

Adverse Selection and Climate Risk: A Response to Ouazad and Khan (2021)

Michael Lacour-Little¹ Andrey Pavlov² Susan Wachter³

October 31, 2022

This research was sponsored by Fannie Mae. The views expressed are the authors' alone and not necessarily those of Fannie Mae nor of the Federal Housing Finance Agency. We also acknowledge helpful comments from Jim Wilcox, Will Doerner, Debra Fuller, Will Larson, Michela Barba, and Gray Kimbrough. We are also grateful to Amine Ouazad and Matt Kahn for sharing their data and computer code.

The data and estimation code for our work is available at https://1sfu-my.sharepoint.com/:f:/g/personal/kdk3_sfu_ca/EuW0cca_d2lKop8BpaQDY5UBoQBMi-Pohlax1FlxijhG0g

¹Senior Director - Economics, Fannie Mae, michael_lacourlittle@fanniemae.com

²Professor of Finance, Beedie School of Business, Simon Fraser University, apavlov@sfu.ca

³Sussman Professor and Professor of Real Estate and Finance, The Wharton School, University of Pennsylvania, wachter@wharton.upenn.edu

Abstract

We identify two critical data-related errors in the work of Ouazad and Kahn (O&K, RFS 2021). Correcting either one of these errors, or both, reverses their original results and eliminates the claimed effect. We document this when using their original data and code. We also perform our entirely independent analysis, which finds no effect as well. However, introducing the two errors into our code generates statistically significant results similar to those O&K report.

The two corrections we implement are to use the correct FHFA conventional loan limits for each county and year and to compare the individual loan amount to that limit accurately. Our results show that there is no statistical evidence that lenders alter their loan origination and securitization behavior in affected areas following a major hurricane. While our findings show that lenders have not transferred climate risk to mortgage insurers in the past, this could occur in the future, especially if climate risk becomes easier to estimate and/or worsens.

1 Introduction

Ouazad and Kahn (O&K, RFS 2021) claim that mortgage lenders change their lending behaviour following a major hurricane. Specifically, O&K claim that lenders originate more loans that are eligible for securitization and also securitize a larger portion of the loans in affected areas. If true, this finding would imply that institutions that securitize mortgages or have exposure to losses from securitized mortgages are subject to adverse selection at the loan origination stage.

We repeat O&K's analysis using both their original data and code as well as our own entirely independent data retrieval and analysis. In doing so, we identify two data processing errors in the original work. Correcting either of these errors reverses the findings reported in O&K's paper. We document this conclusion using both the original data and code and also our own independent analysis. Furthermore, if we introduce the two data processing errors into our code, we find statistically significant results similar to those O&K report.

In other words, the two errors generate significant results both in the original code and in our implementation. Correcting either of the two errors eliminates the significance of the results both in the original code and in our independent implementation.

The Federal Housing Finance Agency (FHFA) specifies precise loan amount limits for a loan to be eligible for securitization through the two Government-Sponsored Enterprises (GSEs), Fannie Mae and Freddie Mac. This limit is known as the "conventional loan limit" and is a statutory requirement.¹ For instance, the conventional limit in Collier County, Florida in 2012 was \$448,500, meaning that only loans originated up to and including this amount were eligible for securitization. These loan limits vary by county and through time.

These precise conventional loan limits generate a well-documented discontinuity of loan orig-

¹Conventional loan limits are readily available at <https://singlefamily.fanniemae.com/originating-underwriting/loan-limits>

inations just above and below the limit. O&K use this discontinuity to find that loans originated just below the limit increase relative to loans originated just above the limit in affected areas following a major hurricane. If correct, such a finding would suggest that loan originators alter their behavior to make more loans eligible for securitization in affected areas.

We document two inaccuracies in O&K's data processing. First, their final data used for model estimation contains only two conventional limits for each year - "standard" and "high-cost". For instance, in 2012, all counties in their data have limits of either \$417,000 or \$625,500. While the standard limit at the time was indeed \$417,000, each "high-cost" county had its own specific limit. For instance, the FHFA limit for Collier County, Florida in 2012 was \$448,500, while the limit in O&K's data for that county was incorrectly set at \$625,500.

The second inaccuracy in O&K's data is in the comparison of each loan amount to the conventional limit. The HMDA loan origination data contains loan amounts rounded to the nearest \$1000, while the FHFA loan limits are reported to the dollar. For instance, the standard conventional loan limit in 2005 was \$359,650. Loans originated near or at the limit that year are reported in the HMDA data as having loan amount of "360" to denote a rounded amount of \$360,000. Thus, a direct comparison of the loan amount (\$360,000) to the conventional limit (\$359,650) results in misclassification of the loan as being above-limit. The following year, 2006, the conventional limit changed to \$417,000, which would not generate any misclassification. In this example, the originations below the limit would appear to have substantially increased (originations above the limit decreased) from 2005 to 2006 even though there was no such increase (decrease). While this is only one specific example, it is noteworthy because five hurricanes occurred in 2005, and the apparent increase in below-limit originations was incorrectly attributed to an effect from these hurricanes rather than to the change in loan classification from 2005 to 2006.

Section 4.3 below details the above issues and quantifies their significance and prevalence

in the data. As a summary, correcting either of the above discrepancies, or both, reverses O&K's original results and shows that there is no evidence lenders altered their origination behavior in affected areas following a major hurricane. Therefore, there is no evidence of unpriced climate risk, as O&K originally claim.

The above empirical findings aside, it is not clear to us that the original identification of treatment and control groups is appropriate. O&K use estimated or actual damage following a hurricane to identify the treated observations. While each storm can revise the way we estimate future risk, those revisions are not confined to the area affected by the particular event. Instead, risk estimate revisions are likely based on certain geographic features and, therefore, extend well outside the affected area. For instance, a particular storm may change how we think about elevation or distance to water, but this change would apply everywhere, not just inside the storm path. Therefore, the treatment observations would be the ones with revised future risk estimate, regardless of their proximity to the particular storm.

Finally, while our findings suggest that lenders do not alter their lending and origination behavior in a statistically discernible way, at least in terms of the original identification method, this could change in the future. As climate-related disasters increase in magnitude and frequency, it is possible that lenders take steps to reduce their exposure to risk areas. One way of doing so is to increase the securitization of loans originated in risk areas. Therefore, we believe that continuous monitoring of lending and securitization in risk areas is prudent.

We discuss the data selection and methodology in detail below. Following this, we present the results of various data selection and estimation methods. Finally, we conclude with an overall summary and suggestions for future monitoring of climate risk-related lending behavior.

2 Literature Review

Kousky et al. (2020) offer an in-depth overview of flood risk and flood insurance effects on the U.S. housing market. They identify several mechanisms that lead to a substantial portion of homeowners who are either under-insured or not at all insured against flood risk. In addition to the devastating effects this can have on individual households, the authors note that uninsured losses may be passed to lenders, servicers, and ultimately to taxpayers. The authors note that while insured property damage improves loan performance, there are several mechanisms that can expose a lender or a loan guarantor to losses, including under-insurance of mortgaged properties and uninsured nearby properties which are then left unrepaired following a disaster.

Ratnadiwakara and Venugopal (2020) also investigate the effects of flooding and find an increase in interest rates and increase in securitizations in affected areas. The authors document a particularly interesting mechanism that drives these findings, namely, that areas subject to disaster events attract less affluent and less creditworthy homebuyers.

Kousky et al. (2020) further offers a comprehensive review of the literature that links home prices to flooding and other climate-related risk. The general finding of this literature is that homes in floodplains and other high-risk areas have lower values than comparable homes outside those areas. (MacDonald et al., 1990; Harrison et al., 2001; Bin et al., 2008; Daniel et al., 2009; Bakkensen and Barrage, 2017; Ortega and Taspınar, 2018; Zhang and Leonard, 2018; Bernstein et al., 2019; Tibbetts and Mooney, 2018; Kusisto, 2018). However, as documented by Bin and Kruse (2006) and Bin et al. (2008), this analysis is often complicated by the attractive amenities that floodplains often provide. This complication makes fully conclusive and robust results difficult to obtain. Nevertheless, taken in aggregate, the evidence is strong that flood risk does have a negative impact on home values. The effect is less clear for other climate-related risks, mostly because the identification of high-risk areas

and homes is harder to define.

