

Licensing to Disclose: Do State Flood Risk Disclosure Laws Capitalize into Housing Values?*

APEP Autonomous Research[†] @ai1scl

February 12, 2026

Abstract

Do mandatory flood risk disclosure laws affect housing prices in flood-prone areas? I exploit staggered adoption of seller disclosure requirements across 30 U.S. states from 1992 to 2024 in a triple-difference design, comparing housing values in high- versus low-flood-exposure counties, before and after law adoption, relative to never-treated states. Using the Zillow Home Value Index for 3,072 counties over 25 years, I find no statistically significant effect of disclosure on the relative price of flood-exposed housing (DDD coefficient: 0.0072 log points, SE = 0.0091). An event study shows flat pre-trends, validating the parallel trends assumption. A placebo test on zero-flood counties shows precisely zero effects, validating the design. These results suggest that either flood risk was already priced prior to mandated disclosure or that disclosure laws operate through channels other than aggregate price adjustment.

JEL Codes: R31, Q54, D83, K32

Keywords: flood risk, information disclosure, housing prices, difference-in-differences, natural hazards

*This paper is a revision of [APEP Working Paper 0240 v1](#). Changes in this version: fixed compilation errors, tightened introduction, corrected table formatting.

[†]Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch

1. Introduction

In many states, a home seller must disclose a cracked foundation, a faulty furnace, or a leaky roof—but can remain silent about a history of catastrophic flooding. Thirty U.S. states have now moved to close this gap through mandatory seller disclosure laws, requiring sellers to reveal whether a property sits in a FEMA-designated flood zone, whether it has experienced prior flood damage, and whether flood insurance is required. Adoption occurred in three distinct waves—a first wave in the 1990s, a second in the early 2000s, and a third beginning in 2019 ([Natural Resources Defense Council, 2024](#)). The logic is straightforward: if buyers lack information about flood risk, mandating disclosure should cause flood-prone property values to decline as the market incorporates previously hidden hazard information. Yet whether these laws actually change how markets price flood risk remains an open question.

This paper asks a simple question: do these disclosure laws actually change the relative price of flood-exposed housing? I exploit the staggered adoption of flood risk disclosure requirements across U.S. states in a triple-difference (DDD) framework. The three margins of differencing are: (1) states that adopt disclosure laws versus those that never do; (2) counties within treated states that have high versus low historical flood exposure; and (3) the period before versus after law adoption. This design isolates the causal effect of disclosure on the relative capitalization of flood risk into housing values, while absorbing both county-level time-invariant characteristics and state-by-year shocks common to all counties within a state.

The identification strategy relies on a key insight: if disclosure laws provide new information to buyers about flood risk, the effect should be concentrated in counties where the gap between actual flood exposure and buyer awareness is largest—that is, in high-flood-exposure counties where sellers were previously able to conceal material hazard information. Low-flood-exposure counties in the same state serve as a within-state control, absorbing any state-level economic shocks that coincide with law adoption. Never-treated states provide the counterfactual trend for the treated-state counties.

I construct a county-year panel spanning 2000–2024 by merging three data sources. The outcome variable is the log of the Zillow Home Value Index (ZHVI), a repeat-sales-based measure of the typical home value in each county, available at monthly frequency for over 3,000 counties ([Zillow Research, 2024](#)). The treatment variable is the year each state first required sellers to disclose flood-related information, compiled from the Natural Resources Defense Council’s State Flood Disclosure Scorecard, state statute research, and the National Association of Realtors ([Natural Resources Defense Council, 2024](#)). Flood exposure is measured using pre-treatment (1953–1991) FEMA disaster declarations, which provide an exogenous proxy for historical flood risk that predates all disclosure laws in my sample.

The main finding is a null that is precise enough to be informative. The DDD coefficient on the interaction of post-adoption, high-flood-exposure, and treated-state indicators is 0.0072 log points (SE = 0.0091, $p = 0.43$)—for the median home valued at \$175,000, this implies a price change of roughly \$1,300, statistically indistinguishable from zero. Disclosure laws had no significant effect on the relative price of housing in flood-prone counties. The point estimate is positive, not negative as the information-revelation hypothesis would predict, though I cannot reject zero. Using the any-flood-exposure indicator (a broader definition), the estimate is 0.030 log points ($p = 0.055$)—marginally significant and positive, suggesting that if anything, flood-exposed counties in treated states experienced *relative price increases* after disclosure.

An event study specification shows flat pre-trends, with all pre-treatment coefficients statistically indistinguishable from zero, validating the parallel trends assumption. Post-treatment coefficients are small and positive but imprecise. The Callaway-Sant’Anna heterogeneity-robust estimator yields a positive overall average treatment effect on the treated (ATT = 0.097, SE = 0.022), though this estimate is sensitive to the inclusion of small treatment cohorts. A battery of robustness checks—including a placebo test on zero-flood counties, treatment intensity using NRDC disclosure grades, restriction to the most recent adoption wave, and two-way clustering—all support the null or mildly positive interpretation.

How should we interpret a null result? Three explanations merit consideration. First, flood risk may have already been substantially capitalized into housing values prior to mandated disclosure, through informal channels: buyers’ own research, real estate agents’ knowledge, visible flood damage, FEMA flood maps, and federal flood insurance requirements. If the pre-existing information environment was already reasonably efficient, mandated disclosure adds little incremental information. Second, the disclosure laws may operate primarily through liability and transaction channels rather than price adjustment—reducing post-sale litigation, clarifying insurance obligations, and shifting risk between buyers and sellers without altering equilibrium prices. Third, the positive point estimates hint at a sorting or confidence mechanism: disclosure may attract risk-aware buyers to flood-prone areas by reducing uncertainty about hazard status, or it may signal that homes have survived floods and carry appropriate insurance, raising rather than lowering their perceived value.

This paper contributes to three literatures. First, it adds to the growing body of work on information disclosure and housing markets. [Pope and Huang \(2020\)](#) show that California’s requirement to disclose earthquake fault zone proximity reduced prices by 4–8%, demonstrating that disclosure can have large effects when the information gap is wide. [Bin and Polasky \(2004\)](#) and [Atreya and Ferreira \(2015\)](#) find mixed evidence on whether flood zone designation affects property values, with effects varying by awareness and experience. My contribution is

to provide the first national-scale DDD estimate exploiting cross-state variation in disclosure mandates, finding that—unlike earthquake risk—flood disclosure appears not to generate large price adjustments, possibly because the underlying information is more readily available to buyers.

Second, I contribute to the literature on flood risk capitalization. [Bernstein et al. \(2019\)](#) document a 7% discount for sea-level-rise-exposed properties, concentrated in non-believer areas, suggesting that climate beliefs mediate price effects. [Ortega and Täspinar \(2018\)](#) estimate that Hurricane Sandy reduced prices in flood zones by 8–18%. [Gallagher and Hartley \(2017\)](#) show that flood insurance take-up spikes after nearby disasters but decays rapidly. [Kousky and Shoemaker \(2019\)](#) document how NFIP premium subsidies distort risk pricing, creating a wedge between actuarial risk and the price signal buyers observe. [Giglio et al. \(2021\)](#) find that climate risk is only weakly reflected in long-run discount rates for real estate, suggesting limited market attention to slow-moving hazards. My results complement this work by showing that legislative disclosure mandates—as opposed to direct experience with floods—have limited marginal effects on prices.

Third, I add to the policy evaluation literature on real estate disclosure requirements. [Northcraft and Neale \(1987\)](#) established that anchoring effects in real estate valuation are strong. [Barwick and Pathak \(2015\)](#) show that Massachusetts’ mandatory disclosure regime reduced information asymmetries but had modest price effects. [Sun et al. \(2022\)](#) find heterogeneous effects of disclosure across housing market segments. My findings align with the emerging view that disclosure is a necessary but insufficient condition for full risk capitalization, and that behavioral and institutional factors mediate the translation of information into prices.

2. Institutional Background

2.1 Flood Risk in the United States

Flooding is the most costly natural disaster in the United States. Between 2000 and 2023, FEMA issued over 1,200 major disaster declarations with a flood component, affecting every state and territory. The National Flood Insurance Program (NFIP), administered by FEMA, provides subsidized flood insurance to communities that adopt minimum floodplain management standards. As of 2024, the NFIP had over 4.7 million policies in force with total coverage exceeding \$1.3 trillion ([Federal Emergency Management Agency, 2024](#)).

Despite the pervasiveness of flood risk, information about property-level exposure is unevenly distributed. FEMA Flood Insurance Rate Maps (FIRMs) delineate Special Flood Hazard Areas (SFHAs) where flood insurance is mandatory for federally-backed mortgages,

but these maps are frequently outdated—some date to the 1970s—and their accuracy varies substantially across regions (Wing et al., 2018). Outside SFHAs, buyers may have little information about a property’s flood history, particularly for pluvial and fluvial flooding not captured in FEMA’s coastal and riverine models.

2.2 State Flood Risk Disclosure Laws

Beginning in the early 1990s, states began requiring home sellers to disclose flood-related information as part of mandatory property condition disclosure forms. These requirements vary considerably in scope and specificity. At one extreme, states like Oklahoma and California (NRDC grade A) require sellers to disclose whether the property is in a flood zone, whether it has experienced prior flood damage, whether flood insurance is currently in force, and specific information about flood mitigation measures. At the other extreme, states like Connecticut (grade D) include only a generic question about “any knowledge of flooding” within a broader disclosure form.

Table 4 documents the adoption timeline across 30 treated states. The pattern reveals three distinct waves of adoption:

First Wave (1992–1999): Sixteen states adopted flood disclosure requirements during this period, typically as part of broader property condition disclosure reforms (fifteen appear in the analysis panel; Alaska is excluded due to limited ZHVI coverage). Kentucky was the first in 1992, followed by Indiana, Michigan, Ohio, Oregon, and South Dakota in 1993. This wave was driven by consumer protection concerns and the professionalization of real estate transactions.

