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

David M. Arseneau and Gazi Kara

Federal Reserve Board

January 18, 2024

Abstract

This paper uses a large-scale redrawing of flood zone maps for the City of New Orleans to identify how banks respond to climate risk in residential mortgage origination. We use geocoding techniques to separate loan-level data on mortgage originations from the largest U.S. banks into treatment versus control properties based on how they were affected by the map changes. We find that banks charge interest rates that are roughly 10 basis points higher for mortgage originations on treated properties that were removed from FEMA flood zones as a result of the map changes and these higher interest rates persist for a period of about two years. These results are robust to a number of different specifications. We also find that the effect is stronger in banks with branches in Orleans and surrounding parishes, suggesting that local knowledge of the mortgage market plays a role in explaining our results.

Keywords: FEMA Maps, Flooding, Mortage lending

PRELIMINARY: PLEASE DO NOT DISTRIBUTE WITHOUT AUTHORS' PERMISSION

*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.

1 Introduction

Climate change is projected to increase the frequency and severity of extreme weather events (U.S. Global Change Research Program, 2017) and, as a result, there is increasing interest in what this implies for the stability of the financial system. One particular area of concern is residential mortgage lending. The U.S. residential real estate market is supported by about $11 trillion mortgage debt market, the largest non-sovereign debt market in the world. The financial crisis of 2007–08 demonstrated that risks in this market can have large spillover effects to the domestic economy and can also have global repercussions. Furthermore, house ownership is the most important source of wealth for most Americans, and hence, changes in house values, such as due to climate related events, have substantial implications for consumption spending, employment and output. Rising sea levels or an increase in the frequency or severity of weather events related to climate change may increase the risk profile of loans made on properties in flood-prone areas. If this risk is not priced properly mortgage borrowers and lenders alike could face exposure to losses that are not well understood.

This paper uses a unique natural experiment to gain in sight into how banks price climate risk through flood insurance in residential mortgage origination. The experiment is a large-scale redrawing of Federal Emergency Management Agency (FEMA) flood maps for the City of New Orleans in 2016. Prior to the redrawing, a large portion of properties in the city were located in special flood hazard areas, which meant that owners of these properties were required to have federal flood insurance if they financed their mortgages with a federally-backed mortgage lender. Following the completion of flood protection infrastructure in the wake of Hurricane Katrina in 2005, city officials successfully lobbied FEMA to reconsider the special flood hazard areas in a large-scale redrawing of the flood maps. The redrawn maps were implemented in late-2016 and resulted in the removal of roughly half of the properties in the city from the special floods zones, meaning that mortgage borrowers for these properties no longer required flood insurance when financing their homes.

We use confidential data on bank mortgage originations in the City of New Orleans and state-of-the-art geo-coding techniques to assess how this policy change affected loan terms at the property level for the largest U.S. banks. Individual properties are divided into treatment properties—those that were in the special flood hazard area prior to the map change, but were removed from the special flood hazard area after the redrawn maps were released—and control properties—those that were will in the special flood hazard area both before and after the revision. We compare interest rates, loan-to-value ratios, and securitization rates for treated versus control properties using a dynamic differences-in-differences specification that controls for various 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 eight basis points higher than for control properties, but this effect is temporary and only lasts for a period of about two years. The magnitude of the interest rate effect is in line with results found in other studies (for example, Nguyen, et. al., forthcoming). These results are robust to a number of different specifications. In addition, we find loan-to-value (LTV) ratios are lower for treated properties, but these findings are not statistically significant likely owing to the large variation of LTV ratios for the mortgages in our sample combined with the relatively small sample size. Finally, do not find any difference in securitization rates, but this likely reflects the fact that the majority of loans in our same are securitized, so the variation in the data is limited.

What drives the higher interest rates for treated properties in our sample? We examine the hypothesis that it reflects bank's knowledge of the local mortgage market. Consistent with this hypothesis, we find that banks with local branches in New Orleans and surrounding parishes increase interest rates more aggressively than banks without local branches. Future work will shed light on whether this local knowledge reflects a better understanding of the underlying risks associated with a treated property (i.e., lack of the insurance requirement for a treated property may make it a riskier loan for the bank). Alternatively, it could be that the higher interest rate reflects the banks understanding that removing a property from the special flood hazard area creates a surplus that accrues to the homeowner as flood insurance is no longer required for that property. Removing the insurance requirement eases the financing burden on the borrower and there is evidence that it boosts the value of the underlying property. Banks with monopoly pricing power owing to local knowledge may try to extract part of this surplus from homeowners through higher interest rates. We will shed additional light on these channels in future drafts of this paper.