The impact of flooding and other risks on individual properties is magnified by the impact of disaster events on entire neighbourhoods. This impact is sometimes due to infrastructure damage, but more often due to a high concentration of homes that were damaged but not repaired. For instance, Kotkin (2014) shows that the number of blighted properties in New Orleans nearly doubled following Hurricane Katrina. Similarly, as already mentioned, Ratnadiwakara and Venugopal (2020) document that home prices fall following a natural disaster and the income distribution of home buyers in the affected areas changes. Earlier studies, such as Masozera et al. (2007) and Vigdor (2008), also document how natural disasters affect different neighbourhoods in very different ways. These findings are important because they suggest that even if a lender or mortgage insurer strictly enforces the requirements for appropriate property insurance for all borrowers, it is still exposed to the risk of neighbourhood decline and the resulting decline in property values.

Regardless of the exact mechanism through which flooding and other natural disasters negatively affect home values, the prevalence of this effect is sufficient grounds for lenders and mortgage insurers to monitor their exposure to such risks. This finding underscores the importance of monitoring potential adverse selection in the underwriting and securitization activities of lenders and insurers as no institution would want to have a more-than-proportional exposure to climate risks.

Beyond the impact on property values, natural disasters impact households in several other ways. For instance, Gallagher and Hartley (2017) investigate how households affected by Hurricane Katrina adjusted their finances. They find that any increases in credit card and other debt following the disaster were short-lived. Long-term households in damaged areas reduced overall borrowing, primarily because they used flood insurance payments to reduce their mortgages. This was particularly the case if the pre-disaster home values were near or below the cost of re-building.

However, a research report by the Urban Institute (2019) documents a very significant financial impact on affected households. It shows that natural disasters have a negative impact on credit scores and mortgage performance and lead to increases in credit card and other high-cost debt. Contrary to Gallagher and Hartley (2017) findings, the Urban Institute (2019) identifies many cases in which these effects persist and worsen over time.

Perhaps a way to reconcile the two studies above is through the work of Kousky et al. (2020) who show that flood insurance acts as intended and protects homeowners and mortgage lenders from the immediate effects of a disaster. However, as discussed above, not all homeowners have insurance and, even if they do, it does not protect them against overall neighbourhood decline. Therefore, it is possible that insured households are protected and either rebuild, or, at the very least, repay their mortgages using insurance payouts. However, uninsured homeowners, and homeowners in neighbourhoods that are severely affected as a whole, suffer substantial negative consequences which in many cases persist over time. On the other hand, rebuilt homes are newer which, in turn, can have an overall positive impact on the neighbourhood.

3 Data and Methodology

In what follows, we describe the data sources, the data definitions, and the estimation methods we utilize. We comment on the reasons for our choices and, when possible, provide alternatives that we have utilized as robustness checks. We further provide comparison to the methods used by O&K both throughout the data and methodology sections and in a separate section below.

3.1 Data Sources

This section describes the data sources and processing that we perform for our independent analysis. We are also grateful to Ouazad and Khan for sharing the data and code for their final model estimation. As stated above, the reversal of the original result occurs both in O&K’s original data and code and in our independent analysis described below.

3.1.1 Loan origination data

Our primary source for mortgage application, origination, and securitization data was collected under the Home Mortgage Disclosure Act (HMDA). This data provides detailed mortgage data at the census tract level, which is the primary level of geographic analysis we use. The specific data is available from the National Archives Catalog (Non-restricted Ultimate Loan Application Registers) for the period 1995 to 2014 and from the Consumer Financial Protection Bureau for the period 2015 – 2020.^{2, 3, 4}

For each reported loan application, this data set includes the lending institution, loan amount, borrower income, race, and ethnicity. The data also contains the census tract of the property, the type of property (1-4 family, manufactured housing, or multifamily), the purpose of the loan (home purchase, home improvement, or refinancing), owner-occupancy status of the property, loan preapproval status, and the outcome of the application (loan denied, loan approved but not accepted, loan withdrawn by the applicant, loan originated, or loan purchased by the lending institution). The data also includes information about the type of institution which purchased the loan and whether it was sold in the same calendar year. The data uses the 1990 census for 1995 to 2002 applications, the 2000 census for 2003

²https://catalog.archives.gov/search?q=*&f.ancestorNaIds=2456161&sort=naIdSort%20asc

³<https://www.consumerfinance.gov/data-research/hmda/historic-data/>

⁴O&K use a combination of final and ultimate LARs in different years. While in our base case we always use the ultimate LARs, we have also replicated O&K’s analysis using the mix they employ.

to 2012 applications, and the 2010 census for applications afterward.

HMDA groups all properties with 1 to 4 units into a single category. This represents a potential challenge because properties with 1, 2, 3, and 4 units have different conventional loan limits. In our analysis, we use the 1-unit conventional limit for all loans in our analysis. We justify this necessary simplification on the basis that only a very small portion of multifamily homes are owner-occupied.

We filter the HMDA records as follows:

- Include only loans in the states from Maine through Texas that are along the Atlantic coast,
- include only conventional loans (i.e., any loan other than FHA, VA, FSA, or RHS),
- include only loans for owner-occupied home purchase,
- include only loans for one to four family housing,
- exclude manufactured housing,
- exclude borrowers who had income exceeding \$5,000,000 or loan-to-income ratios above 4.5 or below 1,⁵
- exclude applications with missing or invalid census tract information, and
- exclude loans with missing income information.

The HMDA data includes many loan outcome categories and classification of approved and securitized loans is subject to some uncertainty. Therefore, we focus our analysis on originated loans, which are clearly identified in the data. Nevertheless, we present securitization

⁵O&K's paper specifies the same exclusion. However, we have been unable to ascertain whether this exclusion has been applied in the implementation or not. Regardless, we perform our analysis both with and without this exclusion.

analysis as well, and we have performed robustness checks using alternative loan outcome classifications.

As discussed in the methodology section below, our analysis relies on the loans originated just above and just below the conforming loan limit. Specifically, we consider loans that were within +/- 5%, +/- 10%, and +/- 20% of the single unit conforming loan limit. This limit varies by year and county and is available from the Federal Housing Finance Agency.

Our main estimation method uses indicator variables that capture loan originations for four years before and after each hurricane. When this window overlaps with the transition in the HMDA data from using the 1990 census to the 2000 census, and from the 2000 census to the 2010 census, we provide consistent definition of affected and non-affected areas. To do so, we match the census tracts in the earlier and the later census by the largest share of population.

The above use of the HMDA data generally follows the methodology of O&K as described in their paper. The biggest difference is that that we focus our analysis on census tracts, while O&K report all of their results using postal zip codes. Nevertheless, as a robustness check, we perform our analysis at the zip code level as well. We convert census tracts to zip codes using the US Census Bureau relationship files.⁶

O&K further report that they merge the HMDA data with the Black Knight McDash data, although we have not been able to ascertain this in their implementation. Regardless, such a merge is not needed for our work as the original HMDA data already allows for the analysis of the approval, origination, and securitization outcomes. Therefore, we have only used the original HMDA data without merging it to other loan origination data sets.

⁶<https://www.census.gov/geographies/reference-files/time-series/geo/relationship-files.2000.html>

3.1.2 Hurricane Wind Path Data

We obtain hurricane wind data from the NOAA HURDAT2 dataset.⁷ This data includes hurricane location, wind speed, and atmospheric pressure data from 1851-2019, sampled every 6 hours. Hurricane size data, presented as wind speed radius by quadrant, is available from 2004.

We apply the following filters and calculations to this data:

- Only hurricane observations since 2004 were considered. We provide results both for the original sample (2004 to 2012) and for an updated sample to 2016. We require data for at least 4 years after a hurricane, so events after 2016 were not considered in the base case. As a robustness check, we have also estimated the model using all events up to 2018, but only with two-year before and after comparison. Table 1 lists the hurricanes we consider in our base case analysis.
- For each hurricane observation, we calculate the distance between the hurricane center and the internal point latitude and longitude of each census tract. Census tracts were flagged as affected if the distance was less than the 64 kt wind radius in the relevant quadrant. Census tract internal point locations were taken from the US Census Bureau Gazetteer files.⁸

We use census tracts both for identification of treatment and for our main analysis. As a robustness check, we also convert census tracts to zip codes, as discussed above. Table 1 lists the hurricanes used in our base case analysis. In one of our robustness specifications, we also utilize hurricanes that occurred in 2017 and 2018: Harvey, Irma, Maria, Florence, and Michael.⁹

⁷<https://www.nhc.noaa.gov/data/>

⁸<https://www.census.gov/geographies/reference-files/time-series/geo/gazetteer-files.html>

⁹O&K use the same events up to and including Hurricane Sandy.

