Does the Salience of Climate-Related Risk Affect Asset Prices?

Cloé Garnache*

Abstract

Using regression discontinuity and difference-in-differences designs to compare prices of California homes across a new wildfire risk zone boundary, I identify null effects in the immediate vicinity of the new risk zone boundary. In contrast, prices of homes newly assigned to a riskier zone drop by 5.7% further away from the new risk zone boundary. These findings suggest that households do not respond to the risk salience triggered by the new risk zone assignment, but instead update their risk beliefs in response to re-zoning in a way that is consistent with the actual wildfire risk.

*Cloé Garnache, cgarnache@ethz.ch, ETH Zürich Center of Economic Research (CER-ETH) and Oslo Metropolitan University HousingLab. I thank Soren Anderson, André K. Anundsen, Lint Barrage, Andy Bernard, Pedro Bordalo, Manudeep Bhuller, Todd Guilfoos, Bård Harstad, Tarjei Havnes, Edwin Leuven, Pierre Mérel, Paolo Piacquadio, Erling Røed Larsen, Kjetil Storesletten, Martin Schmalz, Eva Steiner, Gaute Torsvik, Karen Helene Ulltveit-Moe, Maxence Valentin, and seminar participants at the 2019 Environment and Energy Economics NBER Spring meeting, RFS Real and Private-Value Assets Conference, 2019 ASSA meetings, 2022 North American Meeting of the Urban Economics Association, BI Business School, Copenhagen Business School, University of Maryland, Michigan State University, Montana State University, Norwegian University for Life Sciences, Norwegian School of Economics (NHH), Ohio State University, University of Colorado, Boulder, the University of Oslo, and the Frisch Centre for Economic Research for helpful comments. I thank April Athnos and Sophia Tanner for valuable research assistance. This research received funding from US Forest Service Grant 14-CS-11052007-030 and was part of a project funded by ERC Conservation GA 683031.

I. Introduction

Climate-related risks are difficult to assess and process for non-experts, e.g., due to lack of information, limited attention, or the lack of a model. As a result, regulators may be concerned that households do not correctly account for such risks when choosing in which neighborhood to locate or how much to pay for a home. Risk zoning, e.g., for flood or wildfire, has become a commonly-used policy intervention to provide risk information to households and make the risk more salient for homes located inside the risk zone. Risk zoning has the potential to affect the location of housing developments, and thus in turn the economic costs of natural disasters (Rappaport and Sachs (2003), Kahn (2005)). For example, an estimated 46 million homes and $9.2 trillion in property value are currently located in the wildland-urban interface and are exposed to wildfire risk in the United States (International Association of Wildland Fire (2013), Radeloff, Helmers, Kramer, Mockrin, Alexandre, Bar-Massada, Butsic, Hawbaker, Martinuzzi, Syphard, and Stewart (2018)).

How information about low-probability, high-consequence risks is incorporated into the prices of large, illiquid, widely-held assets is a first-order question in asset pricing. Residential real estate is the largest asset class held by US households and a central component of the financial system through mortgage and insurance markets. A growing share of its value is exposed to climate-related hazard. Understanding whether, and how, households as marginal investors in this market process newly-disclosed risk information, and whether they respond to the *salience* of a discrete risk classification rather than to the underlying continuous risk, has direct implications for housing-market price formation, for the spatial reallocation of development relative to risk boundaries, and for the collateral risk borne by mortgage lenders and insurers. Risk re-zoning events provide a setting to examine whether the salience of climate-related risk drives household risk-taking and asset prices.

This paper estimates the effect of a new climate-related risk zone on residential real estate prices, both at the new zone boundary and further away. Comparing the two estimates separates the effect of risk salience from the effect of risk information.[1] I exploit California's 2007 revision of its wildfire risk zones, which required sellers to disclose a home's risk-zone status to prospective buyers. This disclosure makes wildfire risk more salient for homes inside the newly designated risk zone than for nearby homes just outside the boundary. Throughout, I treat risk salience as discontinuous at the new zone boundary. While some homes are newly assigned to the risk zone, others experience a change in their pre-2007 risk zone level. Using real estate sales transactions for the universe of homes in southern California, I exploit this policy shift combined with the spatial discontinuity in the salience of risk at the new risk zone boundary to identify the effect of risk salience on residential real estate prices.

To identify the effect of the risk zone on residential real estate prices, I rely on three empirical designs. First, a spatial regression discontinuity (RD) design compares homes just inside and just outside the new boundary. I find that the estimated local average treatment effect (LATE) on prices is null. Second, a difference-in-discontinuities (RD-DiD) design compares prices across the boundary before and after the 2007 change: the LATE is again null. Third, a difference-in-differences (DiD) design in the immediate vicinity of the boundary yields a null average treatment effect on the treated (ATT). Under the assumption that mandated disclosure makes the weighting of wildfire risk discontinuous at the boundary, these null results indicate that

[1] I use the term salience as defined by Taylor and Thompson (1982): "the phenomenon that when one's attention is differentially directed to one portion of the environment rather than others, the information contained in that portion will receive disproportionate weighting in subsequent judgments."

households do not respond to the increased salience of risk. The estimates are robust to alternative bandwidths, sample selections, covariates, and fixed effects. The null results could instead mean that households ignore wildfire risk entirely. To test this, I estimate the DiD including homes further from the boundary. The ATT is -5.7% in the preferred specification with home and county-by-year fixed effects. Households are attentive to wildfire risk: they update their beliefs in response to the new zone, but not to the salient boundary itself. Despite this price decline for reclassified homes, mortgage loan-to-value ratios show no corresponding adjustment, implying that lenders' effective collateral exposure grew silently as affected home values fell.

I explore several mechanisms that could explain the coexistence of null LATEs and ATT in the immediate vicinity of the new risk zone boundary and significantly negative ATT on home prices further away. Contrasting these price estimates with high-resolution, objective wildfire risk measures that were not available to households at the time, I find that the actual wildfire risk is continuous in the immediate vicinity of the new risk zone boundary and increases further away, mirroring the price results. These findings suggest that households use the new risk zone to update their risk beliefs and, thus, make their asset pricing decisions consistent with a rational decision-making model.[2]

These findings contribute to the literature on belief formation and asset prices (Brunnermeier, Gollier, and Parker (2007), Brunnermeier, Farhi, Koijen, Krishnamurthy, Ludvigson, Lustig, Nagel, and Piazzesi (2021)) and to the literature on salience (Chetty, Looney, and Kroft (2009), Bordalo, Gennaioli, and Shleifer (2012), Bordalo, Gennaioli, and Shleifer

[2]Insurance premiums, mortgage market conditions, and neighborhood sorting show no significant response to the new risk zone. Sales transaction volume in treated areas shows a marginally significant decline of approximately 13–16% ($p < 0.10$) after the re-zoning, with flat pre-trends (Section E of the Supplementary Material). These results are consistent with Keys and Mulder (2020), who find that climate-risk disclosure reduces transaction volume but leaves mortgage-market conditions, including loan-to-value ratios, unaffected.

(2013), Gabaix (2014)). In the context of a large and risky asset, I find that households do not respond to the increased salience of risk generated by a policy intervention. This contrasts with documented responses to the salience of prices (e.g., Finkelstein (2009)) and of risk (Célérier and Vallée (2017), Frydman and Wang (2020)) involving smaller purchases, and speaks to a prior literature on risk salience and real estate prices.

