

EXPECTING CLIMATE CHANGE: A NATIONWIDE FIELD EXPERIMENT IN THE HOUSING MARKET

Daryl Fairweather*

Redfin

Matthew E. Kahn

University of Southern California & NBER

Robert D. Metcalfe

Columbia University & NBER

Sebastian Sandoval-Olascoaga

Massachusetts Institute of Technology

October 9, 2026

Abstract

Climate change increases property risk in the United States, yet home buyers lack affordable, location-specific risk information. We study the causal effects of providing flood risk information using a field experiment conducted with Redfin. Redfin randomly assigned 17 million users to receive flood risk disclosures for properties they viewed. Flood risk information changed which properties buyers searched for, bid on, and closed on, shifting them toward lower-risk homes without detectable changes in other observable property attributes. Effects were stronger in high-risk areas. Information provision shifted prices and the hedonic equilibrium, suggesting proactive climate adaptation.

*We would like to thank Spencer Banzhaf, Chris Costello, Matthew Eby, Cary Frydman, Justin Gallagher, Ed Glaeser, Andrei Gomberg, Allan Hsiao, Ed Kearns, John List, John Matsusaka, Paulina Oliva, Jeremy Porter, Esteban Rossi-Hansberg, Miguel Urquiola, Oren Ziv, and attendees at NBER Summer Institute EEE, ITAM, MIT, Stanford, UIC, and University of Pittsburgh for their comments. We also thank Robert Huang and Navya Kumar for excellent research assistance. Kahn & Metcalfe gratefully acknowledge funding from the USC Wrigley Institute for Environment and Sustainability Faculty Innovation Award. Disclaimer: Kahn, Metcalfe, and Sandoval-Olascoaga did not receive any funding or payment for the research from Redfin or First Street Foundation. MIT IRB: 2110000488; AEARCTR-0014485.

1 Introduction

Climate change threatens economies and human well-being worldwide. How individuals and firms adjust their behavior in response to new climate-related information will increasingly shape the social costs of carbon emissions (Nordhaus, 1993; Hinkel et al., 2014; Auffhammer, 2018; Kahn, 2020; IPCC, 2021). Adaptation can take many forms, including relocating to safer areas (“higher ground”) or investing in protective measures after location choices are made (Barreca et al., 2016; Deschenes et al., 2017; Ito and Zhang, 2020).

Such adaptive behaviors depend on awareness of high-stakes climate risks embedded in housing markets, particularly flood risk. Yet growing evidence shows that U.S. households substantially underestimate these risks (Bakkensen and Barrage, 2022; Wagner, 2022). As a result, properties in high-flood-risk areas may be significantly overvalued—by as much as \$200 billion in aggregate (Gourevitch et al., 2023)—a distortion expected to grow as climate exposure increases (Marcoux and Wagner, 2025).

Greater awareness of current and future flood risks could reduce both economic and human costs if it induces adaptive behavior. However, there is limited causal evidence on how individuals respond to information about property-specific climate risks, particularly in the context of housing search, purchase decisions, and residential location choice. To explore this, we conducted a rigorous assessment to measure the causal impact of climate change information on searching and buying decisions within the United States national housing market. We present the first large-scale field experimental evidence offering insights into how precise flood risk information influences the behavior of home buyers across the United States. Our nationwide natural field experiment was conducted in late 2020 utilizing the Redfin website and application (“app”), involving the participation of a substantial user base (17 million individuals in total).¹

In an unexpected and unannounced manner, Redfin randomized the assignment of complete property-level flood information at the individual customer level. Randomization occurred at the device level. For each property that a treated user searched for, they were provided with two key pieces of information: a flood risk score ranging from 1 (minimal risk) to 10 (extreme risk), and the predicted likelihood of flooding over a 30-year period. Customers were unable to filter their search results by flood score, and sellers were unaware of this information.

Apart from the flood risk information, all other features of the Redfin search experience remained consistent for both the treated and control groups. The flood risk information had no

¹For home buying in the U.S., using an online search brokerage such as Redfin is extremely common; according to the National Association of Realtors, 97% of U.S. home buyers use the internet for their search.

direct bearing on any search or pricing algorithms. Given Redfin’s significant market share during the experiment, our study generated variation in flood information across approximately 20% of the U.S. internet property market over the experiment’s duration, resulting in nearly 400 million property views. This widespread participation allowed for robust analysis of the effects of flood information on searching and buying behavior within the housing market, and whether the information affected the hedonic equilibrium.

In the natural field experiment, the new detailed flood risk information displayed on Redfin for all properties in the United States was generated by First Street Foundation (FSF) (Bates et al., 2021). FSF was a registered 501(c)(3) non-profit organization (now private) that had a complete U.S. flood risk predictive model. It has been argued that the FSF model and output are more reliable and valid than those provided by FEMA for a number of reasons (Wing et al., 2018, 2022). Unlike FEMA’s flood maps, which primarily serve the requirements of the National Flood Insurance Program and thus do not measure total flood risk for the whole of the US, the FSF flood score is generated by a predictive model that incorporates not only past flood risks but also considers future climate change trends, as well as pluvial flooding and fluvial flooding on small waterways, neither of which FEMA models. FEMA maps were created to cover coastal storm surge areas and main river channels, but neglect to cover heavy precipitation and small waterway flooding (creeks, streams, etc). When those risks are taken into account, the number of properties that should be in FEMA zones is about 17.9 million (versus the 8.0 million currently classified as at risk), which means that while FEMA zones are inland, they just do not take into account all sources of flooding, and those missed sources (precipitation and small waterway flooding) are more likely inland more than on the coast (FirstStreet, 2025).² The FSF predictive model has undergone peer review (Armal et al., 2020; First Street Foundation, 2020; Porter et al., 2021), and ongoing research efforts aim to validate its model predictions by leveraging recent natural disasters as natural experiments.

Our main hypothesis is that the presence of new flood risk information at the property-level will affect consumer behavior and the housing market. We are precise on the populations in our data that would be exposed to the new flood risk news to test this "new news" hypothesis—these are individuals who are searching for properties that were previously seen as safe (in a state of the world before the FSF model and outputs), but now are high risk for flood (because of the FSF output and predictions). By testing how information shapes individuals’

²Moreover, there are some important differences between FSF and the US Government approach. First, FSF’s model incorporates environmental factors to assess recent and future changes in flood risk over a 30-year period. Second, FEMA maps must be approved by local stakeholders, making them subject to political and business influence, whereas FSF models are independent of local lobbying. Third, FSF provides property-level flood scores, offering finer granularity than FEMA. Overall, FSF’s 1-in-100-year hazard layer identifies 1.7 times more at-risk properties than FEMA’s designations (First Street Foundation, 2020). FEMA maps are also outdated for policy and consumer use (Mulder and Kousky, 2023).

expectations, we aim to shed light on the role of information in facilitating forward-looking adaptation to climate change.

The Redfin data allow us to observe the full search process for treated and control users, including their interactions with the Redfin app for each property viewed, both before and during the field experiment. Our partnership with Redfin provides comprehensive insights into user behavior without imposing structure on the search process, as we observe it directly on the platform where treatment is orthogonal to all search and purchase activity. An additional identification advantage is Redfin’s use of in-house real estate agents, which gives us detailed information on bidding behavior for users who engage a Redfin agent, before and during the experiment, for both treated and control groups. This vertical integration across the home search and purchase process allows us to analyze how forward-looking flood risk information affects sorting and market outcomes.³

Additionally, we augment our analysis by linking the Redfin data with market data on all U.S. property transactions and listings during our experimental period and up until 2025. This linking allows us to examine the impact of treatment information disclosure on the hedonic equilibrium of the housing market. By incorporating transaction data, we can assess how the availability of randomized flood risk information influences the overall market dynamics and pricing. All of these data sources provide a novel and comprehensive foundation for our project, enabling us to accurately measure search dynamics, bidding behavior, and the impact of information disclosure on individuals’ climate-risk adaptive behavior.

We report three main sets of results. First, our analysis revealed that the randomized flood risk information had a significant and meaningful impact on users’ search behavior. Specifically, individuals who randomly received the flood risk information were more likely to search and browse properties with lower flood risk compared to the control group. Among those who initially searched for homes with high flood risk, the treatment led to a 12.4% reduction in the flood risk of their searched homes after two months. We found that consumers learned which properties and locations had high flood risk and created a sense of which properties were to be avoided (these homes make up 4% of the US housing stock). The information had little impact on the flood score for searched homes for those searching for relatively lower risk homes before the treatment started.

We also found that when the flood score was very variable in a defined location (i.e., high standard deviation of flood scores within a zip code), the impact of the randomized flood information was significantly large, but the information had no impact when the flood score

³Throughout the duration of the experiment, Redfin had an average of 1,757 lead agents per month across the country (Redfin, 2021). This wealth of agent data is instrumental in enhancing our understanding of bidding patterns and buying behavior, but also how search behavior maps onto bidding and buying behavior, which has not been previously linked in the literature.

was very uniform within a zip code. Together with the above result, this suggests that the impact of the information led people to search geographically more narrowly and overall more efficiently. Finally, we also found that the flood risk information did not push people to search for homes with other climate risky attributes that were not observable to the individual user (such as heat, wind, and wildfire). We predicted that treatment effects from the information should be for the homes that were not previously identified as high risk, but now are due to climate change and better modeling (from FSF). In fact, roughly 2.5% of U.S. homes fall into this category, and this is where we see the largest treatment effects of the flood risk information on search behavior. All of these results suggest that individuals actively and organically adjusted their search parameters and strategies in response to the flood risk information provided by Redfin.⁴

However, it is important to note that we failed to reject several null hypotheses. For instance, treatment effects on search do not vary with the median household income of the ZIP code from which users searched at baseline. Furthermore, we found no predictive impact on the treatment effects of the political affiliation from which the user was searching.

Second, our analysis of the impact of flood risk information on physical engagement with homes and overall bidding behavior revealed several interesting findings. While the randomized flood risk information did not significantly change the overall probability of touring a property, placing an offer, or closing a bid through Redfin (the extensive margin behaviors),⁵ it did have an effect on the bidding behavior at the intensive margin. Specifically, we observed that treated individuals who were initially browsing high flood risk properties were more likely to make offers on properties with approximately 57% lower flood scores compared to their control counterparts. We show that this is not a selection effect but a change in where the user makes an offer. In line with our "new news" hypothesis model, among users whose pre-experiment searches included waterfront but not coastal properties, treated users who had been browsing medium- and high-risk homes went on to bid on properties with 28.2% and 49.8% lower flood scores, respectively. These results suggest that even in-land medium-risk homes (approximately 11% of the U.S. housing stock) are negatively affected by full information on flood risk.

Third, flood-risk information moved outcomes in the broader housing market, not only among Redfin users. Using the full MLS transaction record, we find that after 2021, when this information became publicly accessible, medium- and high-risk homes began to sell at

⁴We also found that geographical areas that had been recorded as flooding in the past year had larger treatment effects of the flood risk information on search behavior.

⁵The information also did not change what type of customer was likely to tour and place an offer through Redfin, so we do not find selection on observables as we move down the funnel (i.e., search for a home → tour a home → make an offer on a home → purchase the home).

growing discounts relative to low-risk homes in the same ZIP code and year (by around 1.7%), implying a substantial revaluation in higher-risk areas. Using exogenous updates in First Street Foundation flood-risk classifications for the same property, we find that reclassification from high to low risk raises the sale–listing price spread by about 4.8%, while reclassification from low to high risk reduces it by about 1.2%.

Moreover, within the experiment itself, we can trace this pricing effect to the hedonic equilibrium directly. We leveraged the exogenous variation created by our randomization, which resulted in some homes having a higher percentage of treated users and others having a lower percentage just through random chance.⁶ We found that when every Redfin user viewing a property is treated (in comparison to no Redfin users treated for that property), property prices dropped by \$8,900 (1.9% of the property price) for not-waterfront homes with high flood risk across the United States. Across the three different hedonic methods that we estimate, we find very high levels of agreement on the size of the effect of the flood risk information on high flood risk homes.

The flood risk information had a tangible effect on property prices, with homes in high flood risk areas experiencing a decrease in value. This value can be construed as the value of the best information available on the expectations of climate change impacting high-risk homes. The ex-ante losses posed by climate change are lower if risk lovers (and maybe those with an edge in upgrading risky homes) are more likely to live in these homes. While we cannot directly observe these attributes of people, our evidence is suggestive of some “reshuffling,” though we note that the estimates behind this interpretation are imprecise. Due to the high risk homes having less competition, we find that prices for medium-risk homes rose by around \$4,000, although this estimate is imprecise. Overall, the information causes less competition on the high risk homes (4% of the homes in the U.S.) and more competition on the medium risk homes (15% of the homes in the U.S.).

Overall, our findings on how flood risk information affected search, engagement, offers, and the housing market provide evidence that the flood risk information was transmitted through the “new news” mechanism. This mechanism influenced the hedonic equilibrium pricing of properties, with homes in high flood risk areas experiencing a reduction in value due to the newly revealed flood risks provided by FSF’s predictive model.

Our research addresses a crucial gap in the literature by conducting a natural field experiment that examines the impact of climate change information on the process of searching for and buying a home with no selection into the field experiment with complete covertness. While previous studies have used observational data or simulations to analyze climate change

⁶One can think of this design as akin to Crépon et al. (2013), but where homes (as opposed to areas) randomly receive from 0 to 100% of the searchers as treatment users.

adaptation (Massetti and Mendelsohn, 2020),⁷ Our experiment leverages a unique opportunity to study the causal effects of flood risk information on consumer behavior.

Existing empirical work examines how climate risk is capitalized into local home prices, with mixed evidence on the degree of capitalization.⁸ However, this literature does not fully address whether market participants are aware of the true risks properties face.⁹ In contrast, we show that individuals avoid high flood-risk properties once provided with new information, consistent with a “new news” hypothesis. The observed effects on search, bidding, and hedonic pricing indicate that flood risk information was not previously known or fully incorporated into market outcomes.

Our research is also related to studies estimating the adaptation costs of coastal flooding, hurricanes, and storms (Seetharam, 2018; Balboni, 2025; Desmet et al., 2021; Fried, 2022; Jia et al., 2026; Bilal and Rossi-Hansberg, 2023; Hsiao, 2023). We complement this literature by adopting a microeconomic perspective on how heterogeneous home buyers sort across properties. Better information about emerging, property-specific risks likely reduces ex-post regret in this matching process. Because housing is a long-term investment, expectations about future climate risks play a central role in search decisions. Results from our natural field experiment indicate that forward-looking risk expectations are critical for evaluating the general equilibrium welfare effects of increased flooding and sea-level rise.¹⁰ By quantifying how buyers trade off current and future climate risks against housing attributes, our study advances understanding of housing market decision-making, an area with limited causal evidence.¹¹

⁷While there are non-experimental papers suggesting that flood maps (Hino and Burke, 2021; Weill, 2026) or flood disclosure requirements (Lee, 2021) may cause a change in property prices, these maps and disclosures are based on the inferior FEMA maps that exclude the new flood-risk science and are not forward-looking. We would not find any treatment effects in our experiment if previous information sets were complete, available, and accessible; hence, our new news interpretation would not exist.

⁸Several studies find weak or partial capitalization of flood risk into property values (Harrison et al., 2001; Bin and Landry, 2013; Beltrán et al., 2018; Bernstein et al., 2019; Murfin and Spiegel, 2020; Baldauf et al., 2020; Gibson and Mullins, 2020; Keys and Mulder, 2020; Giglio et al., 2021; Hino and Burke, 2021).

⁹Recent research shows that most households in high flood-risk areas lack flood insurance (Gallagher, 2014; Kousky et al., 2020; Wagner, 2022), many are not flood-risk literate (Botzen et al., 2015; Royal and Walls, 2019; Kousky and Netusil, 2023), and adaptation investments such as home elevation are underprovided due to long time horizons (Hovekamp and Wagner, 2023). Conell-Price and Mulder (2024) find that owners who underestimate their flood risk do not update beliefs when informed of true risk using FSF data, suggesting motivated reasoning on the supply side and supporting information interventions on the demand side.

¹⁰This finding is consistent with evidence on the adaptation benefits of improved weather forecasts (Rosenzweig and Udry, 2014; Schlenker and Taylor, 2021; Molina and Rudik, forthcoming; Downey et al., 2023; Goldsmith-Pinkham et al., 2023; Shrader, 2023; Shrader et al., 2023; Burlig et al., 2024).

¹¹Greenstone (2017) notes the lack of field experimental applications of Rosen’s hedonic model; this paper directly addresses that gap.

2 The Field Experiment

2.1 Background

The field experiment was conducted through one of the largest full-service real estate brokerages in the United States, Redfin, which had more than 47 million average monthly users in 2021. Redfin pairs agents with individuals buying or selling homes. We first describe the treatment background and then the experimental design.

Redfin allows users to search for properties nationwide. The search process is as follows: (1) the user opens the app and sees a map with nearby homes for sale (Figure A1, center screen; green price labels denote listings); (2) the user moves the map or changes locations to view properties elsewhere; and (3) the user applies filters, including price, bedrooms, bathrooms, home type, property and lot size, time on Redfin, and basic features (e.g., garage or pool). At the time of the experiment, users could not filter on flood score or walk/transit score (Figure A1, left two screens), including those in the treatment group. Registered users could like and save properties, view search history, and restart saved searches (Figure A1, right two screens), in addition to receiving emails. Clicking on a green price label opens the full property listing, and we observe all properties viewed on the app.