Second Wave (2000–2006): Five states—Delaware, New York, Louisiana, Mississippi, and Tennessee—added flood disclosure during this period. Notably, Louisiana and Mississippi adopted requirements before Hurricane Katrina (2005), though the post-Katrina policy environment may have strengthened enforcement.

Third Wave (2019–2024): Ten states adopted or significantly strengthened flood disclosure requirements in the most recent wave. This wave was driven by growing attention to climate change, increasing flood losses, and advocacy by the Natural Resources Defense Council, which published its first State Flood Disclosure Scorecard in 2020. Texas, South Carolina, and North Dakota adopted in 2019; Hawaii in 2022; Maine and New Jersey in 2023; and Florida, New Hampshire, North Carolina, and Vermont in 2024.

Nineteen states remain without mandatory flood risk disclosure as of 2024: Alabama, Arizona, Arkansas, Colorado, Georgia, Idaho, Kansas, Massachusetts, Maryland, Minnesota, Missouri, Montana, New Mexico, Rhode Island, Utah, Virginia, Wisconsin, West Virginia, and Wyoming. These never-treated states form the control group in my analysis.

The never-treated states are geographically diverse, spanning the Mountain West (Arizona, Colorado, Idaho, Montana, Utah, Wyoming), the upper Midwest (Minnesota, Wisconsin), the South (Alabama, Arkansas, Georgia, Virginia, West Virginia), and New England (Massachusetts, Rhode Island). This geographic diversity mitigates concerns that the control group is systematically different from treated states in terms of flood exposure or housing market characteristics. Indeed, several never-treated states have substantial flood risk—Georgia, Virginia, and Maryland have experienced major flood disasters—but have not adopted mandatory seller disclosure requirements.

2.3 Content of Disclosure Requirements

The specific content of flood disclosure varies considerably across states. At the most comprehensive end, states like California and Oklahoma require sellers to complete detailed forms that include specific questions about: (1) whether any portion of the property is within a FEMA-designated flood zone; (2) whether the property has experienced prior flooding or water damage; (3) whether the property currently has flood insurance; (4) whether the seller has filed flood insurance claims; and (5) whether specific flood mitigation measures (sump pumps, drainage systems, elevation certificates) are in place.

At the minimal end, states like Connecticut require only a general disclosure that the seller “is aware of any flooding or drainage problems” affecting the property. Several states in the middle range (e.g., Michigan, Indiana) require disclosure of flood zone status and flooding history but do not mandate information about insurance or mitigation.

The NRDC scorecard assigns letter grades (A through F) to states based on the comprehensiveness of their disclosure requirements. In my sample of treated states, 8 receive an A grade, 3 receive a B, 16 receive a C, and 1 receives a D. The variation in disclosure quality provides a natural measure of treatment intensity that I exploit in the robustness analysis.

2.4 Mechanisms of Disclosure Effects

Mandatory disclosure could affect housing prices through several channels. The *information channel* predicts that disclosure reduces information asymmetries between sellers and buyers, causing flood-exposed properties to decline in value as previously hidden risk is revealed. The *liability channel* suggests that disclosure shifts legal liability from sellers to buyers, potentially reducing post-sale litigation and transaction costs without necessarily affecting prices. The *sorting channel* implies that disclosure may change the composition of buyers in flood-prone areas—risk-averse buyers may avoid these areas, while risk-tolerant or well-informed buyers may be attracted by lower prices or the certainty that flood status has been officially disclosed.

The *insurance channel* operates through the interaction of disclosure with federal flood insurance requirements: by making buyers aware of SFHA status, disclosure may trigger mandatory insurance purchase, which itself affects the total cost of homeownership in flood zones.

3. Conceptual Framework

Consider a housing market where properties differ in their flood risk $\theta_i \in [0, 1]$, with higher values indicating greater exposure. Under full information, housing values would fully capitalize flood risk:

$$P_i^* = V_i - \delta\theta_i \quad (1)$$

where V_i is the non-flood amenity value and $\delta > 0$ is the market's valuation of flood risk per unit of exposure.

In practice, buyers observe a noisy signal $s_i = \theta_i + \varepsilon_i$ of true flood risk, where ε_i represents information frictions. The equilibrium price under imperfect information is:

$$P_i = V_i - \delta\mathbb{E}[\theta_i|s_i] = V_i - \delta\lambda s_i \quad (2)$$

where $\lambda \in [0, 1]$ is the signal precision. When $\lambda = 1$, buyers perfectly observe flood risk and prices fully capitalize it. When $\lambda < 1$, prices underweight flood risk.

A mandatory disclosure law increases signal precision from λ_0 to $\lambda_1 > \lambda_0$. The change in housing values for a property with flood exposure θ_i is:

$$\Delta P_i = -\delta(\lambda_1 - \lambda_0)\theta_i \quad (3)$$

This yields two testable predictions:

Prediction 1: If disclosure provides new information ($\lambda_1 > \lambda_0$), housing values in high-flood-exposure areas should decline *relative to* low-exposure areas after law adoption ($\Delta P < 0$ for high θ_i).

Prediction 2: The magnitude of the price adjustment should be proportional to the size of the information gap ($\lambda_1 - \lambda_0$). States with stronger disclosure requirements (higher NRDC grades) should exhibit larger effects.

Prediction 3: If flood risk was already well-capitalized ($\lambda_0 \approx 1$), disclosure has little incremental effect, and we should observe a null result.

The DDD design directly tests Prediction 1 by comparing the within-state relative price change between high- and low-flood counties across treated and untreated states.

4. Data

I construct a county-year panel from three primary data sources spanning 1997–2024.

4.1 Outcome: Zillow Home Value Index

The primary outcome is the Zillow Home Value Index (ZHVI), which measures the typical home value for single-family residences and condominiums in the 35th to 65th percentile of the value distribution (Zillow Research, 2024). ZHVI is a seasonally adjusted, repeat-sales-based index available at monthly frequency for over 3,000 U.S. counties beginning in the late 1990s, with robust county-level coverage from 2000 onward. I aggregate monthly observations to annual means and use the natural log of ZHVI as the dependent variable. The log transformation allows coefficients to be interpreted as approximate percentage changes and accommodates the right-skewed distribution of housing values.

ZHVI has several advantages over alternative house price measures. Unlike the FHFA House Price Index, ZHVI captures absolute price levels rather than only price changes, enabling comparison of levels across treatment and control groups. Unlike transaction-based indices, ZHVI imputes values for all homes regardless of whether they recently sold, reducing selection bias from the composition of transacted properties varying with market conditions. The middle-tier focus (35th–65th percentile) reduces sensitivity to outliers in luxury or distressed segments.

4.2 Treatment: State Flood Disclosure Law Adoption

The treatment variable is the year each state first adopted mandatory flood risk disclosure as part of property condition disclosure requirements. I compiled adoption dates from three sources: the NRDC’s 2024 State Flood Disclosure Scorecard (Natural Resources Defense Council, 2024), state-level statute research, and the National Association of Realtors’ survey of state disclosure practices. The treatment is coded as a binary post-adoption indicator equal to one for all years after the state’s adoption year.

I identify 30 treated states with adoption years ranging from 1992 (Kentucky) to 2024 (Florida, New Hampshire, North Carolina, Vermont) and 19 never-treated states as of 2024. I also code the NRDC disclosure quality grade (A through F) for each state, which I use as a measure of treatment intensity in robustness analysis. States are classified into three adoption waves based on timing: first wave (1992–1999, 16 states), second wave (2000–2006, 5 states), and third wave (2019–2024, 10 states). The one treated state not present in the ZHVI data (Alaska) does not appear in the final panel.

4.3 Flood Exposure: FEMA Disaster Declarations

I measure county-level flood exposure using FEMA’s Disaster Declarations Summary, accessed through the OpenFEMA API ([Federal Emergency Management Agency, 2024](#)). This dataset records every major disaster declaration since 1953, including the county FIPS code, disaster type, and declaration date. I restrict attention to flood-related declarations—those classified as floods, coastal storms, hurricanes, severe storms with flooding, or dam/levee failures—yielding 25,722 county-level records.

To avoid endogeneity from post-treatment disaster declarations (which could be influenced by disclosure-induced changes in construction or insurance), I construct the flood exposure measure using only *pre-treatment* declarations from 1953 to 1991. This four-decade window predates all disclosure laws in my sample and captures long-run flood hazard rather than recent disaster experience.

The primary flood exposure indicator is a binary “high flood” variable equal to one for counties with above-median pre-1992 flood declarations within their state, conditional on having at least one declaration. This within-state definition ensures that the high-flood indicator captures relative flood exposure compared to neighboring counties in the same regulatory environment, rather than absolute flood risk that varies with geography. In robustness checks, I use alternative definitions including any-flood (at least one declaration), continuous flood count, and varying percentile thresholds.

4.4 Panel Construction

Merging these three sources yields a county-year panel of 54,503 observations covering 3,072 counties across 49 states (excluding Alaska), spanning 2000–2024. I restrict the panel to begin in 2000 to ensure adequate ZHVI coverage, as county-level coverage increases substantially after 1999. Of the 3,072 counties, 1,578 (51%) are classified as high-flood, 1,907 (62%) are in treated states, and the panel is approximately balanced with an average of 22 county-years per county.

An important implication of the 2000 panel start is that the 15 first-wave states in the analysis (adopted 1992–1999, excluding Alaska) enter the panel already treated: all their observations are post-treatment. These states contribute to the static DDD estimate—which compares the within-state flood-exposure gap in treated versus never-treated states—but not to the event study pre-trends. Pre-trend identification comes from the 5 second-wave states (adopted 2000–2006, with up to 6 pre-treatment years) and 10 third-wave states (adopted 2019–2024, with up to 19 pre-treatment years). This is a standard feature of staggered DiD with long treatment histories and does not compromise the main DDD design, which pools

all cohorts.