In terms of related literature, the number of studies focusing on the response of the residential real estate market to emerging climate-related risks is rapidly growing. Within this literature, we view this paper as most closely related to four studies in particular. Sastry (2021) 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, Sila (forthcoming) study the effect of sea level rise on mortgage lending rates at the zip code level using European data. 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 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. Finally, Kahn and Ouzad (forthcoming) find that banks securitize mortgages in areas hit by major hurricanes at a faster rate relative to other mortgages implying that banks effectively transfer climate risk to the government for federally-backed loans. Our results do not have strong implications for securitization, but our data may not be well suited to answer this question because a large number of loans in our dataset are securitized anyway.

The remainder of this paper is organized as follows. The next section discusses the institutional background on FEMA flood maps, offers details on the large-scale redrawing in New Orleans, 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 results and Section 5 discussions potential mechanisms driving the 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 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 that episode for identification in our empirical analysis.

2.1 Flood Insurance, Flood Maps, and the Process to Update the Maps

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 that 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 if they are financed with mortgages from government-backed lenders.

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 and therefore borrowers for these properties whose mortgages come from federally-backed lenders are required to purchase flood insurance.

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.

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; Kousky, Palim and Pan, 2020). 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 important to homeowners, the initial drafts receive considerable scrutiny. There is a formal process in place for dealing with challenges to the maps, which happen 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 (or, Orleans Parrish) 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.

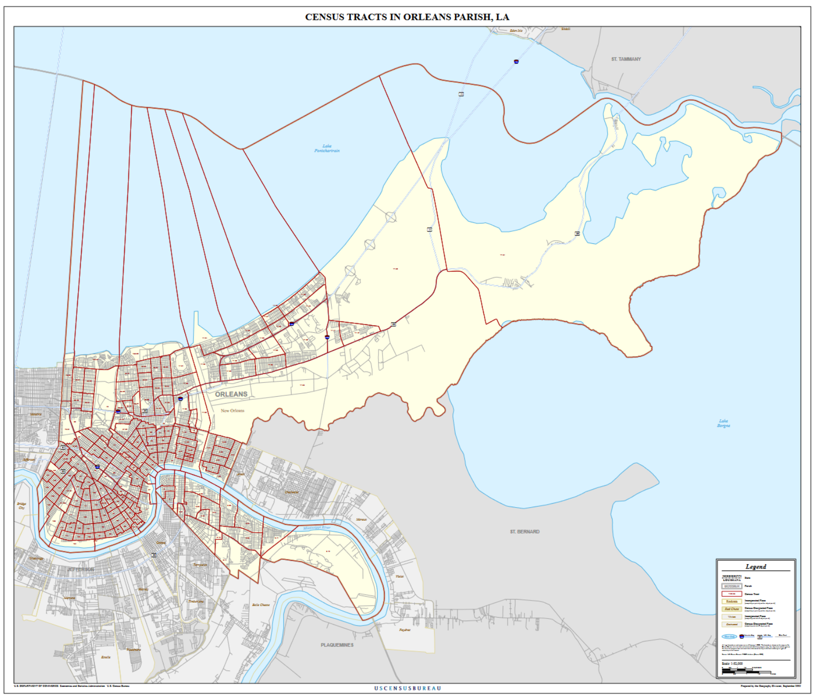

Figure 1: The City of New Orleans, broken out by Census Tracts

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, 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 and took effect on October 1, 2016. Figure 2 shows the old flood maps in the left panel with the SFHA zone highlighted in red. The panel on the right shows the updated maps with the new SFHA highlighted in blue. 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.

Figure 3 layers the old new map on top of the old one. Areas in which the two vintages of maps overlap