For flood risk information, Redfin partnered with the non-profit First Street Foundation (FSF), which developed the Flood Factor Score to predict current and future property-level flood risk.¹² FSF models flooding from fluvial, pluvial, and coastal sources while incorporating current and future environmental conditions at the property level (Bates et al., 2021). The model captures flooding in ungauged or under-gauged areas and regions outside traditional flood models. It relies on a Regionalized Flood Frequency Analysis (RFFA) approach using statistical propensity matching to infer flow parameters for ungauged streams and river reaches based on comparable gauged areas. A key feature is the inclusion of pluvial (rainfall) flooding as a probabilistic risk with associated depths and return periods (First Street Foundation, 2020). The model is peer-reviewed, open-source, and widely regarded as among the best available in the United States.¹³


The Flood Factor Score ranges from 1 to 10, categorized as minimal (1), minor (2), moderate

¹²FSF sells its forecasts to U.S. government agencies and the GSEs and also makes flood risk data for all U.S. residential properties publicly available. It has received funding from donors including the 2040 Foundation, Hightide, and Grantham.

¹³The FSF model was developed through an open science process with contributions from many researchers and ongoing evaluation. Risk scores are reported without confidence intervals, reflecting an active debate about communicating uncertainty to the public. Like all predictive models, it involves Type I and Type II errors, potentially overstating risk in some areas and understating it in others (Bates, 2023), though FSF reduces false positives by using the conservative SSP2-4.5 climate scenario rather than SSP5-8.5.

(3–4), major (5–6), severe (7–8), and extreme (9–10) flood risk. We group these categories into three levels, with minimal and minor classified as low, moderate and major as medium, and severe and extreme as high.¹⁴ The Flood Factor Score reflects the probability of at least one flood over the life of a 30-year mortgage signed today. The score is two-dimensional: higher values indicate greater likelihood and severity of flooding over the next 30 years (Figure A2). It incorporates current and future risk from major flood types and their combinations, including intense rainfall, river overflow, high tides, and coastal storm surge. Scores can vary substantially within neighborhoods due to local factors such as elevation, proximity to water bodies, and flood mitigation infrastructure (First Street Foundation, 2020). This information represents the most objective flood risk signal available to individuals.¹⁵ Figure 1 shows the variation in the properties classified as high risk by FSF and FEMA.

Figure 1: Percentage of all US Properties With High Flood Risk by County



Note: This map displays, by county, the percentage of properties in the United States classified as high risk by both FEMA (zones A, AE, V, or VE) and First Street (flood factor ≥ 7). The risk percentages are divided into equal terciles, as shown in the legend. The tercile cutoffs are 0, 5.7%, 9.2%, and 98.3% for First Street, and 0, 2.3%, 5.6%, and 92.8% for FEMA. Overall, about 4% of properties in the United States are classified as high risk; this share is slightly higher at the county level because properties are more concentrated in low-risk areas.


¹⁴Results using the six Flood Factor categories are reported in the working paper version, here.

¹⁵We do not test whether individuals fully update beliefs in response to the information or assess alternative mechanisms such as attention. We provide indicative evidence that behavioral responses strengthen over time, which is inconsistent with a pure attention mechanism.

2.2 The Design of the Field Experiment

The overall design was a natural field experiment because users were unaware of the experiment, and there was no selection into or out of it (Harrison and List, 2004). In this natural field experiment, Redfin randomly assigned new and existing Redfin users to either the treatment group, in which they were shown a Flood Factor Score section on every on-the-market and off-the-market property listing page they visited, or to a control group without a Flood Score section. Figure 2 presents the experiment experience using a mobile device for both the control and treatment groups within a property page with a flood score of 6 (i.e., major flood risk). The control group (Figure 2(a)) did not see the flood risk section while scrolling a property page, whereas the treatment group (Figure 2(b)) had full access to it.¹⁶

Figure 2: The Visual Experience for the Treatment Group and the Control Groups



The flood risk section was placed after the property's public facts and neighborhood information. When a user clicks on the flood risk section, they view an "about flood risk" paragraph, the colored (i.e., from light green to dark red depending on the risk) Flood Factor Score, and the flood likelihood of the property over time (1, 5, 10, 15, 20, 25, and 30 years).

Redfin users in the treatment group always see the main flood risk number as they scroll down the page for every property that they search for, but they can ignore the more in-depth flood risk probabilities over time by not clicking on the tab (that unfolds the full 30-year flood risk information). They choose their own intensity of engagement with the provided

¹⁶Figure A3 shows the color and labels Redfin used to show a property's flood score within the flood risk section, ranging from 1 to 10.

information. Redfin did not nudge home searchers to focus on this information or advertise the new information feature to consumers (they had to discover it organically, with no help learning about the flood risk). It is included as an additional piece of information about the home. Throughout the searcher’s web-page experience, we observe their full organic search activity (for all users in both treatment and control groups). The flood risk information covered 99.9% of the U.S. housing market (First Street Foundation, 2020).

2.3 Implementation of the Field Experiment

The nationwide natural experiment enrolled 17,017,490 unique users for 12 weeks between October 12th, 2020 and January 3rd, 2021.¹⁷ The number of unique users in the field experiment represented approximately 40% of the average monthly number of unique Redfin users in 2020 (Redfin, 2020). The experiment enrolled both new users browsing the website and app and existing users who had browsed the website and app before entering an experiment arm. Of the enrolled users, 964,152 are registered Redfin users and 648,714 have flood-risk information at baseline (Table A11).¹⁸ We focus on this registered sample because: (i) they have a higher likelihood of actually purchasing a home; (ii) we observe their behavior for a longer time frame (usually from before the experiment, so we know their pre-experiment flood risk search); and (iii) we can guarantee compliance with the assignment mechanism.

Randomization Architecture. Our experiment was implemented as a platform-level randomized controlled trial with randomization at the browser-identifier level. The unit of assignment is a persistent cookie-based identifier associated with a browser or device, which is the platform’s native unit of traffic.

The platform operationalizes enrollment and treatment assignment through two independent, deterministic hash mappings of the browser identifier and the experiment identifier. First, each browser identifier is mapped by a hash function to one of a fixed set of buckets, and the experiment is active for a predetermined fraction of those buckets. In our case, it was set at approximately 40 percent of public traffic and held constant over the 12-week window. This is a per-identifier probability rather than a quota, so enrollment does not depend on the time of day, the week, or how many other identifiers have already been enrolled. Because the mapping is deterministic, enrollment resolves identically for a given identifier on every visit, for returning and newly arriving identifiers alike. Identifiers outside the active buckets receive the default experience and are excluded from all analysis. Second, enrolled identifiers are mapped by an independent hash of the browser identifier and the experiment identifier

¹⁷At this point in time, the macroeconomic environment was between the COVID-19 peaks and a stable economy. We had full access to Redfin data for one week prior to the first day of the experiment.

¹⁸Figure A4 shows the spatial distribution of users by state where they were browsing the most at baseline.

to control or treatment cohorts of equal size.

A browser identifier enters the experiment when it first loads a property-details page on which the Flood Factor section would be served, within the experiment window. Both mappings are evaluated on the server when the page is generated, so the control or treatment experience is determined before the page renders and independently of how the user subsequently interacts with it. Because identifiers reach such pages at different times depending on when and how they browse listings, entry is staggered over the study period (see Figure Figure A5).

2.4 The Unbiasedness of the Treatment Effect Due to Randomization

The design and implementation of our randomization ensure the internal validity of the study by satisfying the three assignment mechanism restrictions and the four exclusion restrictions required to identify an unbiased treatment effect. The assignment mechanism satisfies non-zero probability, as all Redfin users have a positive probability of treatment assignment and the functional form of assignment is controlled; individualism, since each user’s assignment is orthogonal to the assignments and outcomes of other users; and unconfoundedness, as treatment assignment is independent of users’ potential outcomes.

The experiment also satisfies the four exclusion restrictions. First, SUTVA holds for the individual-level search, engagement, and bidding outcomes analyzed in Sections 4.1–4.3: a given user’s browsing and bidding behavior does not depend on the treatment status of other users, and there are no alternative versions of the treatment.¹⁹ We deliberately relax this assumption in Section 5, where the object of interest is a market-level price response that operates precisely through the treatment status of the other bidders on a listing; we discuss the resulting spillovers there. Second, observability is satisfied because users must be observed on the platform to be assigned to treatment or control, ensuring that first-period potential outcomes are always observed. Although attrition may occur over time as users progress through the housing search process, we test for differential attrition and find no evidence that treatment assignment affects outcome observability.

Third, we have complete compliance for registered users, whose treatment status is enforced throughout the experiment; because we cannot guarantee compliance for unregistered users, our primary analysis focuses on registered users, though we also examine effects for the full sample. Given that these conditions hold, the assignment mechanism is statistically independent of all potential outcomes, allowing us to recover an internally valid and unbiased

¹⁹For example, a treated user’s housing search does not affect another user’s search.

estimator of the effect of flood risk information on user behavior.²⁰

2.5 Demand Side Hypotheses Through the Lens of a Hedonic Housing Assignment Model

A parametric closed form hedonic assignment model allows us to concisely state our paper’s core hypotheses. We show and solve the model in Appendix A.1. Our model suggests that flood safety is currently undervalued because of a lack of information on risk, but once that information becomes available (i.e., “new news”), people act on it and value it, and those with a higher preference for safety act on it even more. This hypothesis predicts that the following consumers will be affected by the new information: (a) consumers searching for properties that are in-land (i.e., not coastal) but have high flood risk defined by FSF; (b) consumers searching for properties that are in-land (i.e., not coastal) but close to a waterfront (e.g., river, lake) and have high flood risk defined by FSF; and (c) Consumers searching for properties that are not defined as risky by FEMA and are on the coast.

For populations (a), (b), and (c), the information will revise their flood risk beliefs upward and will search and buy lower flood risk homes, as per propositions 1 and 2. We believe that the central tenet and test of this new news hypothesis is the following: two groups of consumers will be less affected by the FSF flood risk information: (d) consumers searching for properties already FEMA-designated as high risk; and (e) consumers searching for properties with zero to low risk. Population (d) should already be informed about the flood risks because of current information dissemination, and because you need flood insurance to obtain a mortgage in these FEMA areas.²¹ These are the predictions for the new news hypothesis on search and offers.

3 Data and Empirical Strategy

3.1 Data

The analysis combines Redfin user-level browsing and engagement data with property-level transaction records from Multiple Listing Services (MLS).

Property views data. When a user clicks on a property listing, Redfin records a view event containing the user’s anonymized identifiers, the timestamp of the view, the device used, the user’s approximate search location, and key property characteristics such as list price, bedrooms, bathrooms, square footage, year built, and flood-risk score. We organize this dataset

²⁰We also satisfy the four main assumptions for external validity as set out by List (2026).

²¹We recognize that there still might not be full information from the FEMA high flood risk ratings, but there should be less of an effect for this population than populations (a) to (c).

as a user-day panel.

Engagement data. Redfin also records detailed interactions on a property page, such as scrolling, clicking on photos, saving a property, or requesting additional information. These data include timestamps and session-level engagement measures such as time spent on the listing and number of interactions. As with the views data, we aggregate these records to the user-day level.

Touring, bidding, and closing data. We observe later stages of the home-buying process through records of property tours, offers, and completed transactions facilitated through Redfin. Each observation includes anonymized user identifiers, the action date, property identifiers, and offer or closing characteristics.

MLS data. Property-level characteristics and transaction outcomes come from MLS records covering markets where Redfin operates. These records include listing and sale dates, listing and sale prices, property structural characteristics, and geographic identifiers.

FEMA flood zone data. We classify a property as FEMA high risk if it falls within a Special Flood Hazard Area in FEMA’s National Flood Hazard Layer, taking zones A, AE, V, and VE as high risk and matching properties to zones by geographic coordinates. This designation includes inland riverine zones as well as coastal zones.

3.2 Empirical Strategy

Because treatment was randomly assigned, our main specification compares outcomes during the experimental period between treatment and control users while controlling for the baseline value of the corresponding outcome whenever available. We cluster standard errors at the individual level. Our main specification is:

$$y_{it} = \beta_0 + \beta_1 D_i + \beta_2 X_i^{pre} + \beta_3 (D_i * X_i^{pre}) + \beta_4 \bar{y}_i^{pre} + \gamma_{b(i)} + u_{it} \quad (1)$$

Where y_{it} is the average value of the outcome for individual i on day t , $D_i \in \{0, 1\}$ represents treatment assignment, X_i^{pre} is the average flood-risk score of the properties viewed by individual i from one week before the experiment began until the day before that individual entered the experiment, grouped into the low, medium, and high categories. $\bar{y}_i^{pre}$ is the average value of the same outcome over that same user-specific pre-experiment period. The term $\gamma_{b(i)}$ symbolizes indicators for the date on which individual i entered the experiment. This specification tests the “new news” hypothesis that treatment effects are larger for users who were initially viewing higher-risk properties. For outcomes with limited or no pre-experiment observations, such as tours, offers, and closings, we instead estimate difference-in-differences

specifications that compare treatment and control users before and after entry into the experiment.

To account for multiple hypotheses, we use the Romano-Wolf procedure to adjust for the familywise error rate (Romano and Wolf (2010)). We apply this procedure to equation 1 across three baseline flood-risk groups (low, medium, and high) and six subsets of registered users: (1) all registered users, (2) those without a waterfront search at baseline, (3) those with a waterfront search at baseline, (4) those without a FEMA search at baseline, (5) those with a FEMA search at baseline, and (6) those with a waterfront search but no coastal search at baseline. We conduct 1,000 resamples for the search regressions and 3,000 resamples for the offer regressions.²²

3.3 Pre-Experiment Balance

Registered vs. Non-Registered Users. Tables A1 and A2 show that registered and non-registered users differ in baseline browsing patterns. Because registered users are more likely to purchase a home, can be observed for longer, and allow us to verify compliance with treatment assignment, we focus the main analysis of search, engagement, and bidding on this sample. We use the full dataset when examining broader housing-market effects.

Pre-Experiment Balance. Tables A3 and A4 report pre-experiment balance tests between treatment and control groups for registered users; we cannot reject the null in any specification. Tables A5 through A10 report the same balance tests separately by baseline flood-risk category and again show no statistically significant differences between treatment and control users. Finally, Table A11 reports the distribution of registered users across baseline flood-risk categories by treatment status.²³

4 Results

4.1 Home Search Behavior Dynamics

We begin by estimating equation (1) without segmenting by flood risk. Table A12 reports the treatment effects on the number of properties browsed (column 1), average browsed property characteristics (columns 2–5), average flood score (column 6), and the Herfindahl–Hirschman

²²This adjustment is conservative, since the heterogeneity analysis is centered on one main hypothesis: that the effect of treatment varies with baseline flood-risk exposure. We do not additionally adjust across broad outcome families such as search, engagement, and offers, since the relevant sample changes as users move through the home-buying funnel, and there is no natural way to aggregate these outcomes into a single family of tests. See also Viviano et al. (2026) and List et al. (2023).

²³A Pearson test fails to reject equality of the distribution across treatment arms ($\chi^2(5) = 5.1, p = 0.398$).

zip code index of user searches (column 7); Table A13 reports effects on year built and on walk, transit, and bike scores.²⁴ We find no significant effects of treatment on daily browsing behavior.

When we classify users by the average flood-risk score of the properties they viewed before the experiment, treated registered users show a change in search behavior. Figure 3(a) reports estimates of β_3 from equation (1) by baseline average flood-score search category. The largest and statistically significant change occurs among users who, prior to the experiment, browsed properties with high average flood risk. Treated registered users in this group reduced the average flood-risk score of the properties they browsed by 0.21-points, representing -4.9%. To put this effect in perspective, it corresponds to average reductions in long-term flood risk of 8.0%, 10.6%, 6.2%, 4.1%, 3.4%, 2.9%, and 2.7% over time horizons of 1, 5, 10, 15, 20, 25, and 30 years, respectively.


For users browsing in areas with relatively low initial flood risk, the scope for treatment-induced reductions in exposure was limited. We therefore also report effect sizes in standard deviation units. The estimated effects are close to zero for the low- and medium-risk groups, at 0.001 and -0.006 standard deviations, respectively, but larger for the high-risk group, at -0.137 standard deviations.

Figure 3(b) plots event-time estimates for registered users who, on average, viewed high-risk properties before the experiment. For users who, on average, viewed high-risk properties before the experiment, treated users increasingly shifted toward properties with lower flood scores than control users. By week 6, the average flood score of viewed properties was 5.9% lower in the treatment group, and by week 9 it was 12.4% lower, both statistically significant. Our analysis showed no statistically significant impacts in event-time studies of registered users who browsed properties with low and medium risk levels before the experiment commenced (Figures A6 (a) and (b), respectively).

Place-based factors may also influence the impact of flood information. For instance, differences in information access and local knowledge between homeowners and non-local buyers can create a situation where the latter may unwittingly purchase properties with higher environmental risks. In addition, place-based risks in flood-prone areas could further affect the effectiveness of our treatment. For example, people in safe areas may be more likely to adjust their home search when looking for properties in high-risk areas, whereas those already in flood-prone areas may be less responsive to our treatment. To test this, we calculated the difference in the mean, standard deviation, and coefficient of variation of the zip code's flood score between an individual's origin and most search destinations at baseline, and we inter-

²⁴Tables A14 and A15 present the same estimates restricted to users with pre-experiment information.