Name of Hurricane	Year of Hurricane
CHARLEY	2004
FRANCES	2004
IVAN	2004
JEANNE	2004
DENNIS	2005
KATRINA	2005
OPHELIA	2005
RITA	2005
WILMA	2005
DOLLY	2008
GUSTAV	2008
IKE	2008
IRENE	2011
ISAAC	2012
SANDY	2012
MATTHEW	2016

Table 1: **List of hurricanes considered in our analysis.** This table lists the hurricanes we define as disaster events in our base case analysis. We further include the additional hurricanes that occurred in 2017 and 2018 in one of our robustness tests.

3.1.3 Wetlands/Open Water

We use wetlands and open water data from the USGS 2001 National Land Cover Database to identify areas at risk.¹⁰ This is a single GeoTIFF raster image file that contains geographical land cover information, including the locations of wetlands and open water. We calculate the minimum geodesic distance from the internal point latitude and longitude of each 2000 and 2010 census tract to wetland or open water with ArcGIS.

In our analysis, we define a census tract to be near wetland or open water if its internal point is within 1.5 km of the respective geographic feature.¹¹

3.1.4 Coastline

We further define coastline based on the NOAA Medium Resolution Shoreline, which is described by NOAA as a “high-quality, GIS-ready digital vector data set created for general use” and covers “over 75,000 nautical miles of coastline, . . . representing the entire continental United States of America”.¹² We calculated the minimum geodesic distance from the internal point latitude and longitude of each 2000 and 2010 census tract to the shoreline using ArcGIS.¹³

3.1.5 Elevation

Our elevation data comes from the USGS Digital Elevation Model.¹⁴ The model is a seamless 1/3 arc-second resolution dataset, with ground spacing of approximately 10 meters north

¹⁰<https://www.mrlc.gov/data/nlcd-2001-land-cover-conus>

¹¹O&K report performing this identification at the census blockgroup level. It is unclear why this was necessary, as O&K’s statistical analysis is performed at the postal zip code level. We performed our analysis at the census tract, zip code, and county levels.

¹²<https://shoreline.noaa.gov/data/datasheets/medres.html>

¹³In a manner similar to wetlands, O&K use blockgroups to identify subject observations. It is not clear why this was needed, as their statistical analysis is at the postal zip code level.

¹⁴<https://apps.nationalmap.gov/downloader/#/>

and south. Our data is comprised of a series of raster images dated 2019-2021. We obtained GeoTIFF tiles for each of the 18 Atlantic states from the National Map Downloader and imported and merged them in ArcGIS. This allows us to calculate the minimum elevation for each 2000 and 2010 census tract polygon.

3.2 Identification of Treatment Observations

A key step in performing the analysis we describe below is the identification of treatment observations, that is observations that fall in an area that experienced damage. Note that the goal is not to identify damage to specific properties but instead to identify areas (census tracts, zip codes, or counties) that experienced damage from each of the hurricanes we examine.

Except for Hurricane Sandy, there are no publicly available measures of damage at the geography level we require. Therefore, we define treatment areas, meaning areas that were likely subject to damage, using wind speed, distance to water, and elevation data. We identify the treatment areas as follows:

- Center of census tract is within the 64-knot wind radius of each hurricane
- Census tracts with minimum elevation below 3 meters,
- Census tracts within 1.5 km from wetland, open water, or shoreline,
- For the specifications based on zip codes, a zip code was defined as treated if it had at least one census tract that was treated.

There is one instance in which actual damage data is available – Hurricane Sandy. For the regressions that used actual damage data, we treated a census tract if the census tract

reported damage in the FEMA damage survey, which covers block groups with at least ten damaged units.

We further repeat our analysis with the NOAA *Sea, Lake, and Overland Surges from Hurricanes* (SLOSH) data by treating zip codes where the average storm surge height is estimated to be greater than 1 foot and are within the 64 kt wind radius of each hurricane. We have also used the same data to identify treatment areas, as an alternative to the identification described above.

As a measure of robustness, we create several classification tree predictors of damage from Hurricane Sandy using the variables described above - whether the census tract is in the 64 kt hurricane path, the CT's minimum elevation, the CT's distance to wetland, and the CT's distance to shoreline. In some cases, we also incorporate additional variables, such as income per capita, poverty rate measures, and property market characteristics.

Furthermore, we also replicate the main analysis using various sub-samples. For instance, we limit the sample to "large originators", which we define to be institutions with at least 14,000 or 20,000 loan applications in the HMDA data. We also create an alternative filter for "largest originator only".

Finally, in addition to census tracts and zip codes, we use counties as a geographic unit for our analysis. For "counties at 50, 100, 200 mi from center of hurricane", we define as treated all the CTs in counties that meet those distance conditions.

O&K identify a treated observation if more than 40% of the blockgroup surface area within the zip code for that observation had predicted damage. However, blockgroups and census tracts do not aggregate into zip codes in a clean way. Even in a single point in time, many blockgroups and census tracts fall in two zip codes. Furthermore, many 1990, 2000, and 2010 census tracts have overlapping boundaries, further complicating the consistent identification of treated observations, especially when the main unit of geographic analysis is zip codes.

This is one of the reasons to use Census Tract as the geographic unit of analysis, as we have done in our base case estimation. However, we also repeat O&K’s approach as a robustness test.

3.3 Estimation Method

We follow O&K in estimating the following regression model:

$$\begin{aligned}
Outcome_{it} = & \alpha BelowConformingLimit_{ijy(t,d)} + \gamma BelowConformingLimit_{ijy(t,d)} \times Treated_{j(i)} \\
& + \sum_{t=-T}^{+T} \xi_t Treated_{j(i)} \times Time_t \\
& + \sum_{y=1995}^{2016} \xi_y BelowConformingLimit_{ijy(t,d)} \times Year_{y(t)} \\
& + \sum_{t=-T}^{+T} \delta_t BelowConformingLimit_{ijy(t,d)} \times Treated_{j(i)} \times Time_t \\
& + Year_{y(t,d)} + Disaster_d + CT_{j(i)} + \epsilon_{it}
\end{aligned} \tag{1}$$

where i represents the loan application, $j(i)$ represents the geographic region (census tract) of the loan application and $d = 1, 2, \dots, 16$ indexes the hurricanes. The first and third summations run from $t = -4$ years prior to the hurricane to $t = +4$ years after the hurricane, and exclude $t = -1$, the reference year. $Outcome_{it}$ is an indicator variable that specifies whether the application was approved, originated, or securitized, depending on the regression. $BelowConformingLimit_{ijy(t,d)}$ specifies whether the loan amount was at or below the conforming limit for the year and county in which it was requested. $Treated_{j(i)}$ specifies whether the loan was requested in a treated census tract. For this variable, the value was set to 1 for applications in census tracts that met the treatment group conditions at least

once, regardless of which year they were impacted. $Year_{y(t)}$ is a series of indicator variables, representing the year in which the loan was requested. The $Time_t$ variable, in the first and third summations, indicates if the loan was requested at T-4, T-3, ..., T+3, T+4 years since the CT was treated by a hurricane, excluding the reference year T-1. $Disaster_d$ is a series of indicator variables representing each hurricane, specifying whether the loan was in a CT affected by hurricane d . In the last line of the equation, the Year, Disaster, and CT variables were controlled for, but their coefficients were not reported, due to the large number of variables (23,700 CTs, with each CT having its own variable, and 5,100 ZIPs). When necessary, CT was replaced with ZIP. All standard errors are two-way clustered by geographic region and year.

When adding the more recent hurricanes to 2018, we restrict t to -2 to +2, as this allows us to include more recent hurricanes and still consider the post-disaster origination history.

We estimate the regression model above with the linear fixed effects package in “R”. This package “transforms away factors with many levels prior to estimating an OLS regression. This is useful for estimating linear models with multiple group fixed effects, and for estimating linear models which uses factors with many levels as pure control variables.” In other words, this package is specifically designed to estimate linear regression models when using multiple group fixed effects, such as CTs/ZIPs/Counties in our implementation. The package also allows for multi-way clustered standard errors.

4 Empirical results

4.1 Base Case

Table 2 reports the main finding of our analysis using the base-case methodology and sample construction we describe in sections 3.1, 3.2, and 3.3 above. This table uses census tract as the geographic unit of analysis. It reports the results based on approved, originated, and securitized loans using sub-samples that are within 5, 10, and 20 percent of the conventional loan limit for each census tract and each year. The conventional limit for some counties changed mid-year in 2008 and in 2011. In our base case analysis, we have used the limits from the first part of those years, before the change. We have also repeated the analysis using the limits from the second part of those years, and with excluding observations from those two years completely. Neither of those alternative methods materially alter the results.

