Mortgage, Monitoring, and Flood Insurance Disincentive

Zhongchen Hu*

Chinese University of Hong Kong, Shenzhen

May 2023

Abstract

Flooding is the most costly natural disaster in the US. To protect collateral value against flood risk, many mortgage borrowers are thus required by law to maintain flood insurance. However, compliance is loosely enforced and lapsed policies are common. This paper hypothesizes that with a high monitoring cost borne by lenders, credit supply will depend on borrowers' insurance incentives. Exploiting an exogenous premium rise ($266 per year) which disincentivizes some borrowers to buy flood insurance, I show that lenders increase the corresponding mortgage denial rates by 0.8 percentage points (3.54% of the mean). This effect is gigantic, 80 times larger than that of lowering a borrower's annual income by $266. I rule out alternative demand-side explanations and provide evidence to support the mechanism that lenders internalize ex-post monitoring costs into ex-ante restrictions on credit supply.

*I am thankful for the advice of Ashwini Agrawal, Charles Cao, Isaac Hacamo, Claire Yurong Hong, Dirk Jenter, Wei Xiong, and Endong Yang. I am also grateful to seminar participants at University of Macau, Chinese University of Hong Kong-Shenzhen, and Shanghai Advanced Institute of Finance (SAIF) for helpful comments. Address for correspondence: huzhongchen@cuhk.edu.cn

1 Introduction

The US government requires mortgage borrowers—per the National Flood Insurance Reform Act of 1994—to purchase flood insurance if the property securing the loan is in a designated high-risk zone, known as a Special Flood Hazard Area (SFHA). This law is intended to ensure mortgage lenders' financial stability by protecting their collateral value against flooding, which is the most costly natural disaster in the US, having caused losses of over one trillion dollars in the past three decades.[1]

When borrowers fail to maintain sufficient coverage during the life of a loan, the law allows and obligates lenders to force-place flood insurance. In principle, lenders should have incentives to monitor. First, failure to do so would incur a penalty.[2] Second, more importantly, lenders are better off if the collateral is protected against flood risk via borrower-purchased insurance. Third, monitoring is essential in this setting, as US households are known to be very reluctant to acquire flood insurance due to, for example, behavioral biases (Kousky, 2011; Kousky et al., 2018; Hu, 2022).

However, the monitoring costs could be non-trivial. First, flood insurance policies are purchased and renewed annually by homeowners, not by banks. The need for renewal requires repeated checks because mortgage borrowers may buy flood insurance upon loan application but then let it lapse subsequently. Anecdotally, lapsed policies are common and non-compliance is an issue in practice (GAO, 2002; Dixon et al., 2006; FEMA, 2018). Second, the force-placement process, as described in Footnote 2, can be frictional and time-consuming. There have often been cases that banks paid penalties for not fulfilling their obligation (see Section 2.2).

With frictions in monitoring and enforcing flood insurance, this paper hypothesizes that credit supply to flood-prone residents will depend on borrowers' insurance incentives, as insurance actions will implicitly affect collateral value. Given that large asset value is at stake, it is important to understand how banks internalize the issue of households' insurance disincentive into their credit-supply decisions.

[1] The loss estimate is based on the National Oceanic and Atmospheric Administration's records of billion-dollar disasters; flood-related events include flooding, severe storm, and tropical cyclone.

[2] The lender must first issue a notice to the borrower that sufficient flood insurance should be obtained at his or her own expense. The borrower is then allowed 45 days to respond. If the borrower fails to comply within 45 days, the lender must purchase the flood insurance and assess the cost to the borrower. Failing to do so would incur a penalty for the lender of up to $2,000 per violation per loan. For example, in 2020, Citibank received a civil money penalty of $17,998,510 for Flood Act violations.

To test my hypothesis, I exploit an exogenous variation in borrowers' incentives to purchase flood insurance. My natural experiment leverages the passage of the Biggert-Waters Flood Insurance Reform Act of 2012 (BW12), which significantly increases the premiums for many properties in SFHAs (i.e., flood-prone areas where insurance is mandatory). For historical reasons (detailed in Section 2.1), the US government underwrites almost all flood insurance policies in the country through the National Flood Insurance Program (NFIP). Before BW12, highly subsidized premiums had been offered to properties in SFHAs for decades, whereas full-risk premiums were charged to other properties. Aiming to make the NFIP—\$17.75B in debt at the time—more sustainable financially, BW12 was enacted in July 2012 to phase out the subsidies for SFHA properties over five years. Effective January 2013, homeowners with subsidized insurance rates began receiving a 25-percent increase each year until their premiums reflect full-risk rates (see Section 2.3 for details of the law change and NFIP subsidies).

I assume that mortgage applicants in SFHAs, facing rising premiums as a result of BW12, will become more reluctant to acquire and maintain flood insurance, while the impact will be minimal outside SFHAs. This setting allows us to evaluate how banks deal with the federal mandate and how they adjust credit supply in response to changes in borrowers' insurance incentive. Ideally, I would compare mortgage application outcomes between otherwise identical SFHA and non-SFHA borrowers, before and after the law change. However, it is impossible to conduct this analysis at the loan level, as the data on mortgage applications—only available from the Home Mortgage Disclosure Act (HMDA)—does not provide information on a property's precise location or its flood zone designation. Hence, I aggregate my analysis at the level of census tracts. (There are 84,414 census tracts in the US, three times as many as there are ZIP codes.)

Specifically, my difference-in-differences strategy compares the mortgage denial rates between high- and low-SFHA census tracts, before and after the policy shock. Here, the treatment group is defined as census tracts having an above-median fraction of SFHA properties; that is, census tracts most affected by BW12. Figure 1 shows that the number of flood insurance policies in-force in the treated (high-SFHA) tracts has been declining dramatically since the new law, while the trend is fairly smooth in the control (low-SFHA) tracts.[3] This result suggests that the policy shock, BW12, indeed significantly

[3]In Figure 1, the turning point for the high-SFHA tracts appears to be about one year after BW12 was enacted, likely because policies purchased before the new law remained effective for one year.

alters households' incentive to acquire flood insurance in high-SFHA tracts, validating the premise of my empirical design.

The core finding of this paper is that in high-SFHA tracts, the denial rate of mortgage applications goes up by 0.80 percentage points following BW12, a 3.54-percent increase relative to the mean denial rate of 22.6 percent in 2012. In absolute terms, it implies 42,905 more mortgage applications being rejected every year because of the premium rise for flood insurance. This effect is gigantic. To get a sense of its magnitude, if applicants' incomes were lowered by the same amount that the premium rose ($266 per year or $22 per month on average), banks are estimated to increase the denial rate by only 0.01 percentage point. This suggests that my result is not simply driven by banks updating liquidity assessment for cash-constrained borrowers. I will, however, discuss this alternative explanation in more depth later.

The difference-in-differences construct of my analysis should tease out any time-invariant factors related to loan origination. For example, it is not a problem that high-SFHA tracts may have high denial rates due to high exposure to flood risk or that they may have low denial rates because properties in SFHAs (perhaps with water views) tend to be purchased by wealthy and high-credit individuals. However, to assess if there are any pre-existing time-varying factors that could drive my result, I run a dynamic regression. I find parallel pre-trends in denial rates between high- and low-SFHA tracts, supporting the identifying assumption that, absent the shock, credit supply in the treated and control tracts would have evolved in parallel. Figure 2, depicting the raw trends, provides intuitive graphical diagnosis; formal evidence is presented in Section 5.2.

An alternative explanation for my findings is that the increasing mortgage denial rates are driven purely by changes in the composition of borrowers. I provide both formal evidence and suggestive arguments to cast doubt on this alternative demand-side explanation. Formally, the difference-in-differences estimate remains almost identical when I include borrower characteristics, such as race, gender, income, loan amount, and loan-to-income ratio. Alternatively, I use these borrower characteristics as the dependent variable in the regression and obtain insignificant estimates, suggesting that the new law did not substantially alter the marginal borrowers' behaviors regarding whether or not to apply for mortgages in the first place. Consistently, I also find that the number of applications stays constant. Overall, these results suggest that the observed effect of

BW12 on mortgage denial rates is not driven by changes in borrowers composition.

Why did increasing premiums on flood insurance appear to have little impact on the demand for mortgages? Intuitively, home buyers are unlikely to change their minds about needing a mortgage when facing BW12's premium increase, as the price of flood insurance is probably just one marginal consideration in what, for most, is a once-in-a-lifetime situation. In fact, households are generally inattentive to flood risk and flood insurance (Chivers and Flores, 2002; Pope, 2008; Hu, 2022). Moreover, it is also plausible that borrowers are aware of the compliance friction; that is, that they have to buy flood insurance at loan application but can let it lapse. The key idea of this paper is that banks anticipate the high ex-post monitoring costs and hence restrict credit supply ex-ante; as a result, some marginal borrowers who would have obtained mortgages are rejected.

The other alternative explanation—in the absence of monitoring costs and enforcement frictions—is that the observed effect of BW12 on denial rates is just a result of borrowers becoming more cash-constrained due to the extra burden of more expensive flood insurance. I provide a set of evidence to quantify such an "income channel." First, I show that lowering a person's annual income by $1,000 is associated with only a 0.036-percentage-point higher denial rate.[4] Given that the post-BW12 insurance premium increased by $266, it would drive up the denial rate by only 0.01 percentage points through the income channel, drastically less than my finding. Second, if the income channel were the driving explanation, we would expect stronger effects in low-income areas. However, surprisingly, I find the opposite—high-income tracts experience stronger post-BW12 increases in mortgage denial rates. Third, in a similar spirit, I explore the heterogeneity across the income distribution within a given census tract. Again, I find that following the law change, lenders reject more applications from high-income rather than low-income borrowers. Overall, the income channel seems unable to account for the large impact of BW12 on mortgage denial rates.

Next, I provide further supplementary evidence to support my proposed mechanism; namely, that lenders internalize the ex-post cost of monitoring flood insurance compliance into their ex-ante decisions on mortgage origination. First, my hypothesis predicts that the increase in mortgage denial rate will be larger in areas where the compliance issue is

[4]This estimate is obtained from my census-tract-year-level regression with the average annual income as the explanatory variable and the denial rate as the dependent variable. Alternatively, a loan-level regression implies an even smaller estimate of 0.012 percentage points. See Section 5.3 for details.

expected to become more severe following BW12. I assume that the demand for flood insurance is relatively inelastic for households who are actually worried about flood risk (or more generally, global warming). I leverage data from the Yale Climate Opinion survey (Howe et al., 2015) to create subsamples of census tracts with a high or low fraction of global warming believers.[5] In the subsample with high (low) flood insurance elasticity, the mortgage denial rate increases by 1.03 (0.58) percentage points, consistent with my hypothesis that lenders in locations with weak climate risk perception will be more concerned about borrowers not complying with flood insurance requirements.

My second set of supporting evidence is that my result is much stronger for local banks than for national banks. This is in line with the prediction of my proposed mechanism that local banks, which hold a less-diversified portfolio of mortgages, are likely to be more cautious about flood risk. Lack of diversification is particularly important in the context of flood risk because of its correlated and catastrophic nature—if a flood occurs at all, it tends to severely affect a large number of properties in a small region—which can be devastating for small local banks. Specially, I partition the HMDA data into mortgage applications made to local or to national banks, using four different measures (detailed in Section 5.5.2). I find that following BW12, local banks increase their mortgage denial rates two to three times more than national banks.

Third, my hypothesis predicts that home improvement loans, normally unsecured, will not be significantly affected in my natural experiment. For all the findings discussed above, I consider only purchase and refinance mortgages, in accordance with the idea that lenders worry about their collateral values being affected by lapsed flood insurance. Here, I use home improvement loans—not secured by property—as a verification check; I do not find mortgage lenders adjusting denial rates for these loan applications.[6]

Fourth, as Ouazad and Kahn (2021) show that lenders transfer flood risk through securitization to government-sponsored enterprises (GSEs), we expect the effect of BW12 to be weaker for GSE loans because lenders will not worry as much about underinsurance

[5]As the Yale Climate Opinion survey is granular only at the county level, I use the average for all census tracts from the same county. My main measure is the percentage of people who answered "Yes" to the question "Do you think global warming is happening?" My results are robust to using an alternative question "Are you worried about climate change?"

[6]While my hypothesis does not imply an obvious distinction between purchase and refinance mortgages, I examine them separately in Section 5.5.3. My result holds in both cases. While the absolute increases in denial rates are of similar magnitude, in terms of percentage change, the increase is greater for purchase loans because of a smaller baseline denial rate.

with the risk being borne by the GSEs. Consistently, I find that, in census tracts with high (low) prevalence of GSE loans, the mortgage denial rate increases by 0.45 (1.07) percentage points following BW12.

In addition to the accept-or-reject decision (i.e., the extensive margin), I examine two intensive-margin decisions for lenders—loan-to-value (LTV) ratios and interest rates, both conditional on acceptance. As the HMDA data lacks information on these two variables, I use alternative data sources from two GSEs—Fannie Mae and Freddie Mac. However, as they provide less-granular information on property location and have various limitations, my analysis on LTVs and interest rates is admittedly less identified and is therefore to be considered supplementary.

Fannie Mae and Freddie Mac's loan performance data have the advantage of being comprehensive and having detailed contract terms, but there are two drawbacks: (a) apparently, only mortgages purchased or guaranteed by Fannie Mae or Freddie Mac are included and (b) census tracts are unavailable and address information is quite coarse, disclosing only the first three digits of a ZIP code. Hence, I have to aggregate my analysis at the three-digit-ZIP level. To ensure that this is a meaningful aggregation for my application, I first replicate my main extensive-margin result at the ZIP level. Specifically, I transform the tract-level SHFA data to the three-digit-ZIP level, using the crosswalk provided by the US Department of Housing and Urban Development.[7] I show that high-SFHA ZIP codes experience a 0.55-percentage-point increase in mortgage denial rate (using HMDA application data). However, conditional on approval, I find no discernible changes in interest rates or LTVs. This result is consistent with the finding of Ouazad and Kahn (2021) that lenders transfer flood risk through securitization to Fannie Mae and Freddie Mac. Due to data limitations, however, it is unclear whether the intensive-margin result can be generalized to other unsecuritized loans.