Figure 3: Average Flood Score of a Daily Search, Registered Users
Change relative to Control



Note: Vertical lines show 95% confidence intervals. In both panels, each point reports the corresponding estimate of $\beta_1 + \beta_3$ from equation 1. The percentage shown next to each point is based on the corresponding estimate from the same specification re-estimated with the natural log of the outcome as the dependent variable, transformed as $(\exp(\beta_1 + \beta_3) - 1) \times 100$. In panel (a), the x-axis shows users' average search pre-experiment flood-score category. In panel (b), the x-axis shows weeks since entry into the experiment for users who, on average, viewed high-risk properties before the experiment.


acted it with the treatment and the average flood score of search destinations at baseline to test our previous hypothesis.

We find a 12.9% decline in the average flood risk of viewed homes among users who, before the experiment, searched mainly in high-risk areas and in areas whose flood-risk coefficient of variation (CV) exceeded that of their origin ZIP code (Figure 4(b)). To assess whether this pattern is related to differences in the distribution of flood risk between origin and destination ZIP codes, we estimate specifications that interact treatment with the difference in the mean and standard deviation of flood-risk scores between destination and origin ZIP codes, respectively. We find no statistically significant effect for the mean difference (Figure A7), but we do find one for the standard-deviation difference (Figure A8). We also do not find statistically significant effects when allowing effects to vary by baseline ZIP-code median household income (Figure A9).

In addition to showing flood risk categories, our study provided a detailed view of flood risk via a quinquennial cumulative flood probability for properties, as shown in Figure 2. We find a statistically significant effect of the cumulative flood probability in 2020, i.e., a one percentage-point increase is associated with a 0.5% decrease in the average flood score of the properties searched. Estimates for longer horizons are not statistically distinguishable from zero (Figure A10).

Figure 4: CATE on the Average Flood Score of a Daily Search for Registered Users, by Characteristics of Most Searched Destination and Origin Zip Code at Baseline

Change relative to Control




Note: Vertical lines show 95% confidence intervals. In both panels, each point reports the corresponding estimate of $\beta_1 + \beta_3$ from equation 1. The percentage shown next to each point is based on the corresponding estimate from the same specification re-estimated with the natural log of the outcome as the dependent variable, transformed as $(\exp(\beta_1 + \beta_3) - 1) \times 100$. The x-axis shows users' average search pre-experiment flood-score category.

We also examine heterogeneity by whether users' pre-experiment searches included homes that were waterfront (coast, river, or lake), located in FEMA high-risk zones, near the coast, or combinations of these characteristics. We classify users based on the attributes of the homes they viewed before the experiment. Figures 5(a) and (b) show that flood-risk information had its strongest effects among users whose pre-experiment searches included waterfront properties. Panels (c) and (d) show a similar pattern by FEMA high-risk status, with larger effects for users whose initial searches did not include FEMA high-risk properties. Panel (e) points in the same direction, suggesting that the waterfront result is driven by users browsing at baseline properties near water but not on the coast. We also find that this pattern is driven by users whose baseline searches were concentrated in states without floodplain disclosure requirements (Figure A11). These patterns are consistent with the "new news" hypothesis, as treatment appears to matter most for users interested in waterfront properties that are not coastal and are therefore less likely to be flagged as risky by FEMA or disclosed through state floodplain requirements.

We next examine which property attributes users trade off in response to lower flood risk. Table A16 shows no statistically significant changes in bedrooms, bathrooms, square footage, list price, Herfindahl-Hirschman Index (HHI), or walk, transit, and bike scores among users who initially browsed high-risk properties. Similarly, users who initially searched for low-

Figure 5: CATE on the Average Flood Score of a Daily Search for Registered Users
Change relative to Control



Note: Vertical lines show 95% confidence intervals. Each point reports the corresponding estimate of $\beta_1 + \beta_3$ from equation 1. The percentage shown next to each point is based on the corresponding estimate from the same specification re-estimated with the natural log of the outcome as the dependent variable, transformed as $(\exp(\beta_1 + \beta_3) - 1) \times 100$. The x-axis shows users' average search pre-experiment flood-score category. Users who did not browse any waterfront or FEMA Risk property before the experiment are classified as "without" waterfront or FEMA Risk search at baseline. A property is classified as being located on the coast when its geographic coordinates (latitude and longitude) are 200 meters or less from the nearest shoreline.

or medium-risk properties do not exhibit significant changes in these attributes. Finally, we analyzed whether users were willing to accept higher fire, heat, or wind risks in exchange for lower flood risk by analyzing changes in these scores for viewed properties. Neither treated nor control users had access to these risk measures during the experiment. Figure A12 shows no statistically significant differences between the groups.

4.1.1 Nonparametric Conditional Average Treatment Effects

The impact of flood information may vary by individual characteristics. We recognize that Redfin customers may vary in their risk aversion, income, local social capital, and family ties to the area they are searching in. Such searchers may also vary in their trust in the data Redfin is supplying.

In the previous section, we demonstrated heterogeneity in treatment effects by analyzing how these effects varied with users' pre-experiment flood-risk search behavior. However, the impact of our treatment could also vary with other individual characteristics, in addition to baseline flood-risk search behavior. Moreover, the estimator used to calculate the previous CATE, i.e., the estimator presented in equation 1 relies on the linearity assumption of the effect that covariates, X_i , have on the treatment. If these effects were nonlinear, our estimated effects would be biased or would not capture the entire distribution of heterogeneous treatment effects.²⁵

Several results are worth highlighting. Figure A32 shows that the algorithm identified that 76% of conditional average treatment effects were negative for users browsing high-risk properties before the experiment began—albeit only 19.4% were statistically significant at the 90% level. The largest reduction in this group was -0.79 in the average flood risk score. Moreover, Figure A33 also plots how causal forests found that about 22.7% of the treatment effects were negative for those users browsing medium risky properties before the experiment began, but none of them were statistically significant. The most significant reduction was -0.33 in the flood risk score relative to the control group. Finally, for those in the low pre-experiment risk groups, 11.4% were negative and statistically significant treatment effects (Figure A34).

²⁵To account for the possibility of nonlinearity and to incorporate the influence of other baseline characteristics on treatment effects, we use a Generalized Random Forest algorithm known as causal forests (Athey et al., 2019; Athey and Wager, 2021). We further describe this algorithm in section A.3. Figures A32 to A34 show the predicted conditional average treatment effects from causal forests stratified by baseline flood risk categories.

4.2 Engagement

The Redfin experiment represents an “intention to treat.” Each searcher must decide whether she keeps engaging with this information. Everyone has a time budget constraint, and individual home buyers know their home purchase priorities. A home buyer may spend relatively little time engaging with flood risk data, given the many other dimensions of a home’s quality to consider. In this section, we test this hypothesis.

Figure A13 shows that treatment had no statistically significant effect on the likelihood of registering to Redfin during the experiment. Table A17 shows no significant differences in daily website usage, measured by seconds spent, sessions per day, unique home views, or total home views, between treatment and control groups. Further analysis of pre-experiment flood-score searches reveals no statistically significant effects on user behavior. Individuals browsing high-risk properties at baseline did not spend more time on the platform or engage in more sessions on Redfin’s website (Figures A14 (a) and (b)). Similarly, Figures A15 (a) and (b) show no significant differences in the number of unique or total homes viewed per day. These outcomes are defined only on days a user visits the platform. Figures A16 (a) and (b) repeat the analysis on a panel that includes every day from a user’s entry through the end of the experiment, with inactive days entered as zeros. Neither the probability of browsing on a given day nor the number of unique home views differs by treatment in any baseline flood-risk group.


However, we find that the treatment affected engagement with the listings. Figure 6 shows how treatment, stratified by baseline average flood-score search category, affected user engagement with (a) the flood-risk section,²⁶ (b) “favoriting” a property, (c) advancing to the next photo, and (d) viewing similar properties, expressed as a share of all properties a user browsed per day. Users who, at baseline, browsed both low- and high-risk properties had an increased engagement with the flood-risk section, relative to control users (Figure 6 (a)). As well, users browsing, on average, properties with low flood risk pre-experiment “favorited” more properties once the flood score became available, whereas those browsing, on average, high risk properties pre-experiment, “favorited” fewer properties (Figure 6 (b)).

Our findings above suggest that users who have previously browsed high-risk properties adjust their search behavior once they learn about flood risk (as seen in the average flood score search results in the previous section). This finding has important economic content because we reject the hypothesis that informed, risk-loving individuals seek out risky homes. If such “complete information” matching were occurring, we would not expect to observe the facts reported above.

²⁶Engaging with the flood risk refers to reaching the flood risk section for treated individuals and reaching the section where the flood risk would appear for control individuals.

We also examine these engagement metrics across property types based on users' pre-experiment search histories. Figures A17 and A18 in the Appendix report the average treatment effect on engagement with the flood-risk section for users who browsed the waterfront and FEMA-designated high-risk properties at baseline. Consistent with the search results, treated users in these groups are more likely to engage with the flood-risk section.

Figure 6: CATE on the Percentage of Times Registered Users Engaged with a Specific Property's Features per Day
% Times User Engaged with Feature for Listings Viewed per Day



Note: Vertical lines show 95% confidence intervals. Each point reports the corresponding estimate of $\beta_1 + \beta_3$ from equation 1, where baseline engagement is not included, as it is not observed pre-experiment. The x-axis shows users' average search pre-experiment flood-score category. Engagement with a property feature is measured as the share of listings visited on a given day in which the user interacted with that feature (e.g., clicking "favorite," advancing to the next photo, or viewing similar listings). For the flood-risk section, engagement is defined as scrolling to that section for treated users; for control users, it is proxied by scrolling past the page location where the section would have appeared.

4.3 Tours, Offers, and Closes

Redfin provides customers with an integrated house purchase process such that home

buyers can search for a specific home and then work with a Redfin-directed brokerage service. It allows users to tour properties, place bids on specific properties, and close deals. In our experiment, for all users who search for a given property, choose to tour it, and then place a bid on that property, we observe each of these steps in the housing purchasing process. Based on revealed preference logic, we view touring and bidding on a home as important (and costly) evidence that one is responding to information. In this section, we analyze whether treatment information affected the extensive margin decision to tour and make an offer, and, conditional on touring and making an offer, whether they changed the type of home they toured and offered on.

Tours. As shown in Figures A19 (a) and (b), access to flood-risk information does not affect the likelihood of booking or canceling a home tour, regardless of users' baseline flood-score search or the flood score of the property. We find no statistically significant differences between treatment and control groups in these outcomes. Similarly, we do not find a statistically significant difference in the flood scores of homes toured between the two groups.

Offers. Figures 7 (a) and (b) show the probability of placing a house offer and the flood score of properties bid on, stratified by baseline average flood-score search category and relative to control users. Figure 7 (a) shows no statistically significant effect of flood score access on the likelihood of making an offer. In contrast, Figure 7 (b) shows that treated users who browsed high flood-risk properties pre-experiment bid on properties with 3.24 (-57%) lower flood scores than their control counterparts. Although offer probabilities do not differ, there may still be selection among users who place offers. While treatment and control groups are balanced on baseline flood risk, selection on observables remains a concern. We show that the result reflects a change in where users bid rather than who bids. Specifically, we examine balance in pre-experiment search behavior for treatment and control users who ultimately placed an offer, separately for every search category. Appendix Tables A18 through A23 show that we cannot reject equality for most outcomes among registered users who placed an offer, regardless of baseline flood risk.²⁷


Table A25 reports Romano–Wolf adjusted p-values for the offer estimates. The high-risk estimate in Figure 7(b) carries an adjusted p-value of 0.08, and the subgroup estimates in Figure 8 are not significant at conventional levels once adjusted. As discussed in Section 3.2, this correction is conservative in our setting because the offer sample shrinks sharply as users move down the funnel and because the subgroup tests all bear on a single underlying hypothesis. We nonetheless read the offer results as strong suggestive evidence rather than

²⁷Figures A20 (a) and (b) illustrate additional trade-offs used to reduce flood exposure. Figure A20 (a) shows that treated users who browsed high-risk properties pre-experiment were 35 percentage points less likely to place an offer on waterfront properties. Figure A20 (b) shows no statistically significant effect on the square footage of properties bid on.

as independently conclusive, and we place correspondingly more weight on the search estimates, which do survive adjustment (Table A24), and on the market-level evidence in Section 5.

Figure 7: CATE on Offers

Change relative to Control



Note: Vertical lines show 95% confidence intervals. In both panels, each point reports the corresponding estimate of $\beta_1 + \beta_3$ from the difference-in-differences version of equation 1. For these outcomes, we use difference-in-differences specifications due to limited pre-experiment observations. The percentage shown next to each point in panel (b) is based on the corresponding estimate from the same specification re-estimated with the natural log of the outcome as the dependent variable, transformed as $(\exp(\beta_1 + \beta_3) - 1) \times 100$. The x-axis shows users' average search pre-experiment flood-score category. Figure (b) estimates are with city fixed effects of the property.


Our analysis also examines how property type relates to offer behavior. Figures 8 (a)–(e) show that the reduction in the flood scores of offers among high-risk users is driven by waterfront and FEMA-designated high-risk properties, as well as waterfront properties not located on the coast. These patterns further support the “new news” hypothesis.

We examine potential trade-offs associated with shifting high-risk users to lower-risk properties. Among them, the flood score of offered properties declines, and Table A26 shows changes in property size and a higher likelihood of submitting offers in zip codes different from those most frequently browsed at baseline. However, given the sample size, these estimates are imprecise. We also analyze whether these changes affect exposure to other climate risks. Using fire, heat, and wind scores from the First Street Foundation, we find no evidence that users substitute into higher-risk properties (Figure A21). We find statistically significant reductions in heat exposure among users with low and high flood-risk scores at baseline.


We then test whether political ideology influences the average treatment effect of flood-risk information. We estimate conditional average treatment effects based on whether the user's zip code voted for Biden or Trump in the 2020 presidential election. The estimated

Figure 8: CATE on the Flood Score of an Offer for Registered Users
Change relative to Control


(a) Without Waterfront Search at Baseline




(b) With Waterfront Search at Baseline




(c) Without FEMA Risk Search at Baseline



(d) With FEMA Risk Search at Baseline



(e) With Waterfront and Without Coastal Search at Baseline



Note: Vertical lines show 95% confidence intervals. In these panels, each point reports the corresponding estimate of $\beta_1 + \beta_3$ from the difference-in-differences version of equation 1. For these outcomes, we use difference-in-differences specifications due to limited pre-experiment observations. The percentage shown next to each point is based on the corresponding estimate from the same specification re-estimated with the natural log of the outcome as the dependent variable, transformed as $(\exp(\beta_1 + \beta_3) - 1) \times 100$. The x-axis shows users' average search pre-experiment flood-score category. Users who did not browse any waterfront or FEMA Risk property before the experiment are classified as "without" waterfront or FEMA Risk search at baseline. A property is classified as being located on the coast when its geographic coordinates (latitude and longitude) are 200 meters or less from the nearest shoreline. Estimates are with city-fixed effects for the property.

effects are similar in direction and magnitude across Blue and Red zip codes for both search and offer outcomes (Figures A22 (a) and (b)). We also examine whether responses to flood-risk information vary with local flood experience. We define a county flood shock as any flood event occurring in the prior 7 days or in the prior year. The 7-day estimates are not statistically different from zero, while the one-year estimates show statistically significant heterogeneity for users with high risk baseline searches (Figures A23(a)-(d)).

Closing. We also report the probability that an individual closes on a home and the flood score of the property on which they close, stratified by baseline average flood-score search category and relative to the control group. Figure A24(a) shows the effect on the probability of closing across baseline risk groups, while Figure A24(b) shows the effect on the flood score of the property closed on. The pattern in panel (b) is qualitatively similar to that for offers, with the largest reduction among users who browsed high-risk properties at baseline. The closing sample is an order of magnitude smaller than the offer sample and the estimates are correspondingly imprecise, so we do not draw strong conclusions from them.

5 Real Estate Market Price Responses to Property Specific Flood Risk Information

While no home buyer gains utility from owning a home at risk of flooding, the population differs in their willingness to pay for such a home. Higher-income individuals have greater economic capacity to avoid such risky homes. More risk loving people and those with the ability and economic capacity to upgrade their homes will be more likely to bid for risky homes than risk averse people who do not want to invest the time and effort in upgrading a home (Shogren and Stamland, 2002). In a setting where buyers and sellers have complete information about climate risks, these risks will be capitalized into the home's sales price. Hedonic real estate regression techniques can be used to recover the marginal value of the home's attributes. Recent papers have adopted this strategy to estimate flood-risk compensation (Ortega and Taspinar, 2018; Bernstein et al., 2019). Our experiment's individual level variation in access to environmental risk information allows us to take the next step here to investigate how different people engage with such information.

In this section, we study how property-specific revelation of flood risk affects the housing hedonic gradient. An ideal field experiment to answer this question would randomize the flood score at the property level (not at the individual searcher level) and vary the treatment saturation across markets, as in Crépon et al. (2013). This exact design was not possible for our field experiment. We therefore employ three complementary identification strategies, and we present them starting from the market as a whole and working back toward the ex-

periment itself.

Section 5.1 uses the full MLS record from 2015 through August 2025 to ask how the price gap between higher- and lower-risk homes evolved once flood-risk information became broadly available. This is the market-level object of interest, and it is where the magnitudes we would defend come from. Section 5.2 narrows to idiosyncratic post-2021 revisions in the displayed flood score of the same property, holding the physical asset fixed while varying only the information buyers see. Section 5.3 returns to the experiment itself and exploits the random variation in treatment saturation across individual listings; this is the cleanest source of exogenous variation available to us, but by construction it identifies an effect only over the 0–8% of the market our experiment reached, and we treat it as a building block rather than as the headline. Section 5.4 closes by combining these estimates into a back-of-the-envelope calculation of the implied revaluation of transacted properties in higher-risk areas. Table 1 summarizes the three approaches, their identifying assumptions, and the estimates they deliver.