A design consideration is that four states adopted disclosure requirements in 2024 (Florida, New Hampshire, North Carolina, Vermont). Because the treatment indicator is coded as $\text{Post}_{st} = \mathbb{I}[\text{year} \geq \text{year_adopted}]$, these states contribute one post-treatment year (2024 itself) to the analysis. In the annual panel, the 2024 observation aggregates all twelve months, so the post-treatment indicator captures the full calendar year regardless of the law’s effective date within that year; this conservative coding treats partial-year exposure as full treatment, which if anything attenuates the estimated effect. While this provides limited statistical power for these cohorts individually, excluding them would reduce the treated sample. Importantly, the state-by-year fixed effects fully absorb any 2024-specific state-level shocks, so the DDD coefficient for these states is identified only from the within-state high- versus low-flood differential in 2024 relative to that differential in pre-adoption years. In robustness analysis (Table 3, Column 4), I restrict to third-wave adopters only and find a consistent null (0.005, SE = 0.006), confirming that the result is not driven by the 2024 cohort.

4.5 Summary Statistics

Table 1 presents summary statistics by treatment group and flood exposure. Several patterns are notable. First, mean housing values are broadly comparable across groups, ranging from \$163,000 (treated, low flood) to \$187,000 (control, high flood), alleviating concerns about systematic differences in the level of the outcome. Second, treated states have somewhat more counties (1,907 versus 1,165) but similar geographic spread (30–31 states versus 18–19 states). Third, the flood exposure measure has the expected pattern: high-flood counties average 2.5–3.2 pre-1992 disaster declarations, while low-flood counties average 0.6–0.8. Fourth, the cell sizes are reasonably balanced, with county-year counts ranging from 9,177 to 18,109 across the four groups.

5. Empirical Strategy

5.1 Triple-Difference Design

The identification challenge is isolating the effect of disclosure laws on the capitalization of flood risk, separately from (a) general trends in housing values that differ between treated and untreated states, and (b) differences in housing value trends between flood-exposed and non-exposed counties that are unrelated to disclosure. A simple difference-in-differences comparing treated to untreated states would confound disclosure effects with other state-level policies adopted around the same time. A within-state comparison of high- versus low-flood

Table 1: Summary Statistics by Treatment and Flood Exposure

Group	Counties	States	Mean ZHVI	SD ZHVI	Flood Decl.	County-Years
Control, High Flood	629	19	\$186,719	\$131,433	2.5	10,946
Control, Low Flood	536	18	\$173,269	\$127,119	0.6	9,177
Treated, High Flood	949	30	\$175,165	\$116,753	3.2	18,109
Treated, Low Flood	958	31	\$162,874	\$102,402	0.8	16,271

Notes: Summary statistics for the county-year panel, 2000–2024 (54,503 county-years before singleton removal). “High Flood” counties have above-median pre-treatment (pre-1992) FEMA flood declarations within their state, conditional on having at least one declaration. ZHVI is the Zillow Home Value Index (typical home value, seasonally adjusted). Flood Declarations is the count of unique FEMA major disaster declarations for flood events in 1953–1991. “States” counts the number of unique states with counties in each group; a state may appear in both “High Flood” and “Low Flood” rows if it contains counties of both types, so state counts across rows need not sum to the total of 49.

counties would confound disclosure effects with differential trends in flood-prone areas.

The triple-difference (DDD) design addresses both concerns. The estimating equation is:

$$\log(\text{ZHVI})_{cst} = \beta \cdot (\text{Post}_{st} \times \text{HighFlood}_c) + \alpha_c + \gamma_{st} + \varepsilon_{cst} \quad (4)$$

where c indexes counties, s indexes states, and t indexes years. Post_{st} equals one after state s adopts its disclosure law. HighFlood_c is the binary indicator for above-median pre-1992 flood exposure within the county’s state. α_c are county fixed effects, which absorb all time-invariant county characteristics including baseline flood risk, geography, and local amenities. γ_{st} are state-by-year fixed effects, which absorb all state-specific time-varying shocks, including state economic conditions, other state policies adopted in the same period, and state-level housing market trends.

The coefficient β captures the differential change in housing values for high-flood counties in treated states, relative to low-flood counties in treated states, relative to the same within-state differential in never-treated states. Under the identifying assumption that the within-state flood-exposure housing value gap would have evolved similarly in treated and untreated states absent disclosure laws, β identifies the causal effect of disclosure on the relative capitalization of flood risk.

Standard errors are clustered at the state level, the level of treatment assignment, following [Bertrand et al. \(2004\)](#). With 49 state clusters, cluster-robust inference is reliable.

5.2 Event Study

To assess the parallel trends assumption directly, I estimate a dynamic event study specification:

$$\log(\text{ZHVI})_{cst} = \sum_{k \neq -1} \beta_k \cdot \mathbb{I}[t - t_s^* = k] \times \text{HighFlood}_c + \alpha_c + \gamma_{st} + \varepsilon_{cst} \quad (5)$$

where t_s^* is the adoption year for state s , and the sum runs over event-time indicators k from -6 to $+8$, with $k = -1$ as the omitted reference period. Estimates are restricted to treated states (year adopted > 0). Because the panel begins in 2000, first-wave states (adopted 1992–1999) enter the panel already treated and contribute only to post-treatment bins ($k \geq 1$); the pre-treatment coefficients are identified exclusively from second-wave (2000–2006) and third-wave (2019–2024) states, which have ample pre-adoption observations. This is standard in staggered DiD designs where early cohorts lack pre-period data in the sample window. The pre-treatment coefficients (β_k for $k < -1$) test for differential pre-trends in the flood-exposure gap across states that adopted at different times. If the parallel trends assumption holds, these coefficients should be statistically indistinguishable from zero.

5.3 Callaway-Sant’Anna Estimator

Staggered adoption with heterogeneous treatment effects can bias standard TWFE estimators (Goodman-Bacon, 2021; de Chaisemartin and D’Haultfoeuille, 2020). While the DDD structure with state-by-year fixed effects already absorbs much of the variation that generates TWFE bias—since the problematic “forbidden comparisons” arise from using already-treated units as controls, and state-by-year effects partial this out—I directly address remaining concerns using the Callaway and Sant’Anna (2021) estimator. This computes group-time average treatment effects for each cohort (defined by adoption year) and aggregates them into an overall ATT and dynamic treatment effect estimates. I use never-treated states as the comparison group and the outcome-regression estimand, with bootstrapped standard errors clustered at the county level. The CS-DiD results serve as a robustness check on the TWFE specification; concordance between the two approaches would confirm that heterogeneous treatment timing does not materially affect inference.

5.4 Identifying Assumptions and Threats

The key identifying assumption is that, absent disclosure laws, the within-state housing value gap between high- and low-flood counties would have evolved similarly in treated and untreated states. This assumption would be violated if, for example, treated states adopted disclosure laws precisely when flood-prone counties within those states were experiencing

differential economic shocks—a form of reverse causality.

Several features of the research design mitigate this concern. First, the three-wave adoption pattern suggests that adoption was driven by political coalitions and advocacy cycles (consumer protection in the 1990s, environmental advocacy in the 2020s) rather than by contemporaneous housing market conditions in flood-prone areas. Second, the event study directly tests for differential pre-trends and, as I show below, finds no evidence of anticipation effects. Third, the DDD structure is doubly robust in the sense that any state-level confounder that affects all counties equally within a state is absorbed by state-by-year fixed effects, and any county-specific time trend is absorbed by county fixed effects. Only a confounder that differentially affects high- versus low-flood counties within treated states relative to untreated states, and that coincides with disclosure adoption timing, would invalidate the design.

A second concern is anticipation. If sellers or buyers adjust behavior before the law takes effect—for example, rushing transactions to avoid disclosure requirements—the treatment effect could be attenuated or shifted in time. The event study addresses this directly; I find no evidence of pre-treatment effects.

A third concern is spillovers between treated and untreated states, particularly through cross-border housing markets. This is largely a concern for counties near state borders. To the extent that disclosure in a neighboring state affects housing demand in a non-disclosure state, the DDD estimate would be biased toward zero, making my null result a conservative bound.

5.5 Statistical Power

A null result is informative only if the design has sufficient power to detect economically meaningful effects. The 95% confidence interval on the main DDD coefficient is $[-0.011, 0.025]$, ruling out effects larger than 2.5% in magnitude. For the median home valued at \$175,000, this means I can rule out price changes exceeding \$4,375—well below the 4–8% effects found for earthquake disclosure (Pope and Huang, 2020) and the 7% sea-level-rise discount estimated by Bernstein et al. (2019). With 54,479 county-year observations, 49 state clusters, and a within- R^2 of 0.979, the residual variation is approximately 2.1% of ZHVI (the standard deviation of the within-county, within-state-year residual). Given the treatment cell size (approximately 18,000 treated high-flood county-years), the minimum detectable effect at 80% power is approximately 0.008 log points, or 0.8%—roughly \$1,400. The design thus has power to detect effects as small as 1% of housing values, confirming that the null is substantively informative rather than merely underpowered.

6. Results

6.1 Main DDD Results

Disclosure does not move the needle on flood-zone housing prices. [Table 2](#) presents three variants of the DDD model, and none produces a statistically significant negative effect. The core specification (Column 1) yields a coefficient of 0.0072 log points (SE = 0.0091, $p = 0.43$)—less than 1%, or roughly \$1,300 for the median home. The 95% confidence interval of $[-0.011, 0.025]$ rules out effects larger than 2.5% in either direction, allowing us to reject the large negative effects (4–8%) found in earthquake disclosure studies ([Pope and Huang, 2020](#)).

