[9]No other data set on mortgages that we know of contains extensive coverage of both property address and interest rate. For example, the Cotality mortgage data has property addresses but interest rate coverage is sparse, whereas HMDA has interest rate information, but not the property address and only starts in 2018. Similarly, McDash also only has interest rate data.

[10]When cleaning the data, we drop a small number of mortgages for which commercial loan flag is set to 1 or loans with missing loan amount, interest rate or LTV ratio at origination.

[11]We do not include a small number of variable rate mortgages, but our results are robust to including them.

[12]The jumbo loan limit for Orleans Parish was $417,000 from the beginning of our sample up to and including 2016. It increased to $424,100 in 2017, $453,100 in 2018, and $484,350 when our sample ends in 2019.

Table 2: Summary Statistics for Mortgage Data

	All loans, (N = 6,063)					Mean by Group	
		Standard				Control	Treated
Variable	Mean	Deviation	p5	Median	p95	(N = 1,366)	(N = 4,697)
Treated Dummy	0.77	0.42	0.00	1.00	1.00	0.00	1.00***
Interest Rate	4.18	0.57	3.25	4.12	5.12	4.19	4.17
LTV ratio	73.35	18.84	37.00	78.00	98.00	75.46	72.74***
Securitized Loan	0.84	0.37	0.00	1.00	1.00	0.87	0.83***
Credit Score	742	50	652	752	806	743	742
Loan term (months)	331	66	180	360	360	335	329***
Refinance	0.42	0.49	0.00	0.00	1.00	0.40	0.42*
Investment Property	.1	.3	0.00	0.00	1.00	.09	.1
Jumbo Loan	0.10	0.30	0.00	0.00	1.00	0.07	0.11***
Original Loan Amount	$277,851	$190,815	$73,606	$244,000	$601,350	$264,435	$281,753***
Implied House Price	$399,764	$295,147	$112,069	$340,411	$890,707	$360,627	$411,147***
Loan type							
Conv. w/o Mort. Ins.	0.72	0.45	0	1	1	0.70	0.73**
Conv. with Mort. Ins.	0.17	0.37	0.00	0.00	1.00	0.20	0.16***
FHA Residential	0.08	0.28	0.00	0.00	1.00	0.08	0.09
VA Residential	0.02	0.14	0.00	0.00	0.00	0.02	0.02
CRA	0.01	0.09	0.00	0.00	0.00	0.01	0.01

Note: Asterisks indicate a significant difference in mean between treatment and control at the 1% (***), 5% (**), and 10% (*) confidence level.

loans, Veterans Administration (VA) residential loans, and Community Reinvestment Act (CRA) loans.

Our baseline analysis focuses on comparing mortgage characteristics in group one (inside the old flood zone, but outside the new flood zone), which we call the treatment group (4,697 mortgages, 77% of the sample), to group two (inside the old flood zone and remains inside the new flood zone), which we call the control group (1,366 mortgages, 23% of the sample). The last two columns of Table 2 compares mortgage originations that fall into the treatment and control groups, respectively. The unconditional means of the interest rate, borrower credit score, and the investment property flag are quite similar between treated and control groups. Treated properties have an average loan size that is about $17,000 larger relative to control properties and the average loan term is shorter by six months. Consistent with the larger loan size, we find the share of jumbo loans to be a touch higher amongst treated properties, however the average LTV ratio is about 3 percentage points lower (72.74 versus 75.46). The larger loan size and lower LTV ratios for treated properties suggests a meaningful difference in the implied house price valuation. Indeed, for a treated property the average implied house price is $411,147 compared to $360,627 for a control property, a difference of roughly $50,000, which is statistically significant at the 1% confidence level. Finally, mortgages for treated properties are less likely to be securitized and less likely to have a mortgage insurance requirement for a conventional loan.

Table 3: Mortgages: Treatment vs. Control, by Census Tract and Bank

	Full Sample	All Treated	All Control	Mixed
		Census Tracts		
# of Mortgages	6,063	2,140	40	3,883
Treatment	4,697	2,140	0	2,557
Control	1,366	0	40	1,326
# of Tracts	147	72	6	69
		Banks		
# of Mortgages	6,063	6	0	6,057
Treatment	4,697	6	0	4,691
Control	1,366	0	0	1,366
# of Banks	22	3	0	19

As will become clear in the analysis below, it is useful to provide a breakdown of the mortgage sample by census tract and bank. The top half of Table 3 shows that for census tracts, our sample covers mortgages in 147 different census tracts within the City of New Orleans.[13] Of those, 72 contain mortgages on *only treated properties* and 6 contain mortgages on *only control properties* based on mortgages in our final sample. All told, mortgages in these these all treated and all control census tracts make up about one-third of the total sample with the remaining two-thirds in 69 census tracts containing mortgages on properties in census tracts with a mix of both treated and controls. The bottom half of the table shows the same information for banks. Although, most of the 22 banks in our sample originated mortgages for a mix of both treated and control properties, three banks originated a total of six mortgages for only treated properties.

3.2 Methodology

To assess how the map changes affected mortgage lending terms, we estimate the following dynamic difference-in-differences model

$$Y_{ijt} = \alpha_j + \mu_t + \lambda_c + \sum_{\tau=2014}^{2019} \beta_\tau Treat_i \mathbf{1}_{\tau=t} + \delta \mathbf{X_{it}} + \epsilon_{ijt} \quad (1)$$

[13] There were 178 census tracts in New Orleans as of 2019.

where: Y_{ijt} is either the interest rate or the LTV ratio for loan i, by bank j, in time t; $Treat_i$ is a dummy variable that takes on the value of one if the property for mortgage i is in the treatment group; $\mathbf{1}_{\tau=t}$ is a full set of annual time dummies set to 1 if $t = \tau$, where $\tau \in [2014, 2019]$ denotes the annual time periods used in our baseline model; and, $\mathbf{X_{i,j,t}}$ is a vector of controls discussed below. The baseline regression allows for any combination of bank-specific fixed effects, denoted α_j, and time-specific fixed effects at the annual frequency, denoted μ_t, as well as census tract fixed effects, denoted λ_c.[14] Finally, $\epsilon_{i,j,t} \sim \mathcal{N}(0, \sigma^2)$ is an idiosyncratic error assumed to be normally distributed with mean zero and variance σ^2. Standard errors are clustered at the census tract level.

For all dependent variables, the common control variables include: the loan term (in months); the (log) loan amount; FICO credit score at origination; and dummies for refinance loans, loans for investment properties, and jumbo loans. For interest rate regressions, we additionally include a full set of controls for loan type with dummy variables for conventional loans without mortgage insurance, conventional loans with mortgage insurance, VA loans, and CRA loans, leaving FHA loans as the excluded category. We also control for the LTV ratio in these regressions. For LTV regressions, we include a dummy variable for conventional loans without mortgage insurance (leaving all other loan types that typically have higher LTV ratios in the excluded category) and control for the interest rate.