My comparison of LATE and ATT in the immediate vicinity of a boundary contrasts with the quasi-experimental literature that examines how climate-related risk salience affects residential real estate prices using ATTs. For instance, Giglio, Maggiori, Rao, Stroebel, and Weber (2021b) examine the effect of future sea-level-rise risk salience on home prices using variation in a "climate attention index" accounting for the frequency of climate-related language in for-sale listings. This setting does not lend itself to a regression discontinuity design because there is no natural boundary where the salience of risk might be discontinuous, in contrast to my setting with a spatial risk salience discontinuity at the new risk zone boundary. A related strand of literature examines the effect of future sea-level-rise risk on home prices, e.g., Bernstein, Gustafson, and Lewis (2019); Baldauf, Garlappi, and Yannelis (2020); Murfin and Spiegel (2020). In contrast to these studies exploiting variation in risk, my paper focuses on variation in the salience of risk at the new risk zone boundary, holding risk constant. More recently, Letta and Nocito (2025) document that a sharp rise in the salience of volcanic risk near Naples, absent new physical damage, moved local real estate prices, illustrating that salience can be quantitatively important for infrequently-traded assets when the signal is sufficiently large. The literature on climate risk capitalization in housing has grown rapidly; see Contat, Hopkins, Mejia, and Suandi (2024) for a recent survey and Giglio, Kelly, and Stroebel (2021a) for a broader review of climate finance. Other studies use changes in risk maps to estimate ATTs, such as flood zone maps (e.g., Gibson and Mullins (2020), Hino and Burke (2021))[3] or wildfire risk ratings (Donovan, Champ, and Butry (2007)).[4] Closely related to this paper, Ma, Walls, Wibbenmeyer, and Lennon (2024) use a boundary discontinuity design around California wildfire risk zones and find a 4.3% price discount at boundaries. Because their design is a pooled cross-section of sales from 2015–2022, their estimate reflects the value differential between homes that have always been inside the zone and homes that have always been outside it, rather than an effect of being newly reclassified.[5]

The paper relates to the literature examining how risk salience from a recent natural disaster affects decision-makers' choices. For example, recent natural disasters have been shown to affect real estate prices in high-risk areas not directly hit by the event, e.g., for flooding and hurricanes (e.g., Hallstrom and Smith (2005), Addoum, Eichholtz, Steiner, and Yönder (2024), Muller and Hopkins (2019))[6] and wildfires (e.g., McCoy and Walsh (2018)). Dessaint and Matray (2017) find that corporate managers increase their cash holdings after a recent hurricane. Gallagher (2014) shows that households buy temporarily more flood insurance after a flood event is reported in their television media market. Choi, Diffenbaugh, and Burke (2024) similarly find that flood exposure raises insurance adoption, but effects decay within five years. Although a recent natural disaster does not change the risk per se, it may affect decision-makers' risk beliefs,

[3]Flood-map updates differ from my setting in two ways. First, location inside a flood zone triggers a mandatory (if imperfectly enforced) flood insurance requirement, and insurance prices change sharply at the floodplain boundary even where the underlying risk does not. Second, the California risk zone was implemented state-wide at once, whereas flood-map revisions are often staggered and revised after community input.

[4]In contrast to my setting, a home's risk rating is endogenous in Donovan et al. (2007) because households can take risk-mitigating actions (e.g., improve construction materials, defensible space) to affect their rating. This distinction has implications for what is estimated in each paper. See Section VI for further discussion.

[5]One interpretation consistent with both findings is a temporal evolution of wildfire risk salience: Ma et al. (2024)'s sample covers a period following several major California wildfires that raised public awareness of zone status and prompted market participants to filter comparables more explicitly by zone assignment, which would attenuate the boundary-smoothing mechanism that appears operative when the new zone was first implemented in 2007.

[6]In contrast, Baldauf et al. (2020) find that recent hurricanes in nearby, unaffected states have a negligible effect on the prices of homes predicted to be affected by future sea level rise.

for example due to recency effects (Agarwal, Driscoll, Gabaix, and Laibson (2008), Malmendier (2021)).

Section II provides background on the risk zones. Section III describes the data, Section IV the empirical strategy, and Section V the results. Section VI discusses mechanisms, and Section VII concludes.

II. Background on the Fire Risk Zones and Wildfire Risk

A. Policy Motivation

California adopted its first wildfire risk zones in 1985. In November 2007, the state adopted new zones under Title 14 to incorporate improved fire science, data, and mapping techniques (CAL FIRE (2022)). The 2007 zones have two objectives: to inform residents about fire hazard in their area and to impose fire-safe building standards on new construction. (A later map revision, released in November 2022, is beyond the scope of this study.)

B. Risk Zone Design

The wildfire risk zones, referred to by the State of California as Fire Hazard Severity Zones or FHSZs, are produced by the California Department of Forestry and Fire Protection (CAL FIRE); see CAL FIRE (2022) for details. CAL FIRE's model predicts potential fire behavior and the probability of an area burning, using "the best available science and data at the time" (CAL FIRE (2022)). Each area receives a hazard score based on continuous inputs: existing and potential fuel (natural vegetation), terrain characteristics, and typical fire weather. There are three zone levels: moderate, high, and very high.

C. Risk Zone Implementation

The 2007 risk zone revises and, in most places, marginally expands the pre-existing 1985 risk zone, with small contractions in some areas. Typically, the 2007 change in the zone boundary does not exceed a couple of kilometers, as illustrated in Sections A.2 and A.3 of the Supplementary Material. Importantly, because the zone represents overall hazard rates, homeowners cannot take private risk mitigating actions to affect their zone assignment. Sellers are mandated by law to disclose to prospective buyers at the time of sale whether their home lies inside a risk zone and, if so, which one. In practice, after January 1, 2008, being located inside a risk zone (no matter the level) comes with permit requirements for new buildings, or extensions/remodeling of the exterior made to existing buildings, that must comply with California building code Chapter 7a. The implication is that roofs, exterior walls, and decks should use ignition-resistant materials.

The zones were designed not to raise insurance premiums. California regulation prohibits insurers from pricing on risk-zone status, and insurance companies use proprietary, fine-resolution probabilistic risk models that supersede the coarse hazard classes in the zones.[7] Nevertheless, I test empirically in Section E of the Supplementary Material and fail to find an

[7]Based on personal communications with bankers and mortgage brokers conducting business both inside and outside the risk zones and underwriting mortgages to a large number of lenders in my study area, they report that the new risk zone does not influence lending practices. See Section A.1 of the Supplementary Material for additional information.

effect of the new risk zone on insurance premium quotes or on mortgages as measured by rate spread, loan-to-value, and mortgage terms.

D. Wildfire Risk Measures

The first publicly available fine-resolution wildfire risk measures for California were published in 2020 by the U.S. Forest Service (Scott, Gilbertson-Day, Moran, Dillon, Short, and Vogler (2020)). They reflect the wildfire risk that prevails circa 2014 for every 30m$\times$30m pixel in the landscape. I assume that this wildfire risk approximates the overall wildfire risk in my 2000–2015 time period. The two risk measures of interest are (1) the wildfire or burn probability (BP), which represents the annual probability of a wildfire burning in a specific location, and (2) the fire risk to homes or risk to potential structures (RPS), which for every pixel provides the probability that a home would burn if a home existed on that pixel. I convert these annual risk measures into the probability of at least one risky event occurring in the next 30 years to better match the time-horizon associated with real estate purchases.