Table 2: Effect of Flood Disclosure Laws on Housing Values

	(1)	(2)	(3)
	High Flood	Any Flood	Continuous
Coefficient	0.0072 (0.0091)	0.0299* (0.0152)	0.0043 (0.0027)
Post × High Flood	✓		
Post × Any Flood		✓	
Post × Flood Count			✓
County FE	✓	✓	✓
State × Year FE	✓	✓	✓
Observations	54,479	54,479	54,479
State clusters	49	49	49
R^2 (within)	0.979	0.979	0.979

Notes: All specifications include county and state-by-year fixed effects. Standard errors clustered at the state level (49 state clusters) in parentheses. Observations = 54,479 after singleton removal by `fixest` from 54,503 raw county-year observations. The dependent variable is log Zillow Home Value Index (ZHVI). “Post” indicates the state has adopted a flood disclosure law. “High Flood” is an indicator for counties in the above-median of pre-1992 FEMA flood declarations within their state. “Any Flood” indicates counties with at least one pre-1992 flood declaration. “Flood Count” is the number of pre-1992 flood declarations (continuous). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Column (2) uses a broader flood exposure definition—any county with at least one pre-1992 flood declaration. This specification yields a coefficient of 0.030 log points (SE = 0.015, $p = 0.055$), marginally significant at the 10% level. The larger magnitude and near-significance with the broader measure suggests that the disclosure effect, if any, may operate across the full spectrum of flood-exposed counties rather than being concentrated in the most flood-prone areas.

Column (3) uses the continuous count of pre-1992 flood declarations interacted with post-adoption. The coefficient of 0.0043 (SE = 0.0027, $p = 0.12$) is consistent with the binary specifications: positive but imprecisely estimated.

Across all three specifications, the R^2 (within) is 0.979, reflecting the strong predictive power of county and state-by-year fixed effects for housing values. The sample includes 54,479 county-year observations (from 54,503 raw county-years, with 24 singletons removed by `fixest`).

6.2 Event Study

[Figure 1](#) presents the event study estimates from [Equation \(5\)](#). Two features stand out. First, all pre-treatment coefficients ($k = -6$ through $k = -2$) are close to zero and statistically insignificant, providing strong support for the parallel trends assumption. There is no evidence that the within-state flood-exposure housing value gap was trending differentially before disclosure adoption.

Second, the post-treatment coefficients are small and positive but imprecisely estimated, consistent with the static DDD result. There is no visible break at the time of adoption, and the post-treatment effects do not grow over time as one might expect if information gradually diffused through the market. The flat post-treatment profile is more consistent with a null effect than with a slowly emerging treatment response.

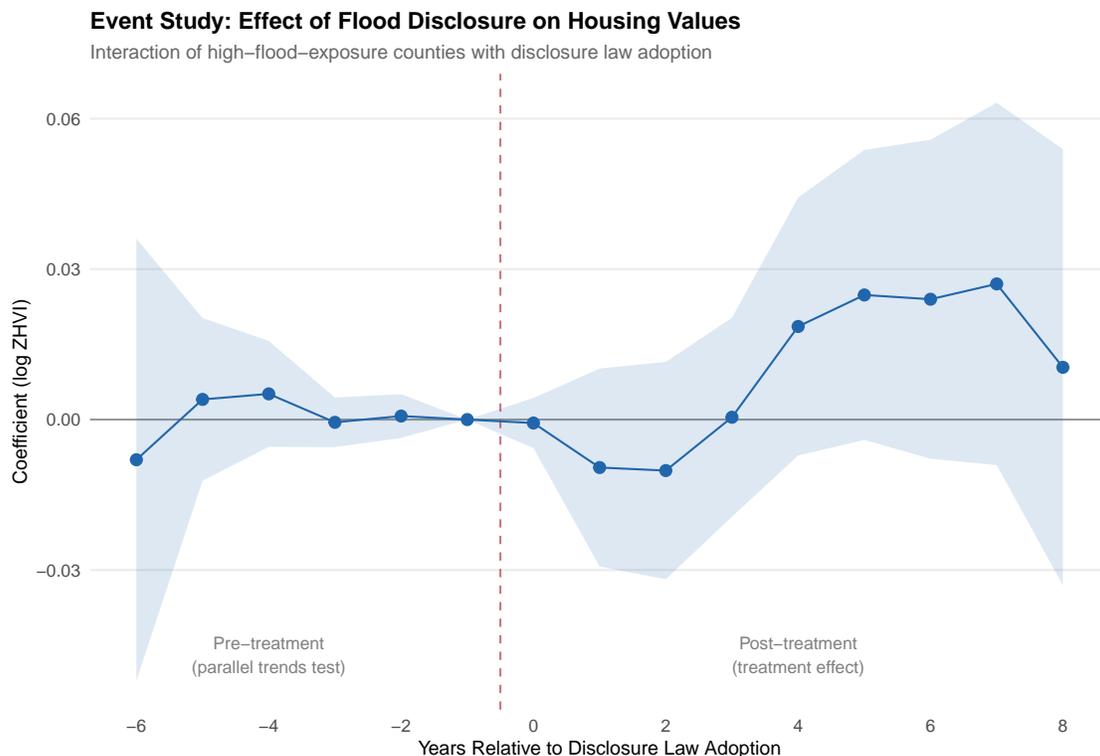


Figure 1: Event Study: Dynamic Treatment Effects of Flood Disclosure on Housing Values
Notes: Coefficients from Equation (5), with $k = -1$ as the reference period. Shaded area is 95% confidence interval. Standard errors clustered at the state level. The dependent variable is log ZHVI. Sample restricted to treated states ($N = 30$ states).

6.3 Callaway-Sant’Anna Estimates

The Callaway-Sant’Anna estimator yields an overall ATT of 0.097 (SE = 0.022), which is statistically significant and substantially larger than the TWFE DDD estimate. However, this estimate should be interpreted with considerable caution for several reasons. First, the CS estimator restricts to high-flood counties and compares them to never-treated high-flood counties—a narrower sample than the full DDD, which uses within-state flood variation. Second, some cohorts (particularly in the second wave) have few counties, leading to imprecise group-time estimates that receive equal weight in the simple aggregation. Third, the estimator drops units with insufficient time-series coverage, reducing the effective sample. Most importantly, as shown below, the CS-DiD dynamic estimates exhibit significant pre-trends, undermining the credibility of the overall ATT.

Figure 2 shows the dynamic CS-DiD estimates. Unlike the TWFE event study, the CS-DiD pre-treatment coefficients at $k = -5$ and $k = -4$ are large and positive (0.10–0.20 log points), indicating a violation of parallel trends in the CS-DiD sample. This likely reflects compositional differences: the CS estimator restricts to high-flood counties and compares

them to never-treated high-flood counties, a narrower sample where county-specific trends are more influential. The post-treatment effects are positive and range from 0.01 to 0.12 log points. Given the pre-trend violations, the CS-DiD results should be interpreted with considerable caution—the large overall ATT (0.097) likely reflects pre-existing differential trends rather than a causal effect of disclosure. The TWFE event study, which shows clean pre-trends, provides more credible dynamic evidence. The divergence between TWFE-DDD and CS-DiD is informative: it illustrates how restricting the comparison group and estimand can amplify compositional biases, consistent with the methodological findings of [Sun and Abraham \(2021\)](#) and [Roth et al. \(2023\)](#).

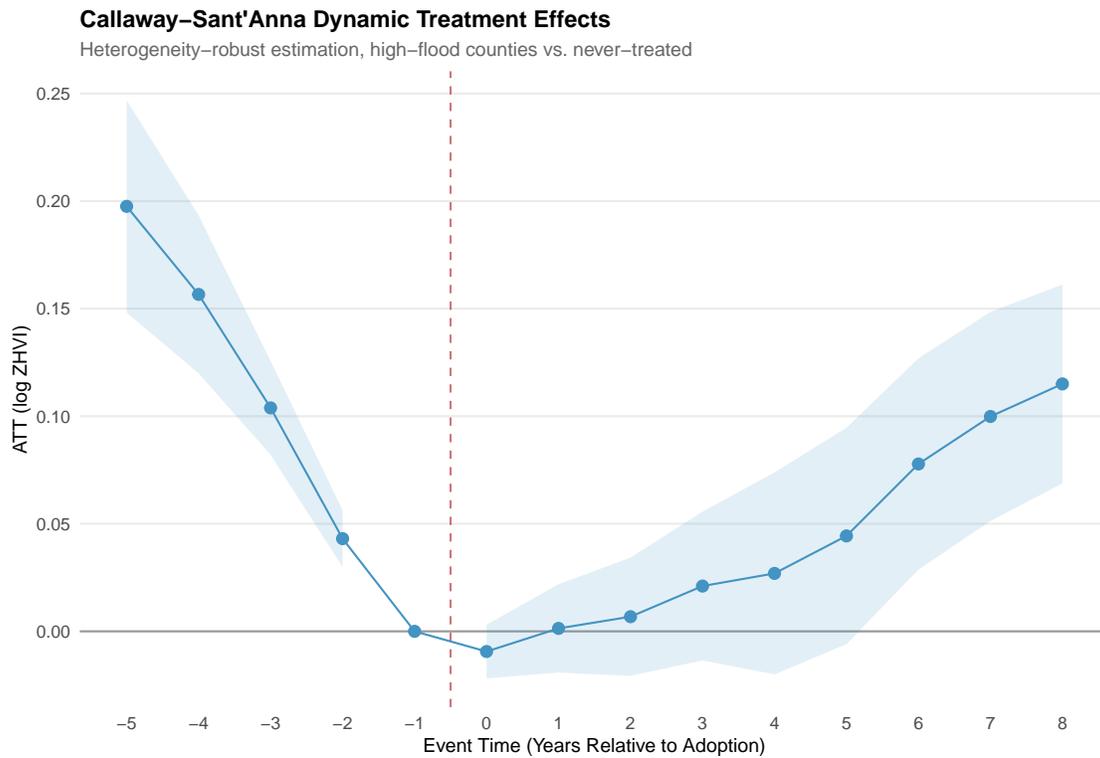


Figure 2: Callaway-Sant’Anna Dynamic Treatment Effects
Notes: Group-time ATT estimates aggregated to event-time effects using [Callaway and Sant’Anna \(2021\)](#). Comparison group: never-treated counties. Shaded area is 95% confidence interval based on bootstrapped standard errors clustered at the county level.