In this dynamic diff-in-diff specification, we interact a full set of annual time dummies ($\mathbf{1}_{\tau=t}$) with the treatment dummy. As a result, the β_τ coefficients estimate the difference in means of the dependent variable between treated and control properties in each time period, after controlling for census tract, bank, loan, and property characteristics. Note that we are not excluding any of the time dummies so the β_τ coefficients are not estimated relative to a particular time period. This choice allows us to be agnostic about the exact timing of the shock. As we discussed above, although the change in flood maps became effective on October 1st, 2016, the new maps were already made public in the first quarter of 2016. This setup also inherently allows testing the parallel trends assumption conditional on observables: we expect the difference in interest rates and LTVs between treated and control properties (captured by β_τ coefficients) to not be statistically different from zero before 2016 when exact map changes were not known.

[14] Specifications which include census tract fixed effects (including our baseline regression) rely entirely on mortgages from tracts that contain a mixture of both treated and control properties for identification. In Section 4.3.3 we explore using demographics rather than fixed effects to control for tract-level developments, allowing us to use information from tracts that are either all treatment or all control as well.

4 Results

We first present our main results on interest rates and LTV ratios for newly originated mortgages on treated properties and then turn to some robustness checks.

4.1 Main Results

Table 5 presents our baseline interest rate results. The columns differ only by the fixed effects used in the estimation. In the first column, we add only year fixed effects. Second column contains both year and census tract fixed effects, while the third column includes year and bank fixed effects. In the fourth column, we include all three sets of fixed effects, namely year, census tract and bank fixed effects, and focus on this more stringent specification as our baseline result. Results show that prior to 2016, interest rates for treatment and control properties, both of which were inside SFHA at the time, were not statistically different from one another. In 2016, the year the change was implemented, newly originated mortgages on treated properties incurred interest rates that were 5 to 6 basis points (bps) higher, on average, than properties in the control group and this difference is statistically significant at the 95% confidence level. However, the effect was short-lived with the point estimate declining to 3 to 5 bps in 2017 and becomes no longer statistically significant in the third and fourth columns. The coefficients remain insignificant across all specifications after 2017. The left panel in Figure 5 plots the coefficients on the interest rate response on an annual basis using 95% confidence intervals based on column 4.

The economic magnitude is meaningful when compared with other studies that examine mortgage pricing for flood risk. For instance, Nguyen, et al. (2022) find that mortgage rates are 10.2 bps higher in zip codes exposed to sea level rise.[15] Another way to think about the economic magnitude is to compare it to traditional risk factors in mortgage origination. Based on our regression results, a 100 unit decrease in FICO score leads to an increase in the interest rate of 20 bps. Cast in this light, our finding of treated properties being charged an additional 6 bps in 2016 is equivalent to 30 unit decrease in the FICO score of an average borrower (= 100/(20/6)). Finally, recalling that one motivation for the map changes was to save home owners money by removing the flood insurance requirement, a simple back-of-the-envelope calculation shows that by raising mortgage

[15] Other studies also find comparable effects on mortgage pricing in various other settings beyond flood risk. For example, Kara and Yook (2023) find that banks charge on average 6.3 bps higher rates on jumbo mortgages when there is a gubernatorial election in their headquarter states, Bhutta, Fuster, and Hizmo (2019) show that when borrowers apply to more than one lender in search of better terms, they obtain 7 bps lower rates, and Buchak et al. (2018) find that non-fintech shadow banks charge on average 2.4 bps higher interest rates relative to traditional banks. Outside of mortgages, Sharma (2025) finds that European firms with exposure to climate-related transition risks are charged a higher business loan rate by 12.7 bps.

interest rates 6 bps banks were able to divert roughly one-quarter of these estimated savings away from households.[16]

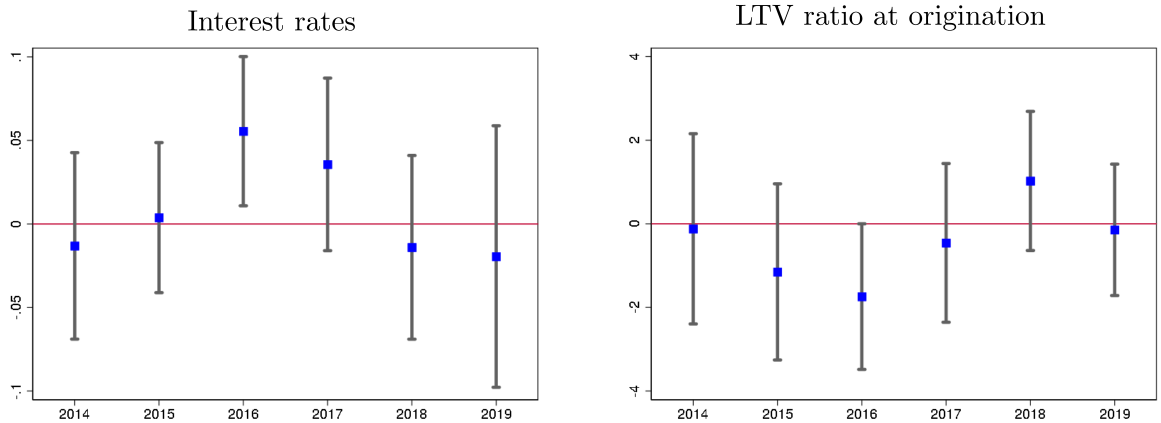

Figure 5: Baseline results: Interest rates and LTV ratios for treated vs control properties

A similar set of results on the response of the LTV ratio at origination are shown in Table 6. In the baseline specification presented in column 4 of Table 6 and shown in the right panel of Figure 5, LTVs on newly originated mortgages for treated properties were lower by 1.7 percentage points relative to control properties in 2016 and this effect is statistically significant at the 95% confidence level. As with the interest rate response, the effect is temporary and appears to dissipate more quickly: across all specifications in Table 6, coefficients on the interaction terms are insignificant after 2016. In terms of the economic significance, the lower LTV ratio corresponds to a higher down payment of roughly \$6,600 for the average treated property in the sample.[17] It is worth noting that the standard errors for the estimated coefficients in the LTV regression shown in the right panel of Figure 5 are generally larger when compared with the interest rate response reported in the left panel across all time periods. Our dynamic diff-in-diff estimation compares the conditional means of the dependent variables (interest rate of LTV ratio) for treated and control properties in every year, and hence, we have a relatively small number of observations. When combined with the large variance for the LTV ratios in our sample, as shown in Table 2, our parameter estimates for the LTV ratio are relatively more imprecise, making it harder to find statistically significant coefficients.

[16] We calculated the cumulative additional interest on a thirty-year fixed rate mortgage that carries 6 bps above the average mortgage for a treated property and compared it to the cumulative insurance premium for NFIP flood insurance in NOLA in 2016 at approximately \$650 per year. The additional interest expenses average roughly 25 percent of the cumulative savings from no longer having to pay the flood insurance premium.

[17] From Table 2, the average LTV ratio for a treated property is 72.74 and the average loan size is \$281,753. Together, these are consistent with an average implied house price of \$387,342. Holding the implied house price constant, lowering the LTV ratio by 1.7 percentage points to 71.04 means the bank is only willing to lend \$275,168, leaving the borrower with a higher downpayment of \$6,585.

