

Oberlin

Digital Commons at Oberlin

Oberlin College Kasper Economics and
Business Working Paper Series

Economics

7-2026

HKC06 - Information versus attention in flood risk disclosure

Dolores Albarracín
University of Pennsylvania

Timothy Hyde
Department of Economics, Oberlin College

Follow this and additional works at: https://digitalcommons.oberlin.edu/economics_wps



Part of the [Economics Commons](#)

Recommended Citation

Albarracín, Dolores and Hyde, Timothy, "HKC06 - Information versus attention in flood risk disclosure" (2026). *Oberlin College Kasper Economics and Business Working Paper Series*. 6.
https://digitalcommons.oberlin.edu/economics_wps/6

This Working Paper is brought to you for free and open access by the Economics at Digital Commons at Oberlin. It has been accepted for inclusion in Oberlin College Kasper Economics and Business Working Paper Series by an authorized administrator of Digital Commons at Oberlin. For more information, please contact mmcfarli@oberlin.edu.



**Department of
Economics**

**OBERLIN COLLEGE KASPER ECONOMICS AND BUSINESS
WORKING PAPER SERIES**

HKC Working Paper No 06
July 2026

Information versus attention in flood risk disclosure

Dolores Albarracín, Timothy Hyde

Information versus attention in flood risk disclosure*

Dolores Albarracín

Timothy Hyde

July 2026

Abstract

Flood risk disclosure is expanding rapidly in real estate markets, but if disclosures move demand by commanding attention rather than by conveying information, their welfare effects are ambiguous. We examine this tension in a discrete choice experiment where 1,498 respondents make hypothetical choices among real Gulf Coast properties under a control condition and four disclosure formats varying in informational content and presentation. Absent disclosure, insurance costs barely influence choices; by contrast, when respondents are shown precise flood insurance estimates, they are willing to pay \$3.34 in annualized housing costs to avoid each annual premium dollar. Using a separate beliefs elicitation, we estimate each format's technical information content and construct a Bayesian benchmark for how much choices should respond. Observed responses exceed this benchmark by a factor of four to six; informational content accounts for at most 30 percent of the response to any format.

Keywords: flood risk, risk disclosure, salience, attention, willingness to pay, stated preferences

JEL Codes: D83, D91, Q54, R31, G22

*Dolores Albarracín: University of Pennsylvania, dalba@upenn.edu. Timothy Hyde: Oberlin College, thye@oberlin.edu. We thank seminar participants at the University of Pennsylvania and the Association of Environmental and Resource Economists annual meeting for their attention and feedback. We also thank Saurabh Bhargava, Nick Muller, Lynn Conell-Price, Philip Mulder, and Sebastian Olascoaga for valuable feedback and input. All remaining errors are our own.

1 Introduction

Disclosure is a key aspect of real estate markets. In the U.S., many states require real estate sellers to complete disclosure forms where they must reveal private information, such as a leaky roof or rotten planks under the deck, that would not otherwise be easily discoverable by the buyers (Nanda 2008). A separate category of disclosures, which we term “reminder” disclosures, convey information that is already publicly available, but may not be top of mind for uninitiated buyers (Pope 2008a). For example, U.S. federal law requires sellers to disclose the *potential* of lead paint in homes built before 1978, even if no lead paint is known to be present (U.S. EPA 1996). These disclosures effectively combine information about the home with a government-mandated reminder about the existence of a hazard. Because these disclosures are based on public information, they can also be disclosed by third parties, such as when online real estate platforms began incorporating natural disaster risk information into listings (Fairweather et al. 2024).

Studies have shown that reminder disclosures have measurable effects on demand, belying standard models that assume that homebuyers have already metabolized all public information (Fairweather et al. 2024; Pope 2008a,b; Ma et al. 2024). Evidence from both the lab and the field shows individuals toggling between neglect of and extreme attention to rare natural disaster risks (Gallagher 2014; Lin 2020; Hu 2022; Hopkins and Hyde 2026; McClelland, Schulze and Coursey 1993). Because reminder disclosures are designed in part to increase the salience of the risk, we hypothesize that any demand response is the combined result of two effects that push the same direction: increased information and increased attention.

We illustrate this phenomenon in a preregistered¹ online experiment that compares different forms of flood risk disclosure in a hypothetical home buying task. We present respondents with mocked up real estate listings for various actual properties recently sold on the Florida Gulf Coast, and experimentally vary the flood risk disclosure accompanying the listing. We compare the effects of a) no disclosure, b) a general reminder about flood risk in Florida, c) FloodFactorTM risk scores on a ten-point scale, d) a range of possible insurance premiums, and e) precise insurance premiums associated with each property.

Our analysis proceeds in three steps of increasing structure. We begin with model-free comparisons of respondent choice behavior under different treatment conditions. Three facts emerge. Absent disclosure, respondents seem indifferent to insurance costs, choosing homes whose flood insurance premiums are statistically indistinguishable from those of the homes they reject. Disclosure induces large, dose-ordered premium avoidance – \$209 of avoided annual premium per choice under a general salience prime, \$269 under an insurance-cost range, \$1,243 under a ten-point flood score, and \$2,458 under exact premiums – partly paid for by accepting higher-rent properties. We

¹See Section 5.3 for more details about our preregistration and subsequent exploratory analyses.

also find disclosure diverts attention from other attributes: it attenuates the control group’s strong sorting toward larger homes across all four formats, including the salience prime, which conveys no property-specific information at all.

We then convert these gaps into dollar valuations with a discrete choice model. Our primary specification is deliberately flexible: we estimate a separate valuation shift for every property under every condition – roughly 400 property-by-condition effects, converted to annualized dollars – imposing nothing about how flood risk enters preferences. These estimates reveal both what respondents responded to, and how. Flood-score respondents devalue properties bin by bin: score bins explain 92 percent of the cross-property variation in their responses, while premium variation *within* bins produces none. Exact-premium respondents respond roughly linearly to the premium itself, with no residual bin structure. We also estimate a more parsimonious model with one estimand per arm: the insurance cost pass-through (ICPT), the marginal willingness to pay in annualized dollars per dollar of annual insurance cost. Relative to a control group whose ICPT is statistically zero, the general reminder raises the ICPT by roughly 0.23, the insurance-cost range by 0.29, the flood score by 1.34, and the precise premium by 3.42; respondents shown precise premiums pay \$3.34 in annualized housing cost to avoid each dollar of annual insurance cost.

Finally, we ask how large these responses should have been using a simple Bayesian benchmark. For a respondent who values insurance costs at par with rental costs and whose attention is unchanged by the display, the ICPT should rise one-for-one with the informational content of the display, measured as the share of cross-property premium variance the display’s signal explains, or equivalent its coefficient of determination R^2 . Using a separate beliefs elicitation in which 1,279 distinct participants guess flood risk from property listings, we estimate this informational content for each format and find that observed responses exceed the benchmark by a factor of four to six: the flood score moved ICPT, our measure of demand, by 1.34 where its information justifies 0.23, and the precise premium moved demand by 3.42 where its information justifies 0.85. Informational content accounts for at most 30 percent of the response to any format. The remainder operates through non-informational channels – attention drawn to the disclosed number, attention drawn away from competing attributes, framing, trust in numeric information – implying that disclosure can move demand well beyond what its informational content justifies.

2 Background

A large literature focuses on the effect of various disclosure policies on real estate markets, including mandates for reminder disclosures pertaining to lead paint (Bae 2012), wildfire risk (Ma et al. 2024), flood risk (Pope 2008*b*), and airport noise (Pope 2008*a*). These studies generally find a demand response to disclosure; respondents exhibit a greater willingness to pay for risk avoidance when the

disclosure comes into effect. Additional studies have found that disclosure of energy consumption and past electric bills can also have a large effect on homebuyer demand (Cassidy 2023; Frondel, Gerster and Vance 2020). Another strand of this literature focuses on the selective disclosure strategies of sellers in cases where certain types of disclosures are not mandated (Bian et al. 2023; Carrillo, Cellini and Green 2013; Haisken-DeNew et al. 2018), but we are not aware of studies of voluntary flood risk disclosure in particular, and questions of strategic behavior are less relevant to the present study.

A separate literature in behavioral economics documents the role of salience in consumer decision-making. Bordalo, Gennaioli and Shleifer (2013) develop a model in which consumers overweight attributes that are salient – contextually distinctive – in the choice environment, generating departures from expected utility maximization that depend on how information is framed and presented. Applied to disclosure, this framework predicts that the same underlying risk information can generate different demand responses depending on whether and how it is made salient at the point of decision. DellaVigna (2009) provides a broader review of inattention and limited consideration in consumer markets. Our paper connects this literature to the disclosure context, showing that salience effects are quantitatively large and dominate the informativeness channel across all disclosure formats we study.

Research on non-mandated disclosure also finds sharp effects. A field experiment surrounding the introduction of flood risk scores from the First Street Foundation (FSF) on the online real estate brokerage Redfin generated highly granular findings on homebuyer response to disclosure (Fairweather et al. 2024). The authors report that Redfin users who were randomly assigned to see flood scores shifted their browsing toward lower risk homes, and went forward to make offers on houses that were much lower risk on average than the control group. In an analogous field experiment conducted on participants who had already purchased homes, Conell-Price and Mulder (2024) find an asymmetric response to disclosure of the FSF flood scores among homeowners, reducing demand for insurance for participants who had previously overestimated their flood risk (but no corresponding increase for participants who had underestimated). Sussman et al. (2021) take an approach similar to the present study, creating real estate listing mockups in an online survey environment and experimentally varying disclosure methods in the context of home energy efficiency. They too find that disclosure increased MWTP for efficiency, although a graphic rating energy efficiency on a horizontal scale was more efficacious than displays of expected energy costs in dollars.

3 Experiment and data sources

3.1 Experimental design

To probe the tension between the information and attention treatments embodied in disclosure, we designed a survey in which respondents are faced with a series of hypothetical home purchase decisions. On each screen, respondents are presented with information on two recently-transacted properties on the Florida Gulf Coast, including a photograph of the property that would be featured on a real estate listing, data on the purchase price, home size, and a programmatically generated satellite view of the property and surrounding neighborhood on Google Maps. These are designed to roughly resemble popular online real estate listing sites like Zillow and Redfin. Figure 1 displays a screenshot from the survey software that shows the layout and format of the information that respondents saw. Each respondent answered ten hypothetical purchase questions covering twenty distinct properties.

While the properties pictured are real properties sold in four counties in Florida during 2021 and 2022, random data were imputed for the sales price, square footage, and number of baths to eliminate any possible correlation between sales price and unobserved characteristics of the properties.² These imputations were drawn randomly each time a property was displayed to a respondent, so any given property was listed with many different sales prices and square footages to different respondents. Meanwhile, the photographs and geographic location of each property were *not* imputed, allowing them to reflect the actual appearance and location of the transacted properties. Year built is a real property characteristic that was not displayed to respondents in any condition; it can influence choices only through visual inference from the listing photographs, and it is uncorrelated with insurance premiums across properties in our sample ($\rho = -0.06$). Table 1 summarizes the source, display status, and experimental variation of every attribute.

Respondents were equally randomized to one of five treatment conditions, which determined what additional information they would see about each property, if any. All respondents saw a property photograph, satellite view, sales price, square footage, number of bedrooms, and number of bathrooms. However, information about flood risk varied across our five conditions. Respondents in the **control** condition did not see any additional information. Respondents in the **salience priming** condition received a general warning about the risk of flood damage to properties in this area at the beginning of the survey.³ However, these respondents had no differentiating information

²Displayed sales prices and square footages were drawn independently for each property display from the empirical distribution of true values across the sampled properties – so displayed attributes are realistic in level and dispersion but statistically detached from the property they accompany; displayed bathroom counts were drawn independently as 2 (80 percent) or 3 (20 percent). All properties selected had three bedrooms to ensure comparability; that information was displayed accurately.