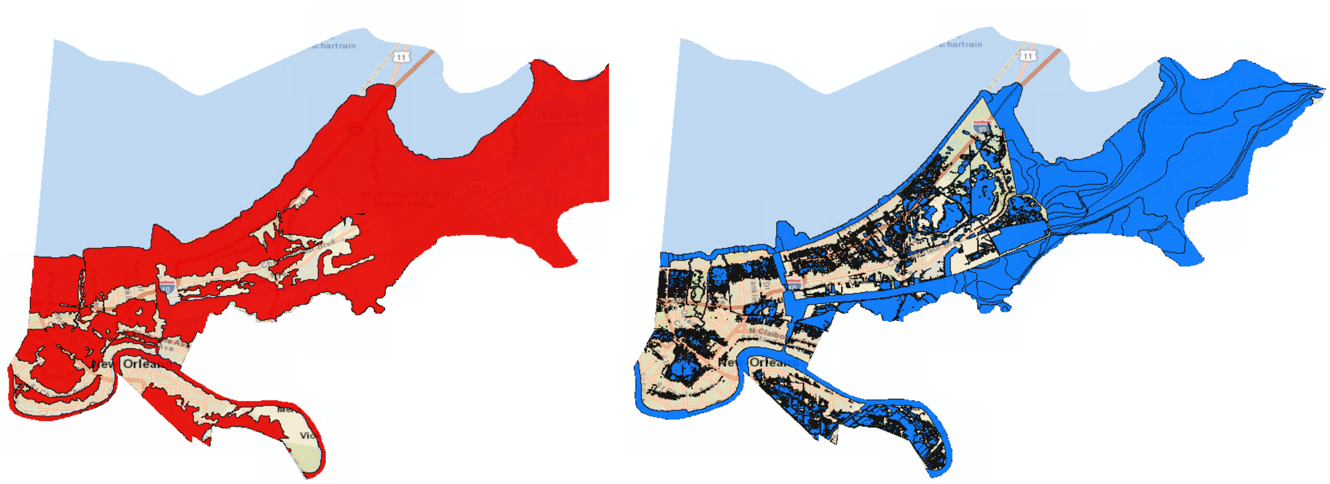

(a) Historic Flood Map (b) Current Flood Map

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

(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 the redrawing resulted in a large scale reduction in the size of the SFHA, no area is colored blue in the new maps.

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 3, 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). 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 SFHA zone.[1]

2.3 Identification Strategy

Under the assumption of local exogeneity of the map changes—meaning that the exact boundary between the red and purple areas in Figure 3 would be random to banks and borrowers—we use the New Orleans flood map revisions to separate individual properties in the Parrish 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 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.

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 in new originations.

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

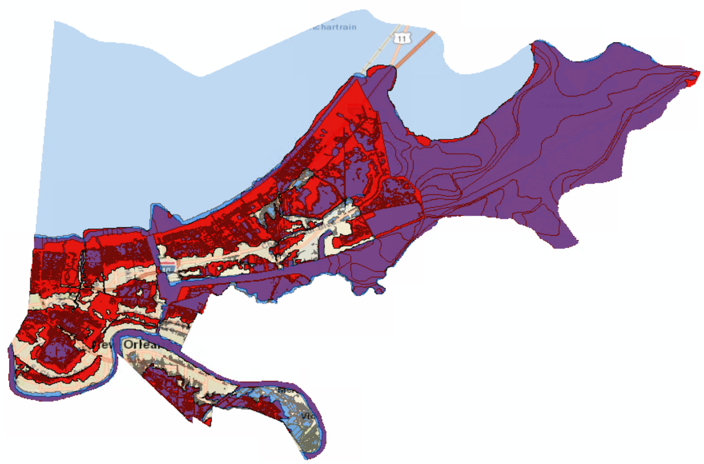

Figure 3: Historic and Current Flood Map Overlay

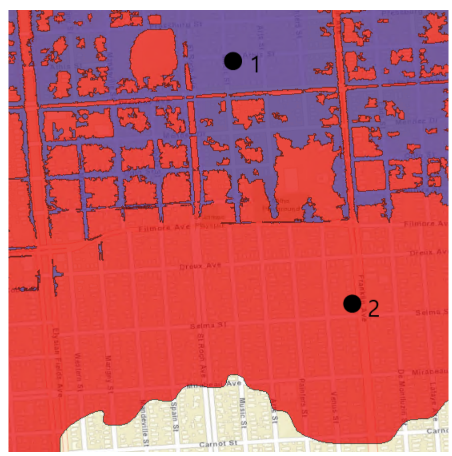

Figure 4: Illustration of treatment versus control properties in Orleans Parrish

3 Data and Methodology

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

3.1 Data

The mortgage data come from the FR-Y14M, which cover first lien residential mortgages in Schedule A as well as home equity loans and home equity lines of credit in Schedule B.[2] These data are collected monthly at the level of the individual loan. Our sample covers the first-lien mortgage loans originated between January

[2] 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

2014 to December 2019. They offer extensive information on every residential mortgage loan or home equity loan or line of credit held on the balance sheet of all banks that submit Y-14 data to the Federal Reserve.[3]

Although the Y14 data cover mortgages for the entire United States, our analysis restricts the sample only to newly originated mortgages within the City of New Orleans. We also restrict the sample to conforming and conventional loans. That means, we exclude jumbo loans that are above the FHFA conforming loan limits, and government sponsored loans, such as FHA and VA loans from our sample. Table 1 below presents some summary statistics for the entire sample as well as the means of the variables for treated and control groups.