4.2 Proximity to the Border

Taken together, our baseline results are consistent with an interpretation that banks viewed mortgages on treated properties as riskier, perhaps owing to the removal of the flood insurance requirement. We provide some additional evidence to support this interpretation by splitting the sample into properties that were located in close proximity to the newly formed border between the new SHFA and non-SFHA regions and those that were farther away. To the extent that our results are driven by the removal of the flood insurance requirement, we would expect to see the strongest effects in treated properties that are in close proximity to the border as these properties are least likely to have experienced a change in true underlying flood risk.

Using geospatial analysis tools in ArcGIS, we measure the distance of each of the treated and control properties to the nearest flood map border. With this distance measure in hand, we construct a new dummy variable, $Close_j$, that is equal to one for properties for which distance to the border is below the median of its distribution. We investigate whether our results are strongest in treated properties close to the border by breaking our sample into two using the $Close_j$ dummy and running our baseline regression with each of these sub-samples.

Results are presented in Table 7 and coefficients for the annual treatment effect from these regressions are depicted in Figures 6 and 7. As shown in the left panel of Figure 6, the statistically significant increase in interest rates for treated properties in 2016 is obtained by properties close to the border, but the results are not statistically significant for properties that are far from the new border, as shown in the right panel. It is also noteworthy that the estimates for properties close to the border are more precise relative to properties far from the border, which have much wider confidence bands.

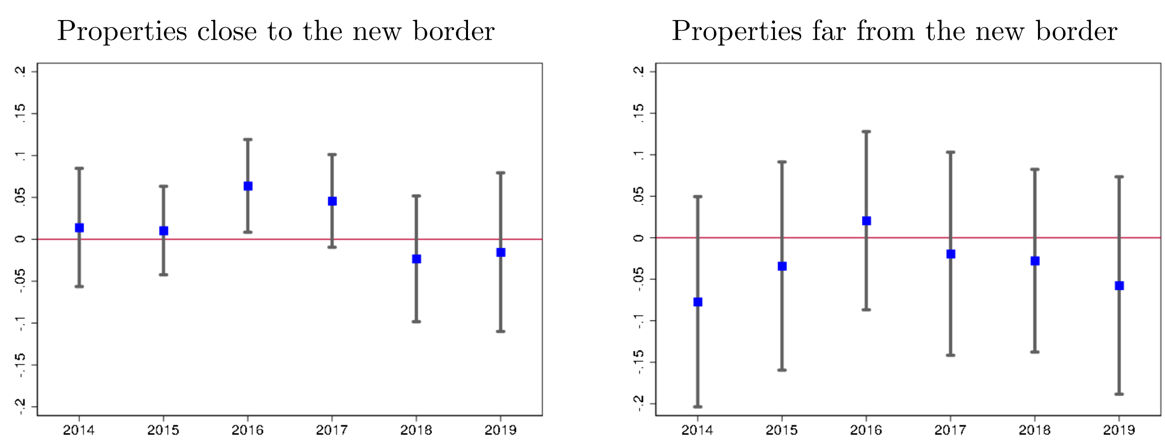

Figure 6: Interest rates for treated vs control properties based on distance to new flood map border

For LTV ratios, the results are insignificant for each half of the sample, as shown in Figure 7. Nevertheless, coefficients for 2016 and 2017 are both negative and have much smaller standard errors for close properties when compared to far properties. We interpret their statistical insignificance as a manifestation of our small sample issue in conjunction with the fact that, as discussed above, LTV ratios have a larger variance in our sample than interest rates. When we further split the sample in half, the statistical power in our dynamic estimation becomes particularly limited due to an even smaller number of control properties in each year, working against finding any significant effect. Therefore, we interpret the LTV results as providing only weak suggestive evidence for the hypothesis posed here.

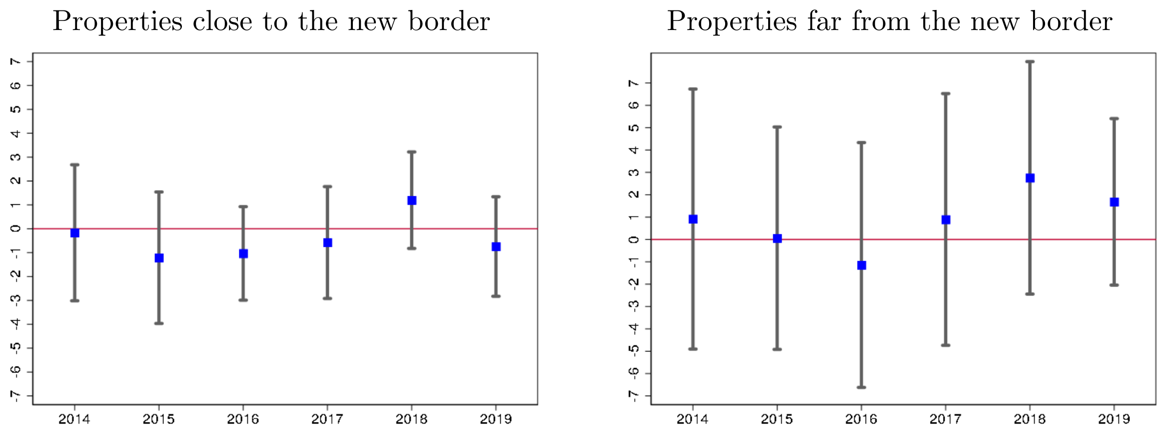

Figure 7: LTV ratios for treated vs control properties based on distance to new flood map border

4.3 Robustness Tests

We present several robustness tests by running our regressions at half-yearly frequency, adding dynamic controls, replacing census tract fixed effects with controls for demographics and income at the census tract level, and directly controlling for flood risk at the property level in the regressions.

4.3.1 Half-yearly Frequency

We first test the robustness of our results at running half-yearly frequency. For this exercise, we include interactions between treatment and half-yearly dummies (as opposed to annual dummies in our baseline regressions). Following our preferred specification, we include bank, census tract and time (half-yearly) fixed effects in these regressions. The results, shown in Figure 8, remain robust to analysis at a higher frequency. The left panel of Figure 8 shows that there was an increase in the interest rate on treated properties in the first and second half of 2016. The effect is insignificant in the first half of 2017 but becomes significant again in the second half and then goes away after 2017 as in our baseline regression with annual dummies.

The right panel shows that LTV results are also largely robust to conducting analysis at the half-yearly frequency. We find a drop in LTV ratio for treated properties at origination in the first half of 2016. This half-year effect is larger in magnitude and statistically more significant than the full year effect shown in our baseline results in the right panel of Figure 5. However, the effect losses significance at conventional levels in the second-half of 2016 despite being negative and meaningful in magnitude at around 2 percentage points. Again, the loss of significance could well be driven by the exacerbated small sample at the half-yearly frequency, as we end up with fewer control properties using higher frequency data, combined with the large variance of LTV ratios. After 2016, the difference in LTV ratios for treated and control properties quickly converge back to zero and remain statistically insignificant in all subsequent time periods.