III. Data

The datasets include home sales transactions, the wildfire risk zone, high-resolution fire risk measures, and neighborhood characteristics. This section describes the main variables, while Section A of the Supplementary Material contains details of the data cleaning, additional data sources to test for sorting and alternative mechanisms, the summary statistics, and illustrations of the spatial distribution of the homes, risk zones, and fire risk measures.

A. Home Sales Transactions

The home sales transactions were obtained from CoreLogic and cover all single family sales (excluding mobile homes) between January 2000 and December 2015 in ZIP codes most likely to be affected by wildfire risk across seven counties in the Los Angeles and San Diego Basins, namely, Santa Barbara, Los Angeles, Orange, Ventura, Riverside, San Bernardino, and San Diego. The key variables are the sale price and sale date, address, and home characteristics such as size, age, number of bedrooms and bathrooms, and presence of a swimming pool. Every home is geocoded and matched to a series of spatial layers, including wildfire perimeters, topography, and fire risk measures. Homes are linked to the American Community Survey (ACS) to obtain demographics at the census tract level.

Starting with a dataset of 2,187,007 unique properties, I focus on owner-occupied properties and arm's-length transactions. All prices are deflated using the Consumer Price Index from the U.S. Bureau of Labor Statistics. In addition, to focus on the sole effect of the new risk zone absent any realization of wildfire risk in the vicinity, I drop a small number of homes located within 10km of a wildfire perimeter that burned within three years prior to a sale.

B. Location Inside the Fire Risk Zone

The geocoded homes are mapped into both the old (pre-2007) and new (post-2007) fire risk zones obtained from CAL FIRE. For each risk zone, there exist three levels: moderate, high, and very high. For each home inside the new risk zone, I calculate its distance to the new risk zone boundaries. For example, if a home is located inside the high risk zone, I calculate the distance to the moderate and to the very high risk zone boundaries. For each home outside the new risk zone, I calculate its distance to the moderate risk zone boundary. The moderate and high risk zones are relatively thin, while the very high risk zone is vast as it extends into the wilderness. In my sample, the median (max) width for the new moderate zone is 1.1km (5.8km), and for the new high risk zone it is 540m (4.1km).

C. Definition of the Treatment and Control Homes

Throughout the paper, I define as control homes those that do not change risk zone level in 2007, while I define as treatment homes those that experience a one-level increase in their risk zone.[8] For example, a treatment home is located inside the moderate risk zone prior to 2007 and inside the high risk zone after 2007. I observe its sale either before or after 2007, or in some cases both before and after 2007. In addition, to improve the comparability of treated and control homes, I restrict observations (1) to come from census tracts that feature both control and treatment homes (this restriction is relaxed in Section C of the Supplementary Material and results are quantitatively similar), and (2) to be located within 1km of the new risk zone boundary.

Few treatment homes move from outside the zone into the moderate zone, so I drop this treatment from the main analysis. I pool the two remaining one-level increases (moderate to high, and high to very high) into a generic one-level increase. Sections A.3 and D of the Supplementary Material report observation counts and DiD estimates for each individual increase. Results are similar.

[8]There exist places where the risk zone "jumps" risk levels, for example from moderate to very high. Throughout the paper, I focus on the sales of homes that experience a one-level increase in risk zone and exclude homes that jump risk levels. Put differently, I restrict the sample to homes that are located inside a risk zone that is one level higher or lower than the adjacent risk zone, i.e., treated and control homes differ in their risk level by exactly one level.

IV. Empirical Strategy

This section lays out the three empirical models. The first two identify the effect of risk salience at the new zone boundary on home prices; the third identifies the effect of the risk information contained in the new zone.

A. Regression Discontinuity Design

I implement a spatial regression discontinuity (RD) design to estimate the local average treatment effect (LATE) of being located inside a riskier wildfire zone on home prices. The spatial RD model is shown in equation (1):

$$\ln p_i = \beta_0 + g_0(x_i) + Z(\tau + g_1(x_i)) + \varepsilon_i, \tag{1}$$

where the dependent variable is the natural log of home i's sale price. The running variable x_i measures the distance of home i to the new risk zone boundary, with cutoff $\bar{x} = 0$. The dummy variable Z takes value 1 for treated homes ($x_i \geq 0$) and 0 for control homes ($x_i < 0$). Treated homes are always located inside a zone that is one level higher after the 2007 policy change than the control homes, which do not change risk zone level. $g_0(x_i)$ and $g_1(x_i)$ are two continuous polynomials of order p. The coefficient τ is the treatment effect, i.e., the change in $\ln p_i$ at the cutoff.

The bandwidth is selected by the MSE-optimal, data-driven procedure of Calonico, Cattaneo, and Titiunik (2014a,b). All specifications use a triangular kernel; the main specification uses a local linear polynomial (Gelman and Imbens (2019)). Following Cattaneo, Titiunik, and

Vazquez-Bare (2020b), throughout the paper I report conventional point estimates, robust bias-corrected standard errors and confidence intervals, and robust p-values. Standard errors are computed using the heteroskedasticity-robust nearest neighbor variance estimator with a minimum of three neighbors.[9] I estimate equation (1) for homes that sell after the new risk zone implementation, both without and with covariates, as including covariates can improve estimation efficiency. Covariates for the RD include: home age, size, number of bedrooms, presence of a swimming pool, distances to the nearest primary road and to the nearest national forest, and location inside the wildland-urban interface.

The strength of the RD design is to compare homes that are located as close as possible to the new risk zone boundary to ensure the treated and control homes are as similar as possible and only differ with respect to their risk zone assignment. The design requires that home characteristics, observed and unobserved, do not change discontinuously at the cutoff. Section B.1 of the Supplementary Material shows that covariates are continuous around the cutoff: re-estimating equation (1) with each covariate as the dependent variable yields no discontinuity.

The RD design also requires that homeowners cannot manipulate their assignment at the cutoff. The risk zones are designed to reflect hazard rates over broad areas such that a particular homeowner cannot take risk mitigating actions, e.g., clearing vegetation around the home, to affect their assignment to the risk zone. Still, I formally test for manipulation of risk zone assignment at the cutoff in Section B.1 of the Supplementary Material. Using the testing

[9] Results are quantitatively similar when clustering standard errors at the census tract level. However, because the boundary crosses *through* census tracts, each tract contains both treated and control homes by design; clustering at that level would group units on opposite sides of the treatment boundary into the same cluster and does not capture the relevant spatial correlation structure. I therefore report heteroskedasticity-robust standard errors throughout.

procedure from Cattaneo, Jansson, and Ma (2020a), I fail to reject the null hypothesis of continuity of the density of the running variable ($p > 0.1$).

Figure 1 panels (a) and (c) depict the local log price means of the treated (right) and control (left) homes, while panels (b) and (d) show the residuals of log prices on covariates in the pre-treatment period (panels a and b) and post-treatment period (panels c and d). Fitting a polynomial of order 4 does not indicate evidence of a discontinuity at the new risk zone boundary in either the pre- or post-treatment period. Panels (e) and (f) illustrate the fire risk to homes and wildfire probability over 30 years, respectively. The risk increases smoothly from left (control) to right (treated homes).