5.1 Market-wide price responses as flood-risk information diffused

We analyze housing transactions using data from the Multiple Listing Service (MLS) from 2015 through August 2025. The dataset covers nearly all residential property sales and includes detailed property characteristics. We complement these data with *First Street Foundation's* (FSF) flood-risk rating available at the time of sale to assess whether the relationship between risk categories and transaction outcomes changed after flood-risk information became publicly available in 2021.²⁸ We estimate the following specification:

$$Y_i = \alpha + \sum_{t \neq 2019} \beta_t D_{t(i)} + \sum_{f \neq 1} \gamma_f F_{f(i)} + \sum'_{t,f} \delta_{tf} D_{t(i)} F_{f(i)} + \mathbf{X}'_i \boldsymbol{\sigma} + \mu_{z(i)} + \lambda_{m(i)} + \varepsilon_i, \quad (2)$$

Where $\sum'_{t,f}$ denotes summation over all (t, f) pairs excluding the baseline cell ($t = 2019, f = 1$). Y_i is the difference between a property's sale price and its listing price. $D_{t(i)}$ are year-of-sale indicators (2015–2025), with 2019 omitted, and $F_{f(i)}$ are indicators for a property's flood-risk

²⁸One salient policy change coincides with this window. FEMA began phasing in Risk Rating 2.0 in October 2021, repricing National Flood Insurance Program premiums toward property-level expected loss and therefore raising premiums disproportionately in higher-risk locations (Mulder and Kousky, 2023). To the extent that expected premium changes are capitalized into sale prices, the post-2021 estimates in this subsection should be read as the joint effect of information dissemination and insurance repricing rather than of information alone. Two features of the evidence suggest repricing is not the whole story. First, the estimated discounts are graded by the First Street flood-factor category, which is not the basis on which NFIP premiums are set. Second, the within-property reclassification estimates in Section 5.2 isolate variation in the *displayed* score while holding the property's physical risk, and therefore its actuarial premium, fixed.


category $f \in \{1, 2, 3\}$, where 1 denotes low risk (omitted), 2 moderate risk, and 3 high risk. The interaction coefficients δ_{tf} capture deviations in outcomes for each risk category and sale year relative to low-risk properties sold in 2019. Thus, $\beta_t + \delta_{tf}$ represents the total change in outcomes for flood-risk category f in year t relative to 2019.

The vector $\mathbf{X}_i$ includes the following property characteristics: number of bedrooms, year built, number of parking spaces, natural logarithm of square feet, and indicators for new construction, garage, and property type. $\mu_{z(i)}$ and $\lambda_{m(i)}$ are ZIP-code and month-of-sale fixed effects, respectively. Standard errors are clustered at the ZIP-code level. As a robustness check, we also estimate a version of equation (2) that absorbs ZIP-by-year fixed effects to examine whether our estimated effects are not driven by ZIP-level or geography-level time trends.

Figure 9(a) and (b) plot the estimated coefficients $\beta_t + \delta_{tf}$ for each flood-risk category relative to 2019 (see *Table A28* for values). Starting in 2021, when flood-risk information became publicly accessible, medium- and high-risk properties began to present declines in their sale-to-list-price spread relative to low-risk properties. By 2025, the decline is steepest for high-risk properties. In Figure 9(a), the sale-to-list-price spread for high-risk properties in 2025 is \$37,783 lower than in 2019. In Figure 9(b), which includes zip-code-by-year fixed effects, high-risk properties have a sale-to-list-price spread that is \$7,647 lower than that of low-risk properties sold in the same zip code and year, around 1.7% of the overall average listing price.

Two pre-2019 interaction coefficients for high-risk properties in column (2) of *Table A28* are statistically distinguishable from zero (\$902 in 2015 and $-\$957$ in 2018). We report them for transparency. They are roughly an order of magnitude smaller than the 2024 and 2025 estimates and do not trend in a common direction, so we do not read them as evidence of a systematic pre-trend, but they do indicate that year-to-year differences of a few hundred dollars between risk categories are within the noise of this specification.

Figure 9: Deviations from 2019 in the (*Sale – Listing Price*) Spread



Note: Vertical lines show 95% confidence intervals. Estimates with zip code fixed effects. Panel (a) has zip code, year, and month fixed effects, whereas panel (b) contains zip code-by-year fixed effects, as well as month fixed effects.

5.2 Exogenous changes to a property's flood scores

To test whether these market responses reflect information rather than compositional effects, we use *idiosyncratic reclassifications* in FSF flood-risk scores after 2021. In a small number of cases, Redfin received updated scores from FSF for technical reasons unrelated to physical risk. Among properties sold more than once, these reclassifications enable comparisons of the same property before and after its flood-risk score was displayed.

We focus on two types of reclassifications: (1) properties sold at least once as low risk and later as high risk, and (2) properties sold at least once as high risk and later as low risk. In total, we had 1,695 unique properties for the former reclassifications and 2,268 for the latter. The corresponding comparison groups are properties that sold at least twice during the period, whose flood-risk score never changed—low-risk properties for (1) and high-risk properties for (2). We estimate the following within-property specification:

$$Y_{it} = \alpha + \beta_1 R_{it} + \mathbf{X}'_{it} \boldsymbol{\sigma} + \mu_i + \kappa_t + \lambda_m + \varepsilon_{it}, \quad (3)$$

Where Y_{it} is the difference between the sale and listing price for property i at sale t . R_{it} is a reclassification indicator taking one at the first sale following a change in the property's displayed flood-risk score and at all later sales, and zero for all earlier sales. For properties whose displayed flood-risk score never changes, it is zero for all sales. The vector $\mathbf{X}_{it}$ includes

the same property characteristics as in equation (2). μ_i , κ_t , and λ_m are property, year-of-sale, and month-of-sale fixed effects, respectively, and standard errors are clustered at the property level.

We estimate two separate specifications, one that includes properties reclassified from high to low risk along with those that remained high risk throughout, and another that includes properties reclassified from low to high risk along with those that remained low risk (see *Table A27* for values). Properties reclassified from high to low risk show an average increase in the sale–listing price spread of about \$22,270 (17238, 27302). When a property’s classification moved from low to high risk, the spread declined by about -\$5,319 (−9933, −705). Both estimates are statistically significant at conventional levels (*Table A27*). Relative to the average listing price in the corresponding samples, these values suggest that moving from high to low or from low to high changes house prices by around 1.2% to 4.85%.

5.3 Exogenous treatment saturation across homes during the field experiment

The two preceding subsections draw on variation that arises outside the experiment. We now return to the experiment itself, which supplies the cleanest exogenous variation we have but over a much smaller slice of the market.

For every home sold during the experimental period, a random fraction of users were assigned to the treatment and control groups. This variation is due to the random assignment of treatment and control groups in small samples. For example, suppose that 50 Redfin home buyers randomly selected in the treatment group chose to click on 14 Elm Street, Belmont, MA 02478 during the experiment (the last three months of 2020). Suppose that during that time, 100 Redfin home buyers chosen at random to be in the control group also looked at that home. This means that 50/150 of the Redfin searchers randomly selected for the experiment knew the home’s flood score during the study period — we will call this the treated fraction.²⁹

Our dependent variable is the difference between the sales price and the list price. This variable reflects the “new news” associated with the home (Bajari et al., 2012). The First Street Foundation flood score treatment information was an unanticipated event that could not be incorporated into the list price. In our econometric model below, the variable, *Frac_Treat*, represents the ratio of the count of people in the treatment group who visited this property divided by the total count of all Redfin searchers who visited this property. We estimate the

²⁹Let’s also remember that Redfin allocated around 41% of its total traffic to this experiment. Thus, by definition, 59% of its monthly traffic did not know about the flood score. In the 14 Elm Street example, the 150 experiment users represent 41% of Redfin’s roughly 366 searchers for this house, so 50/366 of Redfin’s searches for this home knew its flood score. If Redfin has a 20 percent market share, the market-wide exposure to treatment for this home is 50/1,829, or about 3%.

following regression:


$$y_{pt} = \beta_0 + \beta_1 F_p + \beta_2 \text{Frac_Treat}_p + \beta_3 (F_p * \text{Frac_Treat}_p) + \mathbf{X}'_{pt} \boldsymbol{\sigma} + \lambda_z + m_t + u_{pt} \quad (4)$$

where, y_{pt} , represents the dollar spread between the sale price and the property listing price, p , sold in the month, t , of the experimental period. F_p shows the property's flood risk category (i.e., low, medium, and high), whereas, Frac_Treat , represents the percentage of people who viewed the property and were in the treatment group as a fraction of all Redfin viewers. $\mathbf{X}_{pt}$ are characteristics of the property, including total views per property, total bedrooms, new construction, month fixed effects, year built, parking spaces, a garage indicator, and log square footage. λ_z and m_t represent zip code- and month-fixed effects, respectively. u_{pt} are the residuals. Standard errors are clustered at the zip code level.


Figure 10 (a) shows the distribution of the variable, Frac_Treat , i.e., the percentage of people who viewed the property and were in the treatment group as a fraction of all Redfin viewers, which follows a normal distribution. Figure 10 (b) shows estimates of $\beta_2 + \beta_3$ from estimator 4, that is, the change in the spread per one percentage-point increase in treatment intensity within each flood-risk category. We did not find a statistically significant effect when estimating equation 4 on the full sample of listings sold during the experimental period.

Figure 10: Association Between Treatment Exposure Intensity and the (*Sale - Listing Price*) Spread

(a) Distribution of Listing Views Conducted by Treated Users per Listing



(b) All Listings
Change per 1 point increase



Note: For panel (b), vertical lines show 95% confidence intervals. Each point reports the corresponding estimate of $\beta_2 + \beta_3$ from equation 4. The percentage shown next to each point is based on the corresponding estimate from the same specification re-estimated using the natural log of the ratio of sale price to listing price as the dependent variable, transformed as $(\exp(\beta_2 + \beta_3) - 1) \times 100$. The x-axis shows the flood-score category of the property.

However, when we divide the listings into two groups based on their characteristics —being on the waterfront or being in a FEMA risk zone —we observe that our intensity treatment


variable, referred to as `Frac_Treat`, has an impact on the difference between the sale price and the listing price. Figure 11 illustrates how a 1 percentage point increase in the intensity treatment variable affects the spread between the sale price and the listing price for properties that are *not situated* on the waterfront (see Figure 11 (a)), as well as for properties that are *not considered* risky by FEMA (see Figure 11 (b)), stratified by the listing’s flood risk score. For both instances, properties considered *highly* risky by First Street Foundation incurred a price penalty as the percentage of treated users viewing the listing in Redfin increased.

A one percentage-point increase in the share of views conducted by treated users led to negative penalties of -\$89 and -\$55 for highly risky properties not on the waterfront and not considered risky by FEMA, respectively. In other words, going from 0% (pre-experiment beliefs) to 100% (every user on Redfin having the FSF flood score) in our variable of interest leads to a price penalty of -\$8,900 (1.9% of average listing prices) and -\$5,500 (1.2% of average listing prices) under list price among severely risky properties not on the waterfront and not considered risky by FEMA, in that order. These results suggest that the intervention influenced risk expectations for properties that were either not perceived as risky (i.e., not on the waterfront) or not designated as risky by a government institution (i.e., not considered risky by FEMA).

We obtain these hedonic estimates based on a thought experiment in which we increase the percentage of people with flood information from 0% to 100%. However, only roughly 40% of Redfin users were in the experiment, and Redfin had 20% market share at the time, so the range from 0 to 100 is actually 0% to 8% of the market receiving this information. We must make some assumptions about the transition from partial to general equilibrium to assess what would happen if all consumers had the information (Banzhaf, 2021). Extrapolating linearly to full market coverage would imply an effect of roughly \$111,250 on non-waterfront high-flood-risk homes, or 24% of the average listing price. We report this for completeness but treat it as an upper bound rather than an estimate and do not rely on it; for magnitudes we prefer the market-level estimates in Sections 5.1 and 5.2.³⁰

³⁰There are two reasons to read the extrapolation as a bound. First, it projects a within-experiment slope estimated over 0–8% of the market onto full coverage, and there is no reason to expect linearity over that range; if anything, one would expect diminishing marginal effects as the marginal informed buyer becomes less pivotal to the auction outcome. Second, it sits at or above the top of the range implied by the present discounted value of expected future flood damages. FEMA reports an average NFIP flood damage cost of \$66,000 between 2016 and 2022. The “high” Flood Factor band spans 30-year cumulative flood probabilities from roughly 27% to nearly 100%, implying annual hazard rates of about 1.0% and 7.4%. Converting the cumulative probability to an annual hazard as $p = 1 - (1 - \pi_{30})^{1/30}$ and computing $\$66,000 \times p \times \left[1 - (1 + r)^{-30}\right]/r$ gives a present value of roughly \$11,000 to \$14,000 at the bottom of the band and \$75,000 to \$96,000 at the top, depending on whether one discounts at 5% or 3%. This abstracts from multiple flood events over the horizon, from any premium buyers place on avoiding disruption beyond the repair cost, and from expected increases in insurance premia, all of which push the implied penalty up.

Figure 11: CATE of an Increase in Exposure Intensity on the (*Sale - Listing Price*)



Note: Vertical lines show 95% confidence intervals. Each point reports the corresponding estimate of $\beta_2 + \beta_3$ from equation 4. The percentage shown next to each point is based on the corresponding estimate from the same specification re-estimated using the natural log of the ratio of sale price to listing price as the dependent variable, transformed as $(\exp(\beta_2 + \beta_3) - 1) \times 100$. The x-axis shows the flood-score category of the property. Properties not located on the waterfront or not classified as FEMA high risk are labeled as not waterfront or not FEMA risk.

We did not observe a statistically significant effect of an increase in our treatment intensity variable on the difference between the *sale price* and *listing price* for properties located on the waterfront versus those classified as risky by FEMA (as shown in Figure A25), suggesting that flood risk was already priced for those properties. Finally, we have information on whether an investor purchased a property, the time on the market (to examine some elements of allocative efficiency), and the loan value.³¹ In Figure A26, we find that for low flood risk homes, moving from 0 to 100% treated Redfin users led to a 2.5 percentage point increase in the probability of an investor buying the property (over a baseline of 9.76% in the control group), consistent with evidence that investor buyers are more attentive to flood risk than owner-occupiers (Bernstein et al., 2019; Hino and Burke, 2021). We do not find any impact of the treatment intensity variable on the time on the market (Figure A27), the loan value of the property's mortgage (Figure A28), and the listing price of non-FEMA-risky and non-waterfront properties (Figure A29).

Our interpretation of these results is that the sales price represents the outcome of an auction process. In high-flood-risk areas, where this information is publicly available to many searchers, treated individuals are less likely to bid aggressively for the home. The winning bidder for the home will pay less for it. By subtracting the asking price, we standardize the

³¹An investor is defined as any buyer whose name on the listing includes at least one of the following keywords: LLC, Inc, Trust, Corp, Homes; or any buyer whose ownership code on a purchasing deed includes at least one of the following keywords: association, corporate trustee, company, joint venture, or corporate trust.

dependent variable.³²

Robustness Checks. As a robustness check, we repeat the exposure–intensity analysis using listings sold between September 1 and September 30, before the experiment began. For these listings, we reconstruct the same treatment-intensity measure by assigning Redfin users in the experiment to treatment and control groups in the same way as during the study period, and then computing the share of treated users who viewed each property. We used views occurring in the two months prior to the sale date to mirror the construction of the main exposure measure.

Figure A30 (a) shows the distribution of the exposure variable, which is centered and symmetric, similar to the distribution in the experimental period. Figure A30 (b) shows the corresponding estimates of our main regression, finding no statistically significant effects between the exposure measure and the sale–listing price spread. We further use our pre-experiment exposure variable to test whether waterfront listings and FEMA-status listings were differentially affected. As shown in Figures A31 (a) to (d), and in contrast to the experimental period, we do not find statistically significant effects.

5.4 Aggregate implications: a back-of-the-envelope calculation

The three strategies above, summarized in Table 1, establish that flood-risk information moved transaction prices for higher-risk homes. We now put a magnitude on that, using the market-wide estimates of Section 5.1, which are the only ones estimated on the full transaction record rather than on the experimental sample.

To quantify the aggregate market impact of flood-risk information, we combine our estimated price effects with transaction-level sales data by year. For each year and risk category, we estimate the implied dollar revaluation of properties sold in that year using the corresponding interaction coefficient, δ_{tf} , from Figure 9(a), multiplied by the number of transactions in that category and year. We use δ_{tf} rather than $\beta_t + \delta_{tf}$ because the latter also embeds the common movement in the sale–listing spread experienced by low-risk properties in the same year, which reflects market-wide conditions rather than flood-risk information. We also express these effects as percentages of the corresponding average list price within each risk category. This calculation provides a simple back-of-the-envelope estimate of the change in the value of transacted properties in higher flood-risk areas following the public release of

³²Given that this is a short run experiment that the seller was unaware was taking place, the asking price is likely independent of the flood score.

this information.³³

Between 2023 and 2025, our calculations suggest that housing markets in higher-risk areas experienced a revaluation once flood-risk information became widely available. In 2023, the estimated monthly revaluation of transacted properties is approximately \$73 million for medium-risk properties and \$206 million for high-risk properties. These amounts rise to \$107 million and \$314 million in 2024, and to \$137 million and \$363 million, respectively, using data through August 2025. Expressed relative to list prices, these effects correspond to 0.59 percent and 1.88 percent for medium- and high-risk properties in 2023, 0.64 percent and 2.91 percent in 2024, and 0.97 percent and 3.80 percent in 2025, respectively.