Table 1: Summary statistics

						Mean by Group	
Variable	Obs	Mean	Std. Dev.	Min	Max	Treated	Control
Treatment dummy	4,793	0.76	0.43	0	1	1	0
Interest rate	4,793	4.18	0.59	2.25	7.375	4.19	4.16
LTV ratio	4,793	71	19	7	100	70	73
Securitized Loan	4,793	0.77	0.42	0	1	0.78	0.77
Credit Score	4,793	747	48	301	848	746	749
Loan amount (thousand $)	4,793	247	111	14	484	246	251
Loan term (months)	4,793	327	69	60	360	325	332
DTI Back-end ratio	4,793	36	15	0	339	35	36
Refinance	4,793	0.46	0.50	0	1	0.47	0.44
Mortgage Insurance	4,793	0.20	0.40	0	1	0.19	0.23
Investment Property	4,793	0.12	0.32	0	1	0.12	0.11
Fixed rate	4,793	0.97	0.18	0	1	0.97	0.97
Internally Sourced	4,793	0.92	0.27	0	1	0.92	0.91

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 from October 1st, 2016 until the end of the sample (December 2019).

We interact these flood maps with the mortgage data through a process called geocoding. Geocoding is a method of converting the specific property address for a given residential mortgage in our data set into latitude and longitude coordinates so that it can be placed on a map with a high degree of precision. We geocode every property in our sample and place each on an overlay of the current vintage of the flood

[3]Banks with $100 billion or more in consolidated assets are required to submit these data for regulatory purposes.

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 properly that was (1.) inside the old flood zone, but outside the new flood zone; (2.) inside the old flood zone and remains inside the new flood zone; (3.) outside the old flood zone, but inside the new flood zone; and (4.) inside the old flood zone and remains inside the new flood zone. The relatively large fraction of properties falling into the first group is not surprising given the political 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.

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 (3,622 mortgages, 76% of the sample), to group two (inside the old flood zone and remains inside the new flood zone), which we call the control group (1,171 mortgages, 24% of the sample). The last two columns of Table 1 compares mortgage originations that fall into the treatment and control groups, respectively.

3.2 Methodology

We estimate the following dynamic difference-in-differences model

$$Y_{ijt} = \alpha_j + \mu_t + \lambda_c + \sum_{\tau} \beta_\tau Treat_i \mathbf{1}_{\tau=t} + \delta \mathbf{X_{it}} + \epsilon_{ijt} \tag{1}$$

where Y_{ijt} is the interest rate, LTV ratio, or secitization dummy 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 half-yearly time dummy set to 1 if $t = \tau$; and, $\mathbf{X_{i,j,t}}$ is a $(nx1)$ vector of controls discussed below. The baseline regression allows for bank-specific fixed effects, denoted α_j, and time fixed effects at half-yearly frequency, denoted μ_t, and census tract fixed effects, denoted λ_c. 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 (log) loan amount, loan term (in months), back-end DTI ratio and FICO score at origination, and dummies for refinance loans, investment properties, mortgage insurance, fixed-rate loans, and internally sourced loans. Additionally, for interest rate regressions we control for the LTV ratio and securitized loan dummy; for LTV ratio regressions we control for the interest rate and securitization dummy; and for the securitization dummy regressions we control for interest rate and LTV ratio.

In this dynamic diff-in-diff specification, we interact a full set of half-yearly time dummies ($\mathbf{1}_{\tau=t}$) with the treatment dummy. As a result, β_τ coefficients estimate the difference in means of the dependent variable between treated and control properties in each time period, after controlling for loan, bank and property characteristics. Note that, we are not excluding any of the time dummies, and hence, β_τ coefficients are not 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 should expect β_τ coefficients to not be statistically different from zero before 2016 when exact map changes were not known.

4 Results

4.1 Main Results

In the presentation of our main results, the focus is on three margins of adjustment associated with mortgage origination: interest rates, LTV ratios and securitization.

We begin with interest rates. Results are presented in Figure 5 using 90% confidence intervals.[4] Interest rates for properties located within and outside the flood zone were not statistically different prior to 2016 when flood maps were redrawn. Following the change, treated properties suffered rates - on average - 8 basis points higher than properties in the control group in 2016 and 2017.[5] However, the effect appears short-lived as it dissipates after 2018. The economic magnitude may not seem large, but it is quite meaningful when compared with other studies that examine mortgage pricing. For instance, Nguyen et al. (forthcoming) find that mortgage rates are 10.2 basis points higher in zip codes exposed to sea level rise.[6]