³The exact language displayed was as follows: “If you’re new to buying a home in Florida, you should know that

Table 1: Property attributes: source, display, and experimental variation

Attribute	Source	Displayed to respondents	Varies across displays of same property
Displayed, randomly imputed			
Sales price	Imputed (empirical dist.)	All conditions	Yes (each display)
Square footage	Imputed (empirical dist.)	All conditions	Yes (each display)
Number of baths	Imputed (2 w.p. 0.8, 3 w.p. 0.2)	All conditions	Yes (each display)
Displayed, fixed			
Number of beds	Real (all = 3)	All conditions	No
Photograph	Real	All conditions	No
Satellite view	Real	All conditions	No
Displayed only in treatment			
Flood Factor score	Real (First Street)	Flood score arm only	No
Insurance cost range	Real (Redfin)	Insurance range arm only	No
Insurance premium	Real (ReadYourRisk)	Precise insurance arm only	No
Never displayed; observed by researcher and included in estimation			
Year built	Real (MLS)	Never (photo-inferable)	No
Distance to water	Real (FGDL)	Never (map-inferable)	No
Elevation	Real (Google Earth)	Never	No

Each attribute of the mocked-up listings is classified by its source, whether respondents saw it, and whether it varied across displays of the same property. Displayed sales prices and square footages were drawn independently for each display from the empirical distribution of true values across the sampled properties, making them realistic in level and dispersion but orthogonal by construction to true flood risk and to unobserved property quality; the same property therefore appeared with many different prices and sizes to different respondents. Attributes in the final group were never stated to respondents but enter the choice models on the researcher side: year built as a control, distance to water in the polynomial amenity controls, and elevation descriptively. Year built and distance are partially inferable from the photograph and satellite view; this visual channel is the control group's only flood-risk signal, and its informativeness is what the beliefs elicitation measures ($\hat{\lambda}^{\text{ctrl}} = 0.151$). The insurance premium enters the choice models for every respondent regardless of condition; only its *display* varies with treatment.

about the properties. Rather the information displayed during the decision phase of the survey was identical to that displayed to the control group.

The three additional treatment conditions involved property-specific information about flood risk. Respondents in the **insurance range** condition viewed a range of values for annual flood insurance cost, matching estimates included on Redfin listings for each property as of November 2023 (Appendix Figure A4). Respondents in the **flood score** condition saw the Flood Factor for each property, along with a brief verbal characterization of the risk, e.g. “9/10, extreme risk” (Appendix Figure A5). Finally, respondents in the **precise insurance estimate** condition saw a single estimated annual flood insurance cost derived from the ReadYourRisk online calculator (Figure 2; see Section 3.4). These treatment conditions are summarized in Table 2.

3.2 Property selection

Properties were selected for inclusion in the experiment based on several criteria designed to ensure comparability and geographic representativeness. We sought properties that would be plausibly comparable to one another in terms of visible characteristics, while varying meaningfully in their

regular home insurance doesn't cover flood damage. Flood insurance can be very expensive for Florida homeowners, depending on where their property is located. When you make your choices, look closely at where each property is on the map. Think about how much flood insurance might cost and whether that might affect your decision.”

Table 2: Experimental treatment conditions

Treatment condition	Sample size	Median survey	
		completion time	Example information display
Control	302	4m 26s	None
Salience prime	299	4m 43s	None (see footnote 1)
Insurance range	301	4m 36s	“Annual flood insurance cost: \$1358 to \$5099”
Flood score	300	4m 11s	“Flood score: 9/10 (extreme risk)”
Precise insurance estimate	296	4m 40s	“Annual flood insurance cost: \$8727”

Respondents were randomly assigned to each treatment group with equal likelihood. Discrepancies in sample size are attributable to very slight differences in attrition rates. The salience prime condition differed from the control condition in that respondents were provided with a warning about the high risk of flood damage for properties in Florida but no differentiated information about each property. Differences in survey completion time are not jointly significant at the $\alpha = 0.05$ level. See Section 3.1 for more detailed descriptions of each treatment condition.

flood risk. To achieve this, properties were randomly drawn from a list of thousands of properties included in the Zillow ZTRAX database with the following features: located in one of four Florida counties for which MLS data was readily available (Manatee, Sarasota, Charlotte, and Lee), listed as having three bedrooms, sold between January 2021 and August 2022, and sales price between \$300,000 and \$600,000.⁴ The final selection of properties was determined by a stratified random sample that ensured equal representation of the ten Flood Factor scores.

Appendix Figure A6 shows the approximate geographic location of each of the 99 properties randomly selected for inclusion in the survey, along with lines indicating distance and direction to the nearest coastal point of each property.⁵ Descriptive statistics for the resulting sample of properties are reported in Appendix Table A7.

3.3 Sample composition

The survey was fielded to 1,498 U.S. respondents on the Prolific online survey platform in April 2024. Appendix Table A8 displays descriptive statistics for the sample, stratified by treatment condition assignment. Due to the voluntary nature of the online survey approach, the sample is not intended to be representative of the U.S. population at large. Indeed, the sample is more likely to reside in Florida, more likely to be college-educated, less likely to be foreign-born or a non-native English speaker, and more likely to agree with the notion that climate change is occurring than the U.S. population at large (Howe et al. 2015). F-tests for differences in means across treatment groups indicate that the randomization successfully achieved balance on key observable characteristics.

⁴Zillow Transaction and Assessment Dataset (ZTRAX). More information on accessing the data can be found at <http://www.zillow.com/ztrax>. The results and opinions are those of the authors and do not reflect the position of Zillow Group.

⁵We originally collected data on 100 properties, but one was excluded from the survey due to a technical error and not shown to any respondents, and thus effectively dropped from the study.

3.4 Ground-truth flood insurance costs

Throughout the paper, we define a property’s “true” annual flood insurance premium π_i as the estimate produced by the ReadYourRisk online calculator, which implements the rating methodology of the National Flood Insurance Program, also known as NFIP (Cocca 2023a; Horn and Webel 2019). Because π_i serves as both the key regressor in the choice models and the object whose variance defines informational content in Section 6, its reliability as a measure of property-level flood costs deserves discussion.

Beginning in October 2021, FEMA replaced its legacy rating system – which had priced policies since the 1970s using coarse flood-zone maps, so that properties with very different risk profiles within the same zone paid similar premiums – with Risk Rating 2.0, a property-level actuarial methodology (FEMA 2021). Under Risk Rating 2.0, each property’s premium reflects its own modeled expected losses, incorporating distance to water, flood type (riverine, storm surge, and rainfall-driven flooding), ground elevation, first-floor height, and replacement cost, using commercial catastrophe models alongside FEMA’s own data. The reform was explicitly designed to move premiums toward actuarial rates at the property level (GAO 2023).

Official NFIP quotes can be generated only one property at a time, by licensed insurance agents, through FEMA’s access-restricted rating engine; FEMA publishes aggregate premium statistics at the ZIP-code level, but there is no public interface through which a prospective buyer – or a researcher – can observe full-risk premiums at the property level (Cocca 2023a; GAO 2023). This restriction is itself a manifestation of the information friction our experiment studies: a homebuyer comparing listings cannot readily learn what flood insurance will cost before making an offer, which is precisely the gap that third-party disclosures such as FloodFactor scores and listing-site insurance estimates have emerged to fill. The ReadYourRisk calculator is an independent implementation of FEMA’s published Risk Rating 2.0 methodology, reconstructing the rating engine from the public technical documentation and constructing documented workarounds where FEMA’s inputs are not public – notably FEMA’s modified national hydrography layers and the proprietary catastrophe-model flood-depth data underlying the methodology’s river classification (Cocca 2023b).⁶

⁶Two measurement caveats apply. First, Risk Rating 2.0 caps annual premium *increases* for existing policyholders (generally at 18 percent per year), so the premium a current owner actually pays may sit well below the full-risk rate, with most preexisting policies not projected to reach full-risk rates until the late 2030s (GAO 2023). The calculator estimates we use correspond to the full-risk rate a *new* policyholder would face, which is the relevant number for a prospective buyer and the one that reflects underlying risk. Second, because the calculator is a near-replica of the official engine rather than the engine itself, its estimates are susceptible to implementation error relative to an official quote, concentrated in the components built from reconstructed rather than public inputs.

4 Empirical Strategy

Our empirical analysis proceeds in three steps of increasing structure: model-free treatment effects on the characteristics of chosen properties, a discrete choice model that converts choices into a menu-independent preference parameter, and a Bayesian benchmark for how large the resulting responses should have been if information content alone drove demand shifts. This section describes the first two; the benchmark requires the beliefs elicitation data and is introduced together with those data in Section 6.

4.1 Model-free treatment effects

Because condition is assigned randomly at the respondent level and every property attribute is drawn identically across arms – imputed attributes by independent randomization, real attributes by the randomization of respondents over a common property pool – the menus of properties offered to respondents are statistically identical across conditions (Appendix Table A6). Any systematic difference across arms in the properties respondents *chose* is therefore a causal effect of the disclosure treatment, identified by randomization alone.

Our preferred model-free estimator uses the within-pair gap. For each decision, define the gap in attribute k as the chosen property’s value minus the rejected property’s value, $g_k = k_{\text{chosen}} - k_{\text{rejected}}$, and estimate the following specification using OLS with standard errors clustered at the respondent level:

$$g_{k,nc} = \phi_0 + \sum_{t \neq \text{ctrl}} \phi_t \mathbb{I}\{t(n) = t\} + u_{nc}, \quad (1)$$

where respondents are indexed by n , each decision is indexed by c , and disclosure formats (constant within respondent) are indexed by t . Differencing within the pair nets out screen-to-screen variation in menu composition, so the constant ϕ_0 measures how the control group sorts on attribute k in the average head-to-head comparison, and each ϕ_t measures how disclosure format t moderates that sorting. Applied to the insurance premium, ϕ_t is the annual premium avoided per choice, in dollars, caused by format t .

4.2 Choice-model estimates

The gap regressions establish that disclosure moves choices, in raw dollars; a choice model converts those movements into preference parameters and reveals what objects move respondents. Our primary specification imposes as little as possible: a conditional logit with a separate utility shift

for every property under every condition,

$$U_{ni} = \alpha r_{ni} + \delta_i + \sum_{t \neq \text{ctrl}} \delta_i^t \cdot \mathbb{I}\{t(n) = t\} + \mathbf{x}_{ni}' \boldsymbol{\delta} + \varepsilon_{ni}, \quad (2)$$

where:

- δ_i is a property fixed effect, and
- δ_i^t is property i 's utility shift under condition t relative to the same property under control.

Nothing in Equation 2 presumes that flood risk, premiums, or any particular signal drives the response. Relative to the gap regressions, this specification reintroduces the logit's distributional assumption on the error term, but imposes nothing about how flood risk enters utility.