<i>Dependent variable:</i>									
	Approved ±5%	Originated ±5%	Securitized ±5%	Approved ±10%	Originated ±10%	Securitized ±10%	Approved ±20%	Originated ±20%	Securitized ±20%
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
belowLimit:T_minus.4	-0.074*** (0.015)	-0.056*** (0.020)	0.009 (0.044)	-0.049* (0.025)	-0.033 (0.020)	0.018 (0.044)	-0.044* (0.018)	-0.041** (0.019)	0.015 (0.042)
belowLimit:Treatment:T_minus.3	-0.021 (0.018)	-0.016 (0.019)	0.021 (0.035)	-0.021 (0.013)	-0.018 (0.012)	0.018 (0.028)	-0.032*** (0.011)	-0.028** (0.010)	0.039* (0.021)
belowLimit:Treatment:T_minus.2	0.026* (0.014)	0.018 (0.016)	0.073*** (0.016)	0.005 (0.009)	0.002 (0.011)	0.073*** (0.012)	0.019*** (0.005)	0.012 (0.007)	0.067*** (0.012)
belowLimit:Treatment:T_minus.0	-0.0001 (0.008)	-0.012 (0.011)	0.013 (0.033)	-0.012* (0.007)	-0.015 (0.010)	0.016 (0.032)	-0.009 (0.006)	-0.016** (0.007)	0.018 (0.026)
belowLimit:Treatment:T_plus.1	-0.003 (0.007)	0.004 (0.008)	-0.017 (0.025)	0.001 (0.007)	0.004 (0.008)	-0.018 (0.023)	-0.002 (0.007)	0.002 (0.010)	-0.014 (0.024)
belowLimit:Treatment:T_plus.2	-0.006 (0.018)	0.005 (0.024)	-0.030 (0.037)	-0.002 (0.011)	0.013 (0.016)	-0.021 (0.034)	0.012 (0.013)	0.016 (0.013)	-0.005 (0.033)
belowLimit:Treatment:T_plus.3	-0.023 (0.018)	-0.034* (0.018)	-0.009 (0.041)	-0.015 (0.012)	-0.028** (0.011)	-0.001 (0.045)	-0.017 (0.010)	-0.026*** (0.008)	0.016 (0.054)
belowLimit:Treatment:T_plus.4	-0.044*** (0.009)	-0.041*** (0.011)	-0.053* (0.030)	-0.034*** (0.009)	-0.035*** (0.007)	-0.046 (0.030)	-0.025** (0.010)	-0.027*** (0.010)	-0.034 (0.035)
Observations	1,599,285	1,599,285	1,248,152	2,589,696	2,589,696	1,998,895	4,722,448	4,722,448	3,626,309
R ²	0.042	0.046	0.115	0.038	0.043	0.138	0.034	0.039	0.154
Adjusted R ²	0.027	0.031	0.099	0.028	0.033	0.128	0.029	0.034	0.148

Note: *Clustered standard errors in parenthesis.*

* p<0.1; ** p<0.05; *** p<0.01

Table 2: Base Case Results. The table reports the results of our base case estimation. The reported coefficients are the interactions between a loan being below its respective conventional limit with being in a treated area, meaning area with substantial damage, and being approved, originated, or securitized in a specific year with a four-year window around each event. The coefficient estimates of particular interest are those following each hurricane, denoted by T+1 to T+4.

Table 2 reports the results of our base case estimation. The reported coefficients are the interactions between a loan being below its respective conventional limit with being in a treated area, meaning area with substantial damage, and being approved, originated, or securitized in a specific year with a four-year window around each event. The coefficient estimates of particular interest are those following each hurricane, denoted by T+1 to T+4.¹⁵

The estimation uses the complete model specification as described above. The table reports only the interaction effects of interest – Below Limit x Treated at time – 4 years to + 4 years relative to the hurricane that treated each respective observation. In other words, the coefficient estimate we report in the table is for the interaction that captures a loan that is below its conventional loan limit, is in a census tract that was treated, meaning damaged as defined above, and was originated or approved four years prior to the respective hurricane, three years prior, and so on up to and including four years after the hurricane. The interaction term for one year prior to the hurricane is omitted, as required by the indicator variable estimation methods.

The coefficient estimates for the interaction terms reported in Table 2, therefore, capture the difference between originations or approvals above and below the conventional limit for treated census tracts within a four-year window before and after each event. The interaction coefficients for the years after a hurricane are of particular interest because they capture the change in the difference between below and above limit originations in treated areas following a hurricane relative to this difference in areas that were not treated. Positive and significant coefficients for the interaction terms in years t+1 to t+4 would suggest that lenders increase originations below the conventional limit in affected areas following a hurricane. Such a finding, in turn, would suggest that, following an event, at least some of the increased risk of future events is transferred to mortgage insurers.

¹⁵We note that the securitization data are subject to some noise because loans originated in one year but securitized in a subsequent year may not be properly flagged as "securitized" in the HMDA data. Therefore, we base our conclusions primarily on the approval and origination data, and report the securitization results only to be consistent with O&K's work.

As reported in Table 2, of the total of 36 regression coefficient estimates for the years following a hurricane (denoted as $t+1$ to $t+4$) not even one is positive and significant. Statistical significance apart, the coefficients for the year after a hurricane, denoted as $t+1$, are very small and in some cases even negative. For instance, the coefficient for originations in the year following a hurricane for the 5% window around the conventional limit is 0.004. This coefficient is also precisely estimated, with a standard error of 0.008. The 95 percent confidence interval for this coefficient is, therefore, -0.012 to 0.02, making it indistinguishable from zero. The coefficient for approvals is even negative, at -0.003, and is even more precisely estimated with a standard error of just 0.007.

This leads to the conclusion that in our base case analysis, there is no detectable increase in the below-limit originations relative to above-limit originations in affected areas relative to non-affected areas following a hurricane. In other words, there is no evidence that an increase in future risk is transferred to mortgage insurers.

Table 2 contains some coefficient estimates, particularly for the interaction for four years following a hurricane, that are statistically significant. However, these coefficients are unlikely to capture an actual effect, as they occur in isolation with coefficients for one, two, and three years following a hurricane being indistinguishable from zero. Furthermore, all those significant coefficients are negative. Even if taken at face value, these results only further rule out the possibility that below-limit originations increased following a hurricane.

For comparison purposes, we replicate the analogous table from O&K in the Appendix. In addition to the difference in estimated coefficients, our results are based on nearly twice as many observations, even more in the case of 20% distance from the conforming loan limit. It is also not clear how in O&K's analysis the number of observations in the 10% and 20% distance from the conforming limit is nearly identical.

The Appendix also provides the results of our base case implementation using zip codes,

rather than census tracts, as the unit of geographic analysis. Since the original loan data is provided at the census tract level, we see no particular advantage in using zip codes. We only provide these results to facilitate comparison to O&K's work.

4.2 Robustness Analysis

Our base case analysis presented in the previous section incorporates the model and data selection choices that we believe are most appropriate to address the issue at hand. However, the analysis is complex and offers numerous opportunities to implement alternative methods that are also reasonable. Table 3 summarizes the base case and the alternative specifications we have considered.

Method	Base Case	Implemented Alternatives
Unit of geographic analysis	Census tract	<ul style="list-style-type: none"> • Postal Zip Code • County (within 50, 100, or 200 mi from center of hurricane)
HMDA time window	1995-2020	2004-2012, 1995-2016, 2004-2016, 1995-2017, 2004-2017
HMDA data classification	All Ultimate LARs (includes corrections up to 18 months following initial submission)	<ul style="list-style-type: none"> • Mix of Final and Ultimate LARs – Final LARs are a snapshot of the loan register at the end of filing period
Hurricane dates	2004-2016	2004-2012, 2004-2018, 2013-2018, 2004-2017, 2013-2017
Damage identification	Detailed method described above	<ul style="list-style-type: none"> • SLOSH model (1 ft surge) • Own classification tree models based on Sandy actual damage, implemented to the entire sample – Variables included share of CT in wetland, distance to hurricane center, average elevation in CT • Counties < 50, 100, 200 mi from hurricane center • 2-year combined treatment dummies • Limit sample to large originators only (minimums of 5,000 to 20,000 observations after other filters were applied) • Actual damage (only available for Sandy) • Census tracts with more than 3 (or 6) treated observations before and after event • Limit sample to small originators only (maximums of 200 to 500 observations after other filters were applied) • Limit sample to single largest originator

Continued on next page

Method	Base Case	Implemented Alternatives
Data sub-samples	All data	<ul style="list-style-type: none"> • Limit data to flood plains only • Limit data to observations in counties with 1+, 3+, or 5+ treated CTs • Limit data to observations in states with 1+, 3+, or 5+ treated CTs • Exclude applications classified as incomplete or withdrawn before a credit decision was made • Exclude pre-approvals • Exclude observations in treated ZIPs or CTs that occurred outside of T-4 to T+4 window • Remove filters on income or loan-to-income
Dependent variable definitions	<ul style="list-style-type: none"> • Originated <ul style="list-style-type: none"> – action taken = “Loan Originated” vs all others • Approved <ul style="list-style-type: none"> – action taken = “Loan Originated” or “Loan approved but not accepted” vs all others • Securitized <ul style="list-style-type: none"> – purchaser type was Gov’t entities; observations were restricted to loans originated or loans purchased by the institution 	<ul style="list-style-type: none"> • Securitized <ul style="list-style-type: none"> – Purchaser type was not “N/A” • Approved <ul style="list-style-type: none"> – Loans not denied vs loans denied

Table 3: Summary of alternative specifications

With various combinations of alternative implementations and other minor variations, we have estimated well over 100 individual models and data selection choices.