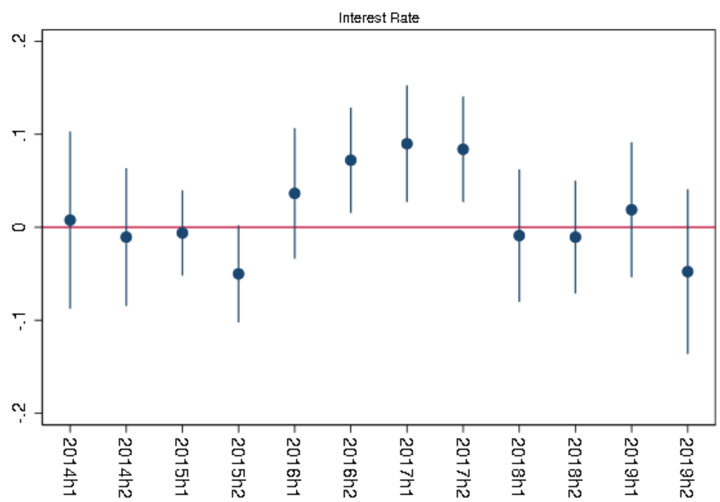

Figure 5: Dynamic baseline results for interest rates

Figure 6 shows that LTV ratios at origination in 2016 and 2017 have all negative coefficients that are lower than other years. However, only one coefficient, in 2016-H2, is statistically significant at the 90% confidence level (this coefficient is insignificant at the 95% confidence level). Moreover, standard errors

[4]These results remain qualitatively same when using 95% confidence bands.

[5]The three statistically significant coefficients in this period are 7, 9, and 8 basis points, respectively.

[6]Other studies also find comparable effects on mortgage pricing in various settings. For example, Kara and Yook (forthcoming) find that banks charge on average 6.3 bp 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 basis points lower rates, and Buchak et al. (2018) find that non-fintech shadow banks charge on average 2.4 basis points higher interest rates relative to traditional banks.

across all time periods are generally larger when compared with the interest rate results. In a dynamic diff-in-diff estimation, we compare the conditional means of dependent variable for treated and control properties in every 6-month period, and hence, a relatively small number of observations combined with the large range of the distribution of the LTV ratio, as shown in table 1, would make it harder to find statistically significant coefficients. As a result, we refrain from concluding that banks do not adjust LTV ratios in response to shrinking flood maps which remove mandatory flood insurance.

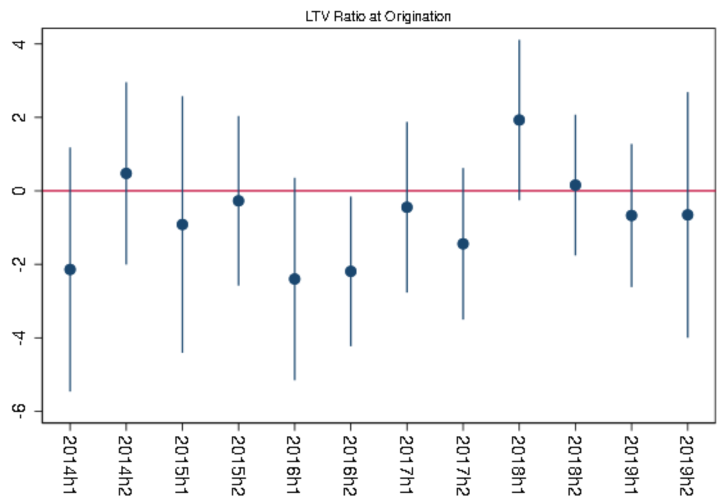

Figure 6: Dynamic baseline results for LTV

We do not find evidence of adjustment in securitization rates after the map changes as shown in Figure 7. Unlike some previous studies (e.g. Kahn and Ouzad, forthcoming), the vast majority of mortgages in our sample are securitized to GSEs, which makes it challenging to detect a differential effect on securitization for treated and control properties. As a result, we do not interpret this result as negating any risk transfer back to government through securitization after flood map changes, as suggested by Kahn and Ouzad (forthcoming). Nevertheless, the fact that majority of our sample are composed of loans securitized with GSEs, still begs the question of why we find interest rate increases in this sample. Banks may want to price flood risk in securitized mortgages for two reasons. First, even within the sample of securitized mortgages, there is residual credit risk because GSEs may ask banks to repurchase these mortgages (put-back risk) if GSEs find evidence for violations of representations and warranties (rep-and-warrants) obligations for securization.[7] After the financial crisis on 2007–08, many large banks were hit by substantial repurchase requests from GSEs and faced associated lawsuits and large fines. Second, banks face residual liquidity risk due to servicing obligations. Banks continue to service the loans in our securitized sample. Servicers for GSE securitized loans are required to advance payments to MBS investors even when mortgages are in delinquency. As

[7] When a loan is sold to GSEs, the originator represents and warrants that the loan meet respective GSE's eligibility and underwriting standards. The conditions for triggering repurchases or "put-backs" include misstatements, misrepresentations, omissions, and data inaccuracies in the loan documents; lack of clear title or first-lien enforceability; lack of compliance with relevant mortgage lending laws and GSE underwriting standards; and ex post triggers related to delinquencies for mortgages that are significantly seasoned (Hartman-Glaser, Stanton and Wallace, 2014).

carrying liquidity has opportunity cost for banks, this obligation may be factored into mortgage prices by banks when they securitize loans with GSEs.