Where the dollar effects prove linear in the premium, they admit a one-parameter summary per arm. Replacing the property-by-condition effects with a premium interaction gives our parsimonious specification:

$$U_{ni} = \alpha r_{ni} + \beta_0 \pi_i + \sum_{t \neq \text{ctrl}} \beta_t \pi_i \cdot \mathbb{I}\{t(n) = t\} + \mathbf{x}_{ni}' \boldsymbol{\delta} + \varepsilon_{ni} \quad (3)$$

where:

- r_{ni} is the annualized purchase price (sales price $\times$ 0.06),
- π_i is the property's true annual flood insurance premium from Risk Rating 2.0 (Section 3.4),
- $\mathbb{I}\{t(n) = t\}$ indicates assignment to disclosure condition $t \in \{\text{salience prime, insurance range, flood score, precise insurance}\}$,
- $\mathbf{x}_{ni}$ is a vector of property and display characteristics (year built, bathrooms, square footage, side-of-screen indicator),
- and ε_{ni} is an i.i.d. Type-I extreme value error.

We estimate both specifications using a conditional logit model, with standard errors clustered at the respondent level; estimation of Equation 2 is detailed in the Appendix (Hainmueller, Hopkins and Yamamoto 2014). Note that the premium π_i enters for *every* respondent, whatever their condition displays: β_0 captures the sensitivity of control-group choices to true insurance costs (as channeled through visual inspection of the listing), while the interaction coefficients β_t capture how each disclosure format changes that sensitivity. Equation 3 is the special case of Equation 2 in which the property-by-condition effects are restricted to lie on lines through the premium, $\delta_i^t = \Delta \beta_t \pi_i$; comparing the two therefore tests the restriction directly.

Crucially, the displayed sales price, square footage, and bathroom count are independently randomly imputed for each respondent from the empirical distribution of property attributes, so displayed price is orthogonal by construction to both flood risk and unobserved property quality. This breaks the price–quality correlation that complicates hedonic estimates of willingness to pay and ensures that α isolates the marginal utility of housing expenditure.

4.3 Definition of insurance cost pass-through

With estimates for Equation 3 in hand, we operationally define **insurance cost pass-through** (ICPT) for condition t as the ratio of insurance cost sensitivity to rental cost sensitivity:

$$\text{ICPT}^t = \frac{\beta_0 + \beta_t}{\alpha}, \quad \text{ICPT}^{\text{ctrl}} = \frac{\beta_0}{\alpha}, \quad (4)$$

measuring willingness to pay in annualized dollars for a one-dollar reduction in annual flood insurance cost. The ICPT is the slope that the flexible dollar effects of Equation 2 trace out against the premium; Section 5 verifies that the two specifications agree.

$\text{ICPT} = 1$ is the **insurance-cost capitalization benchmark**: the response of an agent who values \$1 of annual insurance cost like \$1 of annual rent. We emphasize that this benchmark is descriptive rather than normative. Insurance premiums are not the full economic cost of flood risk – uninsured losses, deductibles, displacement, resale risk, and expected premium growth are all excluded – so a fully informed agent could rationally exhibit an ICPT above 1.⁷

A separate complication contaminates measured ICPT levels. In any sample of real coastal properties, flood risk is positively correlated with waterfront amenity value: high-premium properties tend to sit closer to the water, which makes them both riskier and more attractive.

We respond to this contamination in two ways. First, our primary estimates are treatment effects. Because the amenity term does not depend on t – treatment assignment is independent of property characteristics – it differences out of the treatment effects comparing ICPT across experimental

⁷NFIP premiums price the *insured* layer of flood risk rather than a property’s total flood exposure. Residential coverage is capped by statute at \$250,000 for the structure and \$100,000 for contents – limits unchanged since 1994 and below replacement cost for much of our sample – and NFIP policies exclude displacement and additional living expenses (Horn and Webel 2019; Kousky 2018). Uninsured losses are therefore one reason a fully informed buyer could rationally capitalize more than one dollar of housing cost per dollar of premium. The same institutional facts, however, bound how far above one a rational response can go: if premiums approximate expected insured losses and the insured layer covers a share κ of total expected flood costs, full-cost capitalization implies a pass-through of roughly $1/\kappa$. Even generous accounting for uninsured losses in our sample puts this ceiling well below the responses we estimate under precise disclosure. For a \$450,000 property with \$250,000 of structure coverage, κ below one-half would require most expected losses to fall outside the insured layer; because the majority of flood claims are partial losses well within the coverage limits (Wing et al. 2020), plausible values of κ for this housing stock imply a rational ceiling on the order of 1.5 to 2 – a back-of-envelope bound, but one comfortably below the precise-insurance pass-through of 3.4 we find (4.0 under arm-specific price sensitivity).

arms.

Second, we estimate specifications with property fixed effects, which absorb all time-invariant property-level characteristics. In these specifications β_0 is no longer separately identified, but the interactions β_t remain identified from randomized variation in condition assignment within properties, and we still recover the treatment effects. Reassuringly, as we show in Section 5, the treatment effects we find are nearly identical to those from specifications without fixed effects, confirming that the amenity confound is not driving the size of the treatment effects.

4.4 External validity

While online survey platforms present an opportunity to reach a broad population living across a wide geographic area – a key factor in this study – they also have some important limitations (Snowberg et al. 2021; Peer et al. 2021). Our estimates are stated-preference responses from an online panel, and the levels of willingness to pay should be interpreted accordingly: the hypothetical-bias literature consistently finds that stated valuations exceed revealed valuations (List and Gallet 2001; Murphy et al. 2005). Three considerations nevertheless make the estimates informative for disclosure design. First, randomization holds hypothetical bias constant across arms, so the cross-format contrasts – the central objects of this paper – difference it out to first order; moreover, the response-weight ratios of Section 6 are invariant to any hypothetical discounting that attenuates sensitivity to both dollar attributes proportionally, and the residual concern – differential attention to displayed versus imputed dollars – is measured directly in Appendix Table A3 rather than left as a conjecture. Second, the choice environment mirrors the setting where disclosure policy actually operates: online listings viewed during search, where formats like FloodFactor scores and insurance cost estimates are in fact deployed; the behavioral content of our estimates – respondents selecting the cheaper-to-insure home in 73 percent of comparisons under precise disclosure, versus chance in control – describes browsing-stage behavior, not offer or purchase decisions. Third, the treatments are the actual display formats used by Redfin, FloodFactor, and premium calculators rather than stylized manipulations, so each treatment effect corresponds to a deployable policy. The proper reading of our magnitudes is therefore comparative: precise insurance disclosure moves stated demand roughly fifteen times as much as a general warning and two and a half times as much as a score-based display, even if the levels of all responses would differ with real stakes.

5 Results

5.1 Model-free evidence

Table 3 reports the gap estimates of Equation 1: the effect of each disclosure format on the difference between chosen and rejected properties, attribute by attribute, with the control group’s mean gap as the constant.

Three patterns emerge. First, **control respondents act indifferently to insurance costs**. The average control respondent chose a home whose annual flood insurance premium was only \$40 higher than the home they rejected (SE \$90) – statistically indistinguishable from ignoring premiums entirely, and also consistent with a negative effect cancelled out by the waterfront-amenity correlation discussed in Section 4. Control respondents selected the cheaper-to-insure house in 49.8 percent of head-to-head comparisons, no better than chance (Table 3, Column 6). Absent disclosure, whatever flood risk information is embedded in photographs and satellite views does not translate into premium-avoiding choices in this choice environment.

Second, **disclosure induces large, dose-ordered premium avoidance**. Relative to control, respondents avoided \$209 of annual premium per choice under the salience prime, \$269 under the insurance range, \$1,243 under the flood score, and \$2,458 under precise premiums; the share of decisions selecting the cheaper-to-insure house rises monotonically across arms, from 53 percent under the salience prime to 62 percent under the flood score and 73 percent under exact premiums. The ordering matches the informational richness of the displays exactly. Respondents partly paid for this avoidance in higher home prices: chosen-rent gaps rise by \$732 and \$891 in the flood score and precise insurance conditions, indicating substitution toward safer but more expensive properties.

Third, **disclosure diverts attention from other attributes**. Control respondents sorted strongly toward larger homes, choosing properties 106 square feet larger than those they rejected. Every disclosure format attenuated this sorting, by a pooled 41 square feet ($t = -3.4$), and the attenuation is statistically indistinguishable across the four arms ($F(3, 1470) = 0.64$, $p = 0.59$) – including the salience prime, which conveys no property-specific information and produces almost no premium response. The attentional footprint of disclosure is roughly constant across formats even though its informational payoff varies by an order of magnitude. These estimates are essentially unchanged when controlling for the premium and rent gaps, indicating that the square-footage effect is not a mechanical byproduct of the premium–rent substitution. As we show in Section 5.2, attention to the displayed *price* exhibits a second, distinct pattern: it falls only in the informative arms, in proportion to the premium response.⁸

⁸Chosen-minus-rejected gaps in year built show a smaller, statistically detectable but non-monotone pattern across arms. Year built was not displayed in any condition and is uncorrelated with insurance premiums across properties

Table 3: Treatment effects on chosen properties

	(1) Insurance premium (\$)	(2) Annual rent (\$)	(3) Square footage	(4) Baths	(5) Year built	(6) Chose lower- premium home
Salience prime	−208.5 (127.5)	−31.4 (186.2)	−36.4 (15.7)	−0.016 (0.014)	1.56 (0.62)	0.033 (0.013)
Insurance range	−269.4 (126.3)	119.6 (186.8)	−31.3 (15.2)	−0.004 (0.014)	2.72 (0.61)	0.039 (0.013)
Flood score	−1242.8 (121.9)	731.5 (178.6)	−44.7 (14.7)	−0.010 (0.014)	1.99 (0.59)	0.122 (0.013)
Precise insurance	−2458.0 (126.0)	891.3 (189.0)	−50.3 (15.0)	−0.012 (0.015)	0.80 (0.64)	0.235 (0.014)
Control mean (constant)	40.2 (90.3)	−1859.1 (132.1)	106.1 (10.8)	0.020 (0.010)	0.61 (0.45)	0.498 (0.009)

Columns 1–5: OLS regressions of the chosen-minus-rejected gap in the indicated attribute on condition indicators (Equation 1), one observation per choice (14,710 decisions), with standard errors clustered at the respondent level in parentheses; the constant is the control group’s mean gap and coefficients are treatment effects relative to control. Column 6: the same regression with an indicator for choosing the home with the lower true annual premium as the outcome; the constant is the control group’s share and coefficients are treatment effects, so condition-level shares are 0.531, 0.537, 0.620, and 0.733 respectively. Menus are balanced across arms (Appendix Table A6), so these are causal effects of the disclosure treatments on choices. Square footage effects are statistically indistinguishable across the four treatment arms ($F(3, 1470) = 0.64$, $p = 0.59$) and robust to controlling for the premium and rent gaps.

5.2 Choice-model estimates

Figure 3 presents our central choice-model result: the estimated valuation shift of each of the 99 properties under each condition, relative to the same property under control (Equation 2), plotted against the property’s true annual premium. Three features stand out. First, the fitted slopes are near zero in the salience prime and insurance range arms, steep in the flood score arm, and steepest in the precise insurance arm – the dose-ordering of Section 5.1, now expressed as valuation gradients. Second, the dispersion around the fits differs sharply between the two informative arms in the expected way: the precise insurance cloud is a tight ramp, while the flood score cloud shows the same tilt with far greater scatter, the visual signature of a display that reveals only a coarse bin. Third, no arm’s response is driven by a handful of extreme properties; the gradients run through the full premium distribution.

We also see evidence that responses in demand to flood score displays are sharply responsive to the specific values (see Appendix Table A1). Respondents in the flood score condition saw only the ten-point score, so their valuation shifts should vary between score bins but not within them; precise-premium respondents should show the reverse pattern. Both predictions hold exactly. In the flood score arm, bin fixed effects explain 92 percent of the cross-property variation in the dollar effects ($F(9, 88) = 102.4$) while the within-bin premium slope is -0.04 ($p = 0.79$): properties

($\rho = -0.06$), so this pattern cannot reflect premium inference; we speculate that an intuitive but invalid “newer looks safer” visual heuristic is operating on the listing photographs.