The majority of the above specifications generate results that are similar to our base case, and do not identify any evidence of increased below-limit originations relative to above-limit originations in treated areas relative to non-treated areas following a hurricane. We have identified the following three exceptions to this conclusion: using actual damage from hurricane Sandy, limiting the sample to large originators only and limiting the sample to flood plains only

Actual damage data from hurricanes is generally not available. Therefore, in the base case described above, we estimate the areas with likely damage based on whether they fell within the 64-knot wind area, elevation, and distance to water. We also estimate this using alternative damage measures, such as the SLOSH model described above.

Actual damage is available for Hurricane Sandy. We replicate the analysis using the base model specification but using actual damage rather than predicted damage to identify treated census tracts. Table 4 presents our findings. In this case, the coefficient estimate for the interaction term one year following Sandy, denoted as time $t+1$, is positive (0.043) and statistically significant.

	<i>Dependent variable:</i>					
	Approved ±5%	Originated ±5%	Approved ±10%	Originated ±10%	Approved ±20%	Originated ±20%
	(1)	(2)	(3)	(4)	(5)	(6)
belowLimit:Treatment:T_minus_4	-0.029** (0.012)	-0.040*** (0.014)	-0.045*** (0.011)	-0.042*** (0.012)	-0.050*** (0.005)	-0.051*** (0.007)
belowLimit:Treatment:T_minus_3	-0.154*** (0.020)	-0.056** (0.022)	-0.149*** (0.012)	-0.093*** (0.013)	-0.138*** (0.008)	-0.095*** (0.008)
belowLimit:Treatment:T_minus_2	-0.007 (0.016)	-0.064*** (0.016)	-0.039*** (0.011)	-0.068*** (0.013)	-0.014* (0.008)	-0.043*** (0.011)
belowLimit:Treatment:T_minus_0	0.131*** (0.016)	0.126*** (0.015)	0.016 (0.011)	0.009 (0.014)	-0.011 (0.008)	-0.010 (0.010)
belowLimit:Treatment:T_plus_1	0.042*** (0.014)	0.043** (0.016)	0.036*** (0.012)	0.034** (0.012)	-0.003 (0.007)	-0.006 (0.009)
belowLimit:Treatment:T_plus_2	-0.009 (0.013)	0.026 (0.017)	-0.024** (0.009)	0.010 (0.009)	-0.023*** (0.007)	-0.003 (0.008)
belowLimit:Treatment:T_plus_3	-0.008 (0.011)	0.017 (0.015)	-0.010 (0.010)	0.019* (0.011)	-0.007 (0.008)	0.012 (0.009)
belowLimit:Treatment:T_plus_4	-0.011 (0.013)	-0.015 (0.012)	0.003 (0.009)	0.008 (0.009)	-0.002 (0.006)	-0.012 (0.007)
Observations	1,614,208	1,614,208	2,610,326	2,610,326	4,756,808	4,756,808
R ²	0.042	0.046	0.038	0.042	0.034	0.039
Adjusted R ²	0.027	0.031	0.028	0.033	0.029	0.033

Note: Clustered standard errors in parenthesis.

*p<0.1; **p<0.05; ***p<0.01

Table 4: **Loan Approvals and Originations Around Hurricane Sandy – Actual Damage.** The table reports the estimation results for one of the three cases in which we find positive and significant coefficients following a hurricane. In this case, we restrict our attention only to Hurricane Sandy and use actual damage data as reported by FEMA. Sandy is the only major storm for which such data is publicly available. While the table reports several positive and statistically significant coefficients, these results are not robust to alternative model specifications and are of limited economic significance.

While we report this result as one of the few cases in which we find a positive and significant effect, we cannot assign very much economic significance to it. First, the result is not robust to alternative specifications, such as the use of zip codes as a geographic unit of analysis or other small model alterations. In those alternatives, we find no significant effect.

Moreover, this result was obtained following numerous specification and data selection attempts. Therefore, it could very well be due to chance and have little external validity.

Even if we take the results following Hurricane Sandy at face value, we note that originations or approvals below the limit decrease relative to those above the limit prior to Sandy. For the coefficients following the event to be meaningful, we would need to see no change in the years prior to an event. This is known as the "parallel trends assumption", which is required for the post-event coefficients to be meaningful. In other words, the post-event coefficients reported in Table 4 have limited, if any, economic significance.

The second situation in which we identify a positive and significant result is when we limit the sample to large originators only. For instance, limiting the sample to originators with 20,000 or more observations in the data generates the results reported in Table 5. While the coefficient for one year after a hurricane is not significant, we do find some statistical significance for the coefficient for two years after a hurricane, denoted by $t+2$. We observe a similar pattern when we limit the data to large originators using other cut-off levels.

Once again, the statistical significance of the coefficient for two years after a hurricane does not translate into actual economic significance. First, only the coefficient for two years after an event is consistently positive and significant. All other coefficients, including the coefficient on the year following a hurricane, are not significant, and even change signs across the different specifications. Even if we take the significance in the coefficient for $t+2$ at face value, we cannot conclude that origination behavior changes because none of the coefficients for $t+1$ are significant.

Furthermore, same as the first exception, this result was obtained as part of numerous alternative estimations. As such, it could very well be the result of random chance and cannot be used to derive an economically meaningful relationship.

	<i>Dependent variable:</i>					
	Approved ±5%	Originated ±5%	Approved ±10%	Originated ±10%	Approved ±20%	Originated ±20%
	(1)	(2)	(3)	(4)	(5)	(6)
belowLimit:Treatment:T_minus_4	-0.091* (0.046)	-0.026 (0.034)	-0.059* (0.029)	-0.027 (0.022)	-0.034 (0.026)	-0.020 (0.018)
belowLimit:Treatment:T_minus_3	0.008 (0.031)	0.001 (0.048)	0.018 (0.015)	0.022 (0.014)	-0.012 (0.012)	0.001 (0.012)
belowLimit:Treatment:T_minus_2	0.044 (0.065)	0.009 (0.082)	-0.002 (0.033)	-0.007 (0.030)	0.005 (0.023)	0.004 (0.021)
belowLimit:Treatment:T_minus_0	0.003 (0.021)	-0.005 (0.024)	-0.001 (0.026)	0.005 (0.032)	-0.001 (0.013)	-0.005 (0.015)
belowLimit:Treatment:T_plus_1	0.048 (0.035)	0.019 (0.032)	0.026 (0.025)	0.019 (0.028)	-0.013 (0.024)	-0.009 (0.019)
belowLimit:Treatment:T_plus_2	0.080*** (0.020)	0.090*** (0.026)	0.053** (0.021)	0.067*** (0.021)	0.033* (0.017)	0.041** (0.018)
belowLimit:Treatment:T_plus_3	-0.069 (0.137)	-0.096 (0.114)	0.052 (0.063)	0.016 (0.071)	0.010 (0.049)	-0.008 (0.051)
belowLimit:Treatment:T_plus_4	0.261** (0.121)	0.205 (0.141)	0.046 (0.037)	0.065 (0.038)	0.054** (0.024)	0.060*** (0.015)
Observations	248,386	248,386	585,422	585,422	1,312,987	1,312,987
R ²	0.143	0.126	0.130	0.111	0.103	0.086
Adjusted R ²	0.082	0.064	0.100	0.080	0.087	0.070

Note: Clustered standard errors in parenthesis.

*p<0.1; **p<0.05; ***p<0.01

Table 5: **Loan Approvals and Originations Using Large Originators Only.** This table reports the estimation results for the second of the three cases in which we find positive and significant coefficients following a hurricane. In this case, we restrict the sample to large originators only. While the table reports several positive and statistically significant coefficients for two years following a hurricane, the coefficients for one year following an event are not significant. These results are also not robust to alternative model specifications and, therefore, are of limited economic significance.