[Insert Figure 1 approximately here]

Because one cannot test formally that no unobservables are changing discontinuously at the new risk zone boundary, this motivates the use of a second design leveraging homes that sold before and after the new 2007 risk zone implementation.[10]

B. Difference-in-Discontinuities Design

The difference-in-discontinuities (RD-DiD) model estimates a continuous polynomial on each side of the cutoff ($\bar{x} = 0$) both before and after the policy change, as depicted in equation (2):

$$\begin{aligned} \ln p_{it} &= \beta_0 + g_0(x_{it}) + Z(\beta_1 + g_1(x_{it})) \\ &\quad + Post_t\left[\delta_0 + g_2(x_{it}) + Z(\tau + g_3(x_{it}))\right] + \varepsilon_{it}, \end{aligned} \tag{2}$$

[10]This design follows that presented in Grembi, Nannicini, and Troiano (2016). In their context, one of the observables changes discontinuously at the cutoff.

where t denotes time and $Post_t$ is a dummy taking value 1 in the period following the 2007 change in risk zoning. The treatment effect, τ, captures the change in $\ln p_{it}$ at the cutoff comparing before and after the policy change. In the main specification, the polynomial $g_j(x_{it})$, where $j \in [0, 3]$, is of degree one.

C. Difference-in-Differences Design

I estimate an average treatment effect on the treated (ATT) comparing the prices of homes that experience an increase in their risk zone level with those of homes that remain in the same risk zone. The DiD model is depicted in equation (3):

$$\ln p_{it} = \beta_0 + \beta_{DiD} Treat_i \times Post_t + \beta_1 Post_t + \mathbf{X}_i' \boldsymbol{\beta}_X + \lambda_i + \mu_{gt} + \varepsilon_{it}, \tag{3}$$

where the dependent variable is the natural log of home i's sale price at time t. The treatment group, $Treat_i$, is defined as all homes that are located inside a zone that is one level higher than that prior to the 2007 risk zone implementation. The control group is defined as homes that remain inside the same risk zone after the new 2007 risk zone implementation. $Post_t$ is a dummy variable taking value 1 following the 2007 change in risk zoning. The coefficient β_{DiD} is the ATT. In the simplest specification, $\mathbf{X}_i$ denotes home and neighborhood covariates, while the preferred specification features home fixed effects λ_i to most flexibly control for time-invariant unobservables, such as home and neighborhood quality. Covariates in the DiD include home age and age squared, home size and home size squared, number of bedrooms, number of baths, swimming pool dummy, elevation, slope, dummy for location inside the wildland-urban interface, distances to nearest primary road, to nearest national forest, and to nearest state or local park. To control for temporal shocks, e.g., macro-level housing shocks, and possible heterogeneity in temporal shocks across the region, μ_{gt} denotes year and quarter fixed effects or, more flexibly, county-by-year and quarter fixed effects.

To examine whether the parallel trends assumption holds, I estimate an event-study specification of equation (3) that interacts treatment with calendar-time indicators, taking the map release (2007Q4, or the year 2007 at annual frequency) as the omitted reference period. Figure 2, panel (a), plots the quarterly $Treat_i \times Quarter_t$ coefficients: the pre-period coefficients are small and jointly insignificant (parallel-trends F-test $p = 0.46$), and treated home prices decline steadily after the policy change. Because the quarterly window is short relative to the housing cycle, panel (b) extends the event study at annual frequency over the full 2000–2015 sample window, yielding seven pre-period estimates. The treated–control price gap is flat throughout the 2000–2006 California housing boom: the annual pre-period coefficients alternate in sign, none is individually significant at the 5% level, and they are jointly insignificant at the 5% level (F-test $p = 0.08$).[11] Treated homes thus show no differential exposure to the run-up that preceded the policy. Furthermore, the post-period decline is not confined to the 2009–2010 housing market trough: it persists through the 2012–2015 recovery, which is inconsistent with the decline merely reflecting a greater cyclical sensitivity of treated homes. Supplementary Material Figure B4 shows that annual transaction counts for treated and control homes track each other closely over the full window. Supplementary Material Figure B3 provides additional visual evidence: panel (a) shows the all-sales sample with covariates interacted with a linear year trend (F-test $p = 0.36$) and panel (b) the repeat-sales sample with home fixed effects (F-test $p = 0.12$).

[11] To ensure this marginal p-value is not driven by early-sample noise, a joint F-test restricting the window closer to the treatment year (2001–2006) yields $p = 0.12$, confirming the absence of pre-treatment trends.

[Insert Figure 2 approximately here]

V. Results

A. Effect in the Immediate Vicinity of the New Risk Zone Boundary

1. Regression Discontinuity and Difference-in-Discontinuities Designs

Table 1 contains the RD and RD-DiD estimation results. Using a spatial RD design, the estimated effect of a one-level increase in risk zone on home prices is -6.3% without covariates and -4.0% with covariates (columns 1 and 2, respectively), but not significant at the 10% level. Following Calonico, Cattaneo, and Titiunik (2015), the MSE-optimal bandwidth is 209m and 216m, without and with covariates, respectively. Running a placebo test for the homes that changed risk zone in 2007 but sold prior to 2007 (denoted hereafter 'RD–pre') also yields insignificant estimates: 0.3% without covariates and 1.0% with covariates (columns 3 and 4, respectively). The MSE-optimal bandwidth in this case is 172m and 229m, without and with covariates, respectively.

[Insert Table 1 approximately here]

The RD-DiD design accounts for any pre-existing differences across homes on different sides of the new risk zone boundary. Table 1 indicates RD-DiD estimates of 1.0% without covariates and -1.5% with covariates (columns 5 and 6, respectively). The MSE-optimal bandwidth in this case is 156m and 199m, without and with covariates, respectively.

The RD and RD-DiD estimates are robust to bandwidths from 100m to 800m (Figure 3, panels (a)–(c)).[12] Section C of the Supplementary Material reports further checks: the estimates are robust to the polynomial order, to alternative sample selections (including relaxing the census-tract restriction), to alternative boundary cutoffs, and to estimating periods of housing booms and busts separately. The RD estimates are insignificant in every individual year (Supplementary Material Table C1).

[Insert Figure 3 approximately here]

2. Difference-in-Differences Design

Panel B, columns 1 and 2, in Table 2 show the DiD estimates of the effect of being located inside a one-level riskier zone relative to neighboring control homes that do not change risk zone in the immediate vicinity of the new risk zone boundary. The ATT estimated within 200m of the new risk zone boundary is -1.1% with census-tract, year, and quarter fixed effects plus covariate$\times$year trends (insignificant; column 1) and essentially zero with home and county-by-year fixed effects (insignificant; column 2). These null estimates are consistent with the LATE estimates in Table 1, obtained for a similar bandwidth.

[Insert Table 2 approximately here]

B. Effect over a Broader Area from the New Risk Zone Boundary

Table 2 shows the main DiD estimates of the effect of being located inside a one-level riskier zone relative to neighboring control homes that do not change risk zone over a broader area (Panel A) and further away (Panel B, columns 3 and 4) from the new risk zone boundary.

[12]The RD estimate in the post-period is negative throughout the bandwidth range and approaches but does not cross conventional significance thresholds at wider bandwidths. Wider bandwidths include more homes farther from the boundary where standard errors narrow, yielding a more precise, yet still insignificant, estimate of the price penalty.