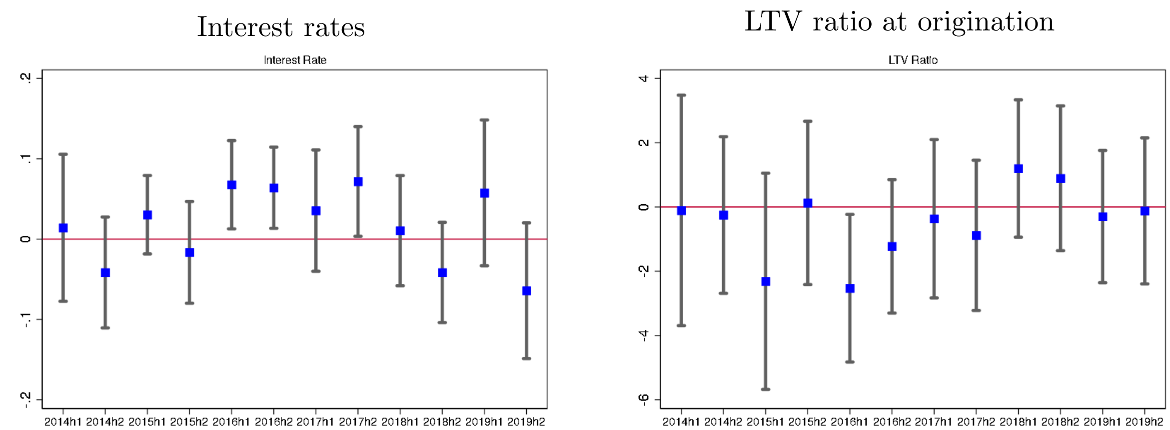

Figure 8: Baseline regressions at the half-yearly frequency

4.3.2 Dynamic Controls

Next, we replace the static controls in our baseline regression with dynamic controls, that is, a specification where all control variables are interacted with annual time dummies to allow for changes in the relationship between outcome variables and controls. To estimate this regression we modify our baseline equation as follows:

$$Y_{ijt} = \alpha_j + \mu_t + \lambda_c + \sum_{\tau=2014}^{2019} \beta_\tau Treat_i \mathbf{1}_{\tau=t} + \sum_{\tau=2014}^{2019} \delta_\tau \mathbf{X_{it}} \mathbf{1}_{\tau=t} + \epsilon_{ijt} \tag{2}$$

Results are presented in the first column of Tables 8 and 9 for interest rate and LTV, respectively. Our main results remain robust. That said, the statistical significance for interest rates drops slightly, particularly for the effect in 2016, as expected because this strategy introduces a large number of control variables and hence reduces the statistical power in estimating the treatment effect given our small sample.

4.3.3 Controlling for Demographics

Our baseline results rely on geographic fixed effects at the census tract level to control for factors affecting the demand for residential mortgages. However, census tract fixed effects negate any information coming from tracts that have either all treated or all control properties, which make up roughly half of our sample (as shown in Table 3 and discussed in Section 3.1). Therefore, the empirical identification in our baseline regression comes from census tracts that have a mixture of both treated and control properties. As an alternative, we ran a specification using data from the U.S. Bureau of the Census to directly control for some demographic differences (i.e., race, education, median income level, etc.) for different census tracts within Orleans Parish. Using tract-level demographic controls rather than fixed effects allows us to draw in information from tracts that contain all treatment or control properties while still controlling for developments in the mortgage market at the tract level. As shown in the second column of Table 8, interest rates for newly originated mortgages on treated properties are still higher in 2016 (the magnitude of the estimated coefficient falls slightly to 5 basis points and remains statistically significant at the 95% confidence level despite slightly lower t-value). Results for the LTV ratio are presented in the second column of Table 9. When we control for demographics, LTV ratios for treated properties in 2016 are even more negative relative the baseline results and the estimated coefficient retains it's statistical significance.

4.3.4 Controlling for Flood Risk

The map changes likely altered the risk profile of treated properties relative to control properties in two potentially offsetting ways. On the one hand, presumably treated properties were removed from the flood zone because FEMA officials viewed them as facing lower flood risk relative to control properties. At the same time, the map changes also implied that mortgages on treated properties no longer required flood insurance, which might increase their riskiness from the lender's point of view. One way to separate these two competing effects is to directly control for flood risk.

To do this, we match-merged our mortgage data set with the Cotality Climate Risk Data on first-floor and ground-floor elevation using the address of each individual property.[18] The left panel in Figure 9 presents a histogram of the ground-floor elevation of all properties in our sample broken out separately by treatment versus control. It shows the ground floor of the average property in our sample sits below sea level and treated properties have, on average, higher ground floor elevation relative to control properties, which is consistent with the removal of these properties from the SFHAs on the new maps. The right panel presents the same distribution but focuses on properties

[18] The sample size declines by about 900 observations as we do not have a perfect merge between the address field in the Y14 and the Cotality dataset.

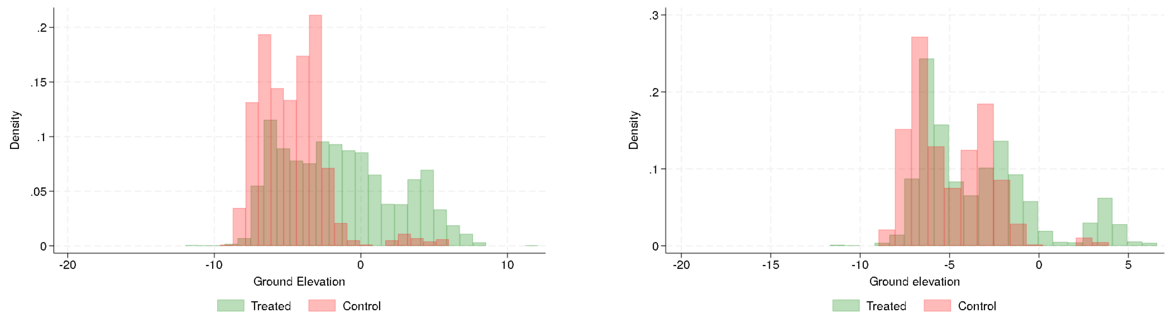

All properties Close properties

Figure 9: Ground floor elevation as a proxy for flood risk at the property level

close to the new flood map borders, where closeness is defined as in Section 4.2. For these close properties, there is a high degree of overlap in the distribution of ground elevation, suggesting underlying flood risks are similar for treated and control properties close to the new map borders.

Given this information, we introduce ground-floor elevation into our baseline regression to directly control for property-level flood risk. As shown in the third column of Table 8, our results for interest rates remain largely unchanged, suggesting that our findings were not driven by underlying differences in flood risk. The estimated coefficient for 2016 increases to 7 basis points and its statistical significance rises somewhat. Results are robust to using other proxies for flood risk as well, including first floor elevation from the Cotality dataset, as shown in the fourth column, and the flood risk factor from the First Street Foundation (not shown).[19] Results for the LTV ratio are presented in the third and fourth columns of Table 9. When we use ground elevation as a control, LTV ratio difference between treated and control properties in 2016 remains negative and comparable in magnitude but loses it's significance at conventional levels. But, when we use first floor elevation as a control, the 2016 effect remains significant, suggesting some mixed evidence.

5 Why are the Results Short-lived?

Our results are temporary, lasting up to two years, suggesting that if banks perceived greater risk for treated properties owing to the possibility that borrowers would no longer voluntarily retain flood insurance, this concern was short-lived. In this section, we use NFIP flood insurance claims from a major flooding event in New Orleans in 2017 to shed some light on this.