This paper contributes to the nascent but fast-growing literature on climate finance (Hong et al., 2020; Giglio et al., 2021), which studies the interactions between climate risk and financial markets. A number of studies examine whether climate risk—in particular, sea level rise—is capitalized into real estate values (Keenan et al., 2018; Bernstein et al., 2019; Baldauf et al., 2020; Murfin and Spiegel, 2020; Addoum et al., 2021; Bakkensen and Barrage, 2021; Giglio et al., 2021; Hino and Burke, 2021). Relatively less is known

[7]The crosswalk data is downloaded from www.huduser.gov/portal/datasets/usps_crosswalk.html. See Wilson and Din (2018) and Din and Wilson (2020) for detailed references.

about the implications of climate risk for the mortgage market, despite its key role in the financial system's stability (Mian and Sufi, 2015), and mixed evidence has been documented. With a few exceptions (Keys and Mulder, 2020; Garbarino and Guin, 2021), most studies suggest that banks do incorporate climate risk into their decision-making (Ouazad and Kahn, 2021; Bakkensen et al., 2022; Blickle et al., 2022; Nguyen et al., 2022; Sastry, 2022). While my result is consistent with the prior findings that banks respond to increasing flood risk, the new insight provided by this paper is that, holding the risk constant, borrowers' willingness to purchase insurance is also an important factor affecting banks. To the best of my knowledge, my paper is the first to examine the impact of flood insurance non-compliance on the mortgage market.

This paper also sheds light on the policy debate on reforming the NFIP; in particular, on the controversy over eliminating subsidies and moving to risk-based rates (Young, 2008; Verchick and Johnson, 2013; Knowles and Kunreuther, 2014; Kousky and Kunreuther, 2014; Lemann, 2015; Nance, 2015; Wriggins, 2015; Horzempa, 2018; Colby and Zipp, 2021). While this strand of the literature typically focuses on the insurance market alone, with opponents of BW12 emphasizing insurance affordability and proponents urging moral hazard and the NFIP's insolvency, my paper is the first to highlight the unintended spillover effect from the insurance market to the mortgage market.

Finally, this paper broadly relates to the literature on how collateral value affects credit availability. Most prior studies focus on corporate borrowing. Benmelech et al. (2005), Benmelech (2009), Benmelech and Bergman (2009, 2011), and Gavazza (2010) exploit variation in the redeployability of assets as a determinant of liquidation value. Cerqueiro et al. (2016) examine a legal reform that exogenously reduced collateral value through the priority structure of claims. On the household side, several papers study the relationship between collateral value uncertainty (measured by house price dispersion) and mortgage credit availability (Lang and Nakamura, 1993; Marschoun, 2000; Jiang and Zhang, 2022). My paper suggests that, in the mortgage market, collateral value is also affected by the borrowers' willingness to acquire flood insurance.

The remainder of the paper is organized as follows. Section 2 presents the institutional background. Section 3 describes the data. Section 4 explains my identification strategy. Section 5 presents the main findings and robustness checks. Section 6 concludes.

2 Institutional Background

2.1 The NFIP: Brief History

In the early 20th century, a small number of private companies offered flood insurance in the US, but the Great Mississippi Flood of 1927 and additional riverine floods in 1928 essentially marked the end of the private flood insurance market (King, 2005). Private insurance companies largely concluded that flood peril was uninsurable, primarily because of the correlated, catastrophic nature of flood damage, the difficulty of determining accurate rates, and the concern of adverse selection (Horn and Webel, 2019). From the 1930s to the 1960s, major flood events were mainly covered by government disaster relief.

In 1968, Congress passed the National Flood Insurance Act and created the National Flood Insurance Program (NFIP), housed under the Federal Emergency Management Agency (FEMA).[8] The program was designed as a partnership between the federal government and local communities. Communities can voluntarily join the program by adopting a floodplain ordinance based on the flood hazard maps provided by FEMA; only residents in participating communities can purchase federal flood insurance. Premiums for these policies vary by flood risk zone, based on the Flood Insurance Rate Maps (FIRM) issued by FEMA. In particular, the high-risk floodplains on the FIRM—those with at least a one-percent probability of flooding in a given year—are referred to as the Special Flood Hazard Areas (SFHAs).

To encourage insurance take-up, the NFIP offered significant subsidies (i.e., discounted premiums that do not reflect actuarially fair pricing) to homeowners, especially to those living in floodplains. FEMA estimates that the subsidized rates are roughly 40 to 45 percent of the full-risk rates (Hayes and Neal, 2011). Two main categories of property pay subsidized rates: (a) *pre-FIRM* properties, built before FEMA published the first FIRM for their community, and (b) *grandfathered* properties, built in compliance with the FIRM that was in effect at the time of construction. Grandfathered properties are allowed to maintain their old lower insurance rates even if they are subsequently remapped into a riskier class. Historically, these heavy government subsidies have given private insurers even less incentive to re-enter the market.[9]

[8] The NFIP was initially housed under the Department of Housing and Urban Development. In 1979, President Carter signed an order creating FEMA, whereupon the NFIP was placed under its aegis.

[9] For a comprehensive overview of the US flood insurance market and the NFIP's history, see Abbott (2008), Michel-Kerjan (2010), Subcommittee (2011), Platt (2012), and Knowles and Kunreuther (2014).

2.2 Mandatory Purchase of Flood Insurance

At the inception of the NFIP, the purchase of flood insurance was completely voluntary. However, because of unexpectedly scant participation by communities and homeowners, Congress amended it in 1973 and in 1994 to require the purchase of flood insurance by many homeowners and to place the onus for ensuring compliance throughout the loan life upon lending institutions. Specifically, with the enactment of the Flood Disaster Protection Act of 1973, flood insurance became compulsory for properties in SFHAs with mortgage loans made or held by federally regulated lending institutions or guaranteed by federal agencies.[10] Two decades later, mortgages purchased by government-sponsored enterprises (such as Fannie Mae and Freddie Mac) were also included, under the National Flood Insurance Reform Act of 1994.

Yet not everyone required to have flood insurance complies. While there is no definitive data on the national compliance rate, evidence suggests that compliance is not universally enforced (GAO, 2002; Kriesel and Landry, 2004; Dixon et al., 2006; Kousky, 2011; Michel-Kerjan et al., 2012; FEMA, 2018). Several studies commissioned by FEMA find that the compliance rate may be less than half in some areas. Banks have often had to pay civil money penalties for violating the "Flood Act."[11] Such examples suggest that the mandatory purchase requirement is not binding in practice, but do not imply that the federal regulators conduct thorough checks on banks. Rather, they acknowledge that large-scale examinations are neither practical nor cost-effective (GAO, 2002).

2.3 The Biggert-Waters Reform

Although the NFIP heavily subsidized the oldest, highest-risk properties (as previously detailed in Section 2.1), FEMA believed that the overall revenues of the program should be enough to cover losses from the "average historical loss year," which was defined as

[10] A federally regulated lending institution is any bank, savings and loan association, credit union, farm credit bank, federal land bank association, production credit association, or similar institution supervised by a federal entity for lending regulation.

[11] Documents of such penalties can be found on the websites of the federal bank regulators—such as the Federal Deposit Insurance Corporation (FDIC), the Federal Reserve Board (FRB), the Office of the Comptroller of the Currency (OCC), the Consumer Financial Protection Bureau (CFPB), and the Office of Thrift Supervision (OTS)—which oversee the lending institutions. To name a few recent examples: Citibank (OCC-regulated) paid a penalty of $17,998,510 in 2020, USAA FSB (OCC-regulated) paid a penalty of $382,500 in 2020, Bryn Mawr Trust (FRB-regulated) paid a penalty of $105,000 in 2021, Park Bank (FDIC-regulated) paid a penalty of $12,841.50 in 2020, and State Farm Bank FSB (OCC-regulated) paid a penalty of $547,200 in 2020.

the mean annual loss over the life of the program.[12] Before 2005, the financial balance of the NFIP—calculated as the cumulative difference between collected premiums and paid claims nationwide—had remained steadily around break-even, as no catastrophic flooding event had yet occurred in the program's history. However, after Hurricane Katrina in 2005, the NFIP's deficit increased from \$1.5 to \$20.8 billion.[13] After Hurricane Sandy in 2012, it was more than \$30 billion in debt to the US Treasury. The program was in crisis.

Such high debt generated broad interest in reform. In July 2012, Congress passed the Biggert-Waters Flood Insurance Reform Act (BW12) to reauthorize and overhaul the NFIP. The underlying principle of the Act was that flood insurance premiums should be risk-based and subsidies should be eliminated to strengthen the NFIP's fiscal soundness.

Under this new legislation (effective January 1, 2013), the key changes to the pricing of flood insurance are as follows. First, for business properties, non-primary residences, and properties that have suffered severe repetitive losses, the historically discounted premiums will increase 25 percent each year until they reflect full-risk rates. Second, subsidies will no longer be offered for policies covering newly purchased properties, lapsed policies, or new policies covering properties for the first time. Third, the grandfathering subsidies will also be eliminated. That means that homeowners will no longer have the option to use risk data from previous flood hazard maps to get lower premiums when their risk designation increases; for these properties, higher rates will be phased in at 20 percent per year until the full-risk price is reached (per up-to-date maps). As the pre-FIRM and grandfathering subsidies were historically concentrated in SFHAs, this change in law affected the high-SFHA census tracts the most.

BW12 quickly faced backlash from some homeowners, local governments, and interest groups. Homeowners complained that the new maps overestimated their risk, that the premium increases were unjustified, and that they could not afford the more expensive flood insurance. In response to the public pressure, in March 2014, Congress passed the Homeowner Flood Insurance Affordability Act (HFIAA), rolling back some of the provisions of BW12.

HFIAA is not a radical departure from BW12 and retains most of its NFIP reforms,

[12] The NFIP was given borrowing authority from the US Treasury to cover higher-than-average loss years. In contrast, an actuarially fair premium set by a private insurance company would factor potential catastrophic loss years into risk-based rates.

[13] Appendix Figure A.1, taken from Michel-Kerjan (2010), depicts the NFIP's financial balance over time. It intuitively shows that the program took a dramatic hit from Hurricane Katrina.

including the 25-percent increase in annual premiums for businesses properties, non-primary residences, and properties with severe repetitive loss. For most primary residences, HFIAA reduced the phaseout rate of the pre-FIRM subsidies to no more than 18 percent per year.[14] The other revision was to reinstate grandfathering. Overall, HFIAA slowed down BW12 rather than repealing it. With the NFIP's deep deficit, the general view was that it seemed inevitable for the program to eliminate subsidies and move to risk-based pricing.

3 Data

This paper uses data from various sources. In this section, I describe details of my sample construction and provide descriptive statistics for the main variables.

3.1 The Home Mortgage Disclosure Act

The main data source is the universe of mortgage applications from HMDA between 2007 and 2016. The law mandates disclosure reports from both depository institutions (such as banks, credit unions, and savings associations) and non-depository institutions that meet the criteria for coverage.[15] The application-level information includes lender identity,[16] year of application, loan amount, borrower income, borrower race and ethnicity, census tract of the house, property type (1–4 family, manufactured housing, or multifamily), loan type (conventional, FHA, VA, or FSA), loan purpose (home purchase, home improvement, or refinancing), occupancy status (owner-occupied or not), and—importantly—the outcome of the application (approved and accepted, approved but not accepted, denied, withdrawn, or closed for incompleteness).

This paper focuses on home purchase and refinancing mortgages for 1–4 family housing and excludes applications that are withdrawn by the applicants or closed for incompleteness. Aggregating at the census-tract-year level, the key outcome variable is the

[14]See Table 4 of Horn and Brown (2018) for a comprehensive summary of the phaseout schemes across different types of property under BW12 and HFIAA.

[15]The primary criteria (among others) are that (a) the institution has an office in a metropolitan statistical area and (b) its total assets exceed the minimum exemption threshold set annually by the regulatory authority (for example, the threshold was $37 and $10 million in 2007 for depository and non-depository institutions, respectively). More details can be found from the official guidelines (see www.ffiec.gov/hmda/pdf/2010guide.pdf).

[16]The unique lender identifier `respondentid` in the HMDA data can be matched to the `RSSDID` used in other datasets of banks, such as the FDIC Summary of Deposits detailed in Section 5.5.2.

application *denial rate*—the number of rejected applications in a given year divided by the total number of applications that receive lender decisions. The covariates include the average borrower income, the average loan amount, the share of white applicants, the share of male applicants, and the average loan-to-income ratio.[17]

3.2 NFIP Data and Definition of Treatment Groups

The NFIP data on flood insurance policies is maintained by FEMA. I obtain more than 50 million transaction-level observations between January 2009 and August 2019. I observe information on the contractual terms, such as policy effective and termination dates, premiums, coverage, and deductibles. I also observe information on the insured property's characteristics, such as census tract, flood zone code (which indicates if the property is in an SFHA or not), elevation, number of floors, and construction date.

As of 2019, the program covers all 50 states and 3,053 (out of 3,143) counties. In an average month of the sample period, there are 5.29 million policies in-force nationally, of which 52.2 percent covers SFHA properties; the program collects $3.32 billion in premiums for $1.26 trillion in coverage. At the individual policy level, the average annual premium is $628 and the average coverage is $238,000.

I calculate the proportion of SFHA policies for each census tract using only pre-BW12 data. The 25th, 50th, and 75th percentiles are 0.05, 0.38, and 0.68, respectively. In my baseline specification, the treated census tracts are defined as having an above-median fraction of SFHA policies, which are, in principle, more affected by BW12, as previously discussed in Section 2.