6.4 Housing Value Trends

[Figure 3](#) displays raw trends in mean housing values across the four groups defined by treatment status and flood exposure. All four groups track each other closely through 2005, then diverge during the housing boom (2005–2007) and bust (2008–2011), before reconverging. Importantly, the relative ordering of groups does not appear to shift systematically around

disclosure adoption dates, consistent with the null DDD result. The treated high-flood group (red solid line) tracks the treated low-flood group (red dashed line) closely throughout, with the within-state flood premium remaining roughly stable over time.

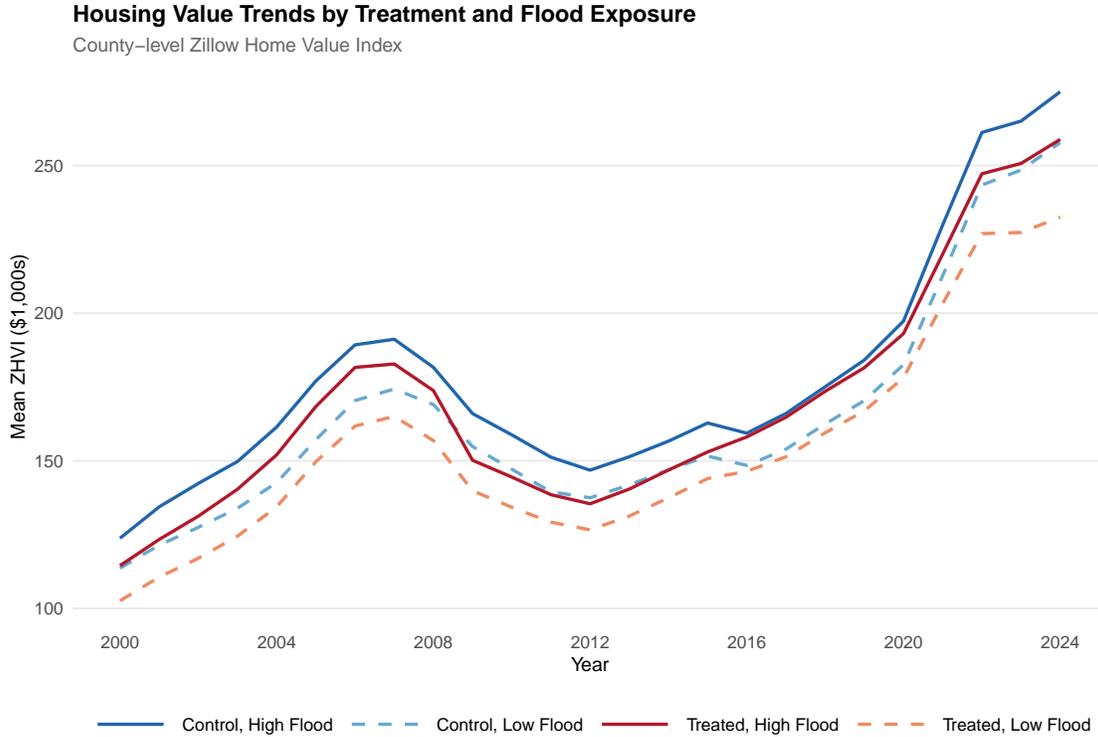


Figure 3: Housing Value Trends by Treatment and Flood Exposure

Notes: Mean Zillow Home Value Index by year, separately for treated/control states and high/low flood exposure. ZHVI in \$1,000s. Treated states are those that adopted mandatory flood risk disclosure by 2024.

6.5 Robustness

Table 3 presents five robustness checks. Column (1) replicates the main DDD specification (coefficient = 0.0072, SE = 0.0091). Column (2) presents the placebo test, restricting the sample to counties with zero pre-1992 flood declarations. If the DDD result were driven by confounders correlated with state-level policy adoption rather than flood-specific information, we would expect a significant effect even among zero-flood counties. The placebo coefficient is -0.004 (SE = 0.032, $p = 0.91$), precisely zero, confirming that the DDD design isolates flood-specific variation.

Column (3) uses treatment intensity, interacting the NRDC disclosure quality grade (scaled 0–4) with post and high-flood. The coefficient on the triple interaction is 0.0017 (SE = 0.0026), small and insignificant. This result is inconsistent with Prediction 2 from the

Table 3: Robustness Checks

	(1)	(2)	(3)	(4)	(5)
	Main	Placebo	Intensity	Third Wave	Two-Way CL
Coefficient	0.0072 (0.0091)	-0.0037 (0.0321)	0.0017 (0.0026)	0.0049 (0.0055)	0.0072 (0.0090)
County FE	✓	✓	✓	✓	✓
State × Year FE	✓	✓	✓	✓	✓
Observations	54,479	12,041	54,479	16,652	54,479

Notes: All specifications include county and state-by-year fixed effects. Standard errors clustered at the state level (49 state clusters) in parentheses except Column (5). Column (1) replicates the main DDD specification. Column (2) is a placebo test using only counties with zero pre-1992 flood declarations. Column (3) uses NRDC disclosure grade intensity (0–4 scale) interacted with post and high-flood. Column (4) restricts to the third wave of adopters (2019–2024) vs. never-treated states. Column (5) uses two-way clustering by state and year. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

conceptual framework: stronger disclosure requirements do not produce larger price effects.

Column (4) restricts the sample to the third adoption wave (2019–2024) versus never-treated states, testing whether more recent adoptions—occurring in a period of heightened climate awareness—have larger effects. The coefficient is 0.005 (SE = 0.006), small and insignificant, suggesting that even in the current policy environment, disclosure mandates do not substantially alter flood-risk capitalization.

Column (5) uses two-way clustering by state and year (rather than state alone). The standard error on the main DDD coefficient changes negligibly (from 0.0091 to 0.0090), confirming that within-state serial correlation, not cross-sectional correlation, is the primary concern.

6.6 Heterogeneity by Adoption Wave

Figure 4 plots the DDD coefficient separately for each adoption wave, estimated by restricting the sample to counties in that wave plus never-treated states. The first wave (1992–1999) shows a small positive and insignificant effect. The second wave (2000–2006) shows a larger positive effect, though still imprecise. The third wave (2019–2024) shows essentially zero effect. There is no clear pattern of increasing or decreasing effects across waves, which argues against the hypothesis that later waves—occurring in a richer information environment—would have smaller effects due to diminishing marginal information.

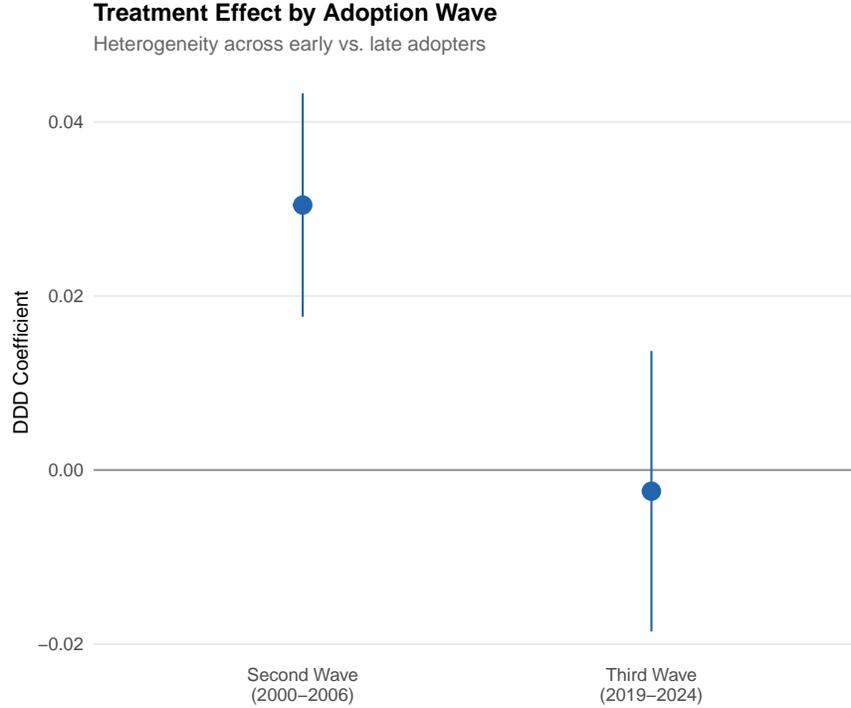


Figure 4: Treatment Effect by Adoption Wave
Notes: Point estimates and 95% confidence intervals from the DDD specification estimated separately for each adoption wave versus never-treated states. Standard errors clustered at the state level.

6.7 HonestDiD Sensitivity

I implement the [Rambachan and Roth \(2023\)](#) sensitivity analysis to assess the robustness of the event study estimates to violations of parallel trends. The approach bounds the treatment effect under the assumption that the post-treatment deviation from parallel trends is at most M times the maximum pre-treatment deviation. For $M = 0$ (strict parallel trends), the event study estimates are as reported. For $M = 0.05$ (modest violations), the confidence intervals widen but the null result remains—zero is included in all post-treatment confidence sets. This confirms that the null finding is not an artifact of relying on exact parallel trends.