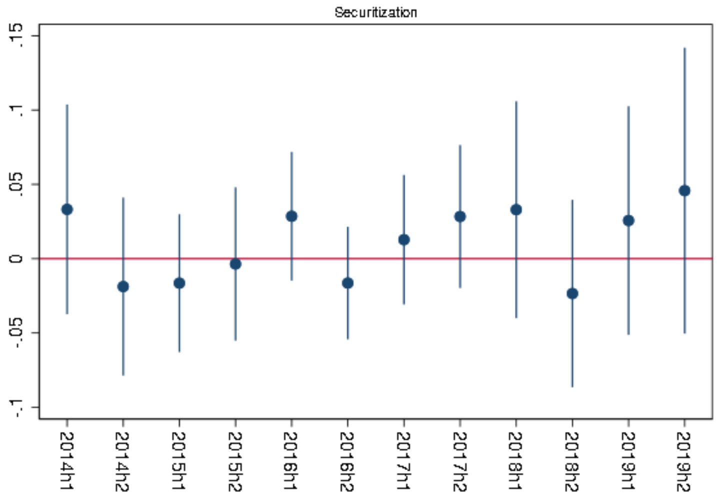

Figure 7: Dynamic baseline results for securitization

How big is the put-back risk? Goodmand and Zhu (2014) show that on average about 9 percent of dollar amount of loans originated and sold to Fannie Mae and Freddie Mac between 2006-2008 were put-back on banks. Outside of this crisis period, the highest put-back rate was observed in 2000 when cumulative put-back rate was about 0.5 percent of the original loan balance for that vintage. Using 40 percent loss rate, the authors calculate that originators experience lifetime losses of less than 20 basis points for mortgages originated in year 2000. Put-back rate was lower in 2010 then it was in 2000, but the traumatic experience of put-backs continued to linger on nation's large banks in 2010s and was documented to be instrumental in their reduced market share in mortgage originations in that decade (Buchak et al, 2018). As of 2013, right before our sample period, the GSEs have successfully closed on repurchases, indemnifications, and negotiated settlements valued in aggregate at \$46.12 billion of direct liability costs to the lenders who securitized with them.[8]

4.2 Robustness Tests

4.2.1 Testing the local randomness assumption for map changes

In this section, we test the validity of our identification assumption regarding local randomness of the map changes by leveraging additional geospatial analysis. One can argue that the strength of the local randomness of the map changes assumption weakens for properties that are further away from the new FEMA SFHA map borders. To address this concern, we check if our results are indeed driven by treated and control properties that are sufficiently close to the new flood map borders. We use geospatial analysis tools in ArcGIS to measure the distance of each of the treated and control properties to the nearest flood map border. Using

[8]See The Complete Guide to Mortgage Buyback Strategies, 5th Edition, Bethesda, MD, Inside Mortgage Finance Publications, 2013.

this distance measure, we construct a new dummy variable, $Close_j$, that is equal to one for properties for which distance to the border is below the 75th percentile of its distribution within each group of properties (that is, for example, we compare a treated properties distance to the border to the 75 percentile of the distribution of the distance measure for the treated properties). We then interact this $Close_j$ dummy with our treatment dummy in a dynamic triple-diff (or difference in difference in differences or DDD) regressions as follows:

$$\begin{aligned} Y_{ijt} &= \alpha_j + \mu_t + \lambda_c + \sum_{\tau} \beta_\tau Treat_i \times \mathbf{1}_{\tau=t} + \delta \mathbf{X_{it}} + \\ & \sum_{\tau} \theta_\tau Treat_i \times Close_j \times \mathbf{1}_{\tau=t} + \sum_{\tau} \eta_\tau Close_j \times \mathbf{1}_{\tau=t} + \epsilon_{ijt} \end{aligned} \tag{2}$$

In this dynamic triple-diff regression β_τ coefficients measure the difference in (conditional) means of the dependent variable between treated and control properties that are far from the border (these are properties for which $Close_j = 0$) for each time period τ, as shown in the left panel of Figure 8. We see that for treated properties that are far from the border, there is no statistically significant increase in interest rates around the map changes in 2016 and 2017 compared to control properties that are also far from the new borders.