[19]With regard to the First Street data, it is worth noting that there is very little variation at the property level in their flood risk factor for properties within Orleans Parish–the overwhelming majority are scored either an 8 or a 9 on a scale of 0 to 10. Nonetheless, what little variation there is suggests that treated properties have lower First Street flood risk factors and is therefore in line with what we see in the Cotality ground floor elevation metric.

5.1 Severe Flooding Event in New Orleans in 2017

In 2017, the year after the map changes were implemented, New Orleans experienced a major flooding event that revealed information about the true underlying flood risk and the extent of insurance coverage for treated properties. The most severe flooding happened in early-August because, despite the post-Katrina infrastructure improvements to levees and flood walls, several critical pumps required to move water from city streets into the levees were either broken or out of service. On August 4, nine inches of rain fell in four hours and overwhelmed the pump system, causing widespread flooding that took over fourteen hours to drain.

This event is informative for our results two reasons. First, it is widely regarded that most assessment of flood risk at the property level is based on observed flooding events, so the 2017 flood event would have provided home owners as well as mortgage lenders with updated information on property flood risk. Second, the flooding revealed information about insurance take-up throughout the city as insured homeowners of flooded properties would have filed flood claims.

5.2 Claims Data from the NFIP

Data for individual claims transactions for all U.S. homeowners that have flood insurance through the NFIP are publicly available from FEMA. These data contain information on the date of the claim, total building damage amount, the flooding event causing the claim, the NFIP flood zone designation of the insured property, the Census block at the time of the loss, whether the building is residential or commercial, as well as extensive additional characteristics of the insured property and associated damages.

Figure 10 plots the number of claims and total damages for residential properties for the City of New Orleans over the full sample for the NFIP claims data from 2010 to 2024. The figure shows that over the first half of period captured by our sample, which starts in 2014 and runs to 2019, there were no NFIP flood claims for the City of New Orleans, including the year of the map change in 2016. The flooding in 2017 was the first major flooding event in the city following the map changes, so it would have been the first opportunity for both homeowners and mortgage lenders to see how treated properties fared and the extent to which flood insurance was retained for borrowers of treated properties.

We interact these claims data with our property-level data set from Cotality that determines the SFHA status for residential properties in New Orleans before and after the map changes (summarized in Table 4). Specifically, merge the property-level data set from Cotality with the NFIP claims data by Census block FIPS code.

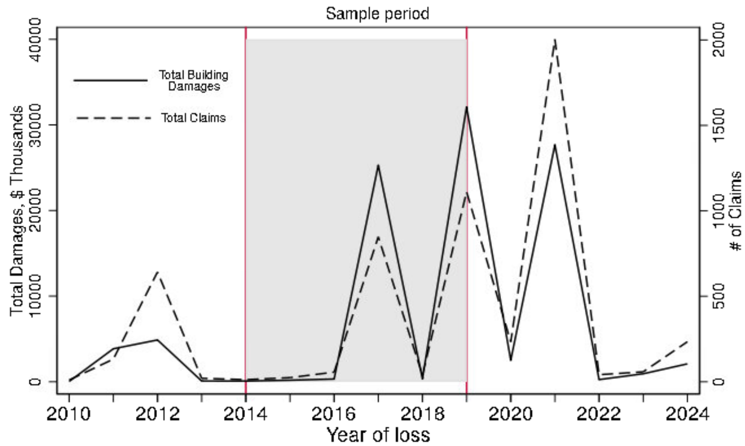

Figure 10: NFIP claims and damages in New Orleans, 2005 to 2024

5.3 Insurance Take-up and Flooding Incidence

To file a claim on on individual property, the homeowner must have flood insurance and must also have experienced damage from the flooding event. So, the geographic profile of the claims data is informative for both insurance take up and flooding incidence. After interacting the claims data with the data on classification of treatment versus control, we aggregate the data to the level of a Census block and, for each block, we calculate that block's share of total NFIP claims as well as the share of treated properties within that block using the total number of properties in the Cotality dataset. If owners of treated properties responded to the map changes by dropping their flood insurance contracts, we would expect to see a negative correlation between the share of treated properties and the share of total claims across all Census blocks. Figure 11 shows this is not the case. Indeed, there is essentially no correlation at all between the two, suggesting that insurance coverage was roughly similar across treated and control properties.

To the extend that insurance coverage is not biased by whether a property is classified as treatment or control, insurance claims should be a decent proxy for flooding incidence.[20] In this case, we define a "heavily treated" Census block as one that is in the upper 75^{th} percentile of the proportion of treated properties within the block. The data are then split into 109 heavily treated blocks and all other 322 blocks and we compare the number of claims and damages across the two.

Conditional on insurance coverage being similar across the two types of properties, if treated properties were removed from the SFHA because they are systematically less risky we would expect

[20] If treatment properties has systematically lower insurance coverage than control properties, the claims data would would be a poor proxy for flood incidence. In this case, claims may proxy well for flood incidence of controls, but not for treatment properties, where we would never observe claims even in the event of flooding owing to lack of insurance.

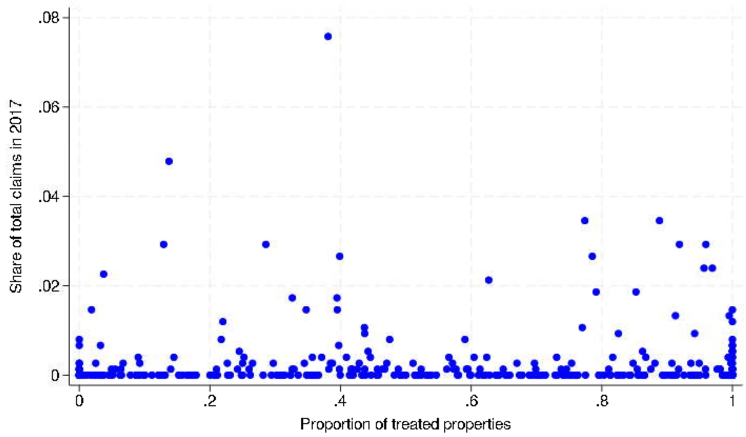

Figure 11: Share of treated properties and share of claims, by Census block, 2017.

to see a lower number of claims and lower damages for mostly treated Census blocks. Table 4 shows this is not the case. The average number of claims in heavily treated Census blocks is broadly similar to all other blocks and the same is true of damages, regardless of whether damages are measured in dollars or as a share of the underlying value of the insured structure.

Table 4: NFIP Claims in 2017, Heavily Treated Blocks vs. All Others

	Heavily Treated (n = 109)		All Other (n = 322)	
Variable	Mean	St. Dev.	Mean	Std. Dev.
Total Claims	2.09	4.82	1.63	5.16
Claims per Property	0.01	0.02	0.03	0.37
Total Damages	$55,704	$158,097	$53,938	$296,188
Damages as Share of Building Value	0.05	0.07	0.03	0.06

We find no discernible difference in either flood risk or insurance take-up for flood insurance between majority treated neighborhoods and others. A plausible explanation for the temporary nature of our results is that banks learned from the flooding event that actual flood risk of treated properties was not materially different from control properties and owners of treated properties retained their flood insurance. This information may have resolved any uncertainty created by the map changes and, hence, banks returned to pricing the two types of mortgages similarly.