Table 1 presents the descriptive statistics at the census-tract level. Over the full sample, the average fraction of SFHA properties is 40 percent. On average, 83 percent of the mortgage applicants are white, 55 percent are male, and 25 percent of the applications are rejected. The average annual income is $100.9 thousand, and the average loan size is $196.6 thousand. Comparing the treated (high-SFHA) and control (low-SFHA) tracts, by construction, the treated tracts have a larger fraction of SFHA properties (68 percent) than the control (11 percent). Their differential exposure to flood risk can partially

[17] For an application with co-applicants, I define variable *white* to take a value of 0 if both co-applicants are non-white, 0.5 if only one of them is white, or 1 if both are white. The variable *male* is defined similarly at the application level. These variables are then averaged at the census-tract level to calculate the share of white and male applicants.

explain the stylized fact that the baseline mortgage denial rate is higher in high-SFHA tracts (26 percent) than in low-SFHA tracts (24 percent). The other explanation is that poorer households tend to live in flood-prone areas. The average applicant income in high-SFHA tracts is \$95.1 thousand, about 10 percent lower than that in low-SFHA tracts (\$106.8 thousand). As long as these intrinsic differences are time-invariant and unrelated to the policy shock, my difference-in-differences strategy will account for them. Sections 5.1 and 5.2 shows evidence supporting this assumption.

3.3 Additional Mortgage Data

As HMDA data lacks information on a loan's contracting structure (in particular, interest rate and loan-to-value ratio), I use other data sources to fill the gaps, at least partially. Below, I describe these supplementary datasets and discuss their advantages and limitation.

The Fannie Mae Single-Family Loan Performance Data and Freddie Mac Single-Family Loan-Level Dataset provide granular information on loan pricing, loan characteristics, and monthly loan performance for conforming loans—a subset of the loans owned or guaranteed by Fannie Mae or Freddie Mac that meet the dollar conforming-loan limit set by the Federal Housing Finance Agency and the funding criteria of Fannie Mae and Freddie Mac. The Fannie Mae dataset includes its single-family, 30-year fixed-rate, fully amortizing, and fully documented conventional mortgages. The Freddie Mac dataset has similar sample selection except that it also includes 15-, 20-, and 40-year fixed-rate mortgages.

The key variables for my analysis are the interest rate and LTV when the loan is originated. These two datasets also provide detailed information on a rich range of borrower and loan characteristics, such as borrower credit score and loan maturity, which are missing in HMDA. However, the disadvantages are at least twofold. First, only conforming loans are included. Second, the property location information is truncated; only the first three digits of the ZIP code are presented. I therefore aggregate the data at the three-digit-ZIP-month level.

4 Empirical Strategy

Leveraging BW12 (detailed in Section 2.3) as an exogenous shock that reduces the incentive of many borrowers living in flood-prone zones to acquire and maintain flood insurance, I use the following difference-in-differences regression to evaluate how lenders adjust their accept-or-reject decisions in response:

$$DenialRate_{it} = \beta_0 + \beta_1 * Treated_i \times Post_t + \beta_2 * Treated_i + \beta_3 * Post_t + \epsilon_{it}. \tag{1}$$

$DenialRate_{it}$ is the denial rate of mortgage applications in census tract i in year t, as defined in Section 3.1. $Treated_i$ is the treatment dummy, as defined in Section 3.2, indicating if census tract i is a high-SFHA area; the treated tracts are expected to be affected the most by BW12, while the control (low-SFHA) tracts are expected to be largely unaffected, consistent with the ex-post evidence shown in Figure 1. $Post_t$ is the post-event dummy that turns on if year t is after the passage of BW12. The interaction $Treated_i \times Post_t$ is the key explanatory variable of interest. Its coefficient β_1 is a difference-in-differences estimate, which measures the differential change in mortgage denial rates for the high-SFHA tracts relative to the low-SFHA tracts, following the premium increase that is disproportionally imposed by BW12 on high-SFHA tracts. To account for serial correlation, I cluster the standard errors at the census-tract level in all specifications.

In order to interpret the estimate of β_1 as the causal effect of borrowers' flood insurance disincentive on credit supply, I must assume that mortgage denial rates in treated and control tracts would have evolved in parallel in the absence of the new law. I test for parallel pre-trends by replacing $Post_t$ in Regression (1) with a sequence of time dummies $\{\mathbb{1}_{t=\tilde{t}}\}$ and running the following dynamic version:

$$DenialRate_{it} = \beta_0 + \sum_{\tilde{t}=2007}^{2016} \beta_{\tilde{t}} * Treated_i \times \mathbb{1}_{t=\tilde{t}} + \beta_2 * Treated_i + \alpha_t + \epsilon_{it}. \tag{2}$$

The coefficients $\{\beta_{\tilde{t}}\}$ on the interaction terms $\{Treated_i \times \mathbb{1}_{t=\tilde{t}}\}$ allow me to examine the patterns in denial rates in the years before and after BW12. $\{\beta_{\tilde{t}}\}_{\tilde{t}<2012}$ correspond to the pre-trends and $\{\beta_{\tilde{t}}\}_{\tilde{t}\geq 2012}$ capture the dynamic effects of borrowers' decreasing incentive to acquire flood insurance. These effects are measured relative to β_{2011}, which is omitted.

In Sections 5.1 and 5.2 below, I present direct evidence for the validity of the parallel pre-trends assumption by showing that the mortgage denial rates for high- and low-SFHA tracts move together before BW12 and diverge only afterward.

One concern with this specification is that BW12 may have altered the behaviors not only of lenders but also of borrowers. As a result, changes in mortgage denial rates after BW12 may be driven by changes in the composition of applicants rather than by lenders adjusting their credit-supply decision-making. I address this concern by estimating specifications that also include borrower characteristics, such as race, gender, income, and loan-to-income ratio. Section 5.2 shows that results from these specifications are not meaningfully different from those of specifications that do not include covariates.

5 Empirical Findings

This section presents the main estimates of the effect of borrowers' weakening incentive to comply with mandatory flood insurance on credit supply. As an initial assessment of the validity of the parallel pre-trends assumption, I begin the section by presenting simple graphical evidence, depicting the raw trends of mortgage denial rates in the high- and low-SFHA tracts. To quantify the causal effect of interest, I then present a series of formal difference-in-differences estimates. Moreover, to support the hypothesized mechanism, I provide supplementary evidence which examines the heterogeneity of my results across locations with different demand elasticity for flood insurance and across different types of lender. Finally, in addition to my main results focusing on the extensive margin of lenders' accept-or-reject decisions, I examine two key intensive-margin decisions for lenders—loan-to-value ratios and interest rates—both conditional on loan origination.

5.1 Graphical Evidence

Figure 2 plots the yearly average mortgage denial rates separately for low-SFHA tracts (dashed line, left y-axis) and high-SFHA tracts (solid line, right y-axis). For visual convenience, the right y-axis is one percentage point higher than the left y-axis, artificially shifting the raw trend for high-SFHA tracts downwards by one percentage point. Consistent with the aggregate descriptive statistics presented in Table 1, in any given year, lenders reject more mortgage applications for properties in high-SFHA tracts than for those in low-SFHA tracts, which likely reflects inherent differences across localities in terms of exposure to flood risk and/or borrower characteristics.

Crucially, the two lines move almost exactly in parallel before 2012 (with a constant gap of roughly one percentage point). However, the trends diverge significantly after 2012, with the denial rate growing much more quickly in high-SFHA tracts than in low-SFHA tracts. Four years after BW12, the gap has widened by more than one percentage point. This result provides strong support for the validity of the parallel pre-trends assumption underlying the difference-in-differences estimates that follow.

5.2 The Effect on Mortgage Denial Rate

To more precisely quantify how lenders adjust their mortgage origination decisions in response to a weakening incentive for borrowers to comply with mandatory flood insurance, Table 2 presents estimates from the difference-in-differences specification given by Regression (1). Column 1 reports estimates from the baseline specification that includes only the $Treated_i$ indicator, the $Post_t$ indicator, and their interaction. The coefficient estimate on $Treated_i$ implies that before BW12, mortgage applications in high-SFHA tracts were more likely to be rejected than those in low-SFHA tracts by 1.20 percentage points on average, consistent with the magnitude of the pre-trends gap shown in Figure 2. The coefficient estimate on $Post_t$ implies that during 2012–2016, the nationwide average mortgage denial rate was 4.89 percentage points lower than it was during 2007–2011, again corroborating the graphical evidence in Figure 2.

As the key result of interest, the coefficient estimate on $Treated_i \times Post_t$ suggests that following BW12, lenders increased the denial rate by 0.80 percentage points for high-SFHA tracts. This estimate is statistically highly significant and economically large—a 3.54-percent increase over the mean denial rate of 22.6 percent in 2012. In Section 5.3 below, I show that this change is gigantic compared to how much lenders would increase the denial rate were applicant income lowered by an amount similar to the premium increase.

In Columns 2 through 5 of Table 2, I add a series of control variables to my baseline Regression (1). I first add several covariates pertaining to borrower characteristics, which can be time-varying and potentially driving my main result; namely, the share of white applicants, the share of male applicants, the average applicant income (in $1,000), and the average loan-to-income ratio. In Column 4, I further add a set of fixed effects for the 50 states. This specification removes the influence of average differences in the levels of denial rates across states and identifies the effect of BW12 by comparing mortgage denial rates for high- and low-SFHA tracts within the same state before and after the law change. Column 5 not only allows for average differences in the levels of denial rates but also allows the aggregate trends to vary across states by interacting state fixed effects with year fixed effects. It is reassuring to see that the key difference-in-differences estimates are all highly significant and relatively stable across specifications. Appendix Table A.1 uses the outcome variable in logarithm and documents findings similar to those in Table 2.

To understand the dynamics of the effect of borrowers' weakening flood insurance incentive on mortgage denial rates, Figure 3 plots estimates from the more-flexible difference-in-differences specification given by Regression (2). Each dot represents the point estimate for the coefficient $\beta_{\tilde{t}}$ on the interaction term $Treated_i \times \mathbb{1}_{t=\tilde{t}}$, which is estimated relative to year 2011 (i.e., the year immediately before BW12 is the omitted category in the regression); the bins around the coefficient estimates represent the 95-percent confidence intervals. As we can see, mortgage denial rates for high-SFHA tracts diverge sharply from the benchmark starting precisely in the year that BW12 is enacted. The effect grows over time. In particular, it has not yet reached its new steady state by the end of 2012. This at least in part reflects the fact that the HMDA data pools both pre- and post-BW12 applications in 2012 and that the provisions of BW12 only become effective in January 2013. However, soon after that, the effect appears to materialize quickly and stabilizes at around 1 percentage point between 2014 and 2016.

Importantly, the trends are statistically indistinguishable in the period prior to BW12 and only begin to diverge afterwards. This supports the parallel pre-trends assumption required for identification in the difference-in-differences research design. In Appendix Figure A.2, I show that the parallel pre-trends also hold in specifications that include the control variables, as in Columns 2 through 5 of Table 2.

5.3 The Income Channel and Ability to Repay

To get a sense of the magnitude of my main result, this section asks how much lenders would increase mortgage denial rates if applicants' incomes were lowered by an amount similar to the BW12 premium increase. This exercise can also help evaluate an alternative explanation; namely, that the effect of BW12 on denial rates is driven by lenders updating liquidity assessment in response to borrowers becoming more cash-constrained due to the extra burden of more expensive flood insurance. This alternative mechanism—referred to here as the "income channel"—could potentially explain my findings in the absence of monitoring costs and frictions in enforcing mandatory flood insurance.

I provide a set of evidence suggesting that the income channel is unlikely to be the driving factor. First, in Table 2, the coefficient estimates of $Income_{it}$—the average annual income (in \$1,000) of mortgage applicants in census tract i in year t—imply that the explanatory power of $Income_{it}$ is much smaller than the main effect. Specifically, lowering the average annual income by \$1,000 is associated with only an increase of 0.034–0.036 percentage points in denial rates. To put that into perspective, we should—knowing that, five years after BW12, the annual premium for SFHA-properties had increased by \$266 on average (see Appendix Figure A.3)—scale the coefficient estimates of $Income_{it}$ from Table 2 by a factor of 0.266. The premium rise caused by BW12 would thus lead—through the income channel—to an increase in denial rates of approximately only 0.01 percentage points, drastically smaller than the main difference-in-differences estimates ranging from 0.47 to 0.81 percentage points.

Second, loan-level regressions confirm the relatively weaker correlation between applicant income and mortgage rejection, compared to the BW12 effect that I document. Appendix Table A.2 regresses the loan-level outcome indicator (equal to 1 if the application is rejected) on the applicant's income and other characteristics. The coefficient estimates of income, stable across all specifications, suggest that if an applicant's annual income decreases by \$1,000, his or her mortgage application is 0.012 percentage points more likely to be rejected. Therefore, my results are unlikely to be driven by the income channel. In other words, if lenders had updated their accept-or-reject decisions only out of concern that the extra expenditure of \$266 on flood insurance would cause some SFHA borrowers to default, we would not have seen such a dramatic increase in denial rates.

Third, if the income channel were the driving explanation, we would expect stronger effects in low-income areas. Moreover, this heterogeneity should be monotonic. To test this prediction, I partition my full sample into quartiles based on the average borrower's income at the census-tract level; I then run the baseline difference-in-differences regression for each subsample. Table 3 presents the results. Inconsistent with the prediction of the alternative explanation, the difference-in-differences estimates are not monotonically decreasing from the bottom to the top income quartile. Quite the opposite, they appear to have a somewhat increasing pattern or an asymmetric inverted U-shape (with high-income tracts experiencing stronger post-BW12 increases in mortgage denial rates).

Finally, in the same spirit as above, I conduct an alternative heterogeneity analysis across the income distribution. Here, instead of partitioning my full sample based on the tract-level mean statistics, I condition on a given census tract and analyze how the denial rates change differently for different income groups. Specifically, in any given census tract and year, I divide all mortgage applications into quartiles according to the applicant's income and then calculate the denial rate for each income group, respectively. The treatment indicator is defined, as before, at the census-tract level. I repeat my main difference-in-differences regression four times, with the outcome variable being the mortgage denial rate for income quartile q in census tract i during year t. As an initial assessment, Figure 4 compares the raw trends between high- and low-SFHA tracts, separately for each income quartile. This graphical evidence provides several insights: (a) the parallel pre-trends assumption holds in all specifications; (b) in any given year, the denial rate is monotonically higher for borrowers with lower income, consistent with intuition; (c) the causal effect of BW12 on mortgage denial rate is not, however, monotonic across income quartiles. More formally, Table 4 presents the regression results, which consolidate the surprising finding that the effect is actually stronger for higher-income borrowers with who should have fewer liquidity concerns.