These figures are the upper end of a defensible range. Column (2) of Table A28, which absorbs ZIP-by-year fixed effects and therefore differences out local price trends, yields substantially smaller differentials for high-risk properties: −\$3,521 in 2023 and −\$7,647 in 2025, implying monthly revaluations of roughly \$55 million and \$90 million, or 0.50 percent and 0.94 percent of list price. We report both because the two specifications bracket the extent to which the estimated discounts reflect property-level flood risk as opposed to the geography in which higher-risk properties are concentrated. The conservative bound is the one we would defend; the wider bound is informative about how much of the raw gap is geographic.

³³For our back-of-the-envelope calculation we use δ_{tf} , the estimated decline in the sale–listing price spread for medium- and high-risk properties *relative to low-risk properties sold in the same year*. From column (1) of Table A28, the fitted values are −\$5,303 (medium), −\$15,288 (high), and −\$2,094 (low) in 2023; −\$6,908, −\$25,017, and −\$3,274 in 2024; and −\$12,651, −\$37,783, and −\$6,780 in 2025. The corresponding differentials are therefore \$3,209 and \$13,194 in 2023, \$3,634 and \$21,743 in 2024, and \$5,871 and \$31,003 in 2025. We combine these with MLS transaction counts of 273,358 medium-risk and 187,425 high-risk properties in 2023, 354,082 and 173,335 in 2024, and 186,793 and 93,782 in 2025. Annual totals are divided by twelve for 2023 and 2024, and by eight for 2025, since the 2025 data cover January through August. To express the effects in percentage terms we divide the differentials by the corresponding average list prices of \$541,779 and \$701,908 for medium- and high-risk properties in 2023, \$566,164 and \$748,421 in 2024, and \$606,100 and \$815,843 in 2025.

Table 1: Three Approaches to Estimating the Price Response to Flood-Risk Information

	(1) Market-wide diffusion Section 5.1, eq. (2)	(2) Within-property reclassification Section 5.2, eq. (3)	(3) Experimental saturation Section 5.3, eq. (4)
Source of variation	Calendar timing of the public release of FSF scores across listing platforms from 2021, interacted with a property’s risk category	Idiosyncratic FSF revisions to the score <i>displayed</i> for a given property, driven by model updates rather than physical risk	Random share of a listing’s Redfin viewers assigned to see the flood score during the 12-week experiment
Sample and period	All MLS residential sales, 2015–August 2025	Repeat sales with a reclassification: 1,695 low→high and 2,268 high→low, plus never-reclassified properties of the same category	Listings sold during the experimental period, October 2020–January 2021
Unit of variation	Risk category × sale year	Property	Listing
Identifying assumption	Absent the information release, the sale–list gap between higher- and lower-risk homes would have followed a common path	Score revisions are orthogonal to the property’s physical risk and to unobserved determinants of the spread	Treated share is random conditional on total views; sellers were unaware, so the list price is independent of the score
Implied change in home value	≈1.7% (ZIP × year specification)	1.2% to 4.85%	1.2% to 1.9% at full saturation
Market the design speaks to	The full U.S. transaction record	Repeat-sale properties whose displayed score changed	The 0–8% of the market the experiment reached
Falsification	Pre-2021 years show no systematic divergence by risk category	Never-reclassified properties in the same risk category serve as controls	Placebo on September 2020 listings, before the experiment, yields a null

Note: The outcome in all three specifications is the dollar spread between the sale price and the listing price. The three approaches are ordered as they appear in the text, from the broadest sample to the cleanest variation. The percentages in the “implied change in home value” row are all expressed relative to the overall average listing price, so they are comparable across columns; the percentages in the back-of-the-envelope calculation of Section 5.4 are instead expressed relative to the average listing price *within* each risk category, and are therefore not directly comparable to this row.

6 Conclusion

A majority of American adults live in owner-occupied housing, often their main asset. Rising greenhouse gas concentrations create new place-based risks for such properties. As real estate platforms integrate detailed climate risk maps, they provide valuable information to home buyers and may accelerate climate adaptation by influencing purchase decisions. Using Redfin’s platform, we studied a randomized field experiment tracking how individuals searched, toured, and bought homes. Flood risk information changed which properties users searched for and bid on, and those shifts were accompanied by measurable changes in transaction prices for high-risk homes. What moved was the composition of the properties users pursued rather than the extensive-margin decisions to tour, bid, or close, which were unaffected. These findings support the “new news” hypothesis: consumers were previously unaware of flood risks but updated their beliefs once given clear, trusted information, advancing

climate adaptation. The three-month experiment affected the price of high-risk properties, and in the years since, as this information has been adopted across the major listing platforms, the gap in sale-to-list price spreads between high- and low-risk homes has widened.

The market is now responding to spatial risk information, showing trust in these tools. Competition among risk modelers and improvement in model accuracy will shape adaptation. As trusted information spreads, developers and insurers will better price and mitigate risks, promoting resilience on the supply side—stimulating a new market for climate-adaptive goods and accelerating overall adaptation.

References

- Armal, S., Porter, J. R., Lingle, B., Chu, Z., Marston, M. L. and Wing, O. E. (2020), ‘Assessing property level economic impacts of climate in the us, new insights and evidence from a comprehensive flood risk assessment tool’, *Climate* **8**(10), 116.
- Athey, S., Tibshirani, J. and Wager, S. (2019), ‘Generalized random forests’, *The Annals of Statistics* **47**(2), 1148–1178.
- Athey, S. and Wager, S. (2021), ‘Policy learning with observational data’, *Econometrica* **89**(1), 133–161.
- Auffhammer, M. (2018), ‘Quantifying economic damages from climate change’, *Journal of Economic Perspectives* **32**(4), 33–52.
- Bajari, P., Fruehwirth, J. C., Timmins, C. et al. (2012), ‘A rational expectations approach to hedonic price regressions with time-varying unobserved product attributes: The price of pollution’, *American Economic Review* **102**(5), 1898–1926.
- Bakkensen, L. A. and Barrage, L. (2022), ‘Going underwater? flood risk belief heterogeneity and coastal home price dynamics’, *The Review of Financial Studies* **35**(8), 3666–3709.
- Balboni, C. (2025), ‘In harm’s way? infrastructure investments and the persistence of coastal cities’, *American Economic Review* **115**(1), 77–116.
- Baldauf, M., Garlappi, L. and Yannelis, C. (2020), ‘Does climate change affect real estate prices? only if you believe in it’, *Review of Financial Studies* **33**(3), 1256–1295.
- Banzhaf, H. S. (2021), ‘Difference-in-differences hedonics’, *Journal of Political Economy* **129**(8), 2385–2414.
- Barreca, A., Clay, K., Deschenes, O., Greenstone, M. and Shapiro, J. S. (2016), ‘Adapting to climate change: The remarkable decline in the us temperature-mortality relationship over the twentieth century’, *Journal of Political Economy* **124**(1), 105–159.

- Bates, P. (2023), 'Fundamental limits to flood inundation modelling', *Nature Water* pp. 1–2.
- Bates, P. D., Quinn, N., Sampson, C., Smith, A., Wing, O., Sosa, J., Savage, J., Olcese, G., Neal, J., Schumann, G. et al. (2021), 'Combined modeling of us fluvial, pluvial, and coastal flood hazard under current and future climates', *Water Resources Research* **57**(2), e2020WR028673.
- Beltrán, A., Maddison, D. and Elliott, R. J. (2018), 'Is flood risk capitalised into property values?', *Ecological Economics* **146**, 668–685.
- Bernstein, A., Gustafson, M. T. and Lewis, R. (2019), 'Disaster on the horizon: The price effect of sea level rise', *Journal of Financial Economics* **134**(2), 253–272.
- Bilal, A. and Rossi-Hansberg, E. (2023), 'Anticipating climate change across the united states'.
- Bin, O. and Landry, C. E. (2013), 'Changes in implicit flood risk premiums: Empirical evidence from the housing market', *Journal of Environmental Economics and Management* **65**(3), 361–376.
- Botzen, W. W., Kunreuther, H. and Michel-Kerjan, E. (2015), 'Divergence between individual perceptions and objective indicators of tail risks: Evidence from floodplain residents in new york city', *Judgment and Decision Making* **10**(4), 365–385.
- Burlig, F., Jina, A., Kelley, E. M., Lane, G. V. and Sahai, H. (2024), Long-range forecasts as climate adaptation: Experimental evidence from developing-country agriculture, Technical report, National Bureau of Economic Research.
- Connell-Price, L. and Mulder, P. (2024), 'Does good news wash out the bad? belief updating and climate risk information', *Belief Updating and Climate Risk Information (March 1, 2024). Consumer Financial Protection Bureau Office of Research Working Paper (24-04)*.
- Crépon, B., Duflo, E., Gurgand, M., Rathelot, R. and Zamora, P. (2013), 'Do labor market policies have displacement effects? evidence from a clustered randomized experiment', *The Quarterly Journal of Economics* **128**(2), 531–580.
- Deschenes, O., Greenstone, M. and Shapiro, J. S. (2017), 'Defensive investments and the demand for air quality: Evidence from the nox budget program', *American Economic Review* **107**(10), 2958–2989.
- Desmet, K., Kopp, R. E., Kulp, S. A., Nagy, D. K., Oppenheimer, M., Rossi-Hansberg, E. and Strauss, B. H. (2021), 'Evaluating the economic cost of coastal flooding', *American Economic Journal: Macroeconomics* **13**(2), 444–86.
- Downey, M., Lind, N. and Shrader, J. G. (2023), 'Adjusting to rain before it falls', *Management Science* **69**(12), 7399–7422.
- First Street Foundation (2020), 'First street foundation flood model: technical methodology document'.
- FirstStreet (2025), High water, high stakes: Fema, flood risk, and the nfip, Technical report, First Street.

- Fried, S. (2022), 'Seawalls and stilts: A quantitative macro study of climate adaptation', *Review of Economic Studies* **89**(6), 3303–3344.
- Gallagher, J. (2014), 'Learning about an infrequent event: evidence from flood insurance take-up in the united states', *American Economic Journal: Applied Economics* pp. 206–233.
- Gibson, M. and Mullins, J. T. (2020), 'Climate risk and beliefs in new york floodplains', *Journal of the Association of Environmental and Resource Economists* **7**(6), 1069–1111.
- Giglio, S., Maggiori, M., Rao, K., Stroebel, J. and Weber, A. (2021), 'Climate change and long-run discount rates: Evidence from real estate', *Review of Financial Studies* **34**(8), 3527–3571.
- Goldsmith-Pinkham, P., Gustafson, M. T., Lewis, R. C. and Schwert, M. (2023), 'Sea-level rise exposure and municipal bond yields', *The Review of Financial Studies* **36**(11), 4588–4635.
- Gourevitch, J. D., Kousky, C., Liao, Y., Nolte, C., Pollack, A. B., Porter, J. R. and Weill, J. A. (2023), 'Unpriced climate risk and the potential consequences of overvaluation in us housing markets', *Nature Climate Change* **13**(3), 250–257.
- Greenstone, M. (2017), 'The continuing impact of sherwin rosen's "hedonic prices and implicit markets: product differentiation in pure competition"', *Journal of Political Economy* **125**(6), 1891–1902.
- Harrison, D., T. Smersh, G. and Schwartz, A. (2001), 'Environmental determinants of housing prices: the impact of flood zone status', *Journal of Real Estate Research* **21**(1-2), 3–20.
- Harrison, G. W. and List, J. A. (2004), 'Field experiments', *Journal of Economic Literature* **42**(4), 1009–1055.
- Hinkel, J., Lincke, D., Vafeidis, A. T., Perrette, M., Nicholls, R. J., Tol, R. S., Marzeion, B., Fettweis, X., Ionescu, C. and Levermann, A. (2014), 'Coastal flood damage and adaptation costs under 21st century sea-level rise', *Proceedings of the National Academy of Sciences* **111**(9), 3292–3297.
- Hino, M. and Burke, M. (2021), 'The effect of information about climate risk on property values', *Proceedings of the National Academy of Sciences* **118**(17), e2003374118.
- Hovekamp, W. P. and Wagner, K. R. H. (2023), 'Efficient adaptation to flood risk', *AEA Papers and Proceedings* **113**, 304–09.
- Hsiao, A. (2023), 'Sea level rise and urban adaptation in jakarta', *Technical Report* .
- IPCC (2021), 'Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change'.
- Ito, K. and Zhang, S. (2020), 'Willingness to pay for clean air: Evidence from air purifier markets in china', *Journal of Political Economy* **128**(5), 1627–1672.

- Jia, R., Ma, X. and Xie, V. W. (2026), 'Expecting floods: Firm entry, employment, and aggregate implications', *American Economic Journal: Macroeconomics* **18**(3), 42–76.
- Kahn, M. E. (2020), 'The climate change adaptation literature', *Review of Environmental Economics and Policy* .
- Keys, B. J. and Mulder, P. (2020), Neglected no more: Housing markets, mortgage lending, and sea level rise, Technical report, National Bureau of Economic Research.
- Kousky, C., Kunreuther, H., LaCour-Little, M. and Wachter, S. (2020), 'Flood risk and the us housing market', *Journal of Housing Research* **29**(sup1), S3–S24.
- Kousky, C. and Netusil, N. R. (2023), 'Flood insurance literacy and flood risk knowledge: Evidence from portland, oregon', *Risk Management and Insurance Review* .
- Lee, S. (2021), 'Adapting to natural disasters through better information: Evidence from the home seller disclosure requirement', *MIT Center for Real Estate Research Paper* (21/17).
- List, J. A. (2026), 'Experimental economics', *University of Chicago Press Economics Books* .
- List, J. A., Shaikh, A. M. and Vayalinkal, A. (2023), 'Multiple testing with covariate adjustment in experimental economics', *Journal of Applied Econometrics* **38**(6), 920–939.
- Marcoux, K. and Wagner, K. R. H. (2025), *Fifty Years of U.S. Natural Disaster Insurance Policy*, Springer Nature Switzerland, pp. 55–79.
- Masseti, E. and Mendelsohn, R. (2020), 'Measuring climate adaptation: Methods and evidence', *Review of Environmental Economics and Policy* .
- Molina, R. and Rudik, I. (forthcoming), 'The social value of hurricane forecasts', *American Economic Journal: Economic Policy* .
- Mulder, P. and Kousky, C. (2023), Risk rating without information provision, in 'AEA Papers and Proceedings', Vol. 113, American Economic Association, pp. 299–303.
- Murfin, J. and Spiegel, M. (2020), 'Is the risk of sea level rise capitalized in residential real estate?', *Review of Financial Studies* **33**(3), 1217–1255.
- Nordhaus, W. D. (1993), 'Reflections on the economics of climate change', *Journal of Economic Perspectives* **7**(4), 11–25.
- Ortega, F. and Taşpınar, S. (2018), 'Rising sea levels and sinking property values: Hurricane sandy and new york's housing market', *Journal of Urban Economics* **106**, 81–100.
- Porter, J. R., Shu, E., Amodeo, M., Hsieh, H., Chu, Z. and Freeman, N. (2021), 'Community flood impacts and infrastructure: Examining national flood impacts using a high precision assessment tool in the united states', *Water* **13**(21), 3125.

Redfin (2020), Redfin Annual Report, Technical report.

URL: https://d1io3yog0oux5.cloudfront.net/_0ae02c41267f78d5f5925aff80ea594/redfin/db/2206/20844/annual-report.pdf

Romano, J. P. and Wolf, M. (2010), ‘Balanced control of generalized error rates’, *The Annals of Statistics* **38**(1), 598–633.

URL: <https://doi.org/10.1214/09-AOS734>

Rosenzweig, M. R. and Udry, C. (2014), ‘Rainfall forecasts, weather, and wages over the agricultural production cycle’, *American Economic Review* **104**(5), 278–283.

Royal, A. and Walls, M. (2019), ‘Flood risk perceptions and insurance choice: Do decisions in the floodplain reflect overoptimism?’, *Risk Analysis* **39**(5), 1088–1104.

Schlenker, W. and Taylor, C. A. (2021), ‘Market expectations of a warming climate’, *Journal of Financial Economics* **142**(2), 627–640.

Seetharam, I. (2018), The indirect effects of hurricanes: Evidence from firm internal networks, Technical report, Working Paper.

Shogren, J. F. and Stamland, T. (2002), ‘Skill and the value of life’, *Journal of Political Economy* **110**(5), 1168–1173.

Shrader, J. (2023), ‘Improving climate damage estimates by accounting for adaptation’, *Available at SSRN* 3212073 .

Shrader, J. G., Bakkensen, L. and Lemoine, D. (2023), Fatal errors: The mortality value of accurate weather forecasts, Technical report, National Bureau of Economic Research.

Viviano, D., Wüthrich, K. and Niehaus, P. (2026), ‘A model of multiple hypothesis testing’, *The Review of Economic Studies* .

Wagner, K. R. (2022), ‘Adaptation and adverse selection in markets for natural disaster insurance’, *American Economic Journal: Economic Policy* **14**(3), 380–421.

Weill, J. A. (2026), ‘Flood risk mapping and the distributional impacts of climate information’, *American Economic Journal: Economic Policy* **18**(3), 275–305.

Wing, O. E., Bates, P. D., Smith, A. M., Sampson, C. C., Johnson, K. A., Fargione, J. and Morefield, P. (2018), ‘Estimates of present and future flood risk in the conterminous united states’, *Environmental Research Letters* **13**(3), 034023.