Coefficients θ_τ measure how this difference additionally changes for properties that are sufficiently close the border. Therefore, the sum of these two coefficients ($\beta_\tau + \theta_\tau$) give us the the difference in (conditional) means of the dependent variable between treated and control properties that are close to the border (these are properties for which $Close_j 4 = 1$). These sums and their associated 90% confidence bands are displayed in the right panel of Figure 8, show that statistically significant increases in interest rates for treated properties compared to control properties that are both close to the new borders. In fact, we see that coefficient in 2017-H2 becomes larger than the baseline estimate and it has lower standard errors. Taken together, this analysis shows that our baseline results are indeed driven by these properties that are sufficiently close to the border, validating our local randomness assumption for map changes.

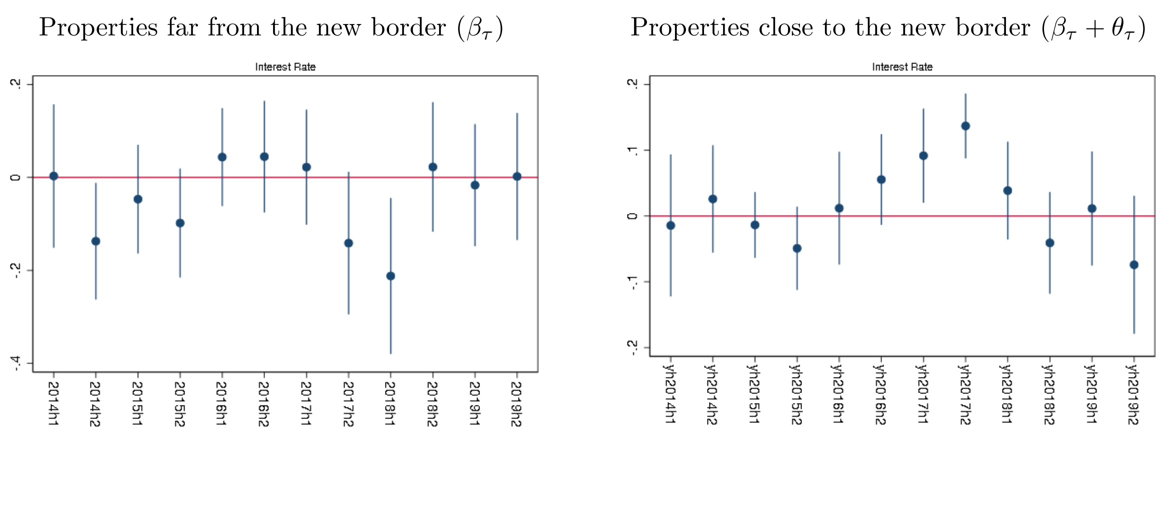

Figure 8: Dynamic triple-diff results to compare banks with and without a branch in NOL

4.2.2 Testing for alternative fixed effect and clustering specifications

First, we progressively add a set of fixed effects to further isolate the effect of flood insurance. We begin by controlling for bank*time fixed effects, thereby controlling for observed and unobserved time-varying heterogeneity across banks, including a bank's health or regulatory constraints. We expand the model by including census tract*time fixed effects in an effort to minimize the effect of time-varying factors related to demand for mortgages at the census tract level. We also experiment with fixed effects at the zip code level instead of the census tract level. All these tests support our earlier findings.

In addition, we interact additional loan characteristics with our time dummies to allow for changes in the relationship between outcome variables and controls. We also allow for rather conservative standard errors, which we double cluster at the census tract and day level. Our main results are virtually the same across these specifications.[9]

5 Inspecting the Mechanism

In this section, we shed light on the mechanism behind our results. Our hypothesis is that banks with superior knowledge about the local market may price climate risk differently than banks with inferior knowledge. We proxy local knowledge with the existence of branches in New Orleans and neighboring parishes using FDIC's Summary of Deposits (SOD) data. We then expand our baseline model with a triple interaction, which informs whether banks that have a branch behave differently than those that don't have a branch.

We estimate the following dynamic triple differences model

$$\begin{aligned} Y_{ijt} = {} & \alpha_j + \mu_t + \lambda_c + \sum_\tau \beta_\tau Treat_i \times \mathbf{1}_{\tau=t} + \delta \mathbf{X_{it}} + \\ & \sum_\tau \theta_\tau Treat_i \times Branch_j \times \mathbf{1}_{\tau=t} + \sum_\tau \eta_\tau Branch_j \times \mathbf{1}_{\tau=t} + \epsilon_{ijt} \end{aligned} \tag{3}$$

Our results are shown in Figure 9. Although we still observe an increase in mortgage rates by banks without a branch, the increase is more pronounced for banks without a branch. Economically, banks without a branch charge rates on average 13 basis points higher than banks without a branch for treated relative to control properties. Similar to our baseline results, we do not find differential effects for LTV ratios and securitization.