Table 4: Treatment effects on insurance cost pass-through

Condition	(1) Pooled conditional logit	(2) Property fixed effects	(3) Per-property slope	(4) Arm-specific price sensitivity
Salience prime	0.231 (0.126)	0.209 (0.117)	0.183 (0.207)	0.200 (0.109)
Insurance range	0.295 (0.127)	0.271 (0.121)	0.270 (0.226)	0.282 (0.124)
Flood score	1.342 (0.134)	1.269 (0.129)	1.463 (0.353)	1.803 (0.255)
Precise insurance	3.415 (0.209)	3.326 (0.190)	2.972 (0.115)	4.018 (0.471)
Control ICPT (level)	−0.074 (0.090)	0 (absorbed)	0 (normalized)	−0.060 (0.075)

Each column estimates the same objects – treatment effects on the insurance cost pass-through (Δ ICPT), in annualized dollars of willingness to pay per dollar of annual insurance cost – under different assumptions. Column 1: the parsimonious conditional logit (Equation 3), our preferred estimates; full coefficients and distance-control variants in Appendix Table A2. Column 2: adding property fixed effects, which absorb all time-invariant amenity characteristics; the control-group level is absorbed and treatment effects are identified from randomized assignment within properties. Column 3: the slope of the roughly 400 nonparametric property-by-condition valuation effects of Equation 2 (Figure 3) regressed on the true premium, arm by arm, imposing no functional form on how flood information enters utility; standard errors unadjusted for first-stage estimation. Column 4: allowing price sensitivity to differ by condition and using each arm’s own rent coefficient as the denominator (Appendix Table A3); informative disclosures reduce price sensitivity by 30–40 percent, so the pooled estimates in column 1 are conservative. Delta-method standard errors clustered at the respondent level in parentheses. Agreement across columns shows the treatment effects are insensitive to amenity controls, functional form, and the denominator convention.

in the same bin whose true premiums differ by thousands of dollars received statistically identical treatment. In the precise insurance arm, the premium slope is -3.03 ($t = -20.4$) while score bins add nothing conditional on the premium ($F(9, 88) = 1.41$, $p = 0.20$).

Pass-through estimates

Table 4 reports the treatment effects on the insurance cost pass-through – the Δ ICPT for each disclosure format – estimated four different ways: the parsimonious conditional logit of Equation 3 (our preferred estimates), the same specification with property fixed effects, the slopes of the per-property valuation effects of Figure 3, and a specification allowing each arm its own price sensitivity. Reading across any row, four methodologically distinct estimates agree; reading down any column, the dose-ordering appears. Full conditional logit coefficients, including the preregistered specification with flexible distance controls (nearly identical), appear in Appendix Table A2.

The pass-through estimates reveal a clear monotone ordering across disclosure formats. In the control condition, respondents show essentially no sensitivity to flood insurance cost differences across properties, with an ICPT of -0.074 (not significantly different from zero). This mirrors the model-free result of Section 5.1: the raw chosen-minus-rejected premium gap in the control group is \$40 and statistically zero. The near-zero control ICPT is consistent with two partially offsetting

forces: respondents have limited ability to visually discriminate flood risk across properties, and flood risk is positively correlated with waterfront amenity value. Both forces tend to push the effective ICPT downward.

Each successive disclosure format generates greater pass-through of insurance costs into home valuations. The general salience prime raises the ICPT by approximately 0.2 relative to control, reflecting the increase in flood risk salience from the Florida-wide warning even absent any property-specific information. The insurance range and flood score conditions produce treatment effects of roughly 0.27-0.30 and 1.27-1.80 respectively. Most strikingly, the precise insurance estimate condition generates a treatment effect between 2.97 and 4.02 across methods – meaning respondents in this condition are willing to pay roughly \$3 or more in annualized rent for a \$1 reduction in annual flood insurance cost, a factor of roughly 3 or more above the capitalization benchmark.

Expressed in sale-price terms, an ICPT of 3.4 implies that a \$5,000 premium differential moves valuations as if prices differed by roughly \$280,000. We do not interpret this level as a market prediction. The dollar magnitude is large in part because respondents are weakly sensitive to the hypothetical displayed price, which enters the denominator of every pass-through ratio; the underlying choice behavior – selecting the lower-premium home in 73 percent of comparisons under precise disclosure, against 50 percent in control – is the model-free content, and Section 7 discusses the interpretation of levels.

A natural concern in any hypothetical-choice experiment is that treatment effects reflect experimenter demand rather than true changes in valuation. The salience prime condition provides a partial bound on this concern. That condition explicitly instructs respondents to think about flood insurance costs when choosing – arguably the most direct signal of the experimenter’s interest anywhere in our design – yet it moves the ICPT by only about 0.2, less than one-tenth of the precise-insurance treatment effect. If conveying the experimenter’s interest in flood risk were the primary driver of responses, the condition that conveys it most explicitly should not produce the smallest effect.

When we allow price sensitivity to vary by condition (Column 4), respondents shown the two informative displays became substantially less sensitive to the displayed price – by roughly 40 percent under the flood score and 30 percent under precise insurance, with no change under the salience prime or insurance range (Appendix Table A3). Combined with the square-footage result of Section 5.1 – which also appears within the choice model, with statistically similar attenuation in all four arms for house size and no movement in bathrooms or year built – there is suggestive evidence that the attentional reallocation has two components: a fixed cost of invoking the flood frame, paid equally by every format, and substitution toward the disclosed number, paid only where there is a number to substitute toward. The elevated pass-through under informative disclosure therefore reflects both overweighting of the disclosed premium and underweighting of

the competing price, and the pooled estimates of column 1, which average the treated arms' reduced price sensitivity with the control group's, are conservative. These estimates are also stable across subsamples, including first versus last choices, excluding the fastest completers, and excluding top-coded-premium properties (Appendix Table A4).

5.3 Pre-registration and deviations

The experiment and its primary analysis were pre-registered on AsPredicted (#171547) on April 19, 2024, before any data were collected. The pre-registration specified the five conditions, a target of 300 respondents per arm, the exclusion of respondents completing the hypothetical purchase task in under 45 seconds, the conditional logit of Equation 3 together with its property controls, and the insurance cost pass-through – estimated as the ratio of the insurance-cost interaction to the annualized rent coefficient – as the primary estimand. The pre-registered exclusion drops 27 of 1,498 respondents, leaving the 1,471 respondents and 14,710 decisions analyzed throughout.

The pre-registration also stated four hypotheses, of which the data support two and reject two. As predicted, the control group does not price flood insurance costs (H1: control pass-through of -0.074 , $p > 0.10$), and every disclosure format shifts pass-through in the predicted direction (H2: significant at the five percent level in three of four arms, and at the ten percent level in the fourth). We also predicted, however, that the flood score would move demand *more* than the dollar-denominated displays (H3), and that the dollar displays and the general reminder would move demand to the capitalization benchmark of one (H4b). Both fail, and they fail informatively. The precise insurance display moves pass-through two and a half times as much as the flood score, reversing the predicted ordering. And no format we predicted would land at the benchmark does: the insurance range moves pass-through to 0.22 and the general reminder to 0.16, far below one, while precise premiums move it to 3.34, far above. What we anticipated as a test of whether disclosure produces calibrated demand instead produced a fifteen-fold spread across formats with nothing well-calibrated – a pattern no valuation of flood costs can explain, since the properties and their true costs are identical across arms and only the display varies.

The analysis that occupies most of this paper was developed in response to that pattern and is exploratory. Specifically, the model-free evidence of Section 5.1, the per-property estimates and display-tracking tests of Section 5.2, the informativeness benchmark of Section 6 and the beliefs elicitation that supports it, the property fixed-effects and arm-specific price-sensitivity specifications, and the subsample robustness of Appendix Table A3 were not pre-registered. We report the pre-registered hypotheses as stated above and flag the exploratory analyses as such.

6 Benchmarking the demand response

To provide context for these normatively large demand responses, we compare these values to the predictions of a simple Bayesian updating model. A natural benchmark is a respondent whose:

- valuation of annual flood insurance premiums at par with annualized rental costs of home purchase, and
- attention to flood risk is unchanged by the display.

For such a respondent there is a simple result: when the informational content of a display rises by x , the ICPT rises by exactly x – information passes through to demand one-for-one. Here, informational content has a familiar meaning: $\lambda^t \in [0, 1]$ is the R^2 of the display’s signal for the true insurance premium – the share of cross-property premium variance the signal explains – with $\lambda^t = 0$ for an uninformative display and $\lambda^t = 1$ for one that fully reveals the premium. The one-for-one rule requires only Bayesian updating and the same conditional logit assumptions as Section 4. The appendix contains a proof sketch of this result. In this exercise, we estimate λ^t for each display using a separate beliefs elicitation dataset and compare observed treatment effects to the benchmark’s prediction.

It is convenient to summarize departures from the benchmark with a single ratio. We define the **response weight** $\gamma^t \equiv \text{ICPT}^t / \lambda^t$: dollars of demand response per dollar of information-justified response, equal to 1 in every arm under the par benchmark. We emphasize that γ^t is a residual by construction: values above 1 establish that the response exceeds what improved discrimination between safe and risky properties can explain, without identifying which non-informational mechanism – heightened attention to the disclosed number, diminished attention to competing attributes such as the displayed price (Appendix Table A3), trust in numeric information, monetary framing, anchoring, inability to balance upfront costs with annual costs – carries the excess. We interpret the excess as attention broadly construed, but lack the ability to identify the individual components in the present study design.

6.1 Measuring the informational content of each display

In addition to the discrete choice experiment data, we fielded a separate beliefs elicitation exercise to a distinct sample of Prolific respondents. Each participant in this sample was shown photographs and satellite images of a subset of the 99 properties and asked to guess the FloodFactor score for each property on a scale of 1-10. Respondents had a small monetary incentive to guess the score accurately. We collected responses from 1,279 unique respondents who evaluate 20 properties

each, yielding a total of 25,580 property-respondent observations.⁹ We translate respondent flood score guesses into dollar-denominated beliefs using the empirical relationship between FloodFactor scores and true premiums in the property sample. These data are used to estimate the information content of the non-disclosure display λ^{ctrl} , and are summarized in Appendix Figure A1.

The beliefs elicitation dataset allows us to estimate λ^t for each condition from outside the choice model. For the control and salience prime conditions – which share the same property-level signal, namely visual appearance and satellite imagery – the slope of a regression of standardized true premiums on standardized mean predicted flood scores across 99 properties is 0.389 (Appendix Figure A1); squaring this correlation gives $\hat{\lambda}^{\text{ctrl}} = 0.151$ (DeGroot 1970) (bootstrap SE = 0.063). For the precise insurance condition, $\lambda^{\text{prec}} = 1$ by construction since the signal is the true premium itself. For the flood score condition, the law of total variance yields $\hat{\lambda}^{\text{score}} = 1 - \bar{s}_{\text{score}}^2 / \hat{\sigma}_0^2 = 0.381$ (bootstrap SE = 0.078), where $\hat{\sigma}_0^2 = 11,174,068$ is the cross-property variance of true premiums and $\bar{s}_{\text{score}}^2$ is the property-weighted mean within-FloodFactor-bin variance. For these two conditions we compute the R^2 of the signal itself, which equals the informativeness of respondents’ beliefs only if respondents fully use the display – know the FloodFactor-to-dollar mapping, trust the stated premium; if they underuse it, the true λ^t is lower and the non-informational shares reported below are understated. Both standard errors come from a joint bootstrap resampling the 99 property clusters (500 replications) – the dominant source of sampling variation, since the estimates condition on a fixed set of 99 properties – which also delivers the covariance between $\hat{\lambda}^{\text{ctrl}}$ and $\hat{\lambda}^{\text{score}}$ used in the confidence intervals below. The insurance range condition is excluded from the benchmark comparison because Redfin range estimates are systematically biased downward relative to true premiums (Appendix Figure A3 and Appendix Table A5), so no credible informational content can be computed for it. Note that the results in Appendix Table A1 provide independent behavioral support for treating the score bin as the flood score arm’s information set: within-bin premium variation, which the display does not reveal, produces no valuation response in that arm.