In Panel A, all homes are located within 600m of the new 2007 risk zone boundary. Columns 1 through 4 represent specifications with increasingly more flexible fixed effects, with estimates ranging from -5.2% to -1.5%. Column 1 features census-tract, year, and quarter fixed effects with home covariates. Column 2 further interacts each covariate with a linear year trend. Columns 3 and 4 feature home fixed effects. The preferred specification (column 4), with home and county-by-year fixed effects to control for regional shocks, yields a price penalty of 1.5%. Panel (d) of Figure 3 plots this estimate across bandwidths from 100m to 800m: it is close to zero at the narrowest bandwidths, which cover the immediate vicinity of the boundary, and settles between 1.2% and 1.5% once the bandwidth exceeds 400m. In Section C of the Supplementary Material, I also show that estimates are robust to dropping every year in the 2000–2015 time period, one year at a time, or to estimating the model separately during periods of housing booms and busts.

To further understand how the ATT estimates vary spatially, I define spatial rings further away from the new risk zone boundary. The ATT estimates for the 200m to 600m ring indicate price penalties of 6.6% and 5.7% ($p < 0.01$; Panel B columns 3 and 4). Estimates are robust to different choices of ring width and positioning of these rings, from relatively close to relatively far away from the new risk zone boundary (200m ring in Panel e and 400m ring in Panel f, Figure 3).[13]

The DiD estimates are robust to a battery of additional checks reported in Section B.2 of the Supplementary Material. First, the Treat$\times$Post coefficient is negative in every window from

[13]Note that the positioning of the 200m-ring right along the new risk zone boundary in Panel (e) corresponds to the immediate vicinity of the new risk zone boundary and thus leads to an insignificant ATT estimate. Furthermore, because moving rings contain fewer sales, in particular further away from the new risk zone boundary (see Supplementary Material Figure A4), Panels (e) and (f) do not extend further than 600m and 800m, respectively.

± 2 to ± 8 quarters around the policy date; it becomes significant at the ± 8-quarter window as the sample grows (Supplementary Material Table B2). Second, a 1:1 nearest-neighbor matching-DiD estimator, which matches treated and control homes on pre-period observables, yields ATTs consistent in sign and similar in magnitude to Table 2, Panel A (Supplementary Material Table B3). Third, replacing the standard control group with type-a controls, homes already carrying the higher hazard rating before 2007, produces a similar spatial gradient as Panel B of Table 2 (Supplementary Material Table B5). Using home, county-by-year, and quarter fixed effects, the 200–600m ring carries a price penalty of -7.0% ($p < 0.01$), while the 0–200m estimate is economically small and statistically insignificant (-0.8%, $p > 0.10$).

VI. Results Interpretation and Potential Mechanisms

This section interprets the joint pattern (null effects at the boundary, a significant price penalty further away) and explores mechanisms. If households respond to the risk information contained in the new zone rather than to the salient label, prices should track the underlying wildfire risk. I therefore investigate how wildfire risk varies over space. I leverage high-resolution (30m$\times$30m) wildfire risk measures (not available to households during my sample period), which I map into each home. Then, I run an RD and OLS design with each risk measure as the dependent variable. Table 3 depicts the effect of being located inside a one-level higher risk zone on the fire risk to homes and the wildfire probability over 30 years. The RD design, Panel A, yields null LATEs on fire risk to homes and wildfire probability over 30 years (p >0.1). In contrast, the OLS approach, Panel B, indicates a 1.8% increase in fire risk to homes over 30 years (p <0.01; columns 1 and 2) and a 2.7% increase in wildfire probability over 30 years for homes located within 600m of the new 2007 risk zone boundary (p <0.01; columns 3 and 4). These findings suggest that the actual wildfire risk is continuous at the new risk zone boundary and increases further away.

[Insert Table 3 approximately here]

The zone revision could in principle affect prices through both a salience channel and a true-risk channel, and these are not perfectly separable. The null LATE at the boundary, where the discrete zone rating flips but the underlying hazard is essentially continuous (Table 3, Panel A), provides the cleanest evidence on the salience channel. The significant DiD further from the boundary, where average wildfire hazard is genuinely higher inside the zone (Panel B), is informative about the risk-information channel. I acknowledge that this separation is not definitive: the zone boundary reflects CAL FIRE's hazard model and may encode information that buyers treat as a relevant continuous signal rather than a purely categorical assignment. The joint RD-and-DiD design provides the best available means to distinguish the two channels given this setting.

In Section E of the Supplementary Material, I explore alternative mechanisms that could influence home prices. Insurance premiums and mortgage market conditions (loan-to-value, mortgage terms, rate spread) show no significant response to the new risk zone. Sales transaction volume in treated areas shows a marginally significant decline of approximately 13–16% ($p < 0.10$) after the re-zoning (Supplementary Material Table E3, columns 1–2), with flat pre-trends confirmed by an event study (Supplementary Material Figure E1). This volume decline is consistent with findings in Keys and Mulder (2020) that transaction volumes can respond to climate risk, although this effect is weak statistically in my setting.

How does the continuous underlying wildfire risk signal find its way into prices despite the discrete zone assignment? Two non-exclusive market channels plausibly explain this. First, the CAL FIRE maps were publicly available. A buyer considering a home just inside the newly-elevated zone could see at a glance that an immediately adjacent home just outside faced nearly identical environmental conditions: the zone boundary is drawn on continuously-varying terrain features (fuel load, slope, wind exposure) that do not change discretely at the line. Second, comparable-based valuation may reinforce this mechanism. Appraisers and buyers draw on recent sales from the surrounding neighborhood, which likely spans homes on both sides of a nearby boundary: a home 50m inside the zone may be routinely compared with homes 50m outside. If the underlying risk is continuous at the line, this aggregation over the set of comparables would smooth away any pure-risk zone penalty at the boundary. I cannot fully isolate these channels or rule out others, and I state the external-validity claim accordingly.

The comparison with Dessaint and Matray (2017) is instructive about what these results do—and do not—imply. Dessaint and Matray (2017) study salience from the *realization* of a nearby hurricane: corporate managers increase cash holdings after a vivid, emotionally salient event, a response operating through recency and affect. The salience mechanism here is different in kind: it arises from a policy change (a re-drawn risk map) rather than a realized disaster, and appeals to an assessment of risk probabilities rather than recency or affect. It is therefore consistent that managers over-react to a salient hurricane *and* that homebuyers may not necessarily over-react to a new map boundary. More broadly, the findings do not imply that salience is unimportant for large assets; rather, they suggest that for infrequently-traded assets and non-event-driven information releases, prices can reflect the continuous underlying risk signal.

The comparison with Donovan et al. (2007) raises a related question of interpretation. In their setting, a home's risk rating is partially determined by household-level mitigation actions (construction materials, defensible space), making the rating *endogenous* to households. Homes that find it cheapest to avoid a high rating do so; the remaining high-rated homes are a selected group. The price response on those homes therefore conflates the information effect with selection of non-mitigating households. In my setting, the 2007 zone represents fire *hazard* at the area level (reflecting fuel load, terrain, and wind patterns) and cannot be changed by household actions. The zone boundary is thus an externally-imposed information signal, and the ATT estimates capture the pure information effect without the selection confound.