In essence, the evidence discussed above all suggests that my findings cannot be explained solely by the income channel. To be clear, I do not intend to refute the income channel, which is certainly a key consideration for lenders. However, in the context of my natural experiment, its magnitude seems second-order and insufficient to account for the large impact of BW12 on mortgage denial rates.

5.4 Changes in Borrowers Composition

The other alternative explanation for my findings is that the increasing mortgage denial rates are driven by changes in the composition of borrowers after the law change, while lenders' decision-making processes stay constant. Below, I provide both suggestive arguments and formal evidence against such an explanation.

First, I argue that home buyers are unlikely to feel less need for mortgages when facing more expensive premiums than before, as flood insurance is probably, for most of them, just one marginal consideration in a once-in-a-lifetime situation. Supporting this view from a broader perspective, the literature has found households generally inattentive to flood risk and flood insurance (Chivers and Flores, 2002; Pope, 2008; Hu, 2022). Second, considering that one can, in practice, purchase flood insurance at loan application but then let it lapse, I consolidate my argument that BW12 is unlikely to change home buyers' decisions on whether or not to apply for a mortgage. Third, if applicant composition does change, intuition suggests it would be the riskier (e.g., cash-constrained) borrowers who drop out or borrow less. The remaining—on average safer—borrowers would then lead to a lower mortgage denial rate in equilibrium, contradicting the stylized fact I document.

More formally, as Table 2 shows, the main difference-in-differences estimate does not change significantly after including borrower characteristics as covariates. This suggests that my main finding is not subsumed by variation in borrower composition. To get a more direct sense of the response from the demand side of mortgage credits, Table 5 uses the same empirical framework given by Regression (1) to quantify the differential change in borrower characteristics for high- versus low-SFHA tracts following BW12. In Column 1, the dependent variable is the number of applications in census tract i in year t, normalized per 1,000 population (the average number of application is 174 and the average population is 4,303). The coefficient estimate of $Treated_i \times Post_t$ is economically minuscule and statistically insignificant, suggesting that the number of applications is unaffected by the law change. The same interpretation applies to the results presented in Columns 2 (share of white applicants), 4 (applicant income), and 5 (loan-to-income ratio). The only exception is the share of male applicants (Column 3). The statistically significant estimate there suggests that the fraction of male applicants decreases by 0.14 percentage points following BW12. Economically, however, the magnitude is small—a change of only 0.26 percent relative to the male share of 54.0 percent in 2012.

5.5 Supporting Heterogeneous Results

In this section, I provide further evidence that the increase in mortgage denial rate following BW12 is driven by lenders internalizing the ex-post cost of monitoring flood insurance compliance into their ex-ante decisions on mortgage origination. To do so, I examine heterogeneity in the magnitude of the denial rate increase along several dimensions in which lenders are expected to respond differently to the law change.

5.5.1 Heterogeneity by Climate Change Belief

My proposed mechanism predicts that as lenders worry about the deterioration in flood insurance compliance, which puts their collateral values at risk, the adjustment to credit supply should be more aggressive in areas where the compliance issue is expected to become more severe following BW12 (that is, where the demand for flood insurance is more elastic).

To get a proxy for flood insurance elasticity, I exploit the heterogeneity in households' belief in climate change. The idea is that the impact of BW12 on borrowers' disincentive to comply with mandatory flood insurance purchases should be weaker for those who do worry about flood risk and, more generally, about global warming. To measure households' climate belief, I leverage data from the Yale Climate Opinion survey (Howe et al., 2015), which is also used in several recent studies of climate risk (Barrage and Furst, 2019; Bernstein et al., 2019; Baldauf et al., 2020; Bakkensen and Barrage, 2021; Goldsmith-Pinkham et al., 2021). The survey contains several questions and provides aggregate statistics at the county level. As the data is not granular to the census-tract level, I need to assume that all census tracts within a county have similar climate beliefs. Following the literature, I use the percentage of people who answered "Yes" to the question "Do you believe global warming is happening?" as my main measure; my results are robust to the other commonly used question, "Are you worried about climate change?" I construct a dummy variable $SCCB_i$ (Strong Climate Change Belief) which indicates if census tract i has an above-median percentage of climate change believers. According to this measure, I divide my sample into subsamples with strong and weak climate change opinions. Thus, the $SCCB = 1$ subsample captures a set of census tracts that seem likely to have low demand elasticity for flood insurance.

I run my baseline difference-in-differences regression in each subsample and present the coefficient estimates in Column 1 of Table 6 for the subsample of weak climate beliefs and Column 2 for the subsample of strong climate beliefs. Comparing the two estimates (1.03 versus 0.58), the negative effect of BW12 on credit supply is apparently much stronger in areas where households are not as worried about climate risk. To test if this difference is statistically significant, I introduce the dummy variable $SCCB_i$ into Regression (1) and interact it with each existing term:

$$\begin{aligned} DenialRate_{it} = \beta_0 + \beta_1 * Treated_i \times Post_t + \beta_2 * Treated_i + \beta_3 * Post_t + \\ \gamma_0 * SCCB_i + \gamma_1 * Treated_i \times Post_t \times SCCB_i + \\ \gamma_2 * Treated_i \times SCCB_i + \gamma_3 * Post_t \times SCCB_i + \epsilon_{it}. \end{aligned} \tag{3}$$

The key coefficients of interest from Regression (3) are β_1 and γ_1, presented in Column 3 of Panel A of Table 6. By construction, the value of β_1 is the same as the one in column 1. The coefficient γ_1 captures the differential response of credit supply to the policy shock of BW12 in tracts with strong climate beliefs relative to tracts with weak climate beliefs; the estimate suggests that the difference is highly significant. In Columns 4 through 6, I repeat the whole exercise using the alternative survey question, "Are you worried about climate change?" The results are robust.

5.5.2 Heterogeneity by Lender Scale

My hypothesis also predicts that local banks will respond to BW12 more aggressively than national banks, since local banks hold less-diversified mortgage portfolios. Lack of diversification is particularly important in the context of flood risk because of this risk's correlated and catastrophic nature; when a flood occurs, it tends to severely affect many properties in a small region, which can be devastating for small local banks. Hence, local mortgage lenders are expected to be more cautious about their collaterals' exposure to flood risk.

I consider a variety of measures to differentiate local from national banks. I first rely on information from my main HMDA data. For each mortgage application, I can observe the lender identity and the property location, which allows me to calculate how many states a given bank covers. Across all banks, the median coverage is two states, but at the loan level, banks with mortgage lending in no more than two states account for less than 7 percent of the total applications. Therefore, in order to have a more balanced sample coverage for local banks, I define them to be those having mortgage business in no more than 10 states (measure 1), which accounts for roughly 25 percent of the HMDA data. Results are robust to various cutoff points, including the median value of two states.

Next, I leverage information about bank branches obtained from the FDIC Summary of Deposits, which can be mapped to the HMDA data using a unique lender identifier called RSSD ID. I consider three additional measures based on the number of branches (measure 2), the fraction of branches that are outside the bank's headquarters state (measure 3), and the number of states in which the bank has a branch (measure 4). The thresholds, which define local banks, are set as 50 branches, 50 percent outside the headquarters state, and branch presence in five states. Based on these thresholds, the sample coverage of local banks in the HMDA data ranges from 25 to 35 percent. Results presented below are robust to other cutoff choices.

To examine whether local and national banks react differently to BW12, I divide my full HMDA sample into (a) applications submitted to local banks and (b) applications submitted to national banks, according to one of the measures described above. Then, for each census tract, I calculate the mortgage denial rate separately for the two lender types. Finally, I run Regression (1) separately for the two outcome variables. Noticing that local banks tend to be less likely to reject mortgage applications, I run the regressions with the denial rates in logarithm, so that the difference-in-differences estimates (representing percentage changes in this specification) are directly comparable.

Table 7 presents the results, which suggest that, following BW12, local banks increase their mortgage denial rates by 8.8–9.9 percent, while national banks respond by only 3.2–5.6 percent. In unreported results (for brevity), I use a regression similar to Regression (3) to show that the difference between local and national banks is statistically highly significant in all specifications. My result is consistent with the finding of Nguyen et al. (2022) that local banks pay closer attention to climate change risk and charge a higher sea-level-rise premium.[18]

5.5.3 Heterogeneity by Loan Purpose

As my hypothesis builds on the idea that lapsed flood insurance leaves collateral value vulnerable to flood risk, it predicts that the effect of BW12 will pertain only to home purchase and refinance loans, for which the property typically serves as collateral. In contrast, home improvement loans, normally unsecured, should not be significantly affected in the context of my natural experiment. As discussed in Section 3.1, relying on the fact

[18] Nguyen et al. (2022) suggest that local banks have superior knowledge about a property's exposure to seal-level-rise risk. In my setting, as the underlying flood risk does not change following the policy shock, I propose a complementary explanation highlighting the diversification consideration of lenders.

that the HMDA data specifies the loan purpose, I restricted my main analysis sample to home purchase and refinance loans (for all the tables seen previously). Here, to test the prediction of my hypothesis, I use home improvement loans as a valuable verification check. Specifically, I calculate the application denial rates for home improvement loans at the census-tract-year level and use the same empirical framework given by Regression (1) to examine how they respond to the policy shock of BW12. I also untie home purchase and refinance loans in this section.

Table 8 presents the results. The outcome variable—loan denial rate—is in percent in Columns 1 through 3 and in logarithm in Columns 4 through 6. Most importantly, Column 1 examines the effect of BW12 on the application outcome for home improvement loans. The difference-in-differences estimate suggests that the denial rate for home improvement loans goes up by only 0.12 percentage points following BW12, which is insignificant both statistically and economically. The effect is about 10 times smaller than that for home purchase mortgages (Column 2) or for refinance mortgages (Column 3).

For the latter two loan purposes, the effects are almost identical in terms of the magnitude of the absolute increase (1.1 percentage points). Note, however, that home purchase mortgages have a much lower baseline rejection rate than refinance mortgages. Therefore, the relative effect of BW12 seems strongest for home purchase loans. This conjecture is confirmed by the results in the last three columns of Table 8 using logarithm.

5.5.4 Heterogeneity by Securitization Prevalence

As Ouazad and Kahn (2021) show that lenders transfer flood risk through securitization to GSEs (such as Fannie Mae and Freddie Mac), we expect the effect of BW12 to be weaker for GSE loans, because lenders will not worry as much about underinsurance and ex-post monitoring since the risk will be borne by the GSEs.

The HMDA data includes a variable called purchaser type, indicating whether an originated loan was then sold to a secondary-market entity. For an average year between 2007 and 2016, there were 8.04 million mortgages originated, 0.76 million applications approved but not accepted, and 2.91 million applications denied. Among the originated loans, 1.62 million and 1.02 million were sold to Fannie Mae and Freddie Mac, respectively.[19]

[19] If a loan is sold to more than one entity, the lending institution reports the entity purchasing the greatest interest, if any. If the institution retains a majority interest, it does not report the sale.

Based on pre-BW12 data, I calculate the ratio between the number of GSE loans and the number of total origination at the census-tract level. The 25th, 50th, and 75th percentiles are 0.23, 0.31, and 0.39, respectively. I then construct subsamples of census tracts with an above- or below-median GSE ratio. The above-median subsample represents areas where securitization is relatively more prevalent (for example, because these areas have more loans satisfying the conforming criteria set by the regulators, Fannie Mae, and Freddie Mac).

I run the difference-in-differences regressions in each subsample and present the coefficient estimates in Table 9. In the baseline specification (Columns 1 and 2), I find that the mortgage denial rate increases by 1.07 (0.45) percentage points following BW12 in census tracts with low (high) prevalence of GSE loans. The results are robust to specifications with covariates and state fixed-effects (Columns 3 to 6).

5.6 The Effect on LTVs and Interest Rates

To complement my main findings on lenders' extensive-margin decision (i.e., accept or reject a loan application), this section examines two intensive-margin decisions—LTV ratios and interest rates, both conditional on acceptance.

As previously discussed in Section 3.3, I have to use alternative data obtained from Fannie Mae and Freddie Mac because HMDA does not collect information on LTV ratios and interest rates. However, the disadvantage is that property addresses are disclosed only at the three-digit-ZIP level (the last two digits of a five-digit ZIP code are truncated). Therefore, in this section, the intensive-margin analysis has to be conducted at the three-digit-ZIP level. The outcome variable of interest is the average LTV ratio or the average interest rate, calculated at the three-digit-ZIP-month level using either Fannie Mae or Freddie Mac data.

To reconstruct the treatment and control groups, I transform the tract-level SHFA data to the three-digit-ZIP level, using the crosswalk provided by the US Department of Housing and Urban Development (see footnote 8). The treated are high-SFHA ZIP codes, which are expected to be more affected by BW12.

Using the same difference-in-differences framework given by Regression (1), I show that high-SFHA ZIP codes experienced a 0.55-percentage-point increase in mortgage denial rate (t-statistic=2.53) following BW12. This result suggests that although the three-digit ZIP codes provide a less-identified setting, it is still a meaningful level of aggregation for my application. Appendix Figure A.4 presents the dynamic effects to support the assumption of parallel pre-trend. There are two points worth noting. First, the magnitude of the difference-in-differences estimate (0.55) is smaller than that in Table 2 (0.80). Second, the standard error (0.22) is greater than that in Table 2 (0.064). These two differences can be explained by the fact that three-digit ZIP codes (893 in total) are much less granular than census tracts (84,414 in total) and thus cannot identify as precisely the most affected population by BW12.

Figure 5 shows that LTV ratios and interest rates did not experience any discernible changes after BW12. This result is consistent with the finding of Ouazad and Kahn (2021) that lenders transfer flood risk through securitization to Fannie Mae and Freddie Mac. However, note that due to data limitation, it is unclear whether the intensive-margin result can be generalized to other unsecuritized loans.