6.2 The benchmark comparison

As established in Section 4, the amenity correlation biases ICPT levels but cancels in treatment effects, so the comparison operates on ΔICPTs anchored at the benchmark rather than at the (contaminated) control. As shown in the appendix, the predicted treatment effect of display t is the information gain itself, $\hat{\lambda}^t - \hat{\lambda}^{\text{ctrl}}$; the treatment effect a Bayesian updater whose attention did not change would exhibit. The response weight is recovered as $\hat{\gamma}^t = (\hat{\lambda}^{\text{ctrl}} + \widehat{\Delta\text{ICPT}}^t) / \hat{\lambda}^t$. The information component of each treatment effect is $\hat{\lambda}^t - \hat{\lambda}^{\text{ctrl}}$ and the non-informational component

⁹Beliefs were elicited for all 100 originally collected properties; all estimates below restrict to the 99 properties shown in the discrete choice experiment (25,295 observations; the excluded property received 285 evaluations).

Table 5: Observed responses versus the Bayesian benchmark

	(1)	(2)	(3)	(4)	(5)
	Information	Expected	Observed	Overstatement	Information
Condition	content	effect	effect	factor	share of
	λ^I	$\lambda^I - \lambda^{\text{ctrl}}$	ΔICPT	(3)/(2)	effect
Control	0.151 ^a	—	—	—	—
Salience prime	0.151 ^a	0	0.231 (0.126)	—	0% ^d
Flood score	0.381 ^b [0.23, 0.53]	0.230 [0.04, 0.42]	1.342 (0.134)	5.8	17% [6%, 29%]
Precise insurance	1 ^c	0.849 [0.73, 0.97]	3.415 (0.209)	4.0	25% [21%, 28%]

Each display is scored by the share of cross-property premium variance it reveals (λ^I), with 95% confidence intervals from a joint property-cluster bootstrap (500 replications) in brackets. A Bayesian respondent valuing insurance and rent dollars at par would exhibit a treatment effect equal to the information gain (column 2); column 3 reports observed treatment effects (Table 4, column 1) with delta-method standard errors clustered at the respondent level. The overstatement factor is the ratio of the observed to the expected effect, reported as a point estimate; the information share is its inverse, the fraction of the observed effect that improved information can explain, and inference on the decomposition operates through its bracketed bounds. Confidence bounds on the shares hold the observed effect fixed and move λ across its interval; nowhere in the confidence region does information account for more than roughly 30 percent of either effect. The beliefs and choice samples are independent, so combined uncertainty sums the bootstrap and delta-method variances.

^aSquared standardized slope of mean guessed flood scores on true premiums across 99 properties; the salience prime inherits the control value because it conveys no property-specific information, so its entire effect is non-informational by construction (^d).

^bLaw of total variance across flood-score bins, $1 - \mathbb{E}[\text{Var}(O \mid \text{score})]/\text{Var}(O)$.

^cBy construction: the disclosed premium is the true premium.

is the remainder, so the two sum to $\widehat{\Delta\text{ICPT}}^I$ by construction. The appendix contains the estimating equations and the confidence-interval methodology.

Table 5 presents the results: the information content of each disclosure method, the demand response relative to the control group predicted by the model, the observed demand response, and the ratio and share that summarize the gap. The flood score disclosure moved demand by 1.342 where its information content justifies a treatment effect of 0.230, and precise insurance moved demand by 3.415 where its information content justifies 0.849. Observed responses exceed information-justified responses by a factor of 5.8 and 4.0 respectively. The salience prime conveys no property-specific information by construction, so its entire effect – 0.231, only marginally significant – is non-informational. Figure 4 displays the same comparison from Columns 2 and 3 of Table 5 graphically.

We note the flood score condition’s near-benchmark ICPT of approximately 1.27 reflects not a well-calibrated demand response to a highly informative disclosure, but rather a coincidental combination of moderate informational content ($\hat{\lambda}^{\text{score}} = 0.381$: FloodFactor bins explain roughly 38% of cross-property premium variance) and an elevated non-informational response ($\hat{\gamma}^{\text{score}} \approx 3.92$) that happen to produce an ICPT near the benchmark. This has important implications for policy: a disclosure format that appears well-calibrated may be getting there through the “wrong” channel. If attention to flood risk were to decay over time or vary across populations, the flood score

ICPT would move away from 1, even holding information fixed. The precise insurance condition, by contrast, exceeds the benchmark robustly – its $\hat{\gamma}^{\text{prec}} = 3.566$ with 95% CI [3.14, 3.99] is well above 1 regardless of assumptions.

The conclusion is also robust to the sampling uncertainty in $\hat{\lambda}$. Holding $\widehat{\Delta\text{ICPT}}$ fixed, moving $\hat{\lambda}^{\text{ctrl}}$ across its 95% confidence interval [0.027, 0.275] bounds the information share of the precise-insurance treatment effect between 21% and 28%; moving $\hat{\lambda}^{\text{score}}$ across its interval [0.228, 0.534] bounds the information share of the flood score effect between 6% and 29%. Nowhere in the confidence region does the information channel account for more than roughly 30% of either treatment effect, so the conclusion that most of the demand response operates through non-informational channels does not depend on the point estimates.

7 Conclusion

This paper studies how different formats for disclosing flood risk in real estate listings affect demand for safer properties. Using a preregistered online experiment with 1,498 Prolific respondents making hypothetical purchase decisions among real Florida Gulf Coast properties, we estimate the insurance cost pass-through (ICPT) – willingness to pay in annualized dollars per dollar of annual flood insurance cost – for five disclosure conditions.

Model-free comparisons of chosen properties, identified because treatment is randomized against balanced menus, anchor the analysis. Absent disclosure, respondents are effectively indifferent to flood insurance costs, choosing homes whose premiums are statistically indistinguishable from the alternatives they reject. Disclosure induces large, dose-ordered premium avoidance – up to \$2,458 per year of avoided premium under exact-premium disclosure – partly paid for in higher rents. It also reallocates attention along two distinct margins: sorting on square footage weakens equally across all four formats regardless of informational content, while sensitivity to the displayed price falls only in the informative arms, by 30 to 40 percent.

Discrete choice model results reveal a clean monotone ordering: each successive format generates a higher ICPT, from near-zero in the control condition, to a modest effect from the general Florida reminder, to near-benchmark demand from the flood score display, to a response several times the benchmark from the precise insurance estimate. The ICPT of approximately 3.34 in the precise insurance condition is our headline finding: respondents who saw precise premium estimates were willing to pay more than three times the capitalization benchmark to avoid flood risk. This estimate uses the rent coefficient pooled across arms and is conservative: with arm-specific price sensitivity the precise-insurance treatment effect rises to roughly 4.0 (Table 4, column 4). No stable valuation of flood costs, whatever its level, can reconcile the responses across arms, since the properties and their true flood costs are identical across arms and only the display varies. This

result is stable across estimation methods, amenity controls, and subsamples (Appendix Table A4).

An interpretive decomposition suggests that improved information accounts for at most one quarter of any treatment effect: none of the general warning response, 17% of the flood score response, and 25% of the precise insurance response – equivalently, respondents moved four to six times as much as the information content of the displays justifies. The remainder operates through non-informational channels – heightened attention to the disclosed number, diminished attention to competing attributes, trust in numerical information, monetary framing – which we interpret as salience broadly construed but cannot separately identify. The flood score condition’s raw ICPT near 1 is doubly coincidental: a genuinely elevated response (amenity-corrected level of at least 1.34, driven by $\hat{\gamma}^{\text{score}} \approx 3.92$ against moderate informativeness $\hat{\lambda}^{\text{score}} = 0.381$) is partially offset by the amenity correlation common to all conditions – it is not a well-calibrated response to flood risk information. This finding has direct implications for disclosure policy: a format that appears well-calibrated may be achieving that appearance through the wrong channel, making it fragile to changes in framing, attention, or context.

Several limitations of this study should inform how the results are interpreted. First, the choices are hypothetical; respondents face no financial consequences for their decisions, which may inflate salience effects relative to real purchase decisions, and which attenuates sensitivity to the displayed price in particular, inflating dollar-denominated willingness-to-pay ratios through the denominator. Second, the sample is drawn from Prolific and skews toward more educated, climate-attentive respondents relative to the U.S. population, which may affect both the baseline salience of flood risk and the responsiveness to different disclosure formats. Third, the interpretive decomposition requires the maintained assumption of a benchmark response weight ($\gamma^{\text{bench}} = 1$) in the control condition; if the control group already exhibits elevated or depressed attention to flood risk, our estimates of γ^t in treatment conditions would be biased accordingly. Fourth, the properties are drawn from four Florida Gulf Coast counties, where flood risk is both high and widely discussed; the effects we document may be larger or smaller in other geographic contexts.

Despite these limitations, the results offer a clear message for the design of natural disaster risk disclosures. Precise numerical estimates of flood insurance costs generate the largest demand response, but that response is dominated by non-informational channels rather than information. Score-based formats like FloodFactor generate a demand response that, while also predominantly non-informational, lands closer to the capitalization benchmark in our sample. Policymakers seeking to promote well-calibrated demand responses to flood risk should consider not only how informative a disclosure format is, but also how strongly it elevates attention to flood risk in the decision context – and whether that elevation is likely to produce responses above or below the desired informational effect.

HYPOTHETICAL CHOICE 1 OF 10

	Property A	Property B
		
		
Price	\$415,000	\$415,000
Sq Ft	1535	1551
Beds	3	3
Baths	3	3

Which property would you be more likely to buy or recommend?

☐ Property A

☐ Property B

Figure 1: Survey screenshot: control condition

The survey screen shows two properties side-by-side, each displaying a property photograph, satellite map view, sales price, square footage, number of bedrooms, and number of bathrooms. Respondents in the control condition saw no additional information about flood risk beyond what was visually apparent from the photographs and satellite imagery.

HYPOTHETICAL CHOICE 2 OF 10




Figure 2: Survey screenshot: precise insurance estimate condition
The survey screen shows two properties side-by-side. In addition to the standard property information, respondents in the precise insurance estimate condition saw a single annual flood insurance cost estimate for each property, derived from the ReadYourRisk.com premium calculator (e.g. “Annual flood insurance cost: \$8,727”).




Figure 3: Per-property valuation effects versus true premium

Each point is the estimated valuation shift of one of the 99 properties under one disclosure condition, relative to the same property under the control condition, in thousands of dollars of annualized rent per year (divide by 0.06 for sale-price equivalents), plotted against the property's true annual flood insurance premium and colored by condition. Estimates come from Equation 2, a conditional logit with property fixed effects fully interacted with condition; the rent coefficient converts utility shifts to dollars. Solid lines are linear fits by condition, estimated on all points; their slopes appear as column 3 of Table 4, where they track the parametric treatment effects. The vertical axis is truncated at $[-15,000, 12,000]$; 71 of the roughly 400 property-condition effects, individually imprecise estimates from thin cells, fall outside the displayed range and are included in the fits. The flood score cloud exhibits substantially greater dispersion around its fit than the precise insurance cloud, reflecting the bin structure of the score display; Appendix Table A1 tests this formally.