VII. Conclusions

This paper asks whether the salience of climate-related risk affects households' large asset pricing decisions. I examine how a new wildfire risk zoning policy in California, whose mandated disclosure makes risk salient at the new zone boundary, affects residential real estate prices. Regression discontinuity, difference-in-discontinuities, and difference-in-differences designs in the immediate vicinity of the boundary yield no evidence of an effect on home prices. A difference-in-differences design including homes further away yields a -5.7% effect. I interpret these estimates as showing that households do not respond to the increased salience of risk, but update their beliefs in response to the risk information contained in the new zone.

Three implications for asset pricing follow from these findings. First, in this market, newly-disclosed risk information is reflected in house prices via the continuous underlying risk signal rather than via the salient discrete classification: the zone level itself does not generate a price penalty at the boundary, but the probability of burning, which is continuous across the boundary and higher on average inside the zone, is priced further inside the zone. Second, the contrast between the null RD result and the significant DiD result would be missed by a design that estimated only an average treatment effect: only the joint use of a discontinuity design and a broader-area design distinguishes a salience effect from an information effect, which has implications for how future studies of climate-risk disclosure are designed. Third, insurance premium quotes did not respond to the re-zoning: California regulation prohibited insurers from pricing on zone status in this period. The price decline therefore does not reflect higher insurance costs, consistent with the risk-information channel. Finally, the price decline came with no corresponding adjustment in mortgage loan-to-value ratios (Supplementary Material Table E3), implying that lenders' effective collateral exposure grew as affected home values fell. This lender non-response echoes Keys and Mulder (2020), who show that mortgage lenders fail to adjust collateral valuations when federal programs absorb or mis-price climate-related risk at below-actuarial rates. These distinctions matter for understanding how climate risk is priced in housing markets and for calibrating the expected effects of disclosure regulation elsewhere.

References

Addoum, J. M.; P. Eichholtz; E. Steiner; and E. Yönder. “Climate Change and Commercial Real Estate: Evidence from Hurricane Sandy.” *Real Estate Economics*, 52 (2024), 687–713.

Agarwal, S.; J. C. Driscoll; X. Gabaix; and D. Laibson. “Learning in the credit card market.” National Bureau of Economic Research (2008).

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

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

Bordalo, P.; N. Gennaioli; and A. Shleifer. “Salience theory of choice under risk.” *Quarterly Journal of Economics*, 127 (2012), 1243–1285.

Bordalo, P.; N. Gennaioli; and A. Shleifer. “Salience and consumer choice.” *Journal of Political Economy*, 121 (2013), 803–843.

Brunnermeier, M.; E. Farhi; R. S. Koijen; A. Krishnamurthy; S. C. Ludvigson; H. Lustig; S. Nagel; and M. Piazzesi. “Perspectives on the future of asset pricing.” *The Review of Financial Studies*, 34 (2021), 2126–2160.

Brunnermeier, M. K.; C. Gollier; and J. A. Parker. “Optimal beliefs, asset prices, and the preference for skewed returns.” *American Economic Review*, 97 (2007), 159–165.

CAL FIRE. “California Department of Forestry and Fire Protection. Fire Hazard Severity Zones.” https://osfm.fire.ca.gov/divisions/

community-wildfire-preparedness-and-mitigation/wildfire-prevention-engineering/fire-hazard-severity-zones/ (2022). Last visited February 3, 2022.

Calonico, S.; M. D. Cattaneo; and R. Titiunik. "Robust data-driven inference in the regression-discontinuity design." *Stata Journal*, 14 (2014a), 909–46.

Calonico, S.; M. D. Cattaneo; and R. Titiunik. "Robust non-parametric confidence intervals for regression-discontinuity designs." *Econometrica*, 82 (2014b), 2295–2326.

Calonico, S.; M. D. Cattaneo; and R. Titiunik. "Optimal data-driven regression discontinuity plots." *Journal of the American Statistical Association*, 110 (2015), 1753–1769.

Cattaneo, M. D.; M. Jansson; and X. Ma. "Simple local polynomial density estimators." *Journal of the American Statistical Association*, 115 (2020a), 1449–1455.

Cattaneo, M. D.; R. Titiunik; and G. Vazquez-Bare. In *Handbook of Research Methods in Political Science and International Relations*. "The Regression Discontinuity Design.", L. Curini, ed. Sage Publications (2020b), chapter 44, 835–857.

Célérier, C., and B. Vallée. "Catering to investors through security design: Headline rate and complexity." *The Quarterly Journal of Economics*, 132 (2017), 1469–1508.

Chetty, R.; A. Looney; and K. Kroft. "Salience and taxation: Theory and evidence." *American Economic Review*, 99 (2009), 1145–77.

Choi, J.; N. S. Diffenbaugh; and M. Burke. "The Effect of Flood Exposure on Insurance Adoption among US Households." *Earth's Future*, 12 (2024), e2023EF004110.

Contat, J.; C. Hopkins; L. Mejia; and M. Suandi. "When Climate Meets Real Estate: A Survey of the Literature." *Real Estate Economics*, 52 (2024), 618–659.

Dessaint, O., and A. Matray. "Do managers overreact to salient risks? Evidence from hurricane strikes." *Journal of Financial Economics*, 126 (2017), 97–121.

Donovan, G. H.; P. A. Champ; and D. T. Butry. "Wildfire risk and housing prices: A case study from Colorado Springs." *Land Economics*, 83 (2007), 217–233.

Finkelstein, A. "EZ-tax: Tax salience and tax rates." *Quarterly Journal of Economics*, 124 (2009), 969–1010.

Frydman, C., and B. Wang. "The impact of salience on investor behavior: Evidence from a natural experiment." *The Journal of Finance*, 75 (2020), 229–276.

Gabaix, X. "A sparsity-based model of bounded rationality." *Quarterly Journal of Economics*, 129 (2014), 1661–1710.

Gallagher, J. "Learning about an infrequent event: Evidence from flood insurance take-up in the United States." *American Economic Journal: Applied Economics*, 6 (2014), 206–233.

Gelman, A., and G. Imbens. "Why high-order polynomials should not be used in regression discontinuity designs." *Journal of Business and Economic Statistics*, 37 (2019), 447–456.

Gibson, M., and J. T. Mullins. "Climate risk and beliefs in New York floodplains." *Journal of the Association of Environmental and Resource Economists*, 7 (2020), 1069–1111.

Giglio, S.; B. Kelly; and J. Stroebel. "Climate Finance." *Annual Review of Financial Economics*, 13 (2021a), 15–36. doi:10.1146/annurev-financial-102620-103311.

Giglio, S.; M. Maggiori; K. Rao; J. Stroebel; and A. Weber. “Climate change and long-run discount rates: Evidence from real estate.” *The Review of Financial Studies*, 34 (2021b), 3527–3571.

Grembi, V.; T. Nannicini; and U. Troiano. “Do fiscal rules matter?” *American Economic Journal: Applied Economics*, 1–30.

Hallstrom, D. G., and V. K. Smith. “Market responses to hurricanes.” *Journal of Environmental Economics and Management*, 50 (2005), 541–561.

Hino, M., and M. Burke. “The effect of information about climate risk on property values.” *Proceedings of the National Academy of Sciences*, 118.

International Association of Wildland Fire. “WUI Fact Sheet.” https://www.frames.gov/catalog/53415 (2013). Last visited January 10, 2018.

Kahn, M. E. “The death toll from natural disasters: The role of income, geography, and institutions.” *Review of Economics and Statistics*, 87 (2005), 271–284.

