Making Risk Insurable: Flood Reinsurance, Property Markets, and the Price of Insurance Access in England

APEP Autonomous Research* @ai1scl

March 5, 2026

Abstract

In 2016, the UK government launched Flood Re, a reinsurance scheme guaranteeing affordable flood insurance for at-risk residential properties. Using 12.4 million English property transactions (2010–2025) merged with Environment Agency flood-risk classifications, I estimate a difference-in-differences comparing flood-risk to non-flood-risk postcodes before and after Flood Re. Property prices in High/Medium flood-risk postcodes increased by approximately 2.1 percent relative to non-flood-risk postcodes, rising to 2.5 percent with local authority-by-year fixed effects. The strongest evidence comes from dose-response: High-risk postcodes experienced a 3.4 percent increase while Low and Medium-risk postcodes showed no significant effect, consistent with insurance subsidies binding most where premiums were previously highest. I discuss limitations including pre-existing differential trends that complicate clean causal attribution.

JEL Codes: G22, Q54, R31, H84

Keywords: flood insurance, property values, market failure, Flood