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 ten basis points more on the interest rates for mortgages originated on properties that were recently removed from the special flood hazard areas. This effect is temporary, lasting up to two years, and is stronger in banks with branches in Orleans and surrounding parishes, suggesting that local knowledge of the mortgage market plays a role.

[9] Robustness test results mentioned in this subsection are available upon request from the authors.

Banks without a branch

(β_τ)

Banks with a branch in NOLA

$(\beta_\tau + \theta_\tau)$

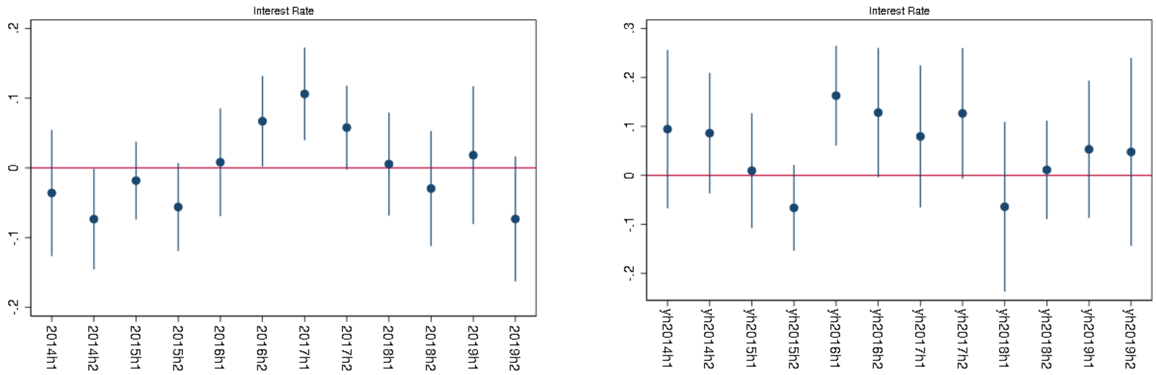

Figure 9: Dynamic triple-diff results to compare banks with and without a branch in NOLA

Future work will examine whether this local knowledge reflects a better understanding of the underlying risk of the property or whether it reflects local monopoly pricing power that is eroded over time due to competition.

References

[1] Blickle, K., and J. Santos, (2022) "Unintended Consequences of "Mandatory" Flood Insurance", *New York Fed Working Paper.*

[2] Bhutta, Fuster, and Hizmo (2019) "Paying Too Much? Borrower Sophistication and Overpayment in the US Mortgage Market", Manuscript.

[3] Buchak, G., Matvos, G., Piskorski, T., and Seru, A. (2018) "Fintech, Regulatory Arbitrage, and the Rise of Shadow Banks", *Journal of Financial Economics*, 130, 453– 83.

[4] Goodman, L. S., Zhu, J. (2014). "Reps and Warrants: Lessons from the GSEs Experience", *The Journal of Fixed Income*, 24(1), 29-41.

[5] Hartman-Glaser, B., Stanton, R., Wallace, N. (2014)."Mortgage Underwriting Standards in the Wake of Quantitative Easing", Working paper, UC Berkeley.

[6] Kara and Yook (*forthcoming*) "Policy Uncertainty and Bank Mortgage Credit", *Journal of Money, Credit, and Banking*

[7] Kahn, M. and A Ouazad, (*forthcoming*) "Mortgate Finance and Climate Change: Securitization Dynamics in the Aftermath of Natural Disasters", *Review of Financial Studies.*

[8] Kousky, C., M. Palim, and Y. Pan (2020). "Flood Damage and Mortgage Credit Risk: A Case Study of Huricane Harvey", *Journal of Housing Research* 29: S86-S120.

[9] Nguyen, D. D., S. Ongena, S. Si, and V. Sila, (*forthcoming*) "Climate Change Risk and the Cost of Mortgage Credit", *Review of Finance.*

[10] Pralle, S., (2019) "Drawing Lines: FEMA and the Politics of Mapping Flood Zones", *Climatic Change*, 152, 227-237.

[11] Sastry, P., (2021) "Who Bears Flood Risk? Evidence from Mortgage Markets in Florida", *MIT Mimeo.*

[12] Wing, O. et al (2018) "Estimates of Present and Future Flood Risk in the Conterminous United States", *Environmental Research Letters* Vol: 13 No:3.