The final exception in which we find some positive and significant coefficient estimates is when we limit the observations to census tracts that overlap with established floodplains. Table 6 reports this analysis, with several different definitions of whether census tract overlaps with a floodplain or not. While most coefficients reported in this table are not statistically significant or are negative, there are two coefficients in the 60 and 70 percent overlap level

that are positive and significant.

The above result is perhaps the closest to generating an overall positive and significant post-treatment effect. Even so, we note that the coefficients in question are only positive and significant with two overlap levels considered. Alternative specifications using lower or higher overlap levels do not generate positive and significant coefficient estimates. Furthermore, the coefficients for two, three, and four post-event years are not significant and/or are negative. Taken together, these results suggest no overall effect, especially in the context of very high sensitivity to the exact definition of what census tracts are included in a floodplain.

One potentially interesting observation is that the coefficient estimates for the year of a hurricane, denoted by $t=0$, are negative and significant for many floodplain overlap definitions. This finding has no implication for the potential adverse selection issue addressed in this paper, but is certainly curious and worthy of a further investigation. We note, however, that the coefficients during the year of the hurricane are identified over very few originations in affected areas, and, therefore, cannot be taken as particularly reliable.

	<i>Proportion of census tract within a floodplain</i>				
	$\geq 0\%$	$\geq 10\%$	$\geq 20\%$	$\geq 40\%$	$\geq 50\%$
	(1)	(2)	(3)	(4)	(5)
belowLimit:Treatment:T_minus_4	-0.056*** (0.020)	-0.067*** (0.019)	-0.082*** (0.020)	-0.09*** (0.027)	-0.106*** (0.023)
belowLimit:Treatment:T_minus_3	-0.016 (0.020)	-0.012 (0.025)	-0.001 (0.028)	0.016 (0.030)	0.045 (0.043)
belowLimit:Treatment:T_minus_2	0.018 (0.016)	0.009 (0.014)	0.011 (0.014)	0.029 (0.024)	0.050 (0.039)
belowLimit:Treatment:T_minus_0	-0.012 (0.011)	-0.025*** (0.007)	-0.034*** (0.010)	-0.041*** (0.010)	-0.041*** (0.012)
belowLimit:Treatment:T_plus_1	0.004 (0.008)	0.011 (0.011)	-0.001 (0.008)	0.013 (0.016)	0.032 (0.022)
belowLimit:Treatment:T_plus_2	0.005 (0.023)	-0.006 (0.023)	-0.011 (0.022)	-0.022 (0.019)	-0.003 (0.027)
belowLimit:Treatment:T_plus_3	-0.034* (0.018)	-0.042** (0.018)	-0.046** (0.018)	-0.051*** (0.024)	-0.044 (0.032)
belowLimit:Treatment:T_plus_4	-0.041*** (0.011)	-0.044*** (0.010)	-0.039*** (0.013)	-0.033* (0.017)	-0.025 (0.019)
Observations	1,599,285	618,828	336,096	159,477	114,854

Continued on next page

	<i>Proportion of census tract within a floodplain</i>				
	$\geq 60\%$	$\geq 70\%$	$\geq 80\%$	$\geq 90\%$	$\geq 95\%$
	(1)	(2)	(3)	(4)	(5)
belowLimit:Treatment:T_minus_4	-0.123*** (0.043)	-0.108 (0.064)	-0.165*** (0.059)	-0.238*** (0.088)	-0.225*** (0.083)
belowLimit:Treatment:T_minus_3	0.008 (0.049)	-0.042 (0.061)	-0.026 (0.065)	-0.076 (0.062)	-0.242 (0.129)
belowLimit:Treatment:T_minus_2	0.064** (0.031)	0.026 (0.037)	0.068 (0.059)	0.183** (0.069)	0.121 (0.107)
belowLimit:Treatment:T_minus_0	-0.041** (0.018)	-0.038* (0.023)	-0.128*** (0.032)	-0.113*** (0.038)	-0.291 (0.172)
belowLimit:Treatment:T_plus_1	0.059*** (0.021)	0.052** (0.022)	0.038 (0.029)	0.041 (0.033)	-0.054 (0.105)
belowLimit:Treatment:T_plus_2	-0.008 (0.036)	-0.014 (0.052)	-0.032 (0.051)	-0.054 (0.077)	-0.097 (0.103)
belowLimit:Treatment:T_plus_3	-0.07** (0.027)	-0.079** (0.037)	-0.062 (0.063)	-0.152* (0.091)	-0.141 (0.093)
belowLimit:Treatment:T_plus_4	-0.015 (0.024)	-0.015 (0.032)	-0.064 (0.042)	-0.071 (0.085)	0.067 (0.183)
Observations	74,788	51,697	30,001	16,615	9,114

Note: Clustered standard errors in parenthesis.

*p<0.1; **p<0.05; ***p<0.01

Table 6: Table 6 reports the estimation results for the last of the three cases in which we find positive and significant coefficients following a hurricane. In this case, we restrict the sample to originations within floodplains only. Our main data source, HMDA, does not identify whether an origination is in a floodplain. Instead, we determine if origination is in a floodplain if the census tract in which it is located overlaps with a floodplain. We report the results for various overlap definitions – from 0 to 95 percent. While the table reports several positive and statistically significant coefficients for certain overlap definitions, particularly in the 60 and 70 percent cases, this significance is not robust to alternative overlap definitions. Therefore, these results are also of limited economic significance.

In short, the three cases in which we identify potentially positive and statistically significant coefficients are likely due to spurious relationships and cannot be taken as an indication of changes in origination behavior following a hurricane. On the contrary, their high sensitivity to model and sample specification is consistent with our base case conclusion of no effect.

Even in the unlikely event that the above statistically significant coefficients are not spurious

and hint at an actual change in lender behavior following a hurricane, such a change would be present in only very specific circumstances. This does not alter our main conclusion that there is no identifiable system-wide effect.

4.3 Comparison to O&K's results

Whenever possible, we have followed the model specification and data selection utilized by O&K. However, we reproduce O&K's original results only when using both their data selection and their model estimation code exactly. Instead, as reported above, our base case analysis suggests no evidence of changes in origination behavior or adverse selection.

While our work follows O&K's overall approach and econometric method, our analysis is distinct in several details, as we have pointed out in the data selection and methodology sections above. Beyond those technical details, our work differs from that of O&K in two important ways.

First, the HMDA data contains loan amounts rounded to the nearest \$1000, while the FHFA loan limits are reported to the dollar. For instance, the standard conventional loan limit in 2005 was \$359,650. Loans originated near or at the limit that year are reported in the HMDA data as having loan amount of "360" to denote a rounded amount of \$360,000. Thus, a direct comparison of the loan amount (\$360,000) to the conventional limit (\$359,650) results in misclassification of the loan as being above-limit while it is at or below the conventional limit in reality. The following year, 2006, the conventional limit changed to \$417,000, which would not generate any misclassification. In this example, the originations below the limit would appear to have substantially increased from 2005 to 2006 even though there was no such increase. While this is only one specific example, it is noteworthy because five hurricanes occurred in 2005, and the apparent increase in below-limit originations could incorrectly be attributed to an effect from these hurricanes rather than to the change in loan classification

between 2005 and 2006.

In the code O&K provided to us, we believe the incorrect comparison occurs on line 80 of file named:

```
2_analysis_hmda/04_make_hmda_event_study_data_set.R
```

This line is as follows:

```
diff_log_loan_amount = log_loan_amount - log(effective_loanlimit/1000)
```

This line compares the loan amount from the HMDA data (which is rounded to the nearest \$1000) to the effective loan limit, which is provided to the dollar. Substituting a loan amount and conventional limit of \$359,650 in the line of code above generates the following expression:

```
diff_log_loan_amount = log(360) - log(359650/1000)
```

This calculation generates a positive number, which in turn incorrectly classifies the loan as being above the conventional limit.

The correct comparison would round the conventional loan limit in the second term above to the nearest \$1000, so that the variable "diff_log_loan_amount" becomes zero and the loan is correctly classified as being at or below the FHA conventional limit.

Figure 1 quantifies the scope of the mis-classification. The figure reports the number of correctly classified and mis-classified originations in the +/- 5% of conventional limit sample. In our analysis, we explicitly handle this discrepancy in the formatting of the source data and thus appropriately classify all loans as above- or below-limit.