Figure 4: Observed responses versus the Bayesian benchmark

For each disclosure format, the gray bar is the treatment effect a Bayesian respondent who values insurance and rent dollars at par would exhibit – the information the display reveals, $\lambda^t - \lambda^{\text{ctrl}}$ – and the crimson bar is the treatment effect we observe (ΔICPT , Table 4, column 1). Brackets report the ratio of observed to justified response. Whiskers are 95% confidence intervals: property-cluster bootstrap for the benchmark effects, delta-method with respondent clustering for the observed effects. The saliency prime conveys no property-specific information, so its justified effect is zero by construction and no ratio is defined; its entire response is non-informational. The insurance range condition is omitted because its displayed values are systematically biased relative to true premiums (Appendix Figure A3 and Appendix Table A5). All values appear in Table 5.

References

- Bae, Hyunhoe.** 2012. “Reducing Environmental Risks by Information Disclosure: Evidence in Residential Lead Paint Disclosure Rule.” *Journal of Policy Analysis and Management*, 31: 404–431.
- Bian, Xun, Justin C. Contat, Bennie D. Waller, and Scott A. Wentland.** 2023. “Why Disclose Less Information? Toward Resolving a Disclosure Puzzle in the Housing Market.” *Journal of Real Estate Finance and Economics*, 66: 443–486.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer.** 2013. “Salience and consumer choice.” *Journal of Political Economy*, 121: 803–843.
- Carrillo, Paul, Stephanie R. Cellini, and Richard K. Green.** 2013. “SCHOOL QUALITY AND INFORMATION DISCLOSURE: EVIDENCE FROM THE HOUSING MARKET.” *Economic Inquiry*, 51: 1809–1828.
- Cassidy, Alecia.** 2023. “How Does Mandatory Energy Efficiency Disclosure Affect Housing Prices?” *Journal of the Association of Environmental and Resource Economists*, 10: 655–686.
- Cocca, Paul.** 2023a. “Advantages of Sprinkles Flood Insurance Premium Estimate Maps over traditional FEMA flood insurance rate maps (FIRMs) and other flood risk maps.” *Read Your Risk, LLC*.
- Cocca, Paul.** 2023b. “What’s Not Transparent about Risk Rating 2.0?” *Read Your Risk, LLC*.
- Conell-Price, Lynn, and Philip Mulder.** 2024. “Does Good News Wash out the Bad? Belief Updating and Climate Risk Information.” *SSRN Electronic Journal*.
- DeGroot, Morris H.** 1970. *Optimal Statistical Decisions*. McGraw-Hill.
- DellaVigna, Stefano.** 2009. “Psychology and Economics: Evidence from the Field.” *Journal of Economic Literature*, 47: 315–72.
- Fairweather, Daryl, Matthew E Kahn, Robert D Metcalfe, and Sebastian Sandoval Olascoaga.** 2024. “Expecting Climate Change: A Nationwide Field Experiment in the Housing Market.” National Bureau of Economic Research Working Paper 33119.
- FEMA.** 2021. “National Flood Insurance Program: Risk Rating 2.0 Methodology and Data Sources.” U.S. Federal Emergency Management Agency.

- Fron del, Manuel, Andreas Gerster, and Colin Vance.** 2020. “The power of mandatory quality disclosure: Evidence from the German housing market.” *Journal of the Association of Environmental and Resource Economists*, 7: 181–208.
- Gallagher, Justin.** 2014. “Learning about an infrequent event: Evidence from flood insurance take-up in the United States.” *American Economic Journal: Applied Economics*, 6: 206–233.
- GAO.** 2023. “New Rate-Setting Methodology Improves Actuarial Soundness but Highlights Need for Broader Program Reform.” U.S. Government Accountability Office.
- Hainmueller, Jens, Daniel J. Hopkins, and Teppei Yamamoto.** 2014. “Causal Inference in Conjoint Analysis: Understanding Multidimensional Choices via Stated Preference Experiments.” *Political Analysis*, 22: 1–30.
- Haiken-DeNew, John, Syed Hasan, Nikhil Jha, and Mathias Sinning.** 2018. “Unawareness and selective disclosure: The effect of school quality information on property prices.” *Journal of Economic Behavior Organization*, 145: 449–464.
- Hopkins, Caroline, and Timothy Hyde.** 2026. “How do risk perceptions change with new information? Evidence from a survey of flood-prone counties.” *Journal of Risk and Uncertainty* 2026, 1–46.
- Horn, Diane P, and Baird Webel.** 2019. “Introduction to the National Flood Insurance Program (NFIP) Diane P. Horn Analyst in Flood Insurance and Emergency Management Introduction to the National Flood Insurance Program (NFIP).” Congressional Research Service.
- Howe, Peter D., Matto Mildenberger, Jennifer R. Marlon, and Anthony Leiserowitz.** 2015. “Geographic variation in opinions on climate change at state and local scales in the USA.” *Nature Climate Change*, 5: 596–603.
- Hu, Zhongchen.** 2022. “Social interactions and households’ flood insurance decisions.” *Journal of Financial Economics*, 144: 414–432.
- Kousky, Carolyn.** 2018. “Financing Flood Losses: A Discussion of the National Flood Insurance Program.” *Risk Management and Insurance Review*, 21: 11–32.
- Lin, Xiao.** 2020. “Feeling Is Believing? Evidence From Earthquake Shaking Experience and Insurance Demand.” *Journal of Risk and Insurance*, 87: 351–380.
- List, John A., and Craig A. Gallet.** 2001. “What experimental protocol influence disparities between actual and hypothetical stated values?” *Environmental and Resource Economics*, 20: 241–254.

- Ma, Lala, Margaret Walls, Matthew Wibbenmeyer, and Connor Lennon.** 2024. “Risk Disclosure and Home Prices: Evidence from California Wildfire Hazard Zones.” *Land Economics*, 100: 6–21.
- McClelland, Gary H., William D Schulze, and Don L. Coursey.** 1993. “Insurance for Low-Probability Hazards: A Bimodal Response to Unlikely Events.” Vol. 1, 95–116. Springer Netherlands.
- Murphy, James J., P. Geoffrey Allen, Thomas H. Stevens, and Darryl Weatherhead.** 2005. “A Meta-analysis of Hypothetical Bias in Stated Preference Valuation.” *Environmental and Resource Economics* 2005 30:3, 30: 313–325.
- Nanda, Anupam.** 2008. “Property condition disclosure law: Why did states mandate ‘Seller tell all’?” *Journal of Real Estate Finance and Economics*, 37: 131–146.
- Peer, Eyal, David Rothschild, Andrew Gordon, Zak Evernden, and Ekaterina Damer.** 2021. “Data quality of platforms and panels for online behavioral research.” *Behavior Research Methods* 2021 54:4, 54: 1643–1662.
- Pope, Jaren C.** 2008a. “Buyer information and the hedonic: The impact of a seller disclosure on the implicit price for airport noise.” *Journal of Urban Economics*, 63: 498–516.
- Pope, Jaren C.** 2008b. “Do Seller Disclosures Affect Property Values? Buyer Information and the Hedonic Model.” *Land Economics*, 84: 551–572.
- Snowberg, By Erik, Leeat Yariv, Marina Agranov, Alessandra Casella, Armin Falk, Guillaume Fréchette, Drew Fudenberg, Johannes Haushoffer, Salvatore Nunnari, and Nichole Szembrot.** 2021. “Testing the Waters: Behavior across Participant Pools.” *American Economic Review*, 111: 687–719.
- Sussman, Reuven, Steven Conrad, Christine Kormos, Celine Park, and Emma Cooper.** 2021. “Context and meaningfulness in energy efficiency labeling: Real estate listings.” *Journal of Environmental Psychology*, 78: 101681.
- U.S. EPA.** 1996. “EPA and HUD Move to Protect Children from Lead-Based Paint Poisoning; Disclosure of Lead-Based Paint Hazards in Housing.” *Fact Sheet*, EPA-747-F-96-002, Revised December 1996.
- Wing, Oliver EJ, Nicholas Pinter, Paul D Bates, and Carolyn Kousky.** 2020. “New insights into US flood vulnerability revealed from flood insurance big data.” *Nature communications*, 11(1): 1444.

Appendix – intended for online publication

Proof sketch: one-for-one information pass-through

This appendix sketches the proof of the one-for-one rule used in Section 6, along with the estimating equations and confidence-interval methodology behind Table 5.

Proposition. *Suppose an agent forms beliefs about a property's flood insurance cost by Bayesian updating from the signal their display provides, values expected insurance costs at par with rental costs, and chooses according to the conditional logit of Section 4. Then $ICPT^t = \lambda^t$, where λ^t is the share of cross-property premium variance explained by the signal. Consequently, a display that raises informational content from $\lambda^{t'}$ to λ^t raises the ICPT by exactly $\lambda^t - \lambda^{t'}$.*

The derivation below establishes the general version $ICPT^t = \gamma^t \lambda^t$, where γ^t is the response weight of Section 6; the proposition is the par case $\gamma^t = 1$.

Setup

Consider agent n choosing between property i and property j under disclosure condition t . Utility is:

$$U_{ni}^t = v_i - \gamma^t \cdot \mathbb{E}[o_i | x_i^t] + \varepsilon_{ni} \quad (5)$$

where v_i captures all visible property characteristics, o_i is the true annual flood insurance premium (the opaque component), x_i^t is the signal received under condition t , $\gamma^t \geq 0$ is the non-informational response weight (interpreted as salience broadly construed), and ε_{ni} is i.i.d. Type-I extreme value.

Bayesian updating

The agent has a prior over o_i with mean μ_0 (common across all properties, normalized to zero in the demeaned framework) and variance σ_0^2 . The signal $x_i^t = o_i + \eta_i^t$ where $\eta_i^t \sim \mathcal{N}(0, s_t^2)$ independently of o_i . Standard Bayesian updating gives:

$$\mathbb{E}[o_i | x_i^t] = \omega^t x_i^t + (1 - \omega^t) \mu_0 = \omega^t x_i^t \quad (6)$$

where the second equality uses $\mu_0 = 0$ (demeaned), and $\omega^t \equiv \sigma_0^2 / (\sigma_0^2 + s_t^2) \in [0, 1]$ is the Bayesian shrinkage factor. Note that $\mu_0 = 0$ is without loss of generality for binary choice: even with a nonzero common prior mean, the prior term cancels in any pairwise utility difference $U_{ni}^t - U_{nj}^t$ and does not affect the conditional logit likelihood.