6 Conclusion

This paper examines how lenders incorporate the ex-post costs of monitoring borrowers' compliance of mandatory flood insurance into their ex-ante loan origination decisions. At the intersection of US flood insurance and mortgage markets, lenders' monitoring is important because flood insurance is ultimately purchased by homeowners, who are known to be inattentive to flood insurance. Lapsed policies will expose housing collaterals to flood risk, which is the most costly natural peril in the US.

To isolate the effect of borrowers' insurance incentives on lenders' loan origination decisions, I exploit the passage of the Biggert-Waters Flood Insurance Reform Act of 2012 (BW12), which significantly increases the premiums (by $266 per year) for many properties in Special Flood Hazard Areas (SFHAs), while leaving other properties unaffected. I show that, following BW12, the denial rate of mortgage applications in high-SFHA tracts increases by 0.80 percentage points relative to low-SFHA tracts. This is a 3.54-percent increase relative to the mean denial rate of 22.6 percent in 2012. Abundant evidence suggests that this result cannot be explained by alternative explanations, such as decreasing disposable income or changes in borrower composition. I show that the effect is stronger for locations with weak climate risk perception, is stronger for local banks, does not exist for unsecured home improvement loans, and is stronger where securitization is less prevalent. All this evidence of heterogeneity supports my hypothesized mechanism that lenders are internalizing the ex-post monitoring cost, resulting in more stringent accept-or-reject decisions.

These results constitute the first direct evidence that borrowers' insurance incentives play an important role in affecting lenders' decision-making. The unintended effect of BW12—namely, that a large number of marginal borrowers lost credit accessibility due to a fairly small rise in insurance premium—highlights important policy implications for regulators.

Figures and Tables

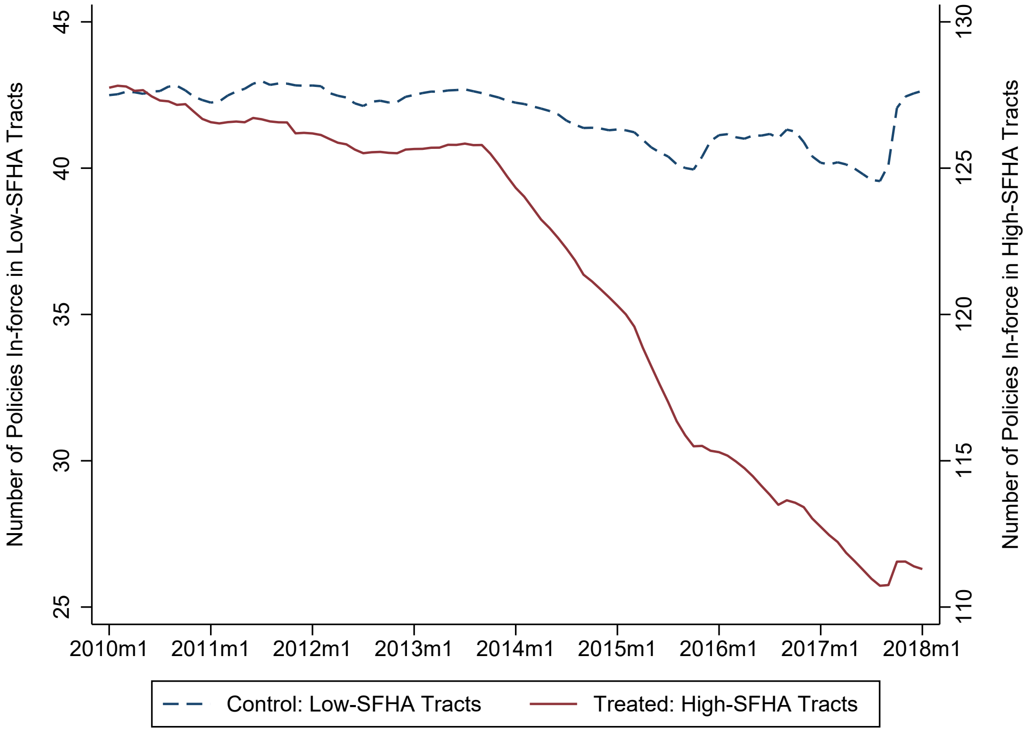

Figure 1. Raw Trends of the Numbers of Flood Insurance Policies In-force

This figure shows the average numbers of flood insurance policies in-force in the low-SFHA (dashed line, left y-axis) and high-SFHA (solid line, right y-axis) census tracts from January 2010 to January 2018.

Figure 2. Raw Trends of Mortgage Denial Rates

This figure shows the average denial rates (in percent) of mortgage applications in the low-SFHA (dashed line, left y-axis) and high-SFHA (solid line, right y-axis) census tracts from 2007 to 2016.

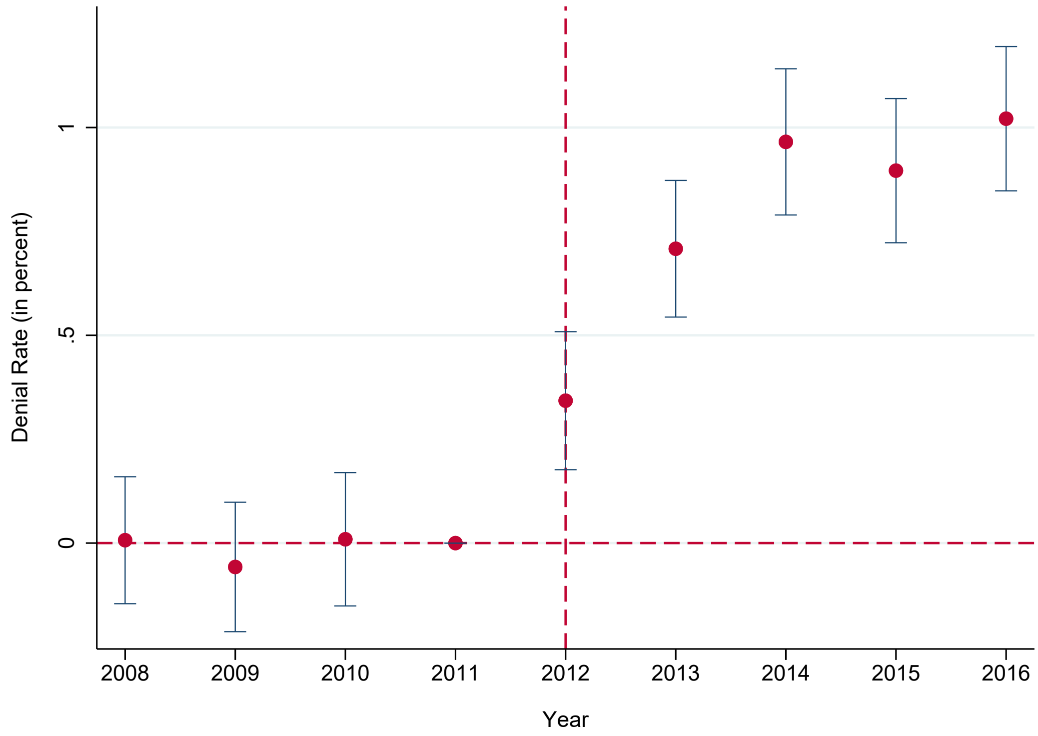

Figure 3. The Effects of Flood Insurance Disincentive on Mortgage Denial Rates

This figure shows the dynamic effects of flood insurance disincentive (caused by BW12) on mortgage denial rates. It plots the coefficient estimates of $\{\beta_{\tilde{t}}\}$ from the regression: $DenialRate_{it} = \beta_0 + \sum_{\tilde{t}=2007}^{2016} \beta_{\tilde{t}} * Treated_i \times \mathbb{1}_{t=\tilde{t}} + \beta_2 * Treated_i + \alpha_t + \epsilon_{it}$. $\{\beta_{\tilde{t}}\}$ are measured relative to β_{2011}, which is omitted. $\{\beta_{\tilde{t}}\}_{\tilde{t}<2012}$ correspond to the pre-trends and $\{\beta_{\tilde{t}}\}_{\tilde{t}\geq 2012}$ capture the dynamic effects of BW12. Standard errors are clustered at the census-tract level. The bands around the coefficient estimates show the 95% confidence intervals.

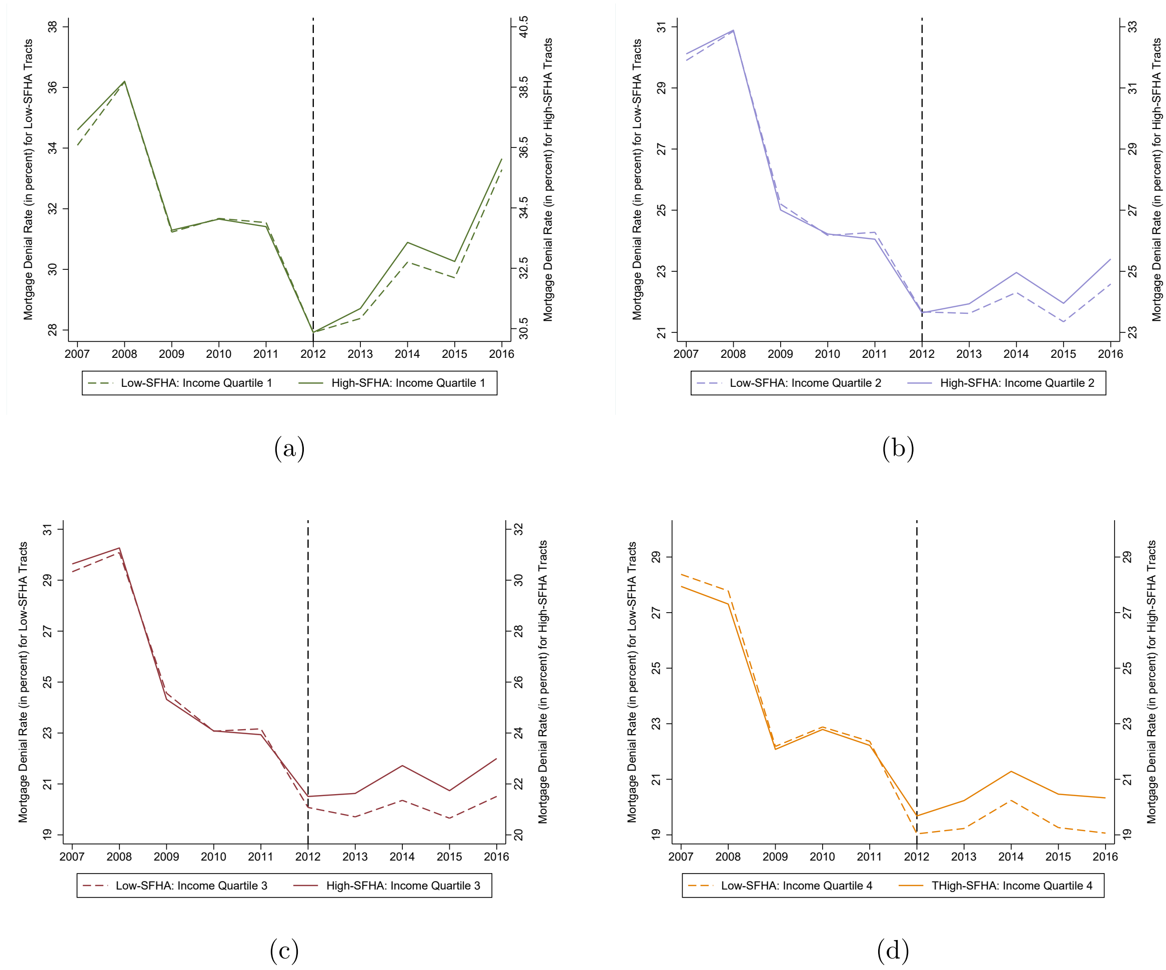

(a) (b) (c) (d)

Figure 4. Raw Trends of Denial Rates by Income Quartiles

This figure shows the average denial rates (in percent) of mortgage applications in four different subsamples constructed according the average applicant's income. Panel (a) shows the lowest income-quartile and Panel (d) shows the highest income-quartile. In each panel, the dashed line (left y-axis) refers to the low-SFHA (i.e., control) census tracts, while the solid line (right y-axis) refers to the high-SFHA (i.e., treated) census tracts.