6 Conclusion

We examine how banks react to property-level flood risk as communicated by FEMA flood maps. A large-scale redrawing of FEMA special flood hazard areas in the City of New Orleans in 2016 led to the removal of the flood insurance requirement for roughly half the properties in the city. Using a dynamic differences-in-differences specification on property-level mortgage data, we find that banks charge nearly six basis points more on the interest rates for mortgages originated on properties that were recently removed from an SFHA. Additionally, we find evidence that LTV ratios for treated properties temporarily declined by about 2 percentage points relative to control properties, suggesting that banks also responded by requiring larger down payments—roughly $6,600 for the average treated property.

We interpret these results as banks pricing in compensation for additional risk associated with dropping the mandatory flood insurance requirement for mortgages on treated properties. This interpretation is supported by evidence suggesting that our main findings are importantly driven by properties that are closest to the new flood map borders where empirical proxies of flood risk are nearly identical across treatment and control properties, leaving the insurance requirement as the main differene between the two.

These effects are temporary, lasting up to two years. An analysis of NFIP flood insurance claims data from a major flooding event in New Orleans in 2017 offers insights into why the effects are short-lived. The claims data show no discernible difference in insurance take-up between treated and control properties, suggesting that many homeowners of treated properties voluntarily retained their flood insurance despite no longer being required to do so. In light of this, the temporary nature of our main results may be due to banks learning about insurance take-up from the 2017 flooding event. As banks observed that there was no material difference in either insurance take-up or flood risk between treatment and control properties, they likely returned to pricing mortgages for both types of properties similarly.

Our study contributes to the growing literature on pricing risks in financial markets by providing evidence of how banks dynamically adjust their lending practices in response to changes in perceived flood risk. While our results show that banks are attentive to flood risk in mortgage origination, the temporary nature of their response highlights the challenges in accurately pricing flooding risks. This underscores the importance of accurate and up-to-date flood risk information, as well as the need for ongoing research into how financial institutions can best incorporate evolving flooding risks into their decision-making processes.

References

[1] Atreya, A., Ferreira, S., & Michel-Kerjan, E. (2015). "What drives households to buy flood insurance? New evidence from Georgia." *Ecological Economics*, 117, 153-161.

[2] Bhutta, N., Fuster, A., & Hizmo, A. (2019). "Paying too much? Borrower sophistication and overpayment in the US mortgage market." Manuscript.

[3] Bin, O., Kruse, J. B., & Landry, C. E. (2008). "Flood hazards, insurance rates, and amenities: Evidence from the coastal housing market." *Journal of Risk and Insurance*, 75(1), 63-82. https://doi.org/10.1111/j.1539-6975.2007.00248.x

[4] Blickle, K., Perry, E., & Santos, J. (2024). "Do mortgage lenders respond to flood risk?" *New York Fed Working Paper.*

[5] Blickle, K., & Santos, J. (2022). "Unintended consequences of 'mandatory' flood insurance." *New York Fed Working Paper.*

[6] Brody, S., Highfield, W., Wilson, M., Lindell, M., & Blessing, R. (2017). "Understanding motivations of coastal residents to voluntarily purchase federal flood insurance." *Journal of Risk Research*, 20(6), 760-775.

[7] Buchak, G., Matvos, G., Piskorski, T., & Seru, A. (2018). "Fintech, regulatory arbitrage, and the rise of shadow banks." *Journal of Financial Economics*, 130, 453-483.

[8] Cannon, C., Gotham, K. F., Lauve-Moon, K., & Powers, B. (2020). "The climate change double whammy: Flood damage and the determinants of flood insurance coverage, the case of post-Katrina New Orleans." *Climate Risk Management*, 27, 100-210.

[9] Horowitz, A. (2020). *Katrina: A history, 1915-2015.* Harvard University Press.

[10] Hung, H. (2009). "The attitude towards flood insurance purchase when respondents' preferences are uncertain: A fuzzy approach." *Journal of Risk Research*, 12(2), 239-258.

[11] Kara, G. I., & Yook, Y. (2023). "Policy uncertainty and bank mortgage credit." *Journal of Money, Credit, and Banking*, 55(4), 783-823.

[12] Kousky, C., Palim, M., & Pan, Y. (2020). "Flood damage and mortgage credit risk: A case study of Hurricane Harvey." *Journal of Housing Research*, 29, S86-S120.

[13] Lave, T., & Lave, L. (1991). "Public perception of the risk of floods: Implications for communication." *Risk Analysis*, 11(2), 255-267.

[14] Li, J., & Landry, C. (2018). "Flood risk, local hazard mitigation, and the community rating system of the National Flood Insurance Program." *Land Economics*, 98(2), 175-198.

[15] Michel-Kerjan, E. (2010). "Catastrophe economics: The National Flood Insurance Program." *Journal of Economic Perspectives*, 24(4), 165-186.

[16] Michel-Kerjan, E., & Kousky, C. (2010). "Come rain or shine: Evidence on flood insurance purchases in Florida." *Journal of Risk and Insurance*, 77(2), 369-397.

[17] Nguyen, D. D., Ongena, S., Si, S., & Sila, V. (2022). "Climate change risk and the cost of mortgage credit." *Review of Finance*, 26(6), 1509-1549.

[18] Ortega, F., & Taşpınar, S. (2018). "Rising sea levels and sinking property values: Hurricane Sandy and New York's housing market." *Journal of Urban Economics*, 106, 81-100. https://doi.org/10.1016/j.jue.2018.06.005

[19] Pralle, S. (2019). "Drawing lines: FEMA and the politics of mapping flood zones." *Climatic Change*, 152, 227-237.

[20] Rhodes, A., & Besbris, M. (2022). *Soaking the middle class: Suburban inequality and recovery from disaster*. Russell Sage Foundation.

[21] Sastry, P. (forthcoming). "Who bears flood risk? Evidence from mortgage markets in Florida." *Review of Financial Studies*.

[22] Savitt, A. (2017). "Insurance as a tool for hazard risk management? An evaluation of the literature." *Natural Hazards*, 86(2), 583-599.

[23] Sharma, B. (2025). "Information about climate transition risk and bank lending." Manuscript.

[24] Shu, E. G., Porter, J. R., Wilson, B., Bauer, M., & Pope, M. L. (2022). "The economic impact of flood zone designations on residential property valuation in Miami-Dade county." *Journal of Risk and Financial Management*, 15(10), 434. https://doi.org/10.3390/jrfm15100434

[25] Weill, J. (2023). "Flood risk mapping and the distributional impacts of climate information." *Finance and Economics Discussion Series 2023-066*. Board of Governors of the Federal Reserve System. https://doi.org/10.17016/FEDS.2023.066

[26] Wing, O., Bates, P. D., Smith, A. M., Sampson, C. C., Johnson, K. A., Fargione, J., & Morefield, P. (2018). "Estimates of present and future flood risk in the conterminous United States." *Environmental Research Letters*, 13(3).