Wing, O. E., Lehman, W., Bates, P. D., Sampson, C. C., Quinn, N., Smith, A. M., Neal, J. C., Porter, J. R. and Kousky, C. (2022), ‘Inequitable patterns of us flood risk in the anthropocene’, *Nature Climate Change* **12**(2), 156–162.

A Appendix

A.1 Demand Side Hypotheses Through the Lens of a Hedonic Housing Assignment Model: Full

A parametric closed form hedonic assignment model allows us to concisely state our paper's core hypotheses. At any point in time, there is a hedonic equilibrium real estate pricing gradient that maps each home's attributes into a market price. In equilibrium, every home is occupied and given market prices nobody wants to trade homes.

To simplify the exposition, we model climate change as a one time change in the distribution of risk across geographic locations. There is a "before" and an "after" period. The assumption that "climate change" is a one time surprise allows us to sidestep the modeling challenge of incorporating expectations into a closed form hedonic assignment model.³⁴

If climate change increases the risk that a given home faces and if a buyer is unaware of this shift, then the buyer may subsequently regret her investment. If climate increases the risk that different homes face and if this knowledge is common knowledge, then in the hedonic equilibrium, those who choose to live in riskier homes will be compensated through lower prices. In our empirical work below, only a subset of buyers are informed about the new risks.

To simplify the algebra below, we focus on one attribute of housing namely flood safety (f). We assume that consumers purchase one property and the numeraire good X . Properties are assumed to vary only with respect to the local flood safety level.³⁵ Consumers only differ by their preference for flood safety (γ). It is assumed to be normally distributed with mean $\bar{\gamma}$ and standard deviation σ_γ .

We assume the supply of properties is exogenous, and their safety follows the distribution $N(\bar{h}, \sigma_h^2)$. We consider a quadratic hedonic pricing function $P(f) = \pi_0 + \pi_1 f + \pi_2 f^2$. Each home buyer's utility is given by:

³⁴We recognize that different people will have different baseline beliefs about emerging climate risks. Some may believe the FEMA risk maps as delineating risks. Others may rely on their past experience in the areas they have lived. Some may be skeptical about whether the risk of extreme events is rising. Those who are most surprised by the "new news" of the FSF nudge and those who become aware of now "known unknowns" are likely to value this forecast information the most. Future hedonic research can follow Severen et al. (2018) by integrating the "forward-looking Ricardian approach" into closed form hedonic assignment models.

³⁵We assume higher flood risk lowers utility but do not parse out the reasons as to why it causes a loss. There are at least five reasons why higher flood risk could cause a utility loss: (1) people think they will die or be physically hurt in the flood; (2) people think that flooding will destroy their home; (3) people think that flood areas will experience rising insurance prices over time; (4) people think they will lose time and resources cleaning up after flood after flood; (5) people think they will spend more money offsetting flood risk to reduce flood damage in high FSF score areas.

$$U_i = Y - P(f) + \gamma_i f := C - \pi_1 f - \pi_2 f^2 + \gamma_i f \quad (5)$$

where Y denotes income, and we collect all constant terms into C . The home buyer chooses the safety level that maximizes their utility. By setting the FOC to 0, each home buyer's chosen safety level is:

$$f_i = \frac{\gamma_i - \pi_1}{2\pi_2} \quad (6)$$

We aggregate f_i to obtain the aggregate demand density for the safety level. Based on our assumption on γ , the demand density follows $N(\frac{\bar{\gamma} - \pi_1}{2\pi_2}, \frac{\sigma_\gamma^2}{4\pi_2^2})$. In equilibrium, the demand density must match the supply density at each safety level. We match the first two moments of the demand density with those of the supply density to solve for π_1 and π_2 , given as follows:

$$\pi_1 = \bar{\gamma} + \frac{\bar{h}\sigma_\gamma}{\sigma_h} \text{ and } \pi_2 = -\frac{\sigma_\gamma}{2\sigma_h} \quad (7)$$

Then the gradient of the hedonic function is $\frac{\partial P}{\partial f} = \bar{\gamma} + \frac{\bar{h}\sigma_\gamma}{\sigma_h} - \frac{\sigma_\gamma}{\sigma_h} f$. The function is concave as consumers have a diminishing marginal willingness to pay for safety.

Proposition 1: Due to climate change, the distribution of flood safety shifts such that; $\bar{h}$ decreases to $\bar{h}'$ and σ_h^2 increases to $\sigma_h'^2$. When home buyers have full information about this new distribution, the price gradient becomes steeper at higher safety levels, in comparison to the pre-climate change world.³⁶

Consumers reevaluate the safety level of each property based on their new information on climate change. In this case, the gradient of the new hedonic function is $\frac{\partial P_{new}}{\partial f} = \bar{\gamma} + \frac{\bar{h}'\sigma_\gamma}{\sigma_h'} - \frac{\sigma_\gamma}{\sigma_h'} f$. Consider a given safety level f_0 . We can calculate the difference in price gradient before and after climate change: $\frac{\partial P_{new}}{\partial f} - \frac{\partial P_{old}}{\partial f} = (\frac{\bar{h}'}{\sigma_h'} - \frac{\bar{h}}{\sigma_h})\sigma_\gamma + (\frac{1}{\sigma_h'} - \frac{1}{\sigma_h})\sigma_\gamma f_0$. Because $\frac{1}{\sigma_h'} - \frac{1}{\sigma_h} > 0$, the difference is an increasing function of f_0 . Note that the first term is negative, which suggests that consumers may be willing to pay less for properties with low safety levels in the full information scenario.

Now suppose only $k\%$ of the population (randomly selected) know the true distribution of safety levels after climate change. We assume the z-score of each property is constant before

³⁶To simplify the discussion, we assume that climate change leads to a one time shift in the distribution of safety.

and after climate change.³⁷ That is, for a property with safety f in the post-climate change world, its pre-climate change safety level is $\bar{h} + \frac{f - \bar{h}'}{\sigma_h'} \sigma_h$. For home buyers without information on the new distribution of safety, their utility is given by:

$$U_i = C' - \pi_1' f - \pi_2' f^2 + \gamma_i (\bar{h} + \frac{f - \bar{h}'}{\sigma_h'} \sigma_h) \quad (8)$$

That is, they are likely to overestimate the safety of a property and derive more utility from living in a riskier place. This implies the following proposition.

Proposition 2: Consider two consumers with the same risk preference, and only consumer A knows the true climate change induced distribution of the safety of properties. Then consumer A bids higher for properties that are safer under climate change and bids lower for the riskier properties.

We now study the equilibrium in the partial information scenario. For consumers with true information, their chosen safety is $f_i^{\text{treat}} = \frac{\gamma_i - \pi_1'}{2\pi_2'}$. For those without true information, by setting the FOC of their utility function to zero, we solve for their choice $f_i^{\text{control}} = \frac{\gamma_i \frac{\sigma_h}{\sigma_h'} - \pi_1'}{2\pi_2'}$. The demand density is the density of $k f_i^{\text{treat}} + (1 - k) f_i^{\text{control}}$. To find the equilibrium, we again match the moments of the demand density with the moments of the supply density $N(\bar{h}', \sigma_h'^2)$. The new hedonic coefficients are given by:

$$\pi_2' = -\frac{\sigma_\gamma}{2\sigma_h'} \sqrt{k^2 + (1 - k)^2 \left(\frac{\sigma_h}{\sigma_h'}\right)^2} \text{ and } \pi_1' = k\bar{\gamma} + (1 - k)\bar{\gamma} \frac{\sigma_h^2}{\sigma_h'^2} - 2\pi_2'\bar{h}' \quad (9)$$

We can compare the price gradient with respect to safety in the partial information post-climate change world with that in the pre-climate change world.

Proposition 3: When a larger fraction of home buyers receive the new information, the increase in the price gradient is larger, especially at high safety levels.

Proposition 1 is a special case of this proposition (i.e. when $k = 1$). Intuitively, when more people know the true risk, more of them are willing to buy safer properties. The price function becomes steeper at high safety levels due to the rising demand.

Our three propositions suggest that safety is currently undervalued due to a lack of information on risk, but once that information on risk becomes available (i.e., new news), people

³⁷This assumption can be relaxed. In general, we need a one-to-one mapping function from pre-climate change safety to post-climate change safety. We choose the function that preserves the z-score for algebraic simplicity.

act on it and value it, and those with a higher preference for safety will act on it even more.³⁸

This “new news” hypothesis allows us to be a bit more precise about who is expected to change their behavior as a result of this information. This hypothesis predicts that the following consumers will be affected by the new information:

- (a) consumers searching for properties that are in-land (i.e., not coastal) but have high flood risk defined by FSF;
- (b) consumers searching for properties that are in-land (i.e., not coastal) but close to a waterfront (e.g., river, lake) and have high flood risk defined by FSF; and
- (c) consumers searching for properties that are not defined risky by FEMA but are on the coast.

For populations (a), (b), and (c), the information will revise their flood risk beliefs upward and will search and buy lower flood risk homes, as per propositions 1 and 2. We believe that the central tenet and test of this new news hypothesis is the following two groups of consumers will be less affected by the FSF flood risk information:

- (d) consumers searching for properties in already FEMA designated as high risk; and
- (e) consumers searching for properties with zero to low risk.

Population (d) should already be informed about the flood risks because of current information dissemination and because you need flood insurance to obtain a mortgage in these FEMA areas.³⁹ These are the predictions for the new news hypothesis on search and offers.

Based on our propositions and predictions, property prices will change their gradient w.r.t. flood risk. The prices on FSF high flood risk but FEMA low/no risk properties will decrease through a lack of demand (either a lower willingness to pay by consumers or a lack of offers overall that will reduce competition and reduce prices) and that the prices on FSF low risk properties will go up because of an increase in demand (either a higher willingness to pay by consumers or an increase in competition for the property that will drive up prices).

We have considered the case when everyone knows the true distribution of property safety, respectively under two states: without and with climate change. Now suppose the true state is that climate change takes place, but only $k\%$ of the population (randomly selected) know the

³⁸We have considered the case where everyone knows the true distribution of property safety, respectively, under two states: without and with climate change. However, we also predict what will happen if there are some people who do not believe in climate change. In Appendix A.1, we discuss this case.

³⁹We recognize that there still might not be full information from the FEMA high flood risk ratings, but there should be less of an effect for this population than populations (a) to (c).

true distribution of safety is $N(\bar{f}', \sigma_f'^2)$. Consumers' utility depends on their perceived safety level and the price of the property. Their perceived safety hinges upon their belief on climate change. Consider those without the true information. They have a $p\%$ probability to believe climate change is not happening and a $(1-p)\%$ to believe the opposite. They seek to maximize their expected utility. Then these consumers face the following maximization problem, where P_2 denotes the price equilibrium function in this economy without full information:

$$\text{Max}_{h_1} pU(h_0, I - P_2(h_1), \alpha) + (1-p)pU(h_1, I - P_2(h_1), \alpha) \quad (10)$$

While consumers with updated information (i.e., know the true distribution, no uncertainty) face the following:

$$\text{Max}_{h_1} U(h_1, I - P_2(h_1), \alpha) \quad (11)$$

Note that from the definition of h_1 and h_0 , we can rewrite $h_0 = \frac{h_1 \sigma_f' + \bar{f}' - \bar{f}}{\sigma_f}$, which would allow us to solve for h_1 in the first case.

The maximization problems can be solved by finding the first order conditions. From the first order conditions, the demand for safety for the two groups of consumers can be written as, where n_1'' and n_0'' are coefficients of $P_2(h)$:

$$h_0 = \frac{(\gamma_1 + \gamma_2 \alpha)(p \frac{\sigma_f'}{\sigma_f} + 1 - p) + p(\frac{\rho \sigma_f'}{\sigma_f^2} - \frac{\omega \pi_1''}{\sigma_f})(\bar{f}' - \bar{f}) + \omega(I - \pi_0'')(p \frac{\sigma_f'}{\sigma_f} + 1 - p) - \theta \pi_1''}{2\omega \pi_1''(p \frac{\sigma_f'}{\sigma_f} + 1 - p) - \rho((\frac{\sigma_f'}{\sigma_f})^2 + 1 - p)} \quad (12)$$

$$h_1 = \frac{\gamma_1 + \gamma_2 \alpha + \omega I - \omega \pi_0'' - \theta \pi_1''}{2\omega \pi_1'' - \rho} \quad (13)$$

A.2 Additional Figures and Tables

Figure A1: The Redfin app




Figure A2: The First Street Foundation flood score matrix calculation




Figure A3: Colors and Labels of Flood Scores Displayed





Figure A4: Users in the Experiment by State Where They Browsed Most at Baseline
% of Registered Users



Note: Registered users with baseline search and flood risk information.

Figure A5: Number of Registered Users Entering the Experiment




Table A1: Registered vs. Non-Registered Users During the Pre-Experiment Phase

	(1) Views	(2) Bedrooms	(3) Bathrooms	(4) Sq. Ft.	(5) List Price	(6) Flood Score	(7) HHI Zip
Registered	1.095*** (0.019)	-0.051*** (0.002)	-5.234 (5.311)	-134.505*** (26.802)	88253.723* (44913.060)	0.020*** (0.002)	-1031.519*** (4.108)
Constant	3.189*** (0.010)	3.472*** (0.001)	9.445 (6.776)	2459.985*** (25.407)	731747.896*** (958.171)	1.736*** (0.001)	8202.836*** (1.431)
Obs.	19,091,234	18,503,273	18,592,940	18,556,087	18,117,389	18,848,928	19,091,234

Note: *** p<0.001, ** p<0.01, * p<0.05. Standard Errors Clustered at the User Level.

Table A2: Registered vs. Non-Registered Users During the Pre-Experiment Phase

	(1) Year Built	(2) Walk Score	(3) Transit Score	(4) Bike Score
Registered	-1.258*** (0.044)	1.040*** (0.045)	0.742*** (0.041)	0.743*** (0.035)
Constant	1974.039*** (0.016)	36.191*** (0.016)	34.221*** (0.015)	41.831*** (0.012)
Obs.	18,558,011	18,257,862	11,321,809	18,090,938

Note: *** p<0.001, ** p<0.01, * p<0.05. Standard Errors Clustered at the User Level.

Table A3: Balance Tests for Registered Users

	(1) Views	(2) Bedrooms	(3) Bathrooms	(4) Sq. Ft.	(5) List Price	(6) Flood Score	(7) HHI Zip
Treatment	0.016 (0.022)	-0.004 (0.003)	0.067 (0.141)	-1.324 (11.835)	55521.173 (88075.090)	-0.002 (0.004)	-0.001 (7.484)
Constant	4.524*** (0.013)	3.423*** (0.002)	2.704*** (0.090)	2319.727*** (8.095)	783014.257*** (20983.674)	1.752*** (0.003)	7046.569*** (5.311)
Obs.	3,756,792	3,702,998	3,700,374	3,686,092	3,695,419	3,721,836	3,756,792

Note: *** p<0.001, ** p<0.01, * p<0.05. Standard Errors Clustered at the User Level.

Table A4: Balance Tests for Registered Users

	(1) Year Built	(2) Walk Score	(3) Transit Score	(4) Bike Score
Treatment	-0.032 (0.079)	-0.004 (0.081)	-0.055 (0.074)	0.043 (0.063)
Constant	1972.939*** (0.055)	37.131*** (0.057)	34.800*** (0.052)	42.475*** (0.045)
Obs.	3,687,767	3,636,595	2,398,056	3,609,231

Note: *** p<0.001, ** p<0.01, * p<0.05. Standard Errors Clustered at the User Level.

Table A5: Balance Tests for Registered Users (Low Flood Score)

	(1) Views	(2) Bedrooms	(3) Bathrooms	(4) Sq. Ft.	(5) List Price	(6) Flood Score	(7) HHI Zip
Treatment	0.029 (0.025)	-0.005 (0.003)	0.085 (0.166)	7.618 (10.503)	66453.369 (102827.156)	0.001 (0.002)	-3.118 (8.168)
Constant	4.588*** (0.014)	3.458*** (0.002)	2.742*** (0.105)	2331.874*** (4.718)	783109.518*** (24460.586)	1.416*** (0.001)	6978.039*** (5.794)
Obs.	3,201,727	3,165,379	3,164,488	3,154,986	3,157,116	3,192,673	3,201,727

Note: *** p<0.001, ** p<0.01, * p<0.05. Standard Errors Clustered at the User Level.

Table A6: Balance Tests for Registered Users (Low Flood Score)

	(1) Year Built	(2) Walk Score	(3) Transit Score	(4) Bike Score
Treatment	-0.005 (0.086)	-0.020 (0.087)	-0.042 (0.079)	0.035 (0.068)
Constant	1972.882*** (0.061)	36.786*** (0.062)	34.306*** (0.056)	41.851*** (0.048)
Obs.	3,156,132	3,116,399	2,076,338	3,093,164

Note: *** p<0.001, ** p<0.01, * p<0.05. Standard Errors Clustered at the User Level.

Table A7: Balance Tests for Registered Users (Medium Flood Score)

	(1) Views	(2) Bedrooms	(3) Bathrooms	(4) Sq. Ft.	(5) List Price	(6) Flood Score	(7) HHI Zip
Treatment	-0.067 (0.051)	-0.007 (0.008)	-0.038 (0.030)	-46.191 (53.746)	-5034.123 (10515.313)	0.007 (0.009)	22.385 (18.659)
Constant	4.150*** (0.043)	3.241*** (0.006)	2.480*** (0.029)	2250.293*** (49.135)	765423.316*** (6953.525)	3.486*** (0.006)	7318.558*** (13.361)
Obs.	496,765	488,893	487,317	484,770	489,071	495,171	496,765

Note: *** p<0.001, ** p<0.01, * p<0.05. Standard Errors Clustered at the User Level.