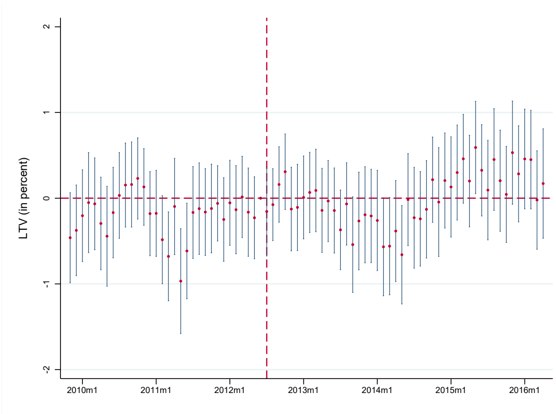

(a) LTV - Fannie Mae Data

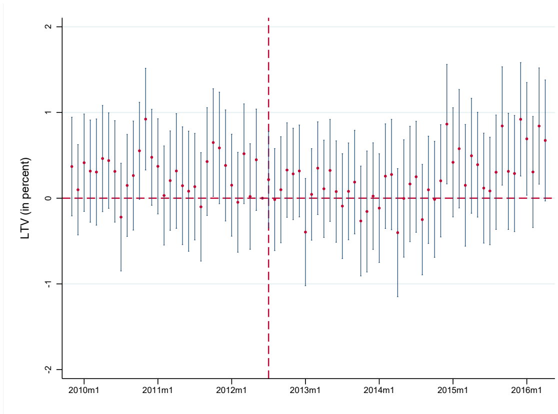

(b) LTV - Freddie Mac Data

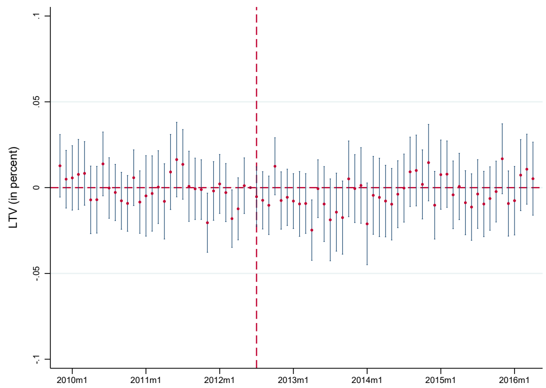

(c) Interest Rate - Fannie Mae Data

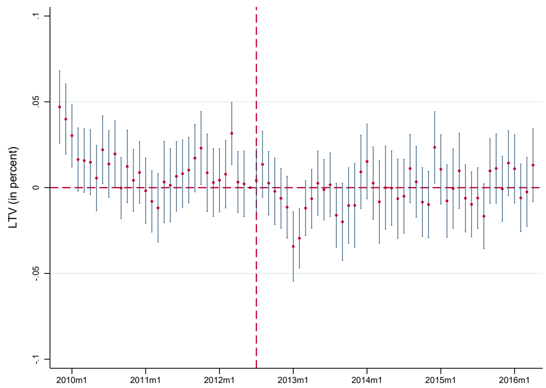

(d) Interest Rate - Freddie Mac Data

Figure 5. LTV and Interest Rate

This figure shows the dynamic effects of BW12 on LTVs and interest rates, both conditional on loan origination. The regression is the same as that described in Figure 3 except that the data here (obtained from Fannie Mae and Freddie Mac) allows for a more granular analysis at the monthly level. Panels (a) and (b) present the results for LTVs, while Panel (c) and (d) present the results for interest rates. Standard errors are clustered at the census-tract level. The bands around the coefficient estimates show the 95% confidence intervals.

Table 1: Descriptive Statistics

This table presents descriptive statistics at the census-tract-year level from 2007 to 2016. The high-SFHA subsample (i.e., the treated group) includes census tracts having an above-median fraction of SFHA policies (using pre-BW12 data). The low-SFHA subsample (i.e., the control group) includes census tracts having a below-median fraction of SFHA policies. *%SFHA* is the fraction of SFHA policies. *Denial rate* is calculated as the number of rejected applications in a given year divided by the total number of applications that receive lender decisions. *%White* is the proportion of white applicants. *%Male* is the proportion of male applicants. *Income ($K)* is the average annual income (in thousand) of mortgage applicants. *Loan ($K)* is the average loan amount (in thousand).

	Full Sample			High-SFHA			Low-SFHA		
	mean	median	s.d.	mean	median	s.d.	mean	median	s.d.
%SFHA	0.40	0.38	0.32	0.68	0.68	0.18	0.11	0.05	0.13
Denial rate	0.25	0.23	0.12	0.26	0.24	0.11	0.24	0.22	0.12
%White	0.83	0.91	0.21	0.86	0.94	0.19	0.80	0.89	0.23
%Male	0.55	0.55	0.07	0.55	0.55	0.07	0.54	0.54	0.07
Income ($K)	100.9	84.2	66.7	95.1	79.9	61.7	106.8	89.4	70.9
Loan ($K)	196.6	159.5	141.2	174.9	144.9	122.8	218.1	179.3	154.4

Table 2: The Effect of Flood Insurance Disincentive on Mortgage Denial Rates

This table shows results from the difference-in-differences regression: $DenialRate_{it} = \beta_0 + \beta_1 * Treated_i \times Post_t + \beta_2 * Treated_i + \beta_3 * Post_t + \gamma * X_{it} + \epsilon_{it}$. $DenialRate_{it}$ is the denial rate of mortgage applications in census tract i in year t. $Treated_i$ is the treatment dummy, indicating if census tract i is a high-SFHA area defined as having an above-median fraction of SFHA properties. $Post_t$ is the post-event dummy that turns on if year t is after the passage of BW12. $White_{it}$ is the proportion of white applicants. $Male_{it}$ is the proportion of male applicants. $Income_{it}$ is the average annual income (in thousand) of applicants. *Loan-to-income*$_{it}$ is the average loan-to-income ratio. Standard errors are clustered at the census-tract level and presented in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	(1)	(2)	(3)	(4)	(5)
$Treated_i \times Post_t$	0.795***	0.778***	0.805***	0.744***	0.471***
	(0.064)	(0.061)	(0.060)	(0.058)	(0.059)
$Post_t$	-4.887***	-4.970***	-4.916***	-4.912***	-2.864***
	(0.046)	(0.044)	(0.043)	(0.042)	(0.871)
$Treated_i$	1.197***	2.445***	1.942***	1.621***	1.774***
	(0.093)	(0.080)	(0.077)	(0.073)	(0.073)
$White_{it}$		-0.225***	-0.217***	-0.226***	-0.222***
		(0.002)	(0.002)	(0.002)	(0.002)
$Male_{it}$		0.121***	0.147***	0.138***	0.127***
		(0.005)	(0.005)	(0.005)	(0.005)
$Income_{it}$			-0.036***	-0.034***	-0.035***
			(0.001)	(0.001)	(0.001)
Loan-to-income$_{it}$			0.072*	0.088*	0.070
			(0.041)	(0.051)	(0.043)
$Constant$	27.330***	38.793***	40.523***	43.245***	41.347***
	(0.069)	(0.293)	(0.308)	(0.382)	(0.419)
State FE	N	N	N	Y	N
State-year FE	N	N	N	N	Y
Observations	544,325	544,055	543,986	543,986	543,986
R-squared	0.041	0.200	0.241	0.302	0.339

Table 3: Heterogeneous Effects in Subsamples based on Income

This table presents the baseline difference-in-differences estimates as in Table 2 in subsamples based on the average mortgage applicant's income. Quartile 1 (4) represents census tracts in the bottom (top) quartile of the income distribution. Standard errors are clustered at the census-tract level and presented in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	Subsample of Census Tracts by Income							
	Quartile 1	Quartile 2	Quartile 3	Quartile 4	Quartile 1	Quartile 2	Quartile 3	Quartile 4
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$Treated_i \times Post_t$	0.775***	1.004***	1.527***	0.682***	0.774***	0.920***	1.337***	0.791***
	(0.140)	(0.115)	(0.116)	(0.112)	(0.131)	(0.106)	(0.109)	(0.109)
$Post_t$	-5.680***	-4.942***	-5.308***	-3.432***	-5.320***	-5.028***	-5.382***	-3.799***
	(0.108)	(0.088)	(0.083)	(0.067)	(0.134)	(0.080)	(0.078)	(0.071)
$Treated_i$	-0.543***	0.009	-0.651***	1.179***	0.663***	1.299***	1.027***	1.420***
	(0.194)	(0.154)	(0.161)	(0.144)	(0.151)	(0.124)	(0.134)	(0.139)
$Constant$	35.426***	28.196***	26.113***	21.440***	52.236***	39.404***	32.576***	16.753***
	(0.156)	(0.122)	(0.122)	(0.092)	(1.518)	(0.523)	(0.656)	(0.811)
$White_{it}$					-0.236***	-0.227***	-0.186***	-0.061***
					(0.003)	(0.003)	(0.004)	(0.005)
$Male_{it}$					0.117***	0.197***	0.185***	0.173***
					(0.009)	(0.008)	(0.009)	(0.011)
$Income_{it}$					-0.093***	-0.043***	-0.039***	-0.003***
					(0.031)	(0.004)	(0.003)	(0.001)
$Loan\text{-}to\text{-}income_{it}$					0.028	0.004	0.768***	0.330***
					(0.043)	(0.034)	(0.042)	(0.097)
Observations	139,903	136,229	134,111	133,965	139,780	136,186	134,061	133,861
R-squared	0.037	0.047	0.054	0.039	0.261	0.231	0.227	0.076

Table 4: Heterogeneous Effects across the Income Distribution within a Census Tract

This table shows results from the difference-in-differences regression: $DenialRate^q_{it} = \beta_0 + \beta_1 * Treated_i \times Post_t + \beta_2 * Treated_i + \beta_3 * Post_t + \epsilon_{it}$, using the full sample. The dependent variable $DenialRate^q_{it}$ is the denial rate of mortgage applications for income quartile q in census tract i in year t. $Treated_i$ and $Post_t$ are defined as in Table 2. Standard errors are clustered at the census-tract level and presented in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	DiD with Y_{it} being $DenialRate^q_{it}$			
	q=1	q=2	q=3	q=4
$Treated_i \times Post_t$	0.286***	0.504***	1.061***	1.283***
	(0.094)	(0.087)	(0.087)	(0.078)
$Post_t$	-3.020***	-4.921***	-5.926***	-5.285***
	(0.067)	(0.061)	(0.062)	(0.056)
$Treated_i$	2.553***	1.968***	0.996***	-0.250***
	(0.114)	(0.106)	(0.107)	(0.095)
$Constant$	32.930***	26.829***	25.988***	24.652***
	(0.082)	(0.078)	(0.079)	(0.071)
Observations	540,067	542,500	541,312	544,325
R-squared	0.013	0.026	0.032	0.029

Table 5: Changes in Borrowers Composition

This table shows results from the difference-in-differences regression: $Y_{it} = \beta_0 + \beta_1 * Treated_i \times Post_t + \beta_2 * Treated_i + \beta_3 * Post_t + \epsilon_{it}$. The dependent variable Y_{it} is the number of mortgage applications in census tract i in year t (Column 1), the proportion of white applicants (Column 2), the proportion of male applicants (Column 3), the average annual income (Column 4), and the average loan-to-income ratio (Column 5). Standard errors are clustered at the census-tract level and presented in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	(1)	(2)	(3)	(4)	(5)
Depedent Variable	Number of Applications	Share of White	Share of Male	Applicant Income	Loan-to-income Ratio
$Treated_i \times Post_t$	0.145	-0.118	-0.144***	0.536	-0.006
	(1.211)	(0.101)	(0.040)	(0.348)	(0.013)
$Post_t$	-8.780***	0.422***	1.540***	2.781***	0.180***
	(1.141)	(0.077)	(0.029)	(0.238)	(0.006)
$Treated_i$	-6.290***	6.040***	0.947***	-11.913***	-0.235***
	(1.209)	(0.197)	(0.045)	(0.579)	(0.008)
$Constant$	50.279***	79.823***	53.441***	105.133***	2.614***
	(1.149)	(0.151)	(0.034)	(0.447)	(0.007)
Observations	544,506	544,066	544,157	544,177	544,177
R-squared	0.003	0.020	0.013	0.008	0.003

Table 6: Heterogeneity by Climate Change Belief

This table presents results of the baseline difference-in-differences estimates as in Table 2 in subsamples with a weak or strong climate change belief measured by the Yale Climate Opinion survey (Howe et al., 2015). Columns 1 to 3 use the question "Do you believe global warming is happening?" Columns 4 to 6 use the question "Are you worried about climate change?" $SCCB_i = 1$ (Strong Climate Change Belief) is a binary variable equal to 1 if census tract i has an above-median percentage of survey respondents answering "Yes" to the question. Columns 1 and 2 run the baseline regression as in Table 2 in respective subsamples constructed according to the "happening" question. Column 3 considers a pooled sample by interacting $SCCB_i$ with the baseline regression: $DenialRate_{it} = \beta_0 + \beta_1 * Treated_i \times Post_t + \beta_2 * Treated_i + \beta_3 * Post_t + \gamma_0 * SCCB_i + \gamma_1 * Treated_i \times Post_t \times SCCB_i + \gamma_2 * Treated_i \times SCCB_i + \gamma_3 * Post_t \times SCCB_i + \epsilon_{it}$. For brevity, only the coefficient estimates of $Treated_i \times Post_t$ and $Treated_i \times Post_t \times SCCB_i$ are presented in this table. Columns 4 to 6 are similar to 1 to 3 but use the "worried" question. Standard errors are clustered at the census-tract level and presented in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	(1)	(2)	(3)	(4)	(5)	(6)
$Treated_i \times Post_t$	1.028***	0.575***	1.028***	0.993***	0.572***	0.993***
	(0.092)	(0.090)	(0.092)	(0.092)	(0.090)	(0.092)
$Treated_i \times Post_t \times SCCB_i$			-0.453***			-0.421***
			(0.129)			(0.128)
Observations	265,082	279,235	544,317	264,349	279,968	544,317
R-squared	0.046	0.038	0.042	0.047	0.036	0.043
Question	Happening	Happening	Happening	Worried	Worried	Worried
Sample	Weak	Strong	Pooled	Weak	Strong	Pooled

Table 7: National Banks Versus Local Banks

This table tests whether local banks respond to BW12 more aggressively than national banks. National and local banks are differentiated using four different measures. Measure 1 relies on the HMDA data: a bank is classified as a local lender if it receives mortgage applications from no more than 10 states. Measures 2 to 4 leverage information about bank branches obtained from the FDIC Summary of Deposits. Measure 2 considers the number of branches. Measure 3 considers the fraction of branches that are outside the bank's headquarters state. Measure 4 considers the number of states in which the bank has a branch. The thresholds, which defined local banks, are set as 50 branches, 50 percent outside the headquarters state, and branch presence in five states. To construct the dependent variable, the full HMDA data is divided into (a) applications submitted to local banks and (b) applications submitted to national banks, according to one of the four measures. The mortgage denial rate is calculated separately for the two lender types. This table presents the baseline difference-in-differences estimate with $log(DenialRate_{it})$ being the dependent variable. Standard errors are clustered at the census-tract level and presented in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	Measure 1		Measure 2	
	National	Local	National	Local
$Treated_i \times Post_t$	0.032***	0.099***	0.054***	0.092***
	(0.002)	(0.005)	(0.003)	(0.004)
Observations	541,130	463,779	527,184	476,120
R-squared	0.055	0.005	0.007	0.028
	Measure 3		**Measure 4**	
	National	Local	National	Local
$Treated_i \times Post_t$	0.055***	0.088***	0.056***	0.097***
	(0.003)	(0.004)	(0.003)	(0.004)
Observations	520,099	502,439	519,393	503,371
R-squared	0.009	0.026	0.007	0.033