7 Results Tables

Table 5: Baseline interest rate regression with alternative fixed-effects specifications

VARIABLES	(1) Interest Rate	(2) Interest Rate	(3) Interest Rate	(4) Interest Rate
2014 x Treated	-0.019	-0.012	-0.018	-0.013
	(0.027)	(0.031)	(0.025)	(0.028)
2015 x Treated	-0.006	-0.001	-0.000	0.004
	(0.019)	(0.025)	(0.017)	(0.023)
2016 x Treated	0.048**	0.055**	0.051**	0.056**
	(0.024)	(0.024)	(0.023)	(0.023)
2017 x Treated	0.042*	0.047*	0.030	0.036
	(0.024)	(0.027)	(0.022)	(0.026)
2018 x Treated	-0.029	-0.022	-0.020	-0.014
	(0.027)	(0.030)	(0.025)	(0.028)
2019 x Treated	-0.025	-0.013	-0.029	-0.020
	(0.040)	(0.039)	(0.040)	(0.040)
LTV ratio	0.003***	0.003***	0.003***	0.003***
	(0.000)	(0.000)	(0.000)	(0.000)
Log loan amount	-0.252***	-0.241***	-0.233***	-0.230***
	(0.011)	(0.016)	(0.010)	(0.015)
Credit score	-0.002***	-0.002***	-0.002***	-0.002***
	(0.000)	(0.000)	(0.000)	(0.000)
Loan term	0.003***	0.003***	0.003***	0.003***
	(0.000)	(0.000)	(0.000)	(0.000)
Refinance	0.064***	0.063***	0.068***	0.068***
	(0.012)	(0.013)	(0.012)	(0.013)
Investment Property	0.531***	0.519***	0.547***	0.535***
	(0.018)	(0.020)	(0.018)	(0.019)
Jumbo loan	0.112***	0.104***	0.084***	0.079***
	(0.022)	(0.023)	(0.022)	(0.023)
Conventional loan without mortgage insurance	0.336***	0.335***	0.362***	0.360***
	(0.028)	(0.029)	(0.025)	(0.026)
Conventional loan with mortgage insurance	0.263***	0.258***	0.293***	0.288***
	(0.027)	(0.027)	(0.024)	(0.024)
VA Residential loan	-0.185***	-0.187***	-0.140***	-0.144***
	(0.044)	(0.046)	(0.041)	(0.043)
CRA loan	-0.659***	-0.682***	-0.637***	-0.654***
	(0.172)	(0.170)	(0.156)	(0.155)
Constant	7.225***	7.107***	6.915***	6.895***
	(0.157)	(0.213)	(0.144)	(0.201)
Observations	6,068	6,066	6,065	6,063
Year Fixed Effects	Yes	Yes	Yes	Yes
Tract Fixed Effects	No	Yes	No	Yes
Bank Fixed Effects	No	No	Yes	Yes
Cluster	Tract	Tract	Tract	Tract
Adjusted R-squared	0.583	0.588	0.599	0.602

Robust standard errors in parentheses, clustered at the census tract level.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 6: Baseline LTV regression with alternative fixed-effects specifications

VARIABLES	(1) LTV Ratio	(2) LTV Ratio	(3) LTV Ratio	(4) LTV Ratio
2014 x Treated	-1.711	-0.173	-1.537	-0.120
	(1.229)	(1.169)	(1.181)	(1.151)
2015 x Treated	-1.977	-1.149	-1.925	-1.151
	(1.229)	(1.074)	(1.228)	(1.066)
2016 x Treated	-2.845***	-1.853**	-2.654**	-1.742**
	(1.081)	(0.899)	(1.064)	(0.881)
2017 x Treated	-1.406	-0.411	-1.396	-0.457
	(1.029)	(0.943)	(1.054)	(0.961)
2018 x Treated	0.043	1.058	0.110	1.025
	(1.007)	(0.876)	(0.993)	(0.842)
2019 x Treated	-1.223	-0.086	-1.289	-0.145
	(0.791)	(0.778)	(0.834)	(0.795)
Interest rate	3.900***	3.280***	3.559***	3.126***
	(0.553)	(0.527)	(0.536)	(0.508)
Log loan amount	7.119***	16.225***	7.059***	16.018***
	(0.673)	(0.797)	(0.652)	(0.796)
Credit score	-0.015***	0.000	-0.016***	-0.000
	(0.005)	(0.004)	(0.005)	(0.004)
Loan term	0.029***	0.019***	0.031***	0.021***
	(0.004)	(0.004)	(0.004)	(0.003)
Refinance	-11.459***	-10.435***	-11.230***	-10.245***
	(0.452)	(0.393)	(0.455)	(0.387)
Investment Property	3.061***	2.378***	3.196***	2.432***
	(0.735)	(0.656)	(0.710)	(0.640)
Jumbo loan	-0.367	-1.539*	-0.802	-1.668*
	(1.097)	(0.907)	(1.038)	(0.862)
Conventional loan without mortgage insurance	-19.450***	-15.645***	-19.528***	-15.834***
	(0.478)	(0.431)	(0.478)	(0.432)
Constant	-9.106	-131.240***	-7.549	-127.898***
	(9.731)	(11.176)	(9.408)	(11.275)
Observations	6,068	6,066	6,065	6,063
Year Fixed Effects	Yes	Yes	Yes	Yes
Tract Fixed Effects	No	Yes	No	Yes
Bank Fixed Effects	No	No	Yes	Yes
Cluster	Tract	Tract	Tract	Tract
Adjusted R-squared	0.515	0.606	0.522	0.609

Robust standard errors in parentheses, clustered at the census tract level.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 7: Distance to the Current Flood Map Border

VARIABLES	Interest Rate Close (< 50th percentile)	Interest Rate Far (> 50th percentile)	LTV Ratio Close (< 50th percentile)	LTV Ratio Far (> 50th percentile)
2014 x Treated	0.014	-0.077	-0.169	0.912
	(0.036)	(0.064)	(1.437)	(2.938)
2015 x Treated	0.010	-0.034	-1.213	0.055
	(0.027)	(0.063)	(1.390)	(2.513)
2016 x Treated	0.064**	0.021	-1.032	-1.144
	(0.028)	(0.054)	(0.988)	(2.766)
2017 x Treated	0.046	-0.019	-0.578	0.894
	(0.028)	(0.062)	(1.182)	(2.844)
2018 x Treated	-0.023	-0.028	1.196	2.756
	(0.038)	(0.056)	(1.020)	(2.627)
2019 x Treated	-0.015	-0.058	-0.743	1.681
	(0.048)	(0.066)	(1.052)	(1.881)
LTV ratio	0.003***	0.003***		
	(0.001)	(0.001)		
Interest rate			3.295***	3.039***
			(0.696)	(0.687)
Log loan amount	-0.234***	-0.221***	16.201***	16.510***
	(0.020)	(0.023)	(1.179)	(1.071)
Credit score	-0.002***	-0.002***	-0.003	0.003
	(0.000)	(0.000)	(0.006)	(0.006)
Loan term	0.003***	0.003***	0.019***	0.022***
	(0.000)	(0.000)	(0.005)	(0.005)
Refinance	0.063***	0.073***	-9.670***	-10.639***
	(0.017)	(0.018)	(0.524)	(0.618)
Investment Property	0.517***	0.554***	2.180*	2.333***
	(0.028)	(0.028)	(1.152)	(0.767)
Jumbo loan	0.117***	0.036	-1.044	-1.839
	(0.031)	(0.029)	(1.015)	(1.384)
Conventional loan without mortgage insurance	0.414***	0.320***	-15.371***	-15.872***
	(0.036)	(0.032)	(0.585)	(0.563)
Conventional loan with mortgage insurance	0.315***	0.271***		
	(0.037)	(0.032)		
VA Residential loan	-0.124*	-0.159***		
	(0.063)	(0.059)		
CRA loan	-0.645***	-0.655***		
	(0.210)	(0.236)		
Constant	6.921***	6.820***	-127.898***	-137.997***
	(0.266)	(0.289)	(15.833)	(14.816)
Observations	3,019	3,020	3,019	3,020
Fixed Effects	Bank, Tract, Time	Bank, Tract, Time	Bank, Tract, Time	Bank, Tract, Time
Cluster	Tract	Tract	Tract	Tract
Adjusted R-Squared	0.598	0.603	0.600	0.621