Table A8: Balance Tests for Registered Users (Medium Flood Score)

	(1) Year Built	(2) Walk Score	(3) Transit Score	(4) Bike Score
Treatment	-0.208 (0.197)	0.019 (0.225)	-0.170 (0.217)	0.070 (0.171)
Constant	1972.633*** (0.139)	39.772*** (0.160)	37.961*** (0.156)	46.543*** (0.122)
Obs.	485,020	477,550	303,114	473,876

Note: *** p<0.001, ** p<0.01, * p<0.05. Standard Errors Clustered at the User Level.

Table A9: Balance Tests for Registered Users (High Flood Score)

	(1) Views	(2) Bedrooms	(3) Bathrooms	(4) Sq. Ft.	(5) List Price	(6) Flood Score	(7) HHI Zip
Treatment	0.022 (0.130)	0.057 (0.032)	0.049 (0.031)	-206.275 (267.535)	39194.587 (46569.980)	-0.020 (0.032)	9.366 (70.725)
Constant	3.812*** (0.085)	2.935*** (0.022)	2.447*** (0.022)	2380.845*** (252.285)	854961.916*** (33718.405)	7.894*** (0.022)	8053.542*** (50.965)
Obs.	34,191	33,382	33,295	33,134	33,716	33,992	34,191

Note: *** p<0.001, ** p<0.01, * p<0.05. Standard Errors Clustered at the User Level.

Table A10: Balance Tests for Registered Users (High Flood Score)

	(1) Year Built	(2) Walk Score	(3) Transit Score	(4) Bike Score
Treatment	0.135 (0.633)	1.438 (0.761)	1.327 (0.707)	0.936* (0.476)
Constant	1978.417*** (0.430)	33.737*** (0.501)	36.600*** (0.506)	45.973*** (0.318)
Obs.	33,193	32,591	14,606	32,271

Note: *** p<0.001, ** p<0.01, * p<0.05. Standard Errors Clustered at the User Level.

Table A11: Pre-Experiment Distribution of Flood Score for Registered Users by Treatment Status

Number of Registered Users; Column Percentages in Parenthesis

Flood Score Category	Control	Treatment	Total
Low	271,954 (83.85)	272,156 (83.90)	544,110 (83.88)
Medium	47,974 (14.79)	47,917 (14.77)	95,891 (14.78)
High	4,411 (1.36)	4,302 (1.33)	8,713 (1.34)
Total	324,339 (100.00)	324,375 (100.00)	648,714 (100.00)

Table A12: ATE for Registered Users on the Characteristics of the Homes Viewed
Change relative to Control

	(1) Views	(2) Bedrooms	(3) Bathrooms	(4) Sq. Ft.	(5) List Price	(6) Flood Score	(7) HHI Zip
Treatment	-0.008 (0.019)	-0.002 (0.002)	-0.152 (0.148)	-684.620 (364.044)	10214.279 (15641.902)	-0.003 (0.002)	-0.980 (5.079)
Constant	3.087*** (0.219)	1.561*** (0.078)	2.773*** (0.145)	2744.734*** (389.449)	742846.963*** (2270.868)	0.800*** (0.004)	3139.932*** (8.689)
Obs.	10,079,929	9,935,289	9,925,205	9,886,396	9,923,585	9,983,156	10,079,929

Note: *** p<0.001, ** p<0.01, * p<0.05. Standard Errors Clustered at the User Level.

Table A13: ATE for Registered Users on the Characteristics of the Homes Viewed
Change relative to Control

	(1) Year Built	(2) Walk Score	(3) Transit Score	(4) Bike Score
Treatment	-0.007 (0.038)	-0.007 (0.033)	-0.041 (0.029)	0.009 (0.026)
Constant	508.244*** (2.305)	7.070*** (0.034)	5.997*** (0.040)	8.592*** (0.039)
Obs.	9,858,362	9,748,579	6,308,618	9,679,443

Note: *** p<0.001, ** p<0.01, * p<0.05. Standard Errors Clustered at the User Level.

**Table A14: ATE for Registered Users on the Characteristics of the Homes Viewed
with Pre-Experiment Information**
Change relative to Control

	(1) Views	(2) Bedrooms	(3) Bathrooms	(4) Sq. Ft.	(5) List Price	(6) Flood Score	(7) HHI Zip
Treatment	-0.011 (0.019)	-0.002 (0.002)	-0.153 (0.149)	-686.918 (365.253)	10435.779 (15698.706)	-0.003 (0.002)	-1.259 (5.092)
Constant	3.066*** (0.221)	1.562*** (0.078)	2.774*** (0.146)	2748.983*** (390.704)	741610.058*** (2276.041)	0.800*** (0.004)	3157.319*** (8.713)
Obs.	10,013,673	9,894,381	9,884,508	9,850,992	9,881,989	9,983,156	10,013,673


Note: *** p<0.001, ** p<0.01, * p<0.05. Standard Errors Clustered at the User Level.

Table A15: ATE for Registered Users on the Characteristics of the Homes Viewed with Pre-Experiment Information
Change relative to Control

	(1) Year Built	(2) Walk Score	(3) Transit Score	(4) Bike Score
Treatment	-0.008 (0.038)	-0.007 (0.033)	-0.041 (0.029)	0.009 (0.026)
Constant	507.771*** (2.309)	7.060*** (0.034)	5.995*** (0.040)	8.584*** (0.039)
Obs.	9,822,267	9,720,841	6,298,423	9,652,044

Note: *** p<0.001, ** p<0.01, * p<0.05. Standard Errors Clustered at the User Level.


Figure A6: Event-time Study on the Average Daily Flood Score of Properties Searched for Registered Users
Change relative to Control



Note: Vertical lines show 95% confidence intervals. Each point reports the corresponding estimate of $\beta_1 + \beta_3$ from equation 1. The percentage shown next to each point is based on the corresponding estimate from the same specification re-estimated with the natural log of the outcome as the dependent variable, transformed as $(\exp(\beta_1 + \beta_3) - 1) \times 100$. The x-axis shows weeks since entry into the experiment for users who, on average, viewed high-risk properties before the experiment.


Figure A7: CATE on the Average Flood Score of a Daily Search for Registered Users, by Characteristics of Most Searched Destination and Origin Zip Code at Baseline

Change relative to Control




Note: Vertical lines show 95% confidence intervals. In both panels, each point reports the corresponding estimate of $\beta_1 + \beta_3$ from equation 1. The percentage shown next to each point is based on the corresponding estimate from the same specification re-estimated with the natural log of the outcome as the dependent variable, transformed as $(\exp(\beta_1 + \beta_3) - 1) \times 100$. The x-axis shows users' average search pre-experiment flood-score category.

Figure A8: CATE on the Average Flood Score of a Daily Search for Registered Users, by Flood Characteristics of the Most Searched Destination and Origin Zip Code at Baseline
Change relative to Control




Note: Vertical lines show 95% confidence intervals. In both panels, each point reports the corresponding estimate of $\beta_1 + \beta_3$ from equation 1. The percentage shown next to each point is based on the corresponding estimate from the same specification re-estimated with the natural log of the outcome as the dependent variable, transformed as $(\exp(\beta_1 + \beta_3) - 1) \times 100$. The x-axis shows users' average search pre-experiment flood-score category.

Figure A9: CATE on the Average Flood Score of a Daily Search for Registered Users, by ZIP-Code Median Household Income
Change relative to Control




Note: Vertical lines show 95% confidence intervals. Each point reports the coefficient on the interaction between treatment, the user's baseline flood-score category, and ZIP-code median household income from equation 1, interacted with median household income, so the point shows how the treatment effect changes with a one-unit increase in income for users in that category. The percentage shown next to each point is based on the same specification re-estimated with the natural log of the outcome as the dependent variable, transformed as in previous figures. ZIP-code median household income is assigned based on the ZIP code where the user conducted most pre-experiment searches and is measured using 2020 Census data.

Figure A10: CATE on the Average Flood Score of a Daily Search for Registered Users, by Cumulative Flood Probability at Baseline
Change relative to Control



Note: Vertical lines show 95% confidence intervals. Each point reports the estimate from equation 1 interacted with the corresponding flood-probability horizon, controlling for the full set of probability measures. The percentage shown next to each point is based on the corresponding estimate from the same specification re-estimated with the natural log of the outcome as the dependent variable and transformed as in previous instances. The x-axis represents each user's average flood probability search before the experiment began.

Figure A11: CATE on the Average Flood Score of a Daily Search for Registered Users with FEMA Risk and Most-Searched State Non-Disclosure Floodplain Status at Baseline
Change relative to Control




Note: Vertical lines show 95% confidence intervals. Each point reports the estimate from equation 1. The percentage shown next to each point is based on the corresponding estimate from the same specification re-estimated with the natural log of the outcome as the dependent variable and transformed as in previous instances. Users who browsed a FEMA Risk property before the experiment are classified as "with" FEMA Risk search at baseline. A state is classified as a non-disclosure floodplain as described in Hino and Burke (2021). The x-axis shows users' average pre-experiment flood-score category.

Table A16: CATE on the Average Outcomes of a Daily Search for Registered Users

	(1) Bedrooms	(2) Bathrooms	(3) Sq. Ft.	(4) List Price	(5) HHI Zip	(6) Walk Score	(7) Transit Score	(8) Bike Score
Low Score	-0.002 (0.003)	-0.161 (0.173)	-799.468 (426.290)	14302.051 (18481.706)	-2.438 (5.510)	0.000 (0.036)	-0.045 (0.031)	0.012 (0.028)
Medium Score	-0.006 (0.005)	-0.117 (0.123)	-37.479 (30.956)	-14107.601 (8792.630)	2.982 (13.753)	-0.067 (0.096)	-0.011 (0.084)	-0.003 (0.073)
High Score	0.015 (0.019)	0.046 (0.030)	9.515 (77.752)	18745.432 (49363.379)	45.668 (53.762)	0.293 (0.390)	-0.052 (0.404)	-0.078 (0.267)


Note: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$. Each cell reports the corresponding estimate of $\beta_1 + \beta_3$ from equation 1 for the specific outcome and average pre-experiment flood-score category.

Figure A12: CATE on the Average Fire, Heat, and Wind Outcomes of a Daily Search for Registered Users
Change relative to Control



Note: Vertical lines show 95% confidence intervals. Each point reports the estimated treatment effect within each baseline flood-risk category, computed as linear combinations of coefficients from equation 1. Outcomes are fire, heat, and wind risk scores. The x-axis shows users' pre-experiment flood-risk categories (low, medium, high).

Figure A13: CATE on the Probability of Platform Registration
Change relative to Control




Note: Vertical lines show 95% confidence intervals. Each point reports the treatment effect on conversion by baseline flood-risk category (low, medium, high), computed as linear combinations from a regression of conversion on treatment interacted with flood-risk categories. The sample includes all registered users in the experiment. Estimates are expressed in percentage points. The x-axis represents each user's average flood probability search before the experiment began.

Table A17: ATE for Registered Users on the Activity on the Platform
Change relative to Control

	(1)	(2)	(3)	(4)
	Total Seconds	Number Sessions	Unique Home Views	Total Home Views
Treatment	3.220 (3.000)	-0.000 (0.002)	0.000 (0.009)	-0.000 (0.011)
Constant	579.530*** (4.768)	1.018*** (0.004)	2.066*** (0.070)	2.761*** (0.069)
Obs.	17,943,086	17,943,086	17,943,086	17,943,086


Note: *** p<0.001, ** p<0.01, * p<0.05. Standard Errors Clustered at the User Level.

Figure A14: CATE on the Time Spent on the Website per Day for Registered Users
Change relative to Control



Note: Vertical lines show 95% confidence intervals. In both panels, each point reports the corresponding estimate of $\beta_1 + \beta_3$ from equation 1. The percentage shown next to each point is based on the corresponding estimate from the same specification re-estimated with the natural log of the outcome as the dependent variable, transformed as $(\exp(\beta_1 + \beta_3) - 1) \times 100$. The x-axis shows users' average search pre-experiment flood-score category.


Figure A15: CATE on the Number of Homes Viewed per Day for Registered Users
Change relative to Control




Note: Vertical lines show 95% confidence intervals. In both panels, each point reports the corresponding estimate of $\beta_1 + \beta_3$ from equation 1. The percentage shown next to each point is based on the corresponding estimate from the same specification re-estimated with the natural log of the outcome as the dependent variable, transformed as $(\exp(\beta_1 + \beta_3) - 1) \times 100$. The x-axis shows users' average search pre-experiment flood-score category.

Figure A16: CATE on Platform Activity per Day for Registered Users, Including Inactive Days
Change relative to Control

(a) Probability of Browsing on a Given Day




(b) Number of Unique Home Views




Note: Vertical lines show 95% confidence intervals. The panel includes every day from a user's entry into when the user did not visit Redfin are entered on which the user did not visit Redfin enter as zero. In both panels, each point reports the treatment effect for the corresponding baseline flood-score category from equation 1, with entry-date fixed effects and standard errors clustered at the user level. Panel (a) is estimated by OLS on an indicator for any listing view that day and is expressed as a change in probability. Panel (b) is estimated by Poisson pseudo-maximum likelihood on the number of unique listings viewed that day and is expressed as the percentage change relative to the control group (incidence rate ratio minus one). The x-axis shows users' average search pre-experiment flood-score category.

Figure A17: CATE on the Percentage of Times Registered Users Engaged with the Flood Risk Section of a Listing per Day
% Times User Engaged with Feature for Listings Viewed per Day

(a) Without Waterfront Search at Baseline




(b) With Waterfront Search at Baseline




Note: Vertical lines show 95% confidence intervals. Each point reports the corresponding estimate of $\beta_1 + \beta_3$ from equation 1, where baseline engagement is not included, as it is not observed pre-experiment. The x-axis shows users' average search pre-experiment flood-score category. Engagement with a property feature is measured as the share of listings visited on a given day in which the user interacted with that feature. For the flood-risk section, engagement is defined as scrolling to that section for treated users; for control users, it is proxied by scrolling past the page location where the section would have appeared. Users who did not browse any waterfront property before the experiment are classified as "without" waterfront search at baseline. On the other hand, users who browsed at least one waterfront property before the experiment are classified as "with" waterfront search at baseline.

Figure A18: CATE on the Percentage of Times Registered Users Engaged with the Flood Risk Section of a Listing per Day
% Times User Engaged with Feature for Listings Viewed per Day

(a) Without FEMA Risk Search at Baseline



(b) With FEMA Risk Search at Baseline



Note: Vertical lines show 95% confidence intervals. Each point reports the corresponding estimate of $\beta_1 + \beta_3$ from equation 1, where baseline engagement is not included, as it is not observed pre-experiment. The x-axis shows users' average search pre-experiment flood-score category. Engagement with a property feature is measured as the share of listings visited on a given day in which the user interacted with that feature. For the flood-risk section, engagement is defined as scrolling to that section for treated users; for control users, it is proxied by scrolling past the page location where the section would have appeared. Users who did not browse any property considered risky by FEMA before the experiment are classified as "without" FEMA risk search at baseline. On the other hand, users who browsed at least one property considered risky by FEMA before the experiment are classified as "with" FEMA risk search at baseline.

Figure A19: CATE on the Probability of Booking a Tour and Canceling a House Tour
Change relative to Control




Table A18: Balance Tests for Registered Users That Eventually Placed an Offer (Low Flood Score)

	(1) Views	(2) Bedrooms	(3) Bathrooms	(4) Sq. Ft.	(5) List Price	(6) Flood Score	(7) HHI Zip
Treatment	1.798 (1.001)	-0.020 (0.032)	-0.007 (0.029)	-34.729 (38.079)	-62713.596 (62957.747)	0.021 (0.017)	-179.117* (85.352)
Constant	7.389*** (0.272)	3.471*** (0.022)	2.509*** (0.019)	2256.972*** (28.375)	708556.499*** (48196.103)	1.398*** (0.011)	6221.737*** (58.266)
Obs.	21,983	21,878	21,866	21,809	21,881	21,975	21,983

Note: *** p<0.001, ** p<0.01, * p<0.05. Standard Errors Clustered at the User Level.

Table A19: Balance Tests for Registered Users That Eventually Placed an Offer (Low Flood Score)

	(1) Year Built	(2) Walk Score	(3) Transit Score	(4) Bike Score
Treatment	0.214 (0.973)	0.521 (0.960)	0.342 (0.794)	1.326 (0.761)
Constant	1973.763*** (0.678)	38.821*** (0.644)	34.358*** (0.533)	43.443*** (0.511)
Obs.	21,830	21,662	16,038	21,558

Note: *** p<0.001, ** p<0.01, * p<0.05. Standard Errors Clustered at the User Level.

**Table A20: Balance Tests for Registered Users That Eventually Placed an Offer
(Medium Flood Score)**

	(1) Views	(2) Bedrooms	(3) Bathrooms	(4) Sq. Ft.	(5) List Price	(6) Flood Score	(7) HHI Zip
Treatment	-0.053 (0.778)	-0.224* (0.098)	-0.146 (0.074)	-170.812 (90.350)	56505.598 (53961.365)	-0.231* (0.097)	124.903 (205.794)
Constant	6.544*** (0.496)	3.404*** (0.076)	2.476*** (0.054)	2127.858*** (63.806)	607466.441*** (35356.561)	3.551*** (0.073)	6599.879*** (128.785)
Obs.	3,107	3,083	3,078	3,072	3,089	3,103	3,107

Note: *** p<0.001, ** p<0.01, * p<0.05. Standard Errors Clustered at the User Level.