Table 8: Heterogeneity by Loan Purpose

This table presents the heterogeneous effect of BW12 on mortgage denial rates across different types of loans: home improvement, home purchase, and refinance. The dependent variable in Columns 1 to 3 is $DenialRate_{it}$, while it is $log(DenialRate_{it})$ in Columns 4 to 6. Standard errors are clustered at the census-tract level and presented in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	(1)	(2)	(3)	(4)	(5)	(6)
$Treated_i \times Post_t$	0.115	1.137***	1.114***	0.006*	0.075***	0.038***
	(0.131)	(0.071)	(0.078)	(0.003)	(0.004)	(0.002)
$Post_t$	-0.423***	-3.673***	-4.119***	-0.022***	-0.230***	-0.128***
	(0.094)	(0.050)	(0.057)	(0.002)	(0.003)	(0.002)
$Treated_i$	-0.043	0.755***	1.059***	-0.013***	0.068***	0.047***
	(0.147)	(0.085)	(0.107)	(0.004)	(0.005)	(0.003)
$Constant$	38.734***	17.683***	32.919***	3.599***	2.698***	3.396***
	(0.109)	(0.063)	(0.080)	(0.003)	(0.003)	(0.002)
Observations	534,230	541,713	543,211	491,430	519,232	539,873
R-squared	0.000	0.020	0.019	0.000	0.029	0.021
Loan purpose	Improvement	Purchase	Refinance	Improvement	Purchase	Refinance
Logarithm	N	N	N	Y	Y	Y

Table 9: GSE Loans and Securitization Prevalence

This table presents the baseline difference-in-differences estimates as in Table 2 in subsamples with a high or low prevalence of GSE loans. The subsamples are constructed based on a variable called purchaser type in the HMDA data, which indicates whether an originated loan was then sold to a secondary-market entity (in particular, the GSEs, such as Fannie Mae and Freddie Mac). For each census tract, the ratio between the number of GSE loans and the number of total origination is calculated using pre-BW12 data. The low (high) GSE-prevalence subsample consists of census tracts with a below- (above-) median ratio. Standard errors are clustered at the census-tract level and presented in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	(1)	(2)	(3)	(4)	(5)	(6)
$Post \times Treated$	1.065***	0.450***	1.041***	0.460***	1.027***	0.236***
	(0.091)	(0.087)	(0.086)	(0.082)	(0.084)	(0.077)
GSE Prevalence	Low	High	Low	High	Low	High
Covariates	N	N	Y	Y	Y	Y
State FE	N	N	N	N	Y	Y
Observations	262,229	281,949	262,028	281,833	262,028	281,833
R-squared	0.034	0.045	0.272	0.141	0.328	0.231

References

E. B. Abbott. Floods, flood insurance, litigation, politics-and catastrophe: The national flood insurance program. *Sea Grant L. & Pol'y J.*, 1:129, 2008.

J. M. Addoum, P. Eichholtz, E. Steiner, and E. Yönder. Climate change and commercial real estate: Evidence from hurricane sandy. *Available at SSRN 3206257*, 2021.

C. Avenancio-Leon and T. Howard. The assessment gap: Racial inequalities in property taxation. *Available at SSRN 3465010*, 2019.

L. Bakkensen, T. Phan, and T.-N. Wong. Leveraging the disagreement on climate change: Theory and evidence. 2022.

L. A. Bakkensen and L. Barrage. Going underwater? flood risk belief heterogeneity and coastal home price dynamics. *The Review of Financial Studies*, 2021.

M. Baldauf, L. Garlappi, and C. Yannelis. Does climate change affect real estate prices? only if you believe in it. *The Review of Financial Studies*, 33(3):1256–1295, 2020.

L. Barrage and J. Furst. Housing investment, sea level rise, and climate change beliefs. *Economics letters*, 177:105–108, 2019.

R. Bartlett, A. Morse, R. Stanton, and N. Wallace. Consumer-lending discrimination in the fintech era. *Journal of Financial Economics*, 143(1):30–56, 2022.

E. Benmelech. Asset salability and debt maturity: Evidence from nineteenth-century american railroads. *The Review of Financial Studies*, 22(4):1545–1584, 2009.

E. Benmelech and N. K. Bergman. Collateral pricing. *Journal of financial Economics*, 91(3): 339–360, 2009.

E. Benmelech and N. K. Bergman. Bankruptcy and the collateral channel. *The Journal of Finance*, 66(2):337–378, 2011.

E. Benmelech, M. J. Garmaise, and T. J. Moskowitz. Do liquidation values affect financial contracts? evidence from commercial loan contracts and zoning regulation. *The Quarterly Journal of Economics*, 120(3):1121–1154, 2005.

A. Bernstein, M. T. Gustafson, and R. Lewis. Disaster on the horizon: The price effect of sea level rise. *Journal of Financial Economics*, 134(2):253–272, 2019.

K. S. Blickle, J. A. Santos, et al. Unintended consequences of" mandatory" flood insurance. Technical report, 2022.

G. Cerqueiro, S. Ongena, and K. Roszbach. Collateralization, bank loan rates, and monitoring. *The Journal of Finance*, 71(3):1295–1322, 2016.

J. Cespedes, Z. Liu, and C. Parra. Almost famous: The short-and long-term effects of housing wealth shocks on career decisions. *Available at SSRN 3491426*, 2019.

J. C. Cespedes, C. R. Parra, and C. Sialm. The effect of principal reduction on household distress: Evidence from mortgage cramdown. Technical report, National Bureau of Economic Research, 2021.

J. Chivers and N. E. Flores. Market failure in information: the national flood insurance program. *Land Economics*, 78(4):515–521, 2002.

S. J. Colby and K. Y. Zipp. Excess vulnerability from subsidized flood insurance: housing market adaptation when premiums equal expected flood damage. *Climate Change Economics*, 12(01): 2050012, 2021.

J. Currie and E. Tekin. Is there a link between foreclosure and health? *American Economic Journal: Economic Policy*, 7(1):63–94, 2015.

A. Din and R. Wilson. Crosswalking zip codes to census geographies. *Cityscape*, 22(1):293–314, 2020.

L. Dixon, N. Clancy, S. A. Seabury, and A. Overton. The national flood insurance program's market penetration rate: Estimates and policy implications. *Santa Monica, CA: RAND Corporation*, 2006.

FEMA. An affordability framework for the national flood insurance program. *The Federal Emergency Management Agency (FEMA)*, 2018.

GAO. Flood insurance: Extent of noncompliance with purchase requirements is unknown. *US Government Accountability Office (GAO)*, 2002.

N. Garbarino and B. Guin. High water, no marks? biased lending after extreme weather. *Journal of Financial Stability*, 54:100874, 2021.

A. Gavazza. Asset liquidity and financial contracts: Evidence from aircraft leases. *Journal of financial Economics*, 95(1):62–84, 2010.

S. Giglio, B. Kelly, and J. Stroebel. Climate finance. *Annual Review of Financial Economics*, 13:15–36, 2021.

P. S. Goldsmith-Pinkham, M. Gustafson, R. Lewis, and M. Schwert. Sea level rise exposure and municipal bond yields. *Jacobs Levy Equity Management Center for Quantitative Financial Research Paper*, 2021.

T. L. Hayes and D. A. Neal. Actuarial rate review: In support of the recommended october 1, 2011, rate and rule changes. *Washington, DC: Federal Emergency Management Agency*, 2011.

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

H. Hong, G. A. Karolyi, and J. A. Scheinkman. Climate finance. *The Review of Financial Studies*, 33(3):1011–1023, 2020.

D. P. Horn and J. T. Brown. *Introduction to the national flood insurance program (nfip).* Congressional Research Service Washington DC, USA, 2018.

D. P. Horn and B. Webel. Private flood insurance and the national flood insurance program. *Congressional Research Service*, 2019.

R. Horzempa. Drowning in debt: the economic failures of the national flood insurance program and why it will never rise above water. *Rev. Banking & Fin. L.*, 38:391, 2018.

P. D. Howe, M. Mildenberger, J. R. Marlon, and A. Leiserowitz. Geographic variation in opinions on climate change at state and local scales in the usa. *Nature Climate Change*, 5 (6):596–603, 2015.

Z. Hu. Social interactions and households' flood insurance decisions. *Journal of Financial Economics*, 144(2):414–432, 2022.

P. Issler, R. Stanton, C. Vergara-Alert, and N. Wallace. Mortgage markets with climate-change risk: Evidence from wildfires in california. *Available at SSRN 3511843*, 2020.

E. X. Jiang and A. L. Zhang. Collateral value uncertainty and mortgage credit provision. *Available at SSRN*, 2022.

J. M. Keenan, T. Hill, and A. Gumber. Climate gentrification: from theory to empiricism in miami-dade county, florida. *Environmental Research Letters*, 13(5):054001, 2018.

A. Kermani and F. Wong. Racial disparities in housing returns. Technical report, National Bureau of Economic Research, 2021.

B. J. Keys and P. Mulder. Neglected no more: Housing markets, mortgage lending, and sea level rise. Technical report, National Bureau of Economic Research, 2020.

R. O. King. Federal flood insurance: The repetitive loss problem. LIBRARY OF CONGRESS WASHINGTON DC CONGRESSIONAL RESEARCH SERVICE, 2005.

S. G. Knowles and H. C. Kunreuther. Troubled waters: The national flood insurance program in historical perspective. *Journal of Policy History*, 26(3):327–353, 2014.

C. Kousky. Understanding the demand for flood insurance. *Natural Hazards Review*, 12(2): 96–110, 2011.

C. Kousky and H. Kunreuther. Addressing affordability in the national flood insurance program. *Journal of Extreme Events*, 1(01):1450001, 2014.

C. Kousky, H. Kunreuther, B. Lingle, and L. Shabman. The emerging private residential flood insurance market in the united states. *Wharton Risk Management and Decision Processes Center*, 2018.

W. Kriesel and C. Landry. Participation in the national flood insurance program: An empirical analysis for coastal properties. *Journal of Risk and Insurance*, 71(3):405–420, 2004.

W. W. Lang and L. I. Nakamura. A model of redlining. *Journal of Urban Economics*, 33(2): 223–234, 1993.

A. Lemann. Rolling back the tide: Toward an individual mandate for flood insurance. *Fordham Environmental Law Review*, 26(2):166–218, 2015.

M. Marschoun. Availability of mortgage loans in volatile real estate markets. *Journal of Urban Economics*, 47(3):443–469, 2000.

A. Mian and A. Sufi. *House of debt: How they (and you) caused the Great Recession, and how we can prevent it from happening again.* University of Chicago Press, 2015.

A. Mian, A. Sufi, and F. Trebbi. Foreclosures, house prices, and the real economy. *The Journal of Finance*, 70(6):2587–2634, 2015.

E. Michel-Kerjan, S. Lemoyne de Forges, and H. Kunreuther. Policy tenure under the us national flood insurance program (nfip). *Risk Analysis: An International Journal*, 32(4): 644–658, 2012.

E. O. Michel-Kerjan. Catastrophe economics: the national flood insurance program. *Journal of economic perspectives*, 24(4):165–86, 2010.

J. Murfin and M. Spiegel. Is the risk of sea level rise capitalized in residential real estate? *The Review of Financial Studies*, 33(3):1217–1255, 2020.

E. Nance. Exploring the impacts of flood insurance reform on vulnerable communities. *International Journal of Disaster Risk Reduction*, 13:20–36, 2015.

D. D. Nguyen, S. Ongena, S. Qi, and V. Sila. Climate change risk and the cost of mortgage credit. *Review of Finance*, 2022.

A. Ouazad and M. E. Kahn. Mortgage finance and climate change: Securitization dynamics in the aftermath of natural disasters. *The Review of Financial Studies*, 2021.

R. H. Platt. *Disasters and democracy: The politics of extreme natural events.* Island Press, 2012.

J. C. Pope. Do seller disclosures affect property values? buyer information and the hedonic model. *Land Economics*, 84(4):551–572, 2008.

P. Sastry. Who bears flood risk? evidence from mortgage markets in florida. 2022.

F. I. Subcommittee. The national flood insurance program: Past, present, and future. *American Academies of Actuaries: Washington, DC, USA*, page 55, 2011.

R. R. Verchick and L. R. Johnson. When retreat is the best option: flood insurance after biggert-waters and other climate change puzzles. *J. Marshall L. Rev.*, 47:695, 2013.

R. Wilson and A. Din. Understanding and enhancing the us department of housing and urban development's zip code crosswalk files. *Cityscape*, 20(2):277–294, 2018.

F. Wong. Mad as hell: Property taxes and financial distress. *Available at SSRN 3645481*, 2020.

J. B. Wriggins. In deep: Dilemmas of federal flood insurance reform. *UC Irvine L. Rev.*, 5: 1443, 2015.

A. T. Young. Replacing incomplete markets with a complete mess: Katrina and the nfip. *International Journal of Social Economics*, 2008.