Keys, B. J., and P. Mulder. “Neglected No More: Housing Markets, Mortgage Lending, and Sea Level Rise.” National Bureau of Economic Research (2020). doi:10.3386/w27930.

Letta, M., and S. Nocito. “Environmental Risk Salience and Real Estate Markets.” (2025). SSRN Working Paper 5436279.

Ma, L.; M. Walls; M. Wibbenmeyer; and C. Lennon. “Risk Disclosure and Home Prices: Evidence from California Wildfire Hazard Zones.” *Land Economics*, 100 (2024), 6–21.

Malmendier, U. “Experience effects in finance: Foundations, applications, and future directions.” *Review of Finance*, 25 (2021), 1339–1363.

McCoy, S. J., and R. P. Walsh. “Wildfire risk, salience, and housing demand.” *Journal of Environmental Economics and Management*, 91 (2018), 203–228.

Muller, N. Z., and C. A. Hopkins. “Hurricane Katrina floods New Jersey: The role of information in the market response to flood risk. NBER Working Paper No. 25984.” (2019).

Murfin, J., and M. Spiegel. “Is the Risk of Sea Level Rise Capitalized in Residential Real Estate?” *The Review of Financial Studies*, 33 (2020), 1217–1255.

Radeloff, V. C.; D. P. Helmers; H. A. Kramer; M. H. Mockrin; P. M. Alexandre; A. Bar-Massada; V. Butsic; T. J. Hawbaker; S. Martinuzzi; A. D. Syphard; and S. I. Stewart. “Rapid growth of the US wildland-urban interface raises wildfire risk.” *Proceedings of the National Academy of Sciences*, 115 (2018), 3314–3319.

Rappaport, J., and J. D. Sachs. “The United States as a coastal nation.” *Journal of Economic Growth*, 8 (2003), 5–46.

Scott, J. H.; J. W. Gilbertson-Day; C. Moran; G. K. Dillon; K. C. Short; and K. C. Vogler. “Wildfire Risk to Communities: Spatial datasets of landscape-wide wildfire risk components for the United States. Fort Collins, CO: Forest Service Research Data Archive. Updated 25 November 2020. https://doi.org/10.2737/RDS-2020-0016.” (2020).

Taylor, S. E., and S. C. Thompson. “Stalking the elusive “vividness” effect.” *Psychological Review*, 89 (1982), 155.

TABLE 1

RD and RD-DiD Estimates of Being Located Inside a One-Level Higher Risk Zone on Log of Home Prices

Table 1 reports RD and RD-DiD estimates of being located inside a one-level higher risk zone on the log of home prices. The spatial cutoff is the new 2007 risk zone boundary. Treated homes ($Z = 1$) are located inside a risk zone that is one level higher than control homes. $Post = 1$ denotes the period after the new risk zone was implemented (Nov. 2007). The MSE-optimal bandwidth is derived using a local linear polynomial (with covariates in columns 2, 4, and 6). Observations are restricted to census tracts featuring both control and treatment homes. Covariates: age, house size, # bedrooms, swimming pool dummy, slope, distances to national forest and to primary road, and dummy for location inside the wildland-urban interface. Columns 1–4: RD in the period after (Post) or prior to (Pre; placebo) the new risk zone implementation. Robust bias-corrected standard errors computed with the nearest neighbor variance estimator; Columns 5–6: Difference-in-discontinuities design leveraging both the post and pre-policy periods. Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

	RD (Post)		RD (Pre)		RD-DiD	
	1	2	3	4	5	6
$Z = 1$	-0.0631	-0.0401	0.0034	0.0099		
	(0.0410)	(0.0267)	(0.0405)	(0.0285)		
Z×Post					0.0101	-0.0153
					(0.0456)	(0.0305)
Observations	3,255	3,169	2,850	3,200	5,543	6,070
Covariates		✓		✓		✓
Bandwidth (m)	209	216	172	229	156	199

TABLE 2

DiD Estimates of a One-Level Increase in Risk Zone on Log of Home Prices

Table 2 reports DiD estimates of a one-level increase in risk zone on the log of home prices. Panel A: homes located within 600m from the new 2007 risk zone boundary. Panel B, columns 1–2: homes located within 200m, columns 3–4: homes located between 200m and 600m from the new 2007 risk zone boundary. Covariates×Year trend: each home covariate is interacted with a linear year trend. Covariates: age and age squared, home size and home size squared, # bedrooms, # baths, swimming pool dummy, elevation, slope, dummy for location inside the wildland-urban interface, distances to nearest primary road, national forest, state or local park. Observations are restricted to census tracts featuring both control and treatment homes. Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

	Panel A: 0–600m			
	1	2	3	4
Treat×Post	-0.0411***	-0.0227***	-0.0522***	-0.0152*
	(0.0080)	(0.0082)	(0.0089)	(0.0085)
Observations	8,978	8,978	6,542	6,542
Adj. R^2	0.85	0.86	0.79	0.82
Quarter FE	✓	✓	✓	✓
Year FE	✓	✓	✓	
County×Year FE				✓
Census tract FE	✓	✓		
Covariates	✓	✓		
Covariates×Year trend		✓		
Home FE			✓	✓
	Panel B: Finer distance bins (rings)			
	0–200m		200–600m	
	1	2	3	4
Treat×Post	-0.0106	0.0000	-0.0660***	-0.0565***
	(0.0097)	(0.0099)	(0.0155)	(0.0172)
Observations	6,074	4,438	2,899	2,104
Adj. R^2	0.86	0.82	0.85	0.84
Quarter FE	✓	✓	✓	✓
Year FE	✓		✓	
County×Year FE		✓		✓
Census tract FE	✓		✓	
Covariates	✓		✓	
Covariates×Year trend	✓		✓	
Home FE		✓		✓

TABLE 3

RD and OLS Estimates of Being Located Inside a One-Level Higher Risk Zone on Fire Risk to Homes and Wildfire Probability over 30 Years

Table 3 reports RD and OLS estimates of being located inside a one-level higher risk zone on the fire risk to homes and the wildfire probability over 30 years. The spatial cutoff is the new 2007 risk zone boundary. Treated homes ($Z = 1$) are located inside a risk zone that is one level higher than control homes. Observations are restricted to census tracts featuring both control and treatment homes. When homes sell multiple times, only one risk measure per home sale is used. Covariates in Panels A and B: slope, distances to nearest primary road and national forest, and dummy for location inside the wildland-urban interface. Panel A: RD estimates. MSE-optimal bandwidth derived using a local linear polynomial (including covariates in columns 2 and 4). The polynomial is of degree one. Robust standard errors computed with the nearest neighbor variance estimator. Panel B: OLS estimates for the model: $Risk_i = \beta_0 + \beta_{OLS} Treat_i + \mathbf{X}_i' \boldsymbol{\beta}_X + \varepsilon_i$, with census tract fixed effects and covariates $\mathbf{X}$. Homes are restricted to being located within 600m from the new 2007 risk zone boundary. Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

	Fire risk to homes		Wildfire probability	
	1	2	3	4
	Panel A: RD estimates			
$Z = 1$	0.0099	0.0134	0.0033	0.0092
	(0.0087)	(0.0087)	(0.0157)	(0.0159)
Covariates		✓		✓
N	4397	4375	3083	2981
	Panel B: OLS estimates			
$Treat$	0.0184***	0.0180***	0.0270***	0.0269***
	(0.0019)	(0.0018)	(0.0027)	(0.0027)
Census tract FE	✓	✓	✓	✓
Covariates		✓		✓
N	5773	5640	5152	5026
R^2_{adj}	0.78	0.84	0.81	0.85