Identification of the ICPT

Substituting the posterior mean into the utility function and taking pairwise differences:

$$U_{ni}^t - U_{nj}^t = (v_i - v_j) - \gamma^t \omega^t (x_i^t - x_j^t) + \varepsilon_{ni} - \varepsilon_{nj} \quad (7)$$

Under logit assumptions, the conditional logit identifies $-\gamma^t \lambda^t$ as the coefficient on the true premium, where $\lambda^t \equiv \text{Cov}(\mathbb{E}[o_i | x_i^t], O_i) / \text{Var}(O_i)$ is the projection slope of posterior beliefs on the true premium (equal to ω^t in the canonical noisy-signal model). The coefficient on annualized rent identifies α . The ICPT is therefore:

$$\text{ICPT}^t = \gamma^t \lambda^t \quad (8)$$

Equivalence of signal and true value as regressors

In practice, we use the true insurance premium O_i as the regressor rather than the signal x_i^t , since O_i is observed but x_i^t is not (respondents see the signal but we do not observe which aspects of the property drove their beliefs). Since $x_i^t = O_i + \eta_i^t$ with $\eta_i^t \perp O_i$, we have:

$$\frac{\text{Cov}(x_i^t, O_i)}{\text{Var}(O_i)} = \frac{\text{Var}(O_i)}{\text{Var}(O_i)} = 1 \quad (9)$$

so the projection of beliefs on O_i has slope λ^t and the omitted belief noise is orthogonal to O_i . In a linear model this produces no attenuation. In the logit, the omitted noise enters the latent index and inflates the effective error variance, which rescales *all* coefficients by a common factor; because the ICPT is a ratio of coefficients, this factor cancels and the ICPT is unaffected. We verify this via simulation across multiple seeds and signal-noise levels: a conditional logit using the true value as regressor recovers $\gamma^t \lambda^t$ (exponent one, not $\gamma^t (\lambda^t)^2$) within sampling error in every case. A previous single-seed simulation suggested an extra λ^t factor, but that was an artifact of holding a single 100-property noise draw fixed across all choice tasks, which induces a spurious finite-sample correlation between the omitted noise and O_i .

Decomposition

Given $\widehat{\text{ICPT}}^t$ from the clogit and $\hat{\lambda}^t$ from the beliefs data, the non-informational response weight is:

$$\hat{\gamma}^t = \frac{\widehat{\text{ICPT}}^t}{\hat{\lambda}^t} \quad (10)$$

with $\hat{\lambda}^t$ in the denominator (exponent one), consistent with the no-attenuation result above.

The benchmark ICPT under the capitalization normalization ($\gamma^{\text{bench}} = 1$) is $\hat{\lambda}^{\text{ctrl}} = 0.151$. Using $\widehat{\Delta\text{ICPT}}^t = \widehat{\text{ICPT}}^t - \widehat{\text{ICPT}}^{\text{ctrl}}$ to eliminate the amenity confound:

$$\text{implied ICPT}^t = \hat{\lambda}^{\text{ctrl}} + \widehat{\Delta\text{ICPT}}^t \quad (11)$$

$$\hat{\gamma}^t = \frac{\text{implied ICPT}^t}{\hat{\lambda}^t} \quad (12)$$

$$\text{Information component} = \hat{\lambda}^t - \hat{\lambda}^{\text{ctrl}} \quad (13)$$

$$\text{Non-informational component} = \widehat{\Delta\text{ICPT}}^t - \text{Information component} \quad (14)$$

The information and non-informational components sum to $\widehat{\Delta\text{ICPT}}^t$ by construction.

Confidence intervals

Standard errors for $\hat{\lambda}^{\text{ctrl}}$ and $\hat{\lambda}^{\text{score}}$ are obtained by a joint bootstrap (500 replications) resampling the 99 property clusters, which also yields their covariance. Property-level resampling is the appropriate choice because both estimands are functions of the property-level distribution of true premiums and beliefs; guesser-level clustering would understate the uncertainty by conditioning on the fixed set of 99 properties. Standard errors for $\widehat{\Delta\text{ICPT}}^t$ come from the clogit delta method with respondent-level clustering. Since the beliefs data and choice data are independent samples, the variance of any derived quantity is the sum of the beliefs-side variance (incorporating the $\hat{\lambda}^{\text{ctrl}}-\hat{\lambda}^{\text{score}}$ covariance) and the choice-side variance. All confidence intervals are computed via the delta method applied to this combined variance.

Estimation of per-property valuation effects

The estimates underlying Figure 3 and Table A1 come from Equation 2, estimated by conditional logit with standard errors clustered at the respondent level. The specification contains 99 property fixed effects and approximately 400 property-by-condition interactions; the likelihood is globally concave and the model converges without difficulty, though a small number of thin property-condition cells (roughly 60 observations each) produce imprecise individual estimates. Results in Table A1 are unchanged when cells with non-computable robust standard errors are excluded. Because Equation 3 is the special case of Equation 2 in which the dollar effects per condition are restricted to lie on a line through the premium, the close agreement between the fitted slopes in Figure 3 and the parametric treatment effects in Table 4 constitutes a direct test of that restriction.

Table A1: Valuation responses track the disclosed object

	Flood score arm	Precise insurance arm
Within-bin premium slope (\$/\$)	−0.043 (0.157) [$p = 0.79$]	−3.027 (0.149) [$p < 0.001$]
Score-bin fixed effects, joint $F(9, 88)$	102.4 [$p < 0.001$]	1.41 [$p = 0.20$]
R^2	0.92	0.88
Properties	99	99

Each column reports an OLS regression of property-level valuation shifts (relative to control, in annualized rent dollars, estimated from Equation 2) on the property's true annual premium, absorbing flood-score-bin fixed effects. In the flood score arm, respondents saw only the ten-point score: bin fixed effects explain nearly all cross-property variation and the residual premium slope is zero. In the precise insurance arm, respondents saw the premium itself: the premium slope is large and precisely estimated, and score bins add no explanatory power conditional on the premium. Estimated valuation shifts are generated regressors; standard errors are unadjusted for first-stage estimation error and the tests should be read as descriptive.




Figure A1: Calibration plot for respondent predictions of Flood Factor

The figure displays average respondent-predicted flood scores against the true FloodFactor score for each of the 99 properties. Small dots show the property-level average prediction from respondents in the beliefs elicitation sample; large filled circles show the mean across all properties with the same true FloodFactor value. The diagonal line represents perfect calibration. Respondent predictions are positively correlated with true flood risk but systematically compressed toward the middle of the scale, consistent with partial Bayesian updating from visual signals. The slope of a standardized regression of true premiums on mean predicted scores across properties is 0.389; squaring this correlation gives our estimate of the visual informativeness weight, $\hat{\lambda}^{\text{ctrl}} = 0.151$, used in Section 6.

Table A2: Conditional logit models

	(1)	(2)	(3)	(4)
Annual rental equivalent of sales price	-9.34e-05*** (3.7e-06)	-9.63e-05*** (3.8e-06)	-0.000102*** (3.9e-06)	-0.000102*** (3.9e-06)
Annual insurance cost	6.9e-06 (8.4e-06)	4.31e-05*** (8.7e-06)		
Ins. cost×I{Condition = Salience prime}	-2.16e-05* (1.17e-05)	-2.11e-05* (1.16e-05)	-2.13e-05* (1.19e-05)	-2.00e-05* (1.19e-05)
Ins. cost×I{Condition = Insurance range}	-2.75e-05** (1.18e-05)	-2.83e-05** (1.19e-05)	-2.76e-05** (1.22e-05)	-2.63e-05** (1.22e-05)
Ins. cost×I{Condition = Flood score}	-0.000125*** (1.16e-05)	-0.000127*** (1.18e-05)	-0.000129*** (1.23e-05)	-0.000128*** (1.23e-05)
Ins. cost×I{Condition = Precise insurance}	-0.000319*** (1.65e-05)	-0.000326*** (1.64e-05)	-0.000338*** (1.63e-05)	-0.000337*** (1.63e-05)
Year built	0.00747*** (0.000808)	0.00285*** (0.000897)	0.0726*** (0.0105)	0.0242** (0.0109)
Number of baths	0.0667** (0.0328)	0.0716** (0.0331)	0.0937*** (0.0339)	0.0938*** (0.0340)
Square footage	0.000622*** (3.89e-05)	0.000640*** (3.96e-05)	0.000678*** (4.12e-05)	0.000678*** (4.11e-05)
Flexible polynomial distance controls		✓		✓
Property-level fixed effects			✓	✓
Observations	29,420	29,420	29,420	29,420
Pseudo R^2	0.123	0.136	0.178	0.178
Akaike information criterion	17907	17636	16973	16970
Bayesian information criterion	17990	17735	17852	17832

Conditional logit models with standard errors clustered at the respondent level in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Annual rental equivalent of sales price calculated by multiplying sales price by 0.06. Annual insurance cost is given by the insurance premium calculator on ReadYourRisk.com (see Section 3.4). Data on number of baths, square footage, and sales price of each property are independently randomly imputed for each respondent in keeping with the overall distribution of property attributes. Year built is a real, non-displayed property characteristic observable to respondents only through listing photographs. Fourth-order polynomial controls for distance from the water are included in Specifications 2 and 4. Property-level fixed effects are included in Specifications 3 and 4 to better control for idiosyncratic features of each property. The insurance cost main effect is absorbed by the property fixed effects in Specifications 3 and 4, since insurance cost varies only at the property level; the interaction coefficients in those columns measure treatment effects relative to the control group directly. Columns (1) and (3) correspond to columns (1) and (2) of Table 4.

Table A3: Conditional logit with condition-specific attribute sensitivity

	(1) Rent interactions	(2) All interactions
Annual rental equivalent of sales price	$-1.118 \times 10^{-4}***$ (8.65×10^{-6})	$-1.135 \times 10^{-4}***$ (8.72×10^{-6})
Rent $\times$ Salience prime	2.36×10^{-6} (1.18×10^{-5})	3.98×10^{-6} (1.19×10^{-5})
Rent $\times$ Insurance range	1.43×10^{-5} (1.17×10^{-5})	1.52×10^{-5} (1.18×10^{-5})
Rent $\times$ Flood score	$4.42 \times 10^{-5}***$ (1.14×10^{-5})	$4.70 \times 10^{-5}***$ (1.15×10^{-5})
Rent $\times$ Precise insurance	$3.40 \times 10^{-5}***$ (1.27×10^{-5})	$3.58 \times 10^{-5}***$ (1.28×10^{-5})
Ins. cost $\times$ Salience prime	$-2.19 \times 10^{-5}*$ (1.17×10^{-5})	$-2.30 \times 10^{-5}*$ (1.17×10^{-5})
Ins. cost $\times$ Insurance range	$-2.75 \times 10^{-5}**$ (1.18×10^{-5})	$-2.85 \times 10^{-5}**$ (1.19×10^{-5})
Ins. cost $\times$ Flood score	$-1.219 \times 10^{-4}***$ (1.17×10^{-5})	$-1.221 \times 10^{-4}***$ (1.17×10^{-5})
Ins. cost $\times$ Precise insurance	$-3.128 \times 10^{-4}***$ (1.69×10^{-5})	$-3.131 \times 10^{-4}***$ (1.70×10^{-5})
Sqft $\times$ Salience prime		$-3.62 \times 10^{-4}***$ (1.31×10^{-4})
Sqft $\times$ Insurance range		$-3.33 \times 10^{-4}***$ (1.29×10^{-4})
Sqft $\times$ Flood score		$-5.11 \times 10^{-4}***$ (1.22×10^{-4})
Sqft $\times$ Precise insurance		$-3.91 \times 10^{-4}***$ (1.34×10^{-4})
Implied Δ ICPT (arm-specific denominator):		
Salience prime	0.200* (0.109)	
Insurance range	0.282** (0.124)	
Flood score	1.803*** (0.255)	
Precise insurance	4.018*** (0.471)	
Rent interactions $\chi^2(4)$ [p]	23.35 [0.000]	25.47 [0.000]
Sqft interactions $\chi^2(4)$ [p]		18.02 [0.001]
Baths interactions $\chi^2(4)$ [p]		0.95 [0.917]
Year-built interactions $\chi^2(4)$ [p]		0.05 [1.000]
Observations	29,420	29,420
Log pseudolikelihood	-8926.6	-8915.2

Specification 1 of Appendix Table A2 augmented with attribute-by-condition interactions. Column 1 interacts the rent coefficient with condition; column 2 interacts rent, square footage, bathrooms, and year built (baths and year-built interactions, individually and jointly null, omitted from the display; insurance cost main effect and non-interacted controls included in both columns but not shown). Standard errors clustered at the respondent level; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Positive rent interactions indicate reduced price sensitivity relative to control: rent sensitivity falls by roughly 40 percent in the flood score condition and 30 percent in the precise insurance condition, and is unchanged in the salience prime and insurance range conditions. Square-footage sensitivity falls in every treatment arm, including the contentless salience prime, by statistically similar amounts, while bathroom and year-built sensitivities are unchanged everywhere – the two-channel pattern discussed in Section 5.2. The implied treatment effects using each arm's own rent coefficient as the denominator appear in Table 4, column 4. The pattern is not a logit-scale artifact: the rent coefficient shrinks while the premium coefficient grows in the same arms, movements in opposite directions that common rescaling cannot produce.