Appendix Figures and Tables

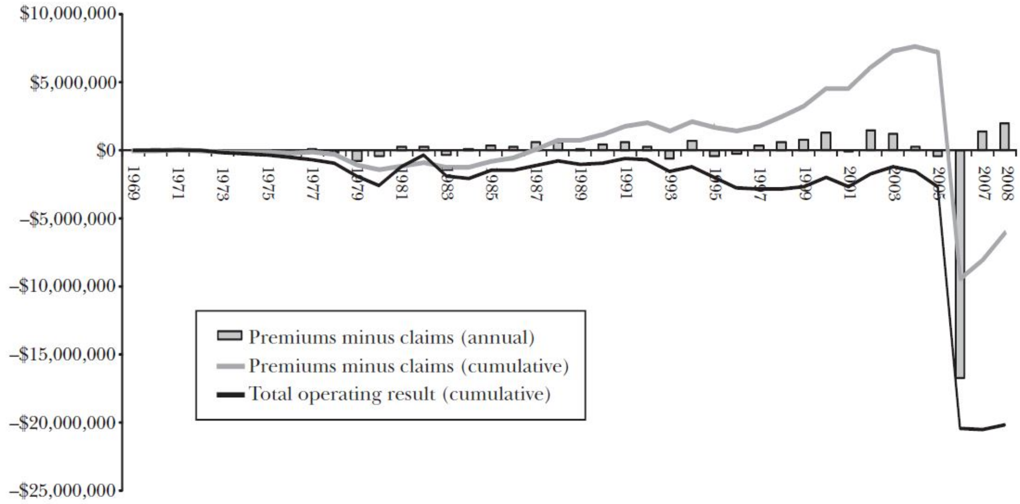

Figure A.1. Financial Operation of the NFIP

This figure is taken from Michel-Kerjan (2010). It depicts the NFIP's financial balance over time. The vertical bars indicate the difference between collected premiums and paid claims nationwide for each year between 1969 and 2008 (in 2008 prices). The light grey line shows the cumulative sum of this difference. The dark grey line further takes into account the operating expenses of the NFIP.

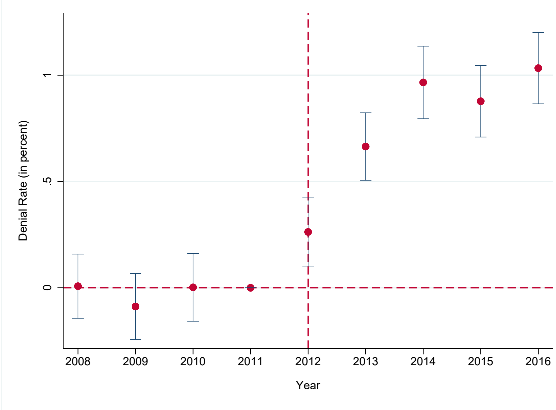

(a)

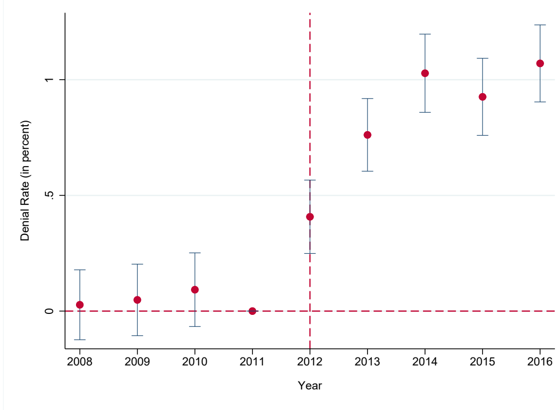

(b)

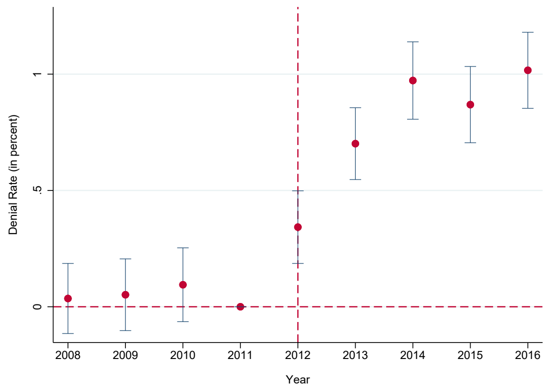

(c)

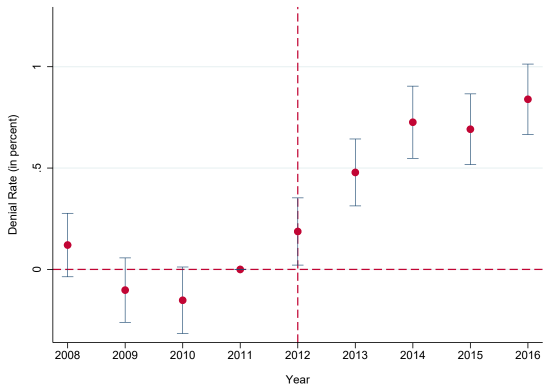

(d)

Figure A.2. Parallel Pre-trends in Alternative Specifications

This figure shows the dynamic effects of BW12 on mortgage denial rates using alternative specifications as discussed in Columns 2 to 5 of Table 2. In general, this figure plots the coefficient estimates of $\{\beta_{\tilde{t}}\}$ from the regression: $DenialRate_{it} = \beta_0 + \sum_{\tilde{t}=2007}^{2016} \beta_{\tilde{t}} * Treated_i \times \mathbb{1}_{t=\tilde{t}} + \beta_2 * Treated_i + \alpha_t + \epsilon_{it}$. Panel (a) includes $White_{it}$ and $Male_{it}$ into the above regression. Panel (b) further includes $Income_{it}$ and $Loan\text{-}to\text{-}income_{it}$. Panel (c) includes state fixed-effects. Panel (d) includes state-year fixed effects. Standard errors are clustered at the census-tract level. The bands around the coefficient estimates show the 95% confidence intervals.

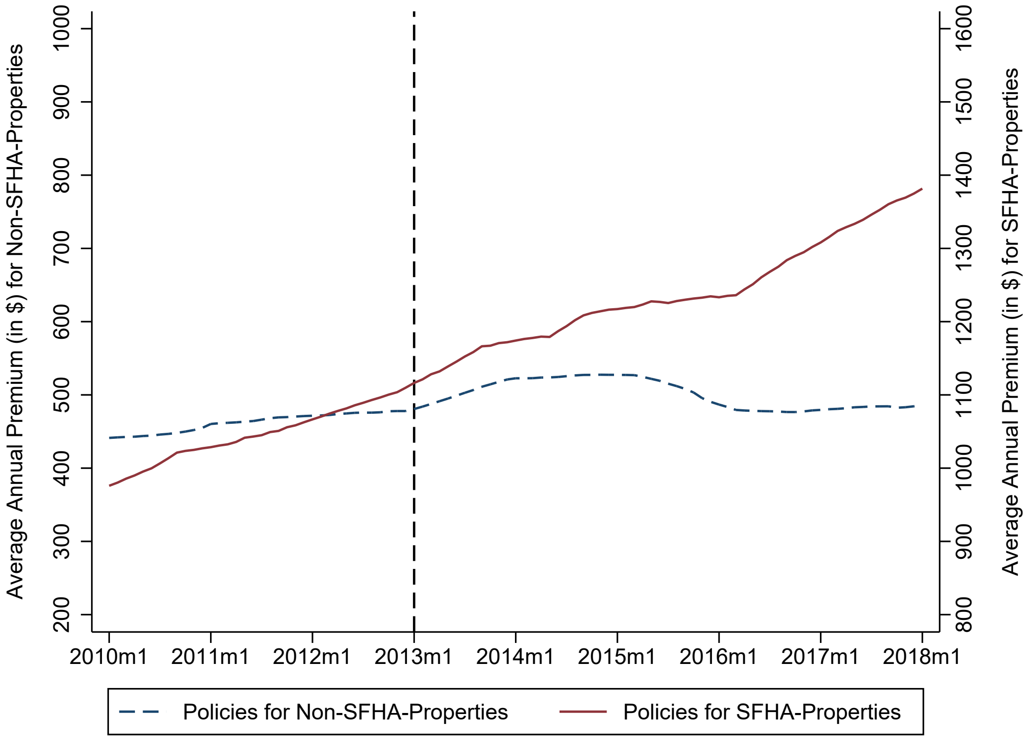

Figure A.3. Raw Trends of the Annual Premiums

This figure shows the average flood insurance premiums over time for non-SFHA properties (dashed line, left y-axis) and SFHA properties (solid line, right y-axis).

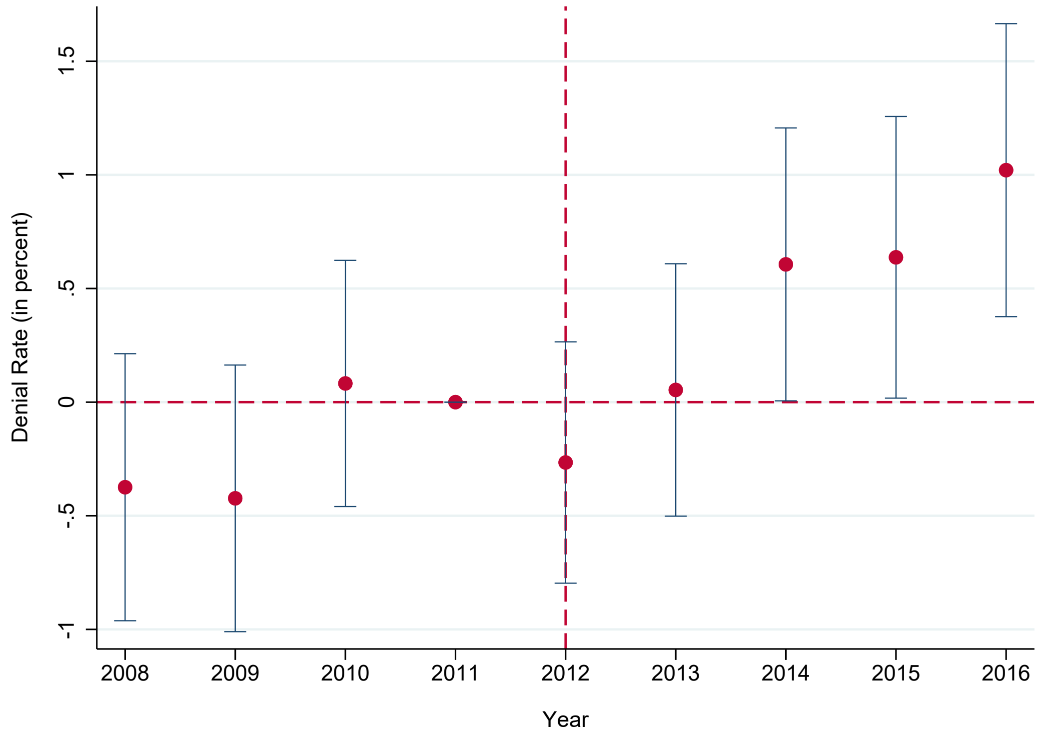

Figure A.4. Difference-in-differences at the three-digit-ZIP level

This figure replicates the main results shown in Figure 3. The only difference here is that the treated and control units are defined at the three-digit-ZIP level. This figure plots the coefficient estimates of $\{\beta_{\tilde{t}}\}$ from the regression: $DenialRate_{it} = \beta_0 + \sum_{\tilde{t}=2007}^{2016} \beta_{\tilde{t}} * Treated_i \times \mathbb{1}_{t=\tilde{t}} + \beta_2 * Treated_i + \alpha_t + \epsilon_{it}$. Standard errors are clustered at the three-digit-ZIP level. The bands around the coefficient estimates show the 95% confidence intervals.

Table A.1: The Effect of Flood Insurance Disincentive on Mortgage Denial Rates

This table shows results from the difference-in-differences regression: $log(DenialRate_{it}) = \beta_0 + \beta_1 * Treated_i \times Post_t + \beta_2 * Treated_i + \beta_3 * Post_t + \gamma * X_{it} + \epsilon_{it}$. All variables are defined same as in Table 2 except that the dependent variable is in logarithm. Standard errors are clustered at the census-tract level and presented in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	(1)	(2)	(3)	(4)	(5)
$Treated_i \times Post_t$	0.039***	0.038***	0.039***	0.037***	0.025***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
$Post_t$	-0.195***	-0.200***	-0.198***	-0.197***	-0.138***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.036)
$Treated_i$	0.059***	0.100***	0.080***	0.067***	0.074***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
$White_{it}$		-0.008***	-0.007***	-0.007***	-0.007***
		(0.000)	(0.000)	(0.000)	(0.000)
$Male_{it}$		0.006***	0.007***	0.006***	0.006***
		(0.000)	(0.000)	(0.000)	(0.000)
$Income_{it}$			-0.001***	-0.001***	-0.002***
			(0.000)	(0.000)	(0.000)
Loan-to-income$_{it}$			0.003***	0.005***	0.004***
			(0.001)	(0.002)	(0.001)
$Constant$	3.214***	3.521***	3.584***	3.666***	3.613***
	(0.002)	(0.009)	(0.009)	(0.012)	(0.013)
State FE	N	N	N	Y	N
State-year FE	N	N	N	N	Y
Observations	542,132	541,999	541,974	541,974	541,974
R-squared	0.050	0.179	0.228	0.302	0.344

Table A.2: Application Outcome and Borrower Characteristics

This table shows results from the loan-application-level regression: $Denied_{it} = \beta_0 + \beta_1 * Income_{it} + \beta_2 * White_{it} + \beta_3 * Male_{it} + \beta_4 * \textit{Loan-to-income}_{it} + \epsilon_{it}$. $Denied_{it}$ is the dummy variable indicating if the mortgage application i filed in year t is rejected by the lender. $Income_{it}$ is the annual income (in thousand) of the applicant. $White_{it}$ is the dummy variable indicating if the race of the applicant is white. $Male_{it}$ is the dummy variable indicating if the applicant is a male. $\textit{Loan-to-income}_{it}$ is the loan-to-income ratio. The coefficient estimates and standard errors are presented in percent. Standard errors are clustered at the census-tract level and presented in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	(1)	(2)	(3)	(4)	(5)
Income	-0.012***	-0.012***	-0.012***	-0.011***	-0.011***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
White		-9.546***	-9.456***	-9.499***	-9.085***
		(0.076)	(0.077)	(0.070)	(0.065)
Male		1.922***	1.919***	1.892***	1.880***
		(0.020)	(0.020)	(0.019)	(0.019)
Loan-to-income			0.251***	0.255***	0.258***
			(0.012)	(0.012)	(0.012)
Constant	24.343***	30.369***	29.552***	30.972***	31.805***
	(0.036)	(0.080)	(0.087)	(0.255)	(0.324)
Observations	97,442,225	86,072,653	86,072,653	86,072,653	86,072,653
R-squared	0.002	0.009	0.011	0.015	0.025
State FE	N	N	N	Y	N
State-year FE	N	N	N	N	Y