**Table A21: Balance Tests for Registered Users That Eventually Placed an Offer
(Medium Flood Score)**

	(1) Year Built	(2) Walk Score	(3) Transit Score	(4) Bike Score
Treatment	0.956 (2.368)	1.092 (2.737)	3.089 (2.383)	0.887 (2.088)
Constant	1970.838*** (1.666)	42.629*** (1.885)	36.817*** (1.532)	49.598*** (1.455)
Obs.	3,062	3,058	2,186	3,044

Note: *** p<0.001, ** p<0.01, * p<0.05. Standard Errors Clustered at the User Level.

**Table A22: Balance Tests for Registered Users That Eventually Placed an Offer
(High Flood Score)**

	(1) Views	(2) Bedrooms	(3) Bathrooms	(4) Sq. Ft.	(5) List Price	(6) Flood Score	(7) HHI Zip
Treatment	0.782 (0.895)	0.801 (0.590)	1.002* (0.403)	1194.595** (412.935)	213546.342* (98026.271)	0.132 (0.569)	-698.900 (1137.249)
Constant	3.719*** (0.579)	2.222*** (0.368)	1.770*** (0.252)	1189.089*** (271.383)	737163.501*** (68200.076)	7.647*** (0.370)	7975.082*** (701.664)
Obs.	209	208	208	208	209	209	209

Note: *** p<0.001, ** p<0.01, * p<0.05. Standard Errors Clustered at the User Level.

Table A23: Balance Tests for Registered Users That Eventually Placed an Offer (High Flood Score)

	(1) Year Built	(2) Walk Score	(3) Transit Score	(4) Bike Score
Treatment	5.634 (7.887)	-8.148 (10.843)	-13.692 (18.859)	4.600 (8.169)
Constant	1977.508*** (5.234)	41.839*** (7.030)	48.643** (12.372)	43.054*** (5.058)
Obs.	208	207	88	207

Note: *** p<0.001, ** p<0.01, * p<0.05. Standard Errors Clustered at the User Level.

Table A24: Romano-Wolf Corrected P-Values for Search Main Results

Figure #	Risk Score at Baseline	Clustered SE P-Value	Romano-Wolf P-Value
3(a)	Low	0.64	0.54
	Medium	0.36	0.14
	High	0.00	0.00
5(a)	Low	0.17	0.01
	Medium	0.71	0.55
	High	0.22	0.03
5(b)	Low	0.33	0.13
	Medium	0.30	0.09
	High	0.04	0.00
5(c)	Low	0.09	0.00
	Medium	0.98	0.97
	High	0.83	0.73
5(d)	Low	0.33	0.13
	Medium	0.16	0.01
	High	0.02	0.00
5(e)	Low	0.16	0.01
	Medium	0.61	0.54
	High	0.07	0.00


Note: Standard Errors Clustered at the User Level for both p-values. 1,000 resamples for the Romano-Wolf p-value.

Table A25: Romano-Wolf Corrected P-Values for Offer Main Results

Figure #	Risk Score at Baseline	Robust SE P-Value	Romano-Wolf P-Value
7(b)	Low	0.80	0.99
	Medium	0.49	0.98
	High	0.00	0.08
8(a)	Low	0.50	0.98
	Medium	0.33	0.96
	High	0.72	0.99
8(b)	Low	0.80	0.99
	Medium	0.94	0.99
	High	0.06	0.56
8(c)	Low	0.76	0.99
	Medium	0.09	0.67
	High	0.77	0.99
8(d)	Low	0.32	0.96
	Medium	0.21	0.91
	High	0.05	0.54
8(e)	Low	0.44	0.98
	Medium	0.03	0.46
	High	0.04	0.50

Note: Robust Standard Errors at the User Level for both p-values. 3,000 resamples for the Romano-Wolf p-value.

Figure A20: CATE on the Characteristics of an Offer
Change relative to Control




Note: Vertical lines show 95% confidence intervals. In both panels, each point reports the corresponding estimate of $\beta_1 + \beta_3$ from equation 1. For these outcomes, we use difference-in-differences specifications due to limited pre-experiment observations. The percentage shown next to each point in panel (b) is based on the corresponding estimate from the same specification re-estimated with the natural log of the outcome as the dependent variable, transformed as $(\exp(\beta_1 + \beta_3) - 1) \times 100$. The x-axis shows users' average search pre-experiment flood-score category. Figure (b) estimates are with city fixed effects of the property.

Table A26: CATE on the Average Outcomes of the Offers
for Registered Users Browsing High Risk Properties at Baseline

	(1) Bedrooms	(2) Bathrooms	(3) Sq. Ft.	(4) List Price	(5) Diff. Zip	(6) Walk Score	(7) Transit Score	(8) Bike Score
High Score	0.154 (0.340)	-0.143 (0.380)	-465.556 (359.451)	-387433.259 (364905.747)	0.082 (0.048)	10.090 (8.481)	-3.593 (7.283)	3.585 (6.966)


Note: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$. Each cell reports the corresponding estimate of $\beta_1 + \beta_3$ from equation 1. For these outcomes, we use difference-in-differences specifications due to limited pre-experiment observations. Standard errors in parenthesis.

Figure A21: CATE on the Average Fire, Heat, and Wind Outcomes of the Offers for Registered Users
Change relative to Control




Note: Vertical lines show 95% confidence intervals. Each point reports the corresponding estimate of $\beta_1 + \beta_3$ from equation 1. For these outcomes, we use difference-in-differences specifications due to limited pre-experiment observations. The x-axis shows users' average search pre-experiment flood-score category.

Figure A22: Flood Risk by Zip Code from Where Users Browsed the Most at Baseline, Stratified by Whether Biden or Trump Won in 2020
Change relative to Control



Note: Vertical lines show 95% confidence intervals. Each point reports the corresponding estimate of $\beta_1 + \beta_3$ from equation 1. For panel (b), we use difference-in-differences specifications due to limited pre-experiment observations. The x-axis shows users' average search pre-experiment flood-score category. We assigned a baseline zip code where users browsed the most at baseline.


Figure A23: CATE on Average Flood Score from Searches Following Flood Events
Change relative to Control



Note: Vertical lines show 95% confidence intervals. Each point reports the corresponding estimate of $\beta_1 + \beta_3$ from equation 1. The percentage shown next to each point is based on the corresponding estimate from the same specification re-estimated with the natural log of the outcome as the dependent variable, transformed as $(\exp(\beta_1 + \beta_3) - 1) \times 100$. For these outcomes, we use difference-in-differences specifications due to limited pre-experiment observations. The x-axis shows users' average search pre-experiment flood-score category. We assigned a baseline county where users browsed the most at baseline.


Figure A24: CATE on Closing

Change relative to Control




Note: Vertical lines show 95% confidence intervals. In both panels, each point reports the corresponding estimate of $\beta_1 + \beta_3$ from equation 1. For these outcomes, we use difference-in-differences specifications due to limited pre-experiment observations. The percentage shown next to each point in panel (b) is based on the corresponding estimate from the same specification re-estimated with the natural log of the outcome as the dependent variable, transformed as $(\exp(\beta_1 + \beta_3) - 1) \times 100$. The x-axis shows users' average search pre-experiment flood-score category. Figure (b) estimates are with city fixed effects of the property.

Figure A25: CATE of an Increase in Exposure Intensity on the (Sale - Listing Price)




Note: Vertical lines show 95% confidence intervals. Each point reports the corresponding estimate of $\beta_2 + \beta_3$ from equation 4. The percentage shown next to each point is based on the corresponding estimate from the same specification re-estimated using the natural log of the ratio of sale price to listing price as the dependent variable, transformed as $(\exp(\beta_2 + \beta_3) - 1) \times 100$. The x-axis shows the flood-score category of the property. Properties located on the waterfront or classified as FEMA high risk are labeled as waterfront or FEMA risk.

**Figure A26: CATE of an Increase in Exposure Intensity
on the Probability a Listing was Bought by an Investor**
Change per 1 point increase



Note: Vertical lines show 95% confidence intervals. Each point reports the corresponding estimate of $\beta_2 + \beta_3$ from equation 4. An investor is defined as any buyer whose name in the listing includes at least one of the following keywords: LLC, Inc, Trust, Corp, or Homes. We also define an investor as any buyer whose ownership code on a purchasing deed includes at least one of the following keywords: association, corporate trustee, company, joint venture, or corporate trust. The x-axis shows the flood-score category of the property.


**Figure A27: CATE of an Increase in Exposure Intensity
on the Time on the Market for All Listings**
Change per 1 point increase



Note: Vertical lines show 95% confidence intervals. Each point reports the corresponding estimate of $\beta_2 + \beta_3$ from equation 4. The percentage shown next to each point is based on the corresponding estimate from the same specification re-estimated with the natural log of the outcome as the dependent variable, transformed as $(\exp(\beta_2 + \beta_3) - 1) \times 100$. The x-axis shows the flood-score category of the property.

Figure A28: CATE of an Increase in Exposure Intensity
on the Loan Value for All Listings


Change per 1 point increase



Note: Vertical lines show 95% confidence intervals. Each point reports the corresponding estimate of $\beta_2 + \beta_3$ from equation 4. The percentage shown next to each point is based on the corresponding estimate from the same specification re-estimated with the natural log of the outcome as the dependent variable, transformed as $(\exp(\beta_2 + \beta_3) - 1) \times 100$. The x-axis shows the flood-score category of the property. Only listings that were bought with a loan.

Figure A29: CATE of an Increase in Exposure Intensity
on the Listing Price for Not Waterfront and Not FEMA Risky Listings


1PP Increase in Exposure Intensity




Note: Vertical lines show 95% confidence intervals. Each point reports the corresponding estimate of $\beta_2 + \beta_3$ from equation 4. The percentage shown next to each point is based on the corresponding estimate from the same specification re-estimated with the natural log of the outcome as the dependent variable, transformed as $(\exp(\beta_2 + \beta_3) - 1) \times 100$. The x-axis shows the flood-score category of the property.

Figure A30: Association Between Treatment Exposure Intensity and the (*Sale - Listing Price*) Spread

(a) Distribution of Listing Views Conducted by Treated Users per Listing




(b) All Listings
1PP Increase in Exposure Intensity



Note: For panel (b), vertical lines show 95% confidence intervals. Each point reports the corresponding estimate of $\beta_2 + \beta_3$ from equation 4. The percentage shown next to each point is based on the corresponding estimate from the same specification re-estimated using the natural log of the ratio of sale price to listing price as the dependent variable, transformed as $(\exp(\beta_2 + \beta_3) - 1) \times 100$. The x-axis shows the flood-score category of the property.

Figure A31: CATE of an Increase in Exposure Intensity on the (*Sale - Listing Price*) Spread
Change per 1 point increase



Note: Vertical lines show 95% confidence intervals. Each point reports the corresponding estimate of $\beta_2 + \beta_3$ from equation 4. The percentage shown next to each point is based on the corresponding estimate from the same specification re-estimated using the natural log of the ratio of sale price to listing price as the dependent variable, transformed as $(\exp(\beta_2 + \beta_3) - 1) \times 100$. The x-axis shows the flood-score category of the property. Properties not located on the waterfront or not classified as FEMA high risk are labeled as not waterfront or not FEMA risk.

Table A27: Effects of Flood-Risk Reclassification on Sale–Listing Price Spread

	(1) High to Low	(2) Low to High
Post reclassification	-5,320** (2,354)	22,270*** (2,567)
Property fixed effects	Yes	Yes
Year fixed effects	Yes	Yes
Month fixed effects	Yes	Yes
Observations	1,970,451	159,859

Note: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$. Standard errors in parentheses.

Table A28: Flood Risk and Sale–Listing Price Spread over Time

	(1) ZIP FE	(2) ZIP × Year FE
2015: Low	3,181*** (126)	
2015: Medium	5,202*** (303)	-298 (219)
2015: High	8,811*** (902)	902* (436)
2016: Low	1,365*** (90)	
2016: Medium	2,245*** (241)	-162 (199)
2016: High	2,483*** (646)	-68 (411)
2017: Low	2,122*** (100)	
2017: Medium	3,006*** (213)	84 (198)
2017: High	711 (527)	-300 (368)
2018: Low	1,623*** (89)	
2018: Medium	2,616*** (237)	168 (205)
2018: High	-621 (523)	-957* (458)
2020: Low	4,695*** (89)	
2020: Medium	4,358*** (232)	-29 (198)
2020: High	827 (687)	-915 (493)
2021: Low	15,537*** (220)	
2021: Medium	15,014*** (336)	-816** (249)
2021: High	14,917*** (921)	-1,443** (501)
2022: Low	9,885*** (201)	
2022: Medium	8,443*** (332)	-994** (331)
2022: High	7,629*** (966)	-1,719* (719)
2023: Low	-2,094*** (207)	
2023: Medium	-5,303*** (433)	-625 (332)
2023: High	-15,288*** (1,085)	-3,521*** (606)
2024: Low	-3,274*** (228)	
2024: Medium	-6,908*** (553)	-800* (352)
2024: High	-25,017*** (1,626)	-6,724*** (878)
2025: Low	-6,780*** (249)	
2025: Medium	-12,651*** (825)	-1,663* (745)
2025: High	-37,783*** (3,344)	-7,647*** (1,347)
ZIP code fixed effects	Yes	Yes
Year fixed effects	Yes	No
ZIP code × year fixed effects	No	Yes
Month fixed effects	Yes	Yes
Observations	33,241,285	33,218,357

Note: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$. Standard errors in parentheses.

A.3 Nonparametric Heterogeneous Treatment Effects


We seek to identify the conditional average treatment effect on the treated (CATE), which seeks to identify differences in treatment effects within the population and how big these differences are with an estimand defined as:

$$\text{CATE} \equiv \tau(X) = \mathbb{E} \left[Y_i^{(1)} - Y_i^{(0)} \mid X_i = x \right] = \mu_1(x) - \mu_0(x) \quad (14)$$

Where $Y_i^{(1)}$ and $Y_i^{(0)}$ are the potential outcomes of outcome, Y , with observed covariates, $X \in \mathbb{R}$, for individual, i . The goal is to identify the CATE or $\tau(X)$ which is the difference in potential outcomes that equates to the difference in the conditional expectation of x : $\mu_1(x) - \mu_0(x)$. Given that our treatment was randomized, we could estimate $\tau(x)$ via estimator 1. However, estimator 1 relies on the linearity assumption of the effect that covariates, X_i , have on the treatment. If these effects were non-linear, our calculated estimates would be biased or wouldn't cover the entire distribution of heterogeneous treatment effects. We rely on causal forests, a Generalized Random Forest algorithm, to relax this assumption and not provide a parametric form (Athey et al., 2019; Athey and Wager, 2021). Causal forests are built to identify how treatment effects vary across users by maximizing the difference in the relationship between our target variable (i.e., our outcome variable) and a specific feature (i.e., our treatment indicator) within other features (i.e., our baseline covariates). This method employs a splitting criterion optimized for detecting splits that reveal treatment effect heterogeneity. The objective is to identify leaf nodes where the treatment effect is constant but distinct from other leaves.


In our analyses, we utilize baseline variables as covariates, X . We use the most searched city, the device most used by the user to browse the web (i.e., mobile or desktop), the number of weeks since the user entered the experiment, and the baseline average flood of all the houses viewed by user, i . We randomly split our data into training (70%) and testing (30%) samples. We used the training sample to fit a causal forest with 50,000 trees grown in the forest and four grown trees on each subsample. After training the causal forest, we predicted conditional treatment effects on the testing sample and a 90% confidence interval around each one. On each graph presented below, we show the percent rank of the conditional average treatment effects with a 90% confidence interval, where a bigger coefficient size tells us that more observations lie within the branch of the calculated conditional average treatment effect.

Figure A32: Causal Forest—High Risk Group
Change relative to Control




Note: The causal forest was trained on 70% of the extreme risk universe, and the plotted effects are calculated on 30% of the rest. We use 50,000 trees to grow the forest. We use baseline variables to fit the model. A bigger size of the circle effects tells us that more observations lie within the branch of the calculated conditional average treatment effect.

Figure A33: Causal Forest—Medium Risk Group
Change relative to Control



Note: For computational reasons, we used 10% of the medium risk universe. From that sample, the causal forest was trained on 70% of the severe risk universe, and the plotted effects are calculated on 30% of the rest. We use 50,000 trees to grow the forest. We use baseline variables to fit the model. A bigger size of the circle effects tells us that more observations lie within the branch of the calculated conditional average treatment effect.

Figure A34: Causal Forest—Low Risk Group
Change relative to Control



Note: For computational reasons, we used 5% of the minimal risk universe. From that sample, the causal forest was trained on 70% of the major risk universe, and the plotted effects are calculated on 30% of the rest. We use 50,000 trees to grow the forest. We use baseline variables to fit the model. A bigger size of the circle effects tells us that more observations lie within the branch of the calculated conditional average treatment effect.

References

- Athey, S., Tibshirani, J. and Wager, S. (2019), 'Generalized random forests', *The Annals of Statistics* **47**(2), 1148–1178.
- Athey, S. and Wager, S. (2021), 'Policy learning with observational data', *Econometrica* **89**(1), 133–161.
- Severen, C., Costello, C. and Deschenes, O. (2018), 'A forward-looking ricardian approach: Do land markets capitalize climate change forecasts?', *Journal of Environmental Economics and Management* **89**, 235–254.