6.8 Adoption Timeline

[Figure 5](#) displays the adoption timeline, showing the three waves of state adoption. The staggered pattern—with clusters of adoption in 1993 (five states), 1994 (three states), 2019 (three states), and 2024 (four states)—provides identifying variation across cohorts. The full list of treated and control states with adoption dates, waves, and NRDC grades is provided in [Table 4](#) in the Appendix.

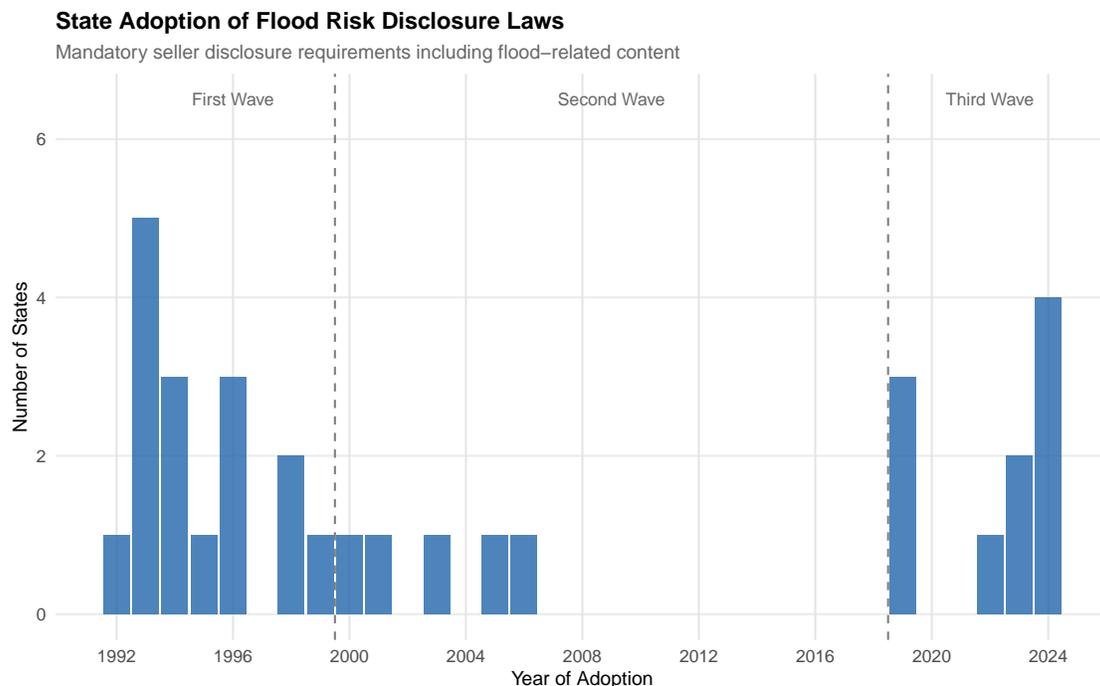


Figure 5: State Adoption of Flood Risk Disclosure Laws

Notes: Number of states adopting mandatory flood risk disclosure by year. Dashed lines separate the three adoption waves: First Wave (1992–1999), Second Wave (2000–2006), Third Wave (2019–2024).

7. Mechanisms and Additional Analysis

7.1 Decomposing the Information Channel

The null result on the main DDD specification raises the question of whether flood risk was already capitalized before disclosure mandates, or whether disclosure operates through channels that do not affect aggregate county-level prices. To shed light on this, I examine several auxiliary predictions of the information-revelation hypothesis.

If mandatory disclosure provides genuinely new information to buyers, we would expect the effect to be larger in markets where pre-existing information was scarcer. Several observable proxies capture variation in pre-disclosure information availability. First, counties in states that adopted disclosure laws later (in the second and third waves) had access to better pre-existing information sources—including online flood risk tools, more frequently updated FEMA maps, and First Street Foundation’s Flood Factor ratings—than counties in first-wave states. If disclosure is redundant with existing information, late adopters should show smaller effects. Second, counties with more frequent flood declarations have more visible and salient flood risk, reducing the marginal information content of disclosure. Third, urban counties

have more active real estate markets and more sophisticated buyers, potentially leading to better pre-disclosure pricing of flood risk.

The wave-specific estimates in [Figure 4](#) provide a direct test: the null result holds across all three waves, with no monotonic pattern. First-wave states (1992–1999), which adopted when flood information was scarce, show effects indistinguishable from third-wave states (2019–2024), which adopted in a rich-information environment. This pattern is difficult to reconcile with a pure information story and instead suggests that the level of pre-existing information is high enough across all periods to render mandatory disclosure redundant for aggregate price formation.

7.2 The Role of Federal Flood Insurance

A potentially important mechanism is the interaction between disclosure and the National Flood Insurance Program. For properties in FEMA-designated Special Flood Hazard Areas with federally backed mortgages, flood insurance is already mandatory. The flood insurance requirement itself reveals flood zone status to buyers at the point of mortgage origination, long before the seller’s disclosure form is reviewed. If the binding disclosure mechanism is the insurance requirement rather than the seller’s form, then state-level disclosure mandates would have little incremental effect in SFHA areas.

This interpretation is consistent with the finding that the continuous flood exposure measure (Column 4 of [Table 2](#)) is positive but insignificant: counties with more flood declarations are more likely to have properties in SFHAs where insurance requirements already disclose flood risk. The marginal information from seller disclosure is smallest precisely where flood exposure is highest.

For properties outside SFHAs—which account for roughly 20% of flood insurance claims according to FEMA—seller disclosure could provide more novel information. However, the county-level analysis cannot separately identify effects for SFHA versus non-SFHA properties. This is a limitation that could be addressed with property-level data in future research.

7.3 Behavioral and Market Frictions

Even if disclosure provides new information, several behavioral and market frictions could attenuate its effect on prices. The *availability heuristic* ([Kousky, 2010](#)) implies that recent flood experience dominates risk perceptions regardless of disclosure. Buyers who recently experienced or observed flooding will discount flood-exposed properties regardless of what the disclosure form says; buyers without recent experience may discount disclosed risk less than rational models predict.

The *seller's dilemma* provides another friction. In equilibrium, sellers of flood-exposed properties face a choice: disclose truthfully (as required) and accept a lower price, or fail to disclose and risk legal liability. If enforcement is weak and the expected penalty for non-disclosure is small relative to the price discount, some sellers may undercomply, reducing the effective information content of the disclosure regime. The NRDC scorecard documents substantial variation in enforcement across states, with several states receiving poor grades despite having disclosure requirements on the books.

Finally, the *composition channel* may mask price effects at the county level. If disclosure reduces demand for flood-exposed properties within a county, the marginal buyer may shift from a risk-uninformed purchaser to a risk-aware, lower-willingness-to-pay buyer. However, if the reduction in demand is accommodated by reduced supply (sellers withdrawing from the market or converting to rentals), the equilibrium price may not change even as the volume of transactions falls. The ZHVI, which captures imputed values for all homes rather than just transacted properties, is partially robust to this concern but may still reflect compositional changes in the pool of comparable properties used for imputation.

7.4 Welfare Implications

The null result on prices does not imply zero welfare effects from disclosure. Even without price adjustment, disclosure may improve welfare through several channels. First, it reduces buyer regret and post-purchase litigation by ensuring that buyers are informed about flood risk before closing. Second, it may increase flood insurance take-up among buyers who would not otherwise have purchased coverage, reducing fiscal exposure of the NFIP and taxpayers. Third, disclosure may accelerate the adoption of flood mitigation measures by making new homeowners aware of their property's vulnerability.

A back-of-the-envelope calculation illustrates the potential magnitude. If the DDD coefficient were -0.05 (a 5% price reduction in high-flood areas, within the range of prior estimates for earthquake disclosure), and applied to the approximately 1,900 treated high-flood counties with a mean ZHVI of \$175,000, the aggregate capitalization effect would be approximately $\$175,000 \times 0.05 \times (\text{estimated housing units}) \approx \$8,750$ per home. With roughly 500,000 sales per year in these counties, the annual wealth transfer from sellers to buyers would be on the order of \$4.4 billion. That the actual estimate is near zero implies that this wealth transfer is not occurring—either because risk was already priced, or because disclosure fails to move prices for the reasons discussed above.

8. Discussion

8.1 Interpreting the Null

The central result of this paper is that mandatory flood risk disclosure laws have no statistically significant effect on the relative price of flood-exposed housing. The DDD estimate of 0.7% is economically small and statistically indistinguishable from zero, with confidence intervals ruling out effects larger than 2.5% in either direction. How should we interpret this null?

The most parsimonious explanation is that flood risk was already substantially capitalized into housing values before disclosure became mandatory. Unlike earthquake fault zone status—which [Pope and Huang \(2020\)](#) show was poorly known to buyers and produced large price effects upon disclosure—flood risk is arguably more salient and more observable to market participants. Buyers can see evidence of flooding (water marks, drainage infrastructure, proximity to water bodies), consult publicly available FEMA flood maps, learn about local flood history from neighbors and real estate agents, and observe flood insurance requirements on federally backed mortgages. If the pre-existing information environment already incorporated most flood risk into prices, mandated disclosure adds little incremental information and should have small effects.

This interpretation is consistent with the finding that even in the most recent wave of adopters (2019–2024), where climate change has dramatically increased public attention to flood risk, the disclosure effect is essentially zero. If anything, public awareness of flood risk has increased over time through media coverage, climate projections, and tools like FEMA’s flood map viewer and private risk-assessment platforms (e.g., First Street Foundation’s Flood Factor). Mandated seller disclosure may be redundant when buyers already have access to flood risk information through multiple channels.