Table A4: Robustness of treatment effects across subsamples

Sample	Δ ICPT flood score	Δ ICPT precise insurance	Respondents	Decisions
Baseline (Table 4, col. 1)	1.342 (0.134)	3.415 (0.209)	1,471	14,710
First five choices only	1.390 (0.198)	3.295 (0.271)	1,471	7,355
Last five choices only	1.305 (0.173)	3.525 (0.286)	1,471	7,355
Excluding fastest quartile of completers	1.337 (0.153)	3.288 (0.232)	1,103	11,030
Excluding top-coded premium properties	1.711 (0.241)	4.073 (0.336)	1,471	8,673

Each row re-estimates the pooled conditional logit (Table 4, column 1) on the indicated subsample and reports the flood score and precise insurance treatment effects with delta-method standard errors in parentheses. First-half versus second-half rows test for learning or fatigue across the ten choice tasks. The completion-speed exclusion drops the fastest quartile by total survey time beyond the attention screens already imposed on the full sample. The final row drops the 23 properties whose ReadYourRisk premium sits at the top-coded maximum of \$8,736; the gradients are, if anything, steeper without them, indicating that the response is not driven by the mass point at the top of the premium distribution.

Table A5: The treatment displays as signals of true insurance cost

Arm	Example display text	Source (vintage)	Correlation with true premium	Mean displayed / mean true
Flood score	“Flood score: 9/10 (extreme risk)”	First Street Foundation (2023)	0.62 ^a	–
Insurance range	“Annual flood insurance cost: \$1358 to \$5099”	Redfin listing estimates (Nov. 2023)	0.18 ^b	0.41
Precise insurance	“Annual flood insurance cost: \$8727”	ReadYourRisk calculator (2023)	1 ^c	1

Each property-specific display scored as a signal of the property’s true annual premium across the 99 properties. The salience prime is omitted: it conveys no property-specific information. Column 4 reports the cross-property correlation between the displayed object and the true premium; column 5 the ratio of the mean displayed dollar value to the mean true premium, a measure of level bias (undefined for the unitless score). ^aSquare root of $\hat{\lambda}^{\text{score}} = 0.381$, the share of premium variance explained by score bins. ^bCorrelation of the Redfin range midpoint with the true premium (endpoint correlations are similar: 0.20 and 0.17); the display is both systematically biased downward in level (Figure A3) and nearly uninformative about the cross-property ranking of premiums, which is why no λ in premium-variance units is computed for this arm. ^cBy construction.




Figure A2: Flood insurance premium by distance to the water

The figure shows a scatterplot of 99 properties relating distance to the water in meters to estimated annual flood insurance premium, along with a smoothed best fit line with bandwidth = 0.8. Property distance to the water is displayed on a log scale; properties closer than 10 m to the water are displayed at 10 m. Property-level premium estimates are based on the premium calculator on ReadYourRisk.com. The figure indicates the generally negative but noisy relationship between flood insurance premiums and distance to the water.

Table A6: Balance of offered menus across treatment arms

	(1) Insurance premium (\$)	(2) Annual rent (\$)	(3) Square footage	(4) Baths	(5) Year built
Salience prime	67.0 (46.2)	92.8 (81.4)	-4.6 (6.7)	-0.012 (0.007)	-0.03 (0.26)
Insurance range	60.0 (46.7)	30.6 (81.0)	-12.0 (6.9)	-0.004 (0.007)	0.29 (0.26)
Flood score	63.3 (46.5)	110.7 (84.1)	-10.9 (6.7)	-0.004 (0.007)	0.09 (0.26)
Precise insurance	30.3 (46.2)	40.4 (81.0)	-5.8 (6.7)	0.003 (0.007)	0.43 (0.27)
Joint $F(4, 1470)$	0.76	0.61	1.03	1.17	1.17
p value	0.55	0.66	0.39	0.32	0.32

Each column is a separate OLS regression of the indicated attribute of *offered* properties (all 29,420 option-level observations, chosen and rejected alike) on condition indicators, with standard errors clustered at the respondent level in parentheses. Because treatment cannot causally affect which properties and imputed attributes were offered, any significant difference would indicate a randomization or imputation failure. Sales price, square footage, and bathrooms are balanced by construction (independent random imputation); insurance premium and year built are real property characteristics whose balance reflects the randomization of treatment assignment. No attribute differs jointly across arms at conventional levels ($p > 0.32$ in every column); with twenty individual coefficients, isolated marginal differences are consistent with chance.




Figure A3: Comparison of Redfin flood insurance ranges with precise estimates

The figure displays the distribution of precise annual flood insurance premium estimates (from ReadYourRisk.com) for properties grouped by the Redfin insurance cost range their listing displayed; Redfin showed only four distinct ranges across the 99 properties, and two of them are nearly identical. Gray dots are individual property estimates; dashed black brackets mark the displayed range itself; filled circles mark column medians, labeled with values and omitted for the single-property range; the premium axis is on a log scale. In the two lowest ranges, which together cover 64 of the 99 properties, the median precise estimate exceeds the *top* of the displayed range – by roughly 18 percent in the first (\$1,072 against a \$905 ceiling, $n = 49$) and by a factor of nine in the second (\$8,723 against \$928, $n = 15$) – even though the two displayed ranges are nearly the same. The middle range's median falls within its bounds, with a heavy right tail of top-coded properties. This downward level bias, together with the near-absence of rank information (Appendix Table A5), motivates our exclusion of the insurance range condition from the benchmark comparison in Section 6.

Table A7: Sample composition: properties

Variable	Data source	Mean	Median	Std dev	Min	Max
Transaction characteristics						
Sales price (\$k)	MLS	404	390	75	303	600
Year of sale	MLS	2021.4	2021	0.5	2021	2022
Structure characteristics						
Square footage	MLS	1766	1694	384	1009	3323
Number of beds	MLS	3.0	3.0	0.0	3.0	3.0
Number of baths	MLS	2.2	2.0	0.4	2.0	4.0
Number of stories	MLS	1.1	1.0	0.3	1.0	3.0
Year built	MLS	1997	2001	16	1950	2022
Hydrological characteristics						
Distance to water (km)	FGDL	2.3	1.2	2.9	0	13.4
Elevation (m)	Google Earth API	2.1	1.8	1.1	0.3	5.5
Flood Factor	First Street Fdn	5.5	6	2.9	1	10
Estimated annual insurance premia (\$)						
ReadYourRisk estimate	ReadYourRisk	3914	2028	3343	448	8736
Redfin (low end estimate)	Redfin listing	796	441	593	435	4719
Redfin (high end estimate)	Redfin listing	2424	928	2081	905	8298
Community rating system discount (%)	FEMA	22	25	6	5	25
County distribution						
Charlotte	13%					
Lee	49%					
Manatee	10%					
Sarasota	28%					

Summary statistics across 99 properties featured in the survey. Included properties were randomly drawn from a pool of transacted properties matching certain geographic, price, and size filters (see Section 3.2 for more details). Most property-level information derived from local multiple listing services (MLS). Distance to water of each property calculated in QGIS using shapefiles from the Florida Geographic Data Library at the University of Florida (FGDL). Flood Factor ratings determined by the First Street Foundation. Precise insurance cost estimates obtained from the calculator on ReadYourRisk.com. Twenty-three properties share the maximum ReadYourRisk estimate of \$8,736, a mass point at the top of the calculator's output for this sample; Appendix Table A4 shows that excluding these properties, if anything, steepens the estimated response gradients.

Table A8: Sample composition: respondents

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Full sample	Control	General reminder	Flood Factor	Insurance range	Insurance estimate	F-test <i>p</i> value
	<i>(n = 1,498)</i>	<i>(n = 302)</i>	<i>(n = 299)</i>	<i>(n = 300)</i>	<i>(n = 301)</i>	<i>(n = 296)</i>	
Age	41 (13)	40 (13)	40 (13)	40 (13)	40 (12)	42 (14)	0.20
Local flood risk (1-5)	1.90 (0.97)	1.93 (0.95)	1.88 (0.96)	1.87 (0.96)	1.93 (1.01)	1.86 (1)	0.83
Climate affirmer	0.85	0.82	0.88	0.86	0.85	0.85	0.50
Female	0.54	0.49	0.60	0.53	0.52	0.55	0.09
Florida resident	0.08	0.07	0.08	0.08	0.07	0.10	0.71
Foreign born	0.06	0.08	0.07	0.05	0.06	0.06	0.82
Nonwhite	0.34	0.33	0.37	0.34	0.35	0.31	0.60
Other native language	0.05	0.05	0.04	0.06	0.04	0.05	0.76

Descriptive statistics for the respondent sample, both in full (Column 1) and by treatment condition (Columns 2–6). Local flood risk refers to the respondent's perception of local flood risk in their home community on a five-point scale, with 5 indicating the greatest level of risk. Climate affirmers were those that agreed with a statement that climate change is occurring from the Yale Program on Climate Change Communication.

HYPOTHETICAL CHOICE 3 OF 10


	Property A	Property B
		
		
Price	\$380,000	\$369,000
Sq Ft	1300	1712
Beds	3	3
Baths	2	3
Annual flood insurance cost	\$1358 to \$5099	\$441 to \$928

Figure A4: Survey screenshot: insurance range condition
The survey screen shows two properties side-by-side. In addition to the standard property information, respondents in the insurance range condition saw the annual flood insurance cost range from Redfin for each property, displayed as a dollar range (e.g. “Annual flood insurance cost: \$1358 to \$5099”).


	Property A	Property B
		
		
Price	\$475,000	\$445,000
Sq Ft	1652	1774
Beds	3	3
Baths	2	3
Flood score	5/10 (major risk)	9/10 (extreme risk)

Figure A5: Survey screenshot: flood score condition

The survey screen shows two properties side-by-side. In addition to the standard property information, respondents in the flood score condition saw the FloodFactor score for each property, displayed as a score out of ten with a verbal risk characterization (e.g. “Flood score: 9/10 (extreme risk)”).



Figure A6: Location of 99 Florida properties

The figure above shows the approximate geographic location of each of the 99 properties randomly selected for inclusion in the survey, along with lines indicating distance and direction to the nearest coastal point of each property. The color of each point reflects the property's Flood Factor, with darker shades indicating higher-risk properties. Properties were randomly drawn from a list of thousands of properties included in the Zillow ZTRAX database with the following features: located in one of four Florida counties for which MLS data were readily available (Manatee, Sarasota, Charlotte, and Lee), listed as having three bedrooms, sold between January 2021 and August 2022, and sales price between \$300,000 and \$600,000.