Robust standard errors in parentheses, clustered at the census tract level.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 8: Robustness Tests for Interest Rate Regression

	(1)	(2)	(3)	(4)
CONTROLS	Dynamic	Demographics	Ground elevation	First-floor elevation
2014 x Treated	-0.001	-0.022	-0.010	-0.016
	(0.025)	(0.025)	(0.030)	(0.030)
2015 x Treated	0.002	-0.004	-0.001	-0.008
	(0.022)	(0.017)	(0.023)	(0.021)
2016 x Treated	0.048**	0.046**	0.068***	0.062**
	(0.024)	(0.023)	(0.023)	(0.024)
2017 x Treated	0.030	0.025	0.025	0.019
	(0.026)	(0.022)	(0.024)	(0.025)
2018 x Treated	-0.009	-0.024	-0.002	-0.009
	(0.028)	(0.026)	(0.031)	(0.030)
2019 x Treated	-0.040	-0.032	-0.007	-0.015
	(0.035)	(0.040)	(0.039)	(0.040)
Ground elevation			-0.006	
			(0.005)	
First-floor elevation				0.001
				(0.003)
Constant	6.984***	7.161***	6.835***	6.854***
	(0.194)	(0.639)	(0.222)	(0.223)
Observations	6,063	6,065	5,190	5,190
Dynamic Controls	Yes	No	No	No
Static Loan Controls	No	Yes	Yes	Yes
Demographic Controls	No	Yes	No	No
Fixed Effects	Bank, Tract, Time	Bank, Time	Bank, Tract, Time	Bank, Tract, Time
Cluster	Tract	Tract	Tract	Tract
Adjusted R-Squared	0.626	0.599	0.592	0.592

Robust standard errors in parentheses, clustered at the census tract level.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 9: Robustness Tests for LTV Regression

	(1)	(2)	(3)	(4)
VARIABLES	Dynamic	Demographics	Ground elevation	First-floor elevation
2014 x Treated	-0.381	-1.203	0.154	-0.181
	(1.071)	(1.248)	(1.329)	(1.305)
2015 x Treated	-1.219	-1.783	-0.961	-1.275
	(1.078)	(1.177)	(1.073)	(1.073)
2016 x Treated	-1.944**	-2.296**	-1.572	-1.943**
	(0.886)	(0.904)	(1.010)	(0.954)
2017 x Treated	-0.397	-1.394	0.345	-0.065
	(0.973)	(1.071)	(1.114)	(1.011)
2018 x Treated	0.596	0.494	0.806	0.453
	(0.852)	(1.000)	(0.942)	(0.848)
2019 x Treated	-0.521	-0.634	-0.330	-0.734
	(0.842)	(0.815)	(0.888)	(0.807)
Ground elevation			-0.178	
			(0.263)	
First-floor elevation				-0.251**
				(0.123)
Constant	-127.360***	-99.615**	-131.176***	-130.856***
	(11.309)	(39.984)	(11.468)	(11.446)
Observations	6,063	6,065	5,190	5,190
Dynamic Controls	Yes	No	No	No
Static Loan Controls	No	Yes	Yes	Yes
Demographic Controls	No	Yes	No	No
Fixed Effects	Bank, Tract, Time	Bank, Time	Bank, Tract, Time	Bank, Tract, Time
Cluster	Tract	Tract	Tract	Tract
Adjusted R-Squared	0.609	0.581	0.614	0.614

Robust standard errors in parentheses, clustered at the census tract level.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

8 Appendix

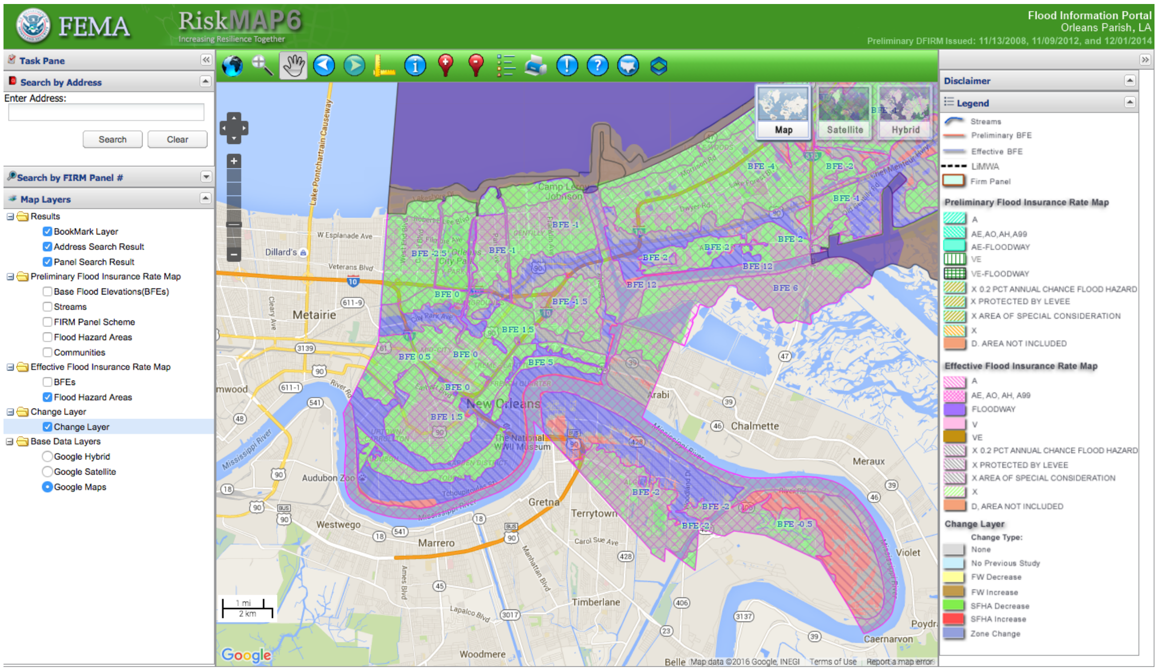

Figure 12: FEMA Website for Flood Zone changes in 2016