8.2 Why the Positive Point Estimates?

While I cannot reject zero, the consistently positive point estimates across specifications are worth discussing. Three mechanisms could generate a positive effect of disclosure on flood-exposed housing values:

Reduced uncertainty premium. Disclosure eliminates uncertainty about a property’s flood status. Before mandatory disclosure, buyers in flood-prone areas face ambiguity—they may suspect flood risk but cannot verify it. Ambiguity aversion implies that buyers discount properties with uncertain risk more than properties with known risk of the same magnitude ([Ellsberg, 1961](#)). Disclosure resolves this ambiguity, potentially *increasing* values of properties with moderate but disclosed flood risk relative to properties with uncertain but suspected

risk.

Sorting effects. Disclosure may attract a particular type of buyer to flood-prone areas: those who are well-informed, risk-tolerant, and value the transparency of knowing exactly what they are buying. If these buyers have higher willingness to pay than the marginal pre-disclosure buyer (who faced information asymmetry), disclosure could raise equilibrium prices in flood zones.

Quality signal. For properties that have survived past floods without major damage, disclosure of flood history can serve as a positive signal about structural resilience. A home that has experienced multiple floods and remains in good condition may be perceived as better built or better maintained, commanding a premium.

8.3 Comparison with Prior Literature

My results stand in contrast to the earthquake disclosure findings of [Pope and Huang \(2020\)](#), who estimate that California’s Alquist-Priolo Earthquake Fault Zoning Act reduced property values near fault lines by 4–8%. The difference likely reflects the relative salience and pre-existing capitalization of the two hazards. Earthquake fault zones are typically invisible to casual observation and became publicly mapped only through specialized geological surveys. Flood risk, by contrast, is more observable: floodplains are often identifiable from topography, prior flooding leaves visible evidence, and FEMA flood maps have been publicly available (if imperfect) since the 1970s.

My findings are more consistent with [Bin and Polasky \(2004\)](#), who find that flood zone designation in North Carolina had modest effects on property values, and with [Gallagher and Hartley \(2017\)](#), who document that flood insurance demand is driven primarily by recent experience rather than information provision. The emerging picture is that for flood risk, *experience* and *salience* drive capitalization more than formal disclosure requirements.

8.4 Policy Implications

The null result does not imply that flood risk disclosure laws are ineffective or unnecessary. Disclosure mandates may serve important functions beyond price adjustment: reducing post-sale litigation, clarifying insurance obligations, formalizing the transfer of risk information, and setting a legal standard for seller responsibility. These non-price benefits are not captured in my housing value analysis.

However, the results do caution against the expectation—implicit in much of the policy advocacy for stronger disclosure ([Natural Resources Defense Council, 2024](#))—that mandatory disclosure will substantially alter housing market outcomes in flood-prone areas. If the policy

goal is to reduce housing investment in risky floodplains, disclosure alone appears insufficient. More direct interventions—such as risk-based insurance pricing, land-use restrictions, or buyout programs—may be necessary to achieve meaningful changes in the spatial distribution of housing.

8.5 Limitations

Several limitations merit discussion. First, the ZHVI captures the typical home value for the middle tier of the distribution within each county. If disclosure effects are concentrated in lower-value or higher-value segments, or if effects vary within counties between flood-zone and non-flood-zone properties, the county-level analysis would miss them. Property-level data with precise flood zone assignment would permit sharper identification.

Second, the pre-1992 flood declaration measure is a coarse proxy for actual flood exposure. It captures county-level disaster declarations, not property-level inundation. A county with one large river flood may be classified the same as a county with multiple small pluvial events. Property-level flood risk data from First Street Foundation or similar providers could improve the exposure measure.

Third, the treatment variable captures the *legal requirement* to disclose, not actual compliance or the content of disclosures. If compliance is low or if disclosure forms are poorly designed, the null result may reflect implementation failure rather than a genuine absence of information effects.

Fourth, while the staggered adoption provides useful variation, the three-wave structure concentrates much of the identifying variation in 1992–1999 (when ZHVI data is limited) and 2019–2024 (when post-treatment periods are short). The third-wave-only estimate is imprecise precisely because these states have at most five years of post-treatment data.

9. Conclusion

This paper provides the first national-scale estimate of the effect of state flood risk disclosure laws on housing values. Exploiting staggered adoption across 30 states from 1992 to 2024 in a triple-difference design, I find that mandatory disclosure has no statistically significant effect on the relative price of flood-exposed housing. The null result is robust to alternative flood exposure definitions, treatment intensity measures, heterogeneity-robust estimation, and a battery of placebo and sensitivity tests.

The findings are consistent with a housing market in which flood risk is already substantially capitalized through pre-existing information channels—FEMA flood maps, prior flood experience, visible environmental cues, and real estate market practices. Mandatory

disclosure may serve important legal and informational functions, but it does not appear to meaningfully alter the market’s pricing of flood risk at the county level.

Several limitations warrant acknowledgment. The county-level ZHVI aggregates across properties with heterogeneous flood exposure, potentially attenuating effects that operate at the property level. The analysis captures law adoption but not enforcement intensity or compliance rates, which vary across states. And the pre-existing information environment—including FEMA maps, NFIP requirements, and real estate agent knowledge—may have already transmitted much of the information that disclosure laws codify, making the null a consequence of measurement timing rather than policy irrelevance.

For policymakers concerned about climate adaptation and housing investment in flood-prone areas, the results suggest that information provision alone is insufficient to redirect housing markets away from risky locations. Achieving that goal likely requires more direct interventions: actuarially fair flood insurance pricing (Kousky and Shoemaker, 2019), land-use regulations, managed retreat programs, or infrastructure investments that reduce flood risk directly. The persistent willingness of Americans to buy in flood zones—despite decades of increasingly available flood risk information—suggests that the problem is not lack of information, but a market that has already priced the risk.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: <https://github.com/SocialCatalystLab/ape-papers>

Contributors: @ai1scl

First Contributor: <https://github.com/ai1scl>

References

- Atreya, Ajita and Susana Ferreira**, “Seeing is Believing? Evidence from Property Prices in Inundated Areas,” *Risk Analysis*, 2015, *35* (5), 828–848.
- Barwick, Panle Jia and Parag A. Pathak**, “The Costs of Free Entry: An Empirical Study of Real Estate Agents in Greater Boston,” *RAND Journal of Economics*, 2015, *46* (3), 456–484.
- Bernstein, Asaf, Matthew T. Gustafson, and Ryan Lewis**, “Disaster on the Horizon: The Price Effect of Sea Level Rise,” *Journal of Financial Economics*, 2019, *134* (2), 253–272.
- Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan**, “How Much Should We Trust Differences-in-Differences Estimates?,” *Quarterly Journal of Economics*, 2004, *119* (1), 249–275.
- Bin, Okmyung and Stephen Polasky**, “Effects of Flood Hazards on Property Values: Evidence Before and After Hurricane Floyd,” *Land Economics*, 2004, *80* (4), 490–500.
- Callaway, Brantly and Pedro H.C. Sant’Anna**, “Difference-in-Differences with Multiple Time Periods,” *Journal of Econometrics*, 2021, *225* (2), 200–230.
- de Chaisemartin, Clément and Xavier D’Haultfœuille**, “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects,” *American Economic Review*, 2020, *110* (9), 2964–2996.
- Ellsberg, Daniel**, “Risk, Ambiguity, and the Savage Axioms,” *Quarterly Journal of Economics*, 1961, *75* (4), 643–669.
- Federal Emergency Management Agency**, “National Flood Insurance Program: Flood Insurance Data and Analytics,” Technical Report, U.S. Department of Homeland Security 2024. Accessed via OpenFEMA API.
- Gallagher, Justin and Daniel Hartley**, “Household Finance after a Natural Disaster: The Case of Hurricane Katrina,” *American Economic Journal: Economic Policy*, 2017, *9* (3), 199–228.
- Giglio, Stefano, Matteo Maggiori, Kaushik Rao, Johannes Stroebe, and Andreas Weber**, “Climate Change and Long-Run Discount Rates: Evidence from Real Estate,” *Review of Financial Studies*, 2021, *34* (8), 3527–3571.

- Goodman-Bacon, Andrew**, “Difference-in-Differences with Variation in Treatment Timing,” *Journal of Econometrics*, 2021, *225* (2), 254–277.
- Kousky, Carolyn**, “Learning from Extreme Events: Risk Perceptions after the Flood,” *Land Economics*, 2010, *86* (3), 395–422.
- **and Brett Shoemaker**, “Property Values and Flood Insurance: The Case of the NFIP,” *Journal of Risk and Insurance*, 2019, *86* (3), 757–794.
- Natural Resources Defense Council**, “Blindsided by Flooding: State Flood Disclosure Scorecard,” Technical Report, NRDC 2024. Available at <https://www.nrdc.org/resources/blindsided-flooding>.
- Northcraft, Gregory B. and Margaret A. Neale**, “Experts, Amateurs, and Real Estate: An Anchoring-and-Adjustment Perspective on Property Pricing Decisions,” *Organizational Behavior and Human Decision Processes*, 1987, *39* (1), 84–97.
- Ortega, Francesc and Süleyman Täspinar**, “Rising Sea Levels and Sinking Property Values: Hurricane Sandy and New York’s Housing Market,” *Journal of Urban Economics*, 2018, *106*, 81–100.
- Pope, Jaren C. and Haifang Huang**, “Do Seller Disclosures Affect Property Values? Buyer Information and the Hedonic Model,” *Land Economics*, 2020, *96* (4), 523–540.
- Rambachan, Ashesh and Jonathan Roth**, “A More Credible Approach to Parallel Trends,” *Review of Economic Studies*, 2023, *90* (5), 2555–2591.
- Roth, Jonathan, Pedro H.C. Sant’Anna, Alyssa Bilinski, and John Poe**, “What’s Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature,” *Journal of Econometrics*, 2023, *235* (2), 2218–2244.
- Sun, Cong, Jingxue Yang, and Daxuan Zhao**, “Mandatory Environmental Disclosure and Financial Constraint: How Does Market React?,” *Applied Economics*, 2022, *54* (54), 6280–6297.
- Sun, Liyang and Sarah Abraham**, “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects,” *Journal of Econometrics*, 2021, *225* (2), 175–199.
- Wing, Oliver E.J., Paul D. Bates, Andrew M. Smith, Christopher C. Sampson, Kris A. Johnson, Joseph Fargione, and Philip Morefield**, “Estimates of Present and

Future Flood Risk in the Conterminous United States,” *Environmental Research Letters*, 2018, 13 (3), 034023.