FIGURE 1

RD Local Means, Price Residuals, and Wildfire Risk at the New Risk Zone Boundary

Panels (a)–(d) show log home prices by distance to the 2007 risk zone boundary, before (a–b) and after (c–d) the re-zoning: local means in (a) and (c), and residuals from a regression of log prices on covariates (age, house size, number of bedrooms, swimming pool dummy, distances to nearest primary road and national forest, and wildland-urban interface dummy) in (b) and (d). Panels (e) and (f) show the fire risk to homes and the wildfire probability. Homes to the right of the cutoff are inside a risk zone one level higher than homes to the left. The number of bins is data-driven (Calonico et al. (2015)). Vertical bars denote 90% confidence intervals. A polynomial of degree four is fitted on each side. Observations are restricted to census tracts featuring both control and treatment homes.

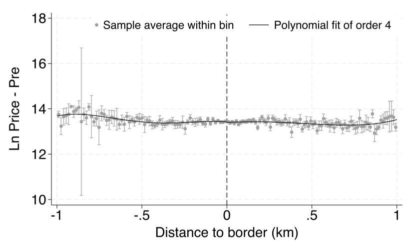

(a) Local means of log home prices (pre-2007)

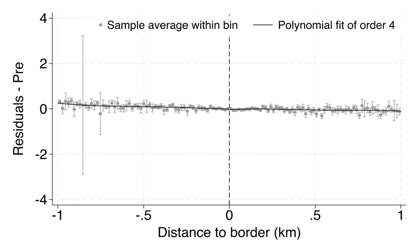

(b) Log home price residuals (pre-2007)

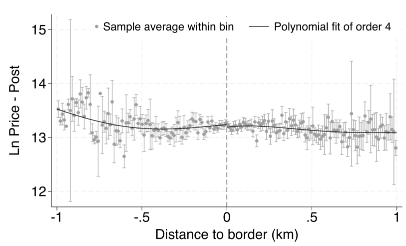

(c) Local means of log home prices (post-2007)

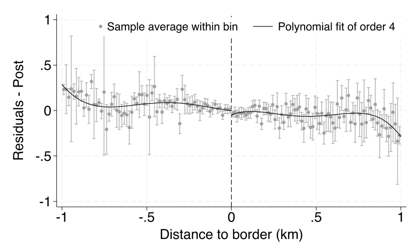

(d) Log home price residuals (post-2007)

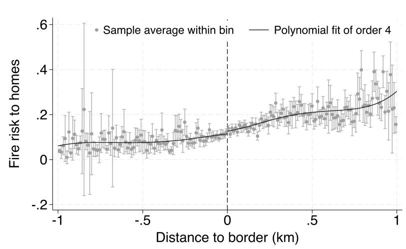

(e) Fire risk to homes

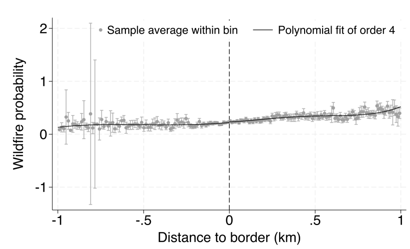

(f) Wildfire probability

FIGURE 2

Event-Study Estimates of Assignment to a One-Level Higher Risk Zone on Log Home Prices

Event-study estimates of the effect of assignment to a one-level higher wildfire risk zone on log home prices. Each marker is the coefficient on the interaction of the treatment indicator with a calendar-time indicator from a difference-in-differences regression on all sales within 600m of the new 2007 risk zone boundary, with census-tract, quarter, and year fixed effects and covariates. Panel (a): quarterly event time; the omitted reference period is 2007Q4. The pre-period coefficients are jointly insignificant (parallel-trends F-test $p = 0.46$). Panel (b): annual event time over the 2000–2015 sample window; the omitted reference period is 2007. The pre-period coefficients are individually and jointly insignificant at the 5% level (parallel-trends F-test $p = 0.08$). Covariates: age and age squared, home size and home size squared, # bedrooms, # baths, swimming pool dummy, elevation, slope, dummy for location inside the wildland-urban interface, distances to nearest primary road, national forest, state or local park. Vertical bars depict 95% confidence intervals based on heteroskedasticity-robust standard errors. Observations are restricted to census tracts featuring both control and treatment homes.

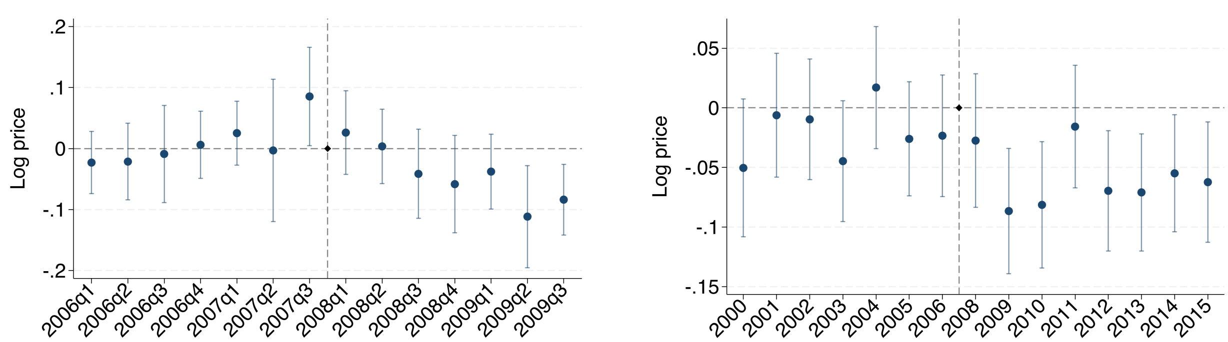

(a) Quarterly, 2006Q1–2009Q3

(b) Annual, 2000–2015 sample window

FIGURE 3

Robustness of RD, RD-DiD, and DiD Estimates to Bandwidth and Ring Choices

Regression discontinuity (RD), difference-in-discontinuities (RD-DiD) and difference-in-differences (DiD) estimates of the effect on log home prices of being located inside a one-level higher risk zone. Panels (a)–(c): RD in the post-policy period, RD in the pre-policy period, and RD-DiD estimates, respectively, with bandwidth ranging from 100m to 800m with 50m intervals. Regressions include covariates (age, size, number of bedrooms, presence of a swimming pool, distances to the nearest primary road and to the nearest national forest, and dummy for location inside the wildland-urban interface). The vertical blue lines mark half (left) and twice (right) the MSE-optimal bandwidth. Panels (d)–(f): DiD estimates for regressions that include home, county-by-year, and quarter fixed effects, and robust standard errors. Panel (d): bandwidth ranging from 100m to 800m with 50m intervals. Panel (e): 200m-ring ranging from 0m-200m to 400m-600m with 50m intervals. Panel (f): 400m-ring ranging from 0m-400m to 200m-800m with 50m intervals. In all panels, the central line depicts the point estimates; the dashed green lines depict the 95 percent confidence interval. Observations are restricted to census tracts featuring both control and treatment homes.