The extent of this mis-classification is also evident in Figure 5 provided by O&K. Below we reproduce this figure from the original paper (left panel) and also provide an equivalent




Figure 1: The figure depicts the number of correctly classified and mis-classified originations from the +/-5% sample.

year	state_code	county_code	county_name	cbsa_number	state	Replication Limit (single-unit)	O&K Limit
2012	9	1	FAIRFIELD	14860	CT	601450	625500
2012	12	21	COLLIER	34940	FL	448500	625500
2012	12	87	MONROE	28580	FL	529000	625500
2012	13	133	GREENE	99999	GA	515200	625500
2012	24	3	ANNE ARUNDEL	12580	MD	494500	625500
2012	24	5	BALTIMORE	12580	MD	494500	625500
2012	24	9	CALVERT	47900	MD	625500	625500
2012	24	13	CARROLL	12580	MD	494500	625500
2012	24	17	CHARLES	47900	MD	625500	625500
2012	24	21	FREDERICK	47900	MD	625500	625500
2012	12	23	COLUMBIA	29380	FL	417000	417000

Table 7: The table provides an example of several counties which had a variety of conforming limits in 2012. In O&K's analysis, these counties have the same incorrect limit of \$625,500.

figure that correctly classifies all originations (right panel). The original figure (left) has an unusually high observation count just above the limit. Since almost no loans are ever originated just above the conventional limit, the count above the limit should be very low. Our replication (right) shows that indeed this count is very low.

The second point of substantive departure from O&K's methodology is that we use the conventional loan limits as provided by the FHFA. Those limits differ by year and county. Table 7 provides an example of several high-cost counties who had a variety of conforming limits in 2012. In O&K's analysis, all these counties have the same incorrect conforming limit of \$625,500, which is simply 50% higher than the standard limit of \$417,000.




Figure 2: The figure reproduces O&K's Figure 5 (left) and our replication (right). The original figure shows a high count of originations above the limit, which is not realistic. Our replication shows that the originations just above the limit are very low, as expected.

In our work, we match each origination to the conventional limit applicable to it at the time of approval, origination, or securitization. This, in turn, provides for the correct classification of loans as above- or below-conventional limit.

The above two differences are important. Using *both* the incorrect two-category conventional limits *and* the incorrect comparison procedure that O&K use results in positive and significant coefficients very similar to what they report. This occurs both in our own data and in the data O&K use. However, using *either* the correct FHFA conventional limits *or* the correct rounding procedure, *or* both, eliminates the significance of the results both in our own data and in O&K's data.

To summarize the impact of the rounding and conventional limits we replicate O&K's code as closely as possible, including the use of zip codes and limiting the sample to the original events. Table 8 reports the results of this analysis using incorrect and correct rounding and incorrect and correct conventional loan limits. The table reports the coefficient estimates and their significance for originations for time T+1 at the +/-5% sample. The first row reports the coefficients using the incorrect rounding procedure that O&K use with various high-cost conventional limits. For those counties, O&K assign a limit of 1.5 times the standard limit. We repeat the analysis using this limit as well as limits of 1.3 times, 1.4 times, 1.6 times and 1.7 times the standard limit. The last column represents the coefficient estimate using the correct limits.

The second row of Table 8 reports the coefficients using a correct rounding procedure, again for various incorrect limits and the correct limit. The bottom right coefficient of 0.004 is the outcome of correct rounding and correct conventional limits, but retaining all other aspects of O&K's approach and data.

The first row, third column coefficient of 0.044 matches exactly the T+1 coefficient for originated loans reported in Table 2 of O&K, which is also replicated in the appendix here.

The last coefficient on the second row, 0.004, is analogous to the coefficient for T+1 for originations in our base case reported in Table 2 above. ¹⁶

Table 8 shows that using the incorrect rounding *and* incorrect high-cost limits generates the positive and significant coefficient that O&K report. This occurs for a wide range of incorrect high-cost conventional limits, in addition to the limit O&K use. Using the correct conventional limits (last column) *or* the correct rounding procedure *or both* eliminates the significance of the coefficient of interest. ¹⁷

	<i>High Cost County Conventional Loan Limit</i>					
	1.3x	1.4x	1.5x (O&K)	1.6x	1.7x	Correct Limits (FHFA)
Incorrect Rounding (O&K)	0.031** (0.012)	0.033*** (0.007)	0.044** (0.016)	0.023*** (0.007)	0.023* (0.011)	0.026* (0.015)
Correct Rounding	0.000 (0.016)	0.005 (0.009)	0.021 (0.022)	-0.001 (0.009)	-0.006 (0.016)	0.004 (0.018)

*p<0.1; **p<0.05; ***p<0.01

Table 8: The table reports the outcome of matching O&K’s analysis as closely as possible, except for the conventional limits and the rounding procedure. The first row, third column coefficient of 0.044 matches O&K’s analogous coefficient exactly. The bottom right coefficient is analogous to our base case reported above, and uses both the correct limits and the correct rounding

5 Conclusion

We document that the claims presented O&K are based on two incorrect treatments of the data. Correcting either of the two inaccuracies, or both, reverses the original O&K results. We find this both in O&K’s original data and code and in our own independent implementation.

¹⁶While the bottom right coefficient of 0.004 is very close to the analogous coefficient reported in Table 2 above, they are not identical because in Table 2 we present our best approach, while in Table 8 we have attempted to replicate the O&K procedure as closely as possible in all aspects other than the rounding and the conventional limits.

¹⁷The coefficient estimate when using incorrect rounding but correct conventional limits is significant at the 10% level in this particular implementation. However, small changes in the specification removes even this level of significance.

This finding means that we have no evidence that lenders change their origination behavior in affected areas following a major hurricane. Therefore, adverse selection in the securitization process is not detectable in the data.

While we find no evidence of risk transfer or adverse selection in the data, the hypothesized risk transfer and adverse selection could occur in the future. Increased frequency and magnitude of extreme climate events makes the assessment of their risk easier going forward and increases the incentives to transfer it to other market participants. Therefore, there is a clear need to design and implement mechanisms and procedures that monitor for potential risk transfer going forward.

6 References

- Bakkensen, L. A. & Barrage, L. (2017), Flood risk belief heterogeneity and coastal home price dynamics: Going under water?, Technical report, National Bureau of Economic Research.
- Bernstein, A., Gustafson, M., & Lewis, R. (2018). Disaster on the horizon: The price effect of sea level rise. SSRN. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3073842
- Bin, O., & Kruse, J. B. (2006). Real estate market response to coastal flood hazards. *Natural Hazards Review*, 7(4), 137–144.
- Bin, O., Crawford, T. W., Kruse, J. B., & Landry, C. E. (2008). Viewscapes and flood hazard: Coastal housing market response to amenities and risk. *Land Economics*, 84(3), 434–448.
- Bin, O., Kruse, J. B., & Landry, C. E. (2008). Flood hazards, insurance rates, and amenities: Evidence from the coastal housing market. *Journal of Risk and Insurance*, 75(1), 63–82.
- Daniel, V. E., Florax, R. J., & Rietveld, P. (2009). Flooding risk and housing values: An economic assessment of environmental hazard. *Ecological Economics*, 69(2), 355–365.
- Gallagher, J., & Hartley, D. (2017). Household finance after a natural disaster: The case of Hurricane Katrina. *American Economic Journal: Economic Policy*, 9(3), 199–228.
- Harrison, D., T. Smersh, G., & Schwartz, A. (2001). Environmental determinants of housing prices: The impact of flood zone status. *Journal of Real Estate Research*, 21(1-2), 3–20.
- Kousky, C., Palim, M., & Pan, Y. (2020). Flood risk and mortgage credit risk: A case study of Hurricane Harvey. *The Journal of Housing Research*, 29(S1).
- Kousky, C., Kunreuther, H., LaCour-Little, M., & Wachter, S. (2020). Flood Risk and the

U.S. Housing Market. *Journal of Housing Research*, 29:sup1, S3-S24.

Kotkin, J. (2014). *Sustaining prosperity: A long term vision for the New Orleans region*. Greater New Orleans, Inc.

Kusisto, L. (2018, August 25). *Flooding risk knocks \$7 billion off home values, study finds*. *The Wall Street Journal*.

MacDonald, D. N., White, H. L., Taube, P. M., & Huth, W. L. (1990). *Flood hazard pricing and insurance premium differentials: Evidence from the housing market*. *Journal of Risk and Insurance*, 654–663.

Masozera, M., Bailey, M., & Kerchner, C. (2007). *Distribution of impacts of natural disasters across income groups: A case study of New Orleans*. *Ecological Economics*, Elsevier, vol. 63(2-3). 299-306.

Ortega, F., & Taspınar, S. (2018). *Rising sea levels and sinking property values: Hurricane Sandy and New York's housing market*. *Journal of Urban Economics*, 106, 81–100.

Ouazad, A. & Kahn, M. (2021). *Mortgage Finance and Climate Change: Securitization Dynamics in the Aftermath of Natural Disasters*. *The Review of Financial Studies*, 35(8), 3617–3665. <https://doi.org/10.1093/rfs/hhab124>

Mortgage Finance and Climate Change: Securitization Dynamics in the Aftermath of Natural Disasters (June 2021). NBER Working Paper No. w26322, Available at SSRN: <https://ssrn.com/abstract>

Ratnadiwakara, D., & Venugopal, B. (2020). *Do Areas Affected by Flood Disasters Attract Lower-Income and Less Creditworthy Homeowners?*, *Journal of Housing Research*, 29:sup1, S121-S143.

Tibbetts, J. & Mooney, C. (2018, August 20). *Sea level rise is eroding home value, and owners might not even know it*. *The Washington Post*.

Urban Institute. (2019, April 11). Insult to Injury: Natural Disasters and Residents' Financial Health.

Vigdor, J. (2008). The Economic Aftermath of Hurricane Katrina. *Journal of Economic Perspectives*, 22 (4): 135-54.

Zhang, L. & Leonard, T. (2018), 'Flood hazards impact on neighborhood house prices', *The Journal of Real Estate Finance and Economics*, 1-19.

7 Appendix

For ease of reference, in this Appendix we reproduce the main results table from O&K's paper.

Table 2: Impact of Billion-Dollar Events on Approvals, Originations, and Securitization Probabilities

This table presents the estimates of the impact of billion-dollar events on the discontinuity in mortgages' approval rates, origination rates, and in securitization conditional on origination. Mortgages with amounts in the $\pm 20\%$, $\pm 10\%$, and $\pm 5\%$ window of the conforming loan limit are considered in every year and every area between 1995 and 2017 inclusive. Pre- and post-treatment indicator variables estimated in the -4 to $+4$ period. The conforming loan limit is determined annually and differs between high cost and general counties. Standard errors 2-way clustered at the ZIP and year level. The unit of observation is the mortgage application in columns (1)-(6), and origination in columns (7)-(9). The control group is the set of mortgages in Zips of Atlantic states and states of the Gulf of Mexico.

	Dependent variable:								
	±20%	Approved ±10%	±5%	±20%	Originated ±10%	±5%	±20%	Securitized ±10%	±5%
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Below Limit _{it} × Treated _{t=-4}	0.003 (0.014)	0.001 (0.014)	0.008 (0.019)	0.022 (0.026)	0.020 (0.026)	0.028 (0.033)	0.041** (0.019)	0.040** (0.018)	0.071*** (0.022)
Below Limit _{it} × Treated _{t=-3}	0.015 (0.010)	0.014 (0.009)	0.016 (0.020)	0.025 (0.018)	0.025 (0.019)	0.045 (0.035)	-0.002 (0.025)	-0.005 (0.026)	-0.016 (0.026)
Below Limit _{it} × Treated _{t=-2}	-0.002 (0.006)	-0.002 (0.006)	0.003 (0.010)	-0.011 (0.010)	-0.010 (0.010)	0.001 (0.015)	-0.017 (0.024)	-0.018 (0.026)	-0.027 (0.024)
Below Limit _{it} × Treated _{t=+0}	0.003 (0.014)	0.001 (0.013)	0.006 (0.014)	0.007 (0.018)	0.005 (0.017)	0.017 (0.020)	0.006 (0.021)	0.005 (0.021)	0.013 (0.029)
Below Limit _{it} × Treated _{t=+1}	0.025*** (0.008)	0.024*** (0.008)	0.030** (0.011)	0.024** (0.011)	0.024** (0.011)	0.044** (0.016)	0.018 (0.019)	0.017 (0.019)	0.042* (0.021)
Below Limit _{it} × Treated _{t=+2}	0.039*** (0.011)	0.040*** (0.011)	0.063*** (0.015)	0.037** (0.017)	0.037** (0.017)	0.075*** (0.019)	0.045* (0.024)	0.046* (0.024)	0.086*** (0.027)
Below Limit _{it} × Treated _{t=+3}	0.061*** (0.016)	0.062*** (0.016)	0.073*** (0.030)	0.057** (0.021)	0.059** (0.021)	0.080** (0.037)	0.095*** (0.029)	0.097*** (0.029)	0.120** (0.043)
Below Limit _{it} × Treated _{t=+4}	0.021 (0.024)	0.019 (0.024)	0.009 (0.028)	0.007 (0.022)	0.007 (0.022)	0.002 (0.036)	0.154** (0.067)	0.155** (0.066)	0.193*** (0.064)
Other Controls	Treated _t , for $t = -4, \dots, +4$, Year f.e., Disaster f.e., ZIP f.e., Below Limit _{it} × Year _{it} , See specification 1								
Observations	1,345,012	1,310,397	803,424	1,345,012	1,310,397	803,424	1,500,360†	1,461,539†	900,765†
R ²	0.066	0.066	0.069	0.070	0.072	0.249	0.250	0.229	0.229
Adjusted R ²	0.061	0.061	0.064	0.064	0.064	0.246	0.246	0.223	0.223

Note: *p<0.1; **p<0.05; ***p<0.01

†: securitizations of originated mortgages occur for both mortgages originated in the current year and for mortgages originated in previous years. The larger number of observations (originated mortgages) in columns (7)-(9) reflects this.

We further replicate our base case results provided in Table 2 above using zip codes as the unit of analysis. While we see no advantage to using zip codes, we provide the table to facilitate comparison to O&K's results.

<i>Dependent variable:</i>									
	Approved ±5%	Originated ±5%	Securitized ±5%	Approved ±10%	Originated ±10%	Securitized ±10%	Approved ±20%	Originated ±20%	Securitized ±20%
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
belowLimit:Treatment:T_minus_4	-0.054*** (0.010)	-0.044*** (0.012)	0.010 (0.047)	-0.034*** (0.014)	-0.025*** (0.010)	0.013 (0.053)	-0.031** (0.012)	-0.032** (0.013)	0.010 (0.057)
belowLimit:Treatment:T_minus_3	-0.025* (0.014)	-0.012 (0.013)	0.022 (0.035)	-0.025 (0.020)	-0.017 (0.015)	0.027 (0.030)	-0.027 (0.018)	-0.025* (0.014)	0.047* (0.025)
belowLimit:Treatment:T_minus_2	0.025*** (0.008)	0.028*** (0.009)	0.104*** (0.018)	0.006 (0.008)	0.010 (0.010)	0.099*** (0.016)	0.012** (0.005)	0.011* (0.006)	0.091*** (0.016)
belowLimit:Treatment:T_minus_0	0.009 (0.008)	0.006 (0.010)	0.039 (0.029)	-0.003 (0.006)	-0.003 (0.007)	0.042 (0.031)	-0.002 (0.004)	-0.006 (0.005)	0.036 (0.027)
belowLimit:Treatment:T_plus_1	-0.010 (0.010)	-0.004 (0.015)	0.0002 (0.023)	0.008 (0.010)	0.010 (0.013)	-0.003 (0.020)	0.003 (0.007)	0.003 (0.011)	-0.0004 (0.022)
belowLimit:Treatment:T_plus_2	-0.002 (0.017)	0.002 (0.020)	-0.039 (0.047)	0.002 (0.012)	0.011 (0.017)	-0.026 (0.043)	0.011 (0.013)	0.014 (0.013)	-0.005 (0.042)
belowLimit:Treatment:T_plus_3	-0.022 (0.013)	-0.032** (0.015)	-0.017 (0.044)	-0.017* (0.009)	-0.027** (0.012)	-0.003 (0.048)	-0.022*** (0.007)	-0.031*** (0.009)	0.016 (0.057)
belowLimit:Treatment:T_plus_4	-0.035*** (0.009)	-0.036*** (0.011)	-0.052 (0.036)	-0.033** (0.012)	-0.037*** (0.011)	-0.046 (0.038)	-0.022* (0.011)	-0.026* (0.013)	-0.028 (0.043)
Observations	1,614,208	1,614,208	1,260,073	2,610,326	2,610,326	2,015,333	4,756,808	4,756,808	3,653,617
R ²	0.030	0.034	0.102	0.029	0.033	0.128	0.028	0.033	0.146
Adjusted R ²	0.026	0.029	0.096	0.026	0.031	0.125	0.026	0.031	0.145

Note: *p<0.1; **p<0.05; ***p<0.01

Table 9: This table repeats Table 2 above, except using zip code as the unit of geographic analysis. This generates no positive and significant coefficients in the T+1 to T+4 periods, just like our base case implementation using census tracts.