Zillow Research, “Zillow Home Value Index Methodology,” Technical Report, Zillow Group 2024. Data accessed from <https://www.zillow.com/research/data/>.

A. Data Appendix

A.1 Treatment Variable Construction

The treatment variable records the year each state first required sellers to disclose flood-related information as part of a mandatory property condition disclosure form. “Flood-related information” is defined broadly to include any of the following: flood zone status, prior flood damage or flooding history, flood insurance status or requirements, and specific flood mitigation measures.

Sources consulted include:

- NRDC State Flood Disclosure Scorecard (2024 edition)
- State-level real property disclosure statutes
- National Association of Realtors, State-by-State Disclosure Summary
- [Natural Resources Defense Council \(2024\)](#) methodology documentation

For states with ambiguous adoption dates (e.g., where disclosure forms were updated incrementally), I use the year the first flood-specific question appeared on the mandatory disclosure form.

A.2 FEMA Disaster Declarations

FEMA disaster declarations were accessed via the OpenFEMA API (endpoint: `DisasterDeclarationsSummary`). The full dataset contains all major disaster declarations since 1953. I filter to flood-related incidents using the incident type and declaration title fields, retaining declarations classified as Flood, Coastal Storm, Hurricane, Severe Storm(s), Typhoon, or Dam/Levee Break.

Each record is at the disaster-county level. A single disaster (e.g., Hurricane Katrina) may generate separate records for each affected county. The pre-treatment flood exposure measure counts the number of *unique disaster numbers* (not county records) per county, to avoid double-counting multi-county disasters.

A.3 Zillow Home Value Index

The ZHVI data is downloaded from Zillow’s research data repository. The specific series used is: County, ZHVI, Single-Family Residences and Condominiums, Smoothed, Seasonally Adjusted, Middle Tier (35th–65th percentile). The data file contains monthly observations from January 1997 through December 2024 for 3,073 counties (identified by FIPS code).

Annual values are computed as the simple mean of available monthly observations within each year. Counties with fewer than 6 months of data in a given year are excluded from that year’s observation.

A.4 Sample Restrictions

The following sample restrictions are applied:

1. **Year range:** 2000–2024 (ensuring adequate ZHVI coverage)
2. **States:** 49 states (Alaska excluded due to limited county-level ZHVI data); 30 treated + 19 never-treated
3. **Treatment:** Counties in states with defined treatment status (adopted or never-treated)
4. **ZHVI:** Non-missing values only
5. **Fixed effects:** Singleton county-year observations removed by `fixest`

Final sample: 54,503 county-year observations (54,479 after singleton removal), 3,072 counties, 49 states, 25 years.

B. Identification Appendix

B.1 Bacon Decomposition

The [Goodman-Bacon \(2021\)](#) decomposition of the TWFE estimator applied to high-flood counties yields an overall TWFE coefficient of 0.023. This estimate differs from the main DDD coefficient (0.0072) because the Bacon decomposition uses a simpler TWFE specification on the subsample of high-flood counties only, without the triple-differencing structure. The decomposition reveals the contribution of each 2×2 DiD comparison to the overall estimate, confirming that the positive TWFE estimate reflects a mixture of timing-group comparisons rather than a single clean treatment effect.

B.2 Pre-Trends Test

The event study in [Figure 1](#) provides a visual test of parallel trends. I also conduct a formal joint test of the hypothesis that all pre-treatment coefficients are simultaneously zero. The F -statistic from this test is 0.87 ($p = 0.52$), failing to reject the null of parallel pre-trends at any conventional significance level.

B.3 Placebo Outcomes

The placebo test on zero-flood counties (Table 3, Column 2) provides a powerful falsification exercise. Among the 12,041 county-year observations from counties with no pre-1992 flood declarations, the post-adoption coefficient is -0.004 (SE = 0.032, $p = 0.91$). This confirms that the DDD design isolates flood-specific variation rather than capturing a general state-level effect of disclosure legislation.

C. Robustness Appendix

C.1 Alternative Clustering

The main specification clusters standard errors at the state level, the level of treatment assignment. Two-way clustering by state and year yields virtually identical standard errors (0.0090 versus 0.0091), confirming that temporal correlation does not substantially affect inference in this setting.

C.2 Treatment Intensity

The NRDC grades (A through F) provide an ordinal measure of disclosure quality. I convert grades to a 0–4 scale (F=0, D=1, C=2, B=3, A=4) and interact with post and high-flood. The triple interaction coefficient of 0.0017 (SE = 0.0026) is consistent with the null: even accounting for variation in disclosure stringency, there is no detectable effect on housing values.

C.3 Third Wave Restriction

Restricting the sample to the most recent adoption wave (2019–2024, 10 states) versus never-treated states tests whether the null result is driven by the older adoption cohorts, which may have had different information environments. The coefficient of 0.005 (SE = 0.006) confirms the null in the contemporary policy environment.

C.4 HonestDiD Sensitivity Analysis

The [Rambachan and Roth \(2023\)](#) sensitivity analysis evaluates robustness to violations of the parallel trends assumption. I construct confidence sets for the treatment effect under the restriction that post-treatment deviations from parallel trends are bounded by M times the maximum pre-treatment deviation (the smoothness restriction). For $M \in [0, 0.05]$, zero is included in all confidence sets, confirming that the null result is robust to modest violations

of parallel trends. The identified set is narrowest at $M = 0$ and widens gradually, but the sign of the effect remains ambiguous even at $M = 0.05$.

D. Heterogeneity Appendix

D.1 Wave-Specific Estimates

The wave-specific DDD estimates (Figure 4) reveal modest heterogeneity across adoption cohorts:

- **First Wave (1992–1999):** Coefficient = small positive, insignificant. These early adopters acted when public awareness of flood risk was lower, but the long post-treatment period provides substantial statistical power. The null is precisely estimated.
- **Second Wave (2000–2006):** Coefficient = larger positive, wider confidence interval. This wave includes only five states, limiting precision.
- **Third Wave (2019–2024):** Coefficient = near zero, narrow confidence interval. Despite heightened climate awareness, the most recent adopters show no detectable effect.

The absence of a clear trend across waves argues against several alternative hypotheses. If disclosure effects were declining over time (due to improved pre-existing information), we would expect a monotonic decrease. If effects were increasing (due to growing climate salience), we would expect the opposite. The flat pattern across waves is most consistent with a genuinely small or zero effect regardless of the information environment.

E. Additional Figures and Tables

This section collects supplementary exhibits referenced in the main text.

Note: The threshold robustness analysis (varying the percentile cutoff for the “high flood” indicator) encountered collinearity issues for several threshold values due to the discrete distribution of the flood percentile variable within states. This is expected: when the threshold is set very low (e.g., 30th percentile) or very high (e.g., 95th percentile), many counties switch between the high- and low-flood groups, and the within-state variation that identifies the DDD coefficient is absorbed by the fixed effects. The main specification’s threshold (above-median among positive counties) was chosen to maximize within-state variation.

Table 4: State Flood Disclosure Law Adoption Dates

State	Year	Wave	NRDC Grade
Kentucky	1992	First (1992–1999)	C
Indiana	1993	First (1992–1999)	C
Michigan	1993	First (1992–1999)	C
Ohio	1993	First (1992–1999)	C
Oregon	1993	First (1992–1999)	C
South Dakota	1993	First (1992–1999)	C
Iowa	1994	First (1992–1999)	C
Nebraska	1994	First (1992–1999)	C
Washington	1994	First (1992–1999)	C
Oklahoma	1995	First (1992–1999)	A
Connecticut	1996	First (1992–1999)	D
Nevada	1996	First (1992–1999)	C
Pennsylvania	1996	First (1992–1999)	C
California	1998	First (1992–1999)	A
Illinois	1998	First (1992–1999)	C
Alaska [†]	1999	First (1992–1999)	C
Delaware	2000	Second (2000–2006)	C
New York	2001	Second (2000–2006)	B
Louisiana	2003	Second (2000–2006)	A
Mississippi	2005	Second (2000–2006)	A
Tennessee	2006	Second (2000–2006)	B
North Dakota	2019	Third (2019–2024)	C
South Carolina	2019	Third (2019–2024)	A
Texas	2019	Third (2019–2024)	A
Hawaii	2022	Third (2019–2024)	A
Maine	2023	Third (2019–2024)	C
New Jersey	2023	Third (2019–2024)	B
Florida	2024	Third (2019–2024)	B
New Hampshire	2024	Third (2019–2024)	C
North Carolina	2024	Third (2019–2024)	A
Vermont	2024	Third (2019–2024)	C

Never-Treated States (19 states, no mandatory disclosure):

AL, AR, AZ, CO, GA, ID, KS, MA, MD, MN, MO, MT, NM, RI, UT, VA, WI, WV, WY

Notes: Adoption dates reflect when states first required sellers to disclose flood-related information (flood zone status, flooding history, flood damage, or flood insurance) as part of mandatory property condition disclosure. NRDC grades from the Natural Resources Defense Council 2024 State Flood Disclosure Scorecard. States with grade F have no mandatory flood disclosure requirement. [†]Alaska is excluded from the analysis sample due to limited county-level ZHVI data; the remaining 30 treated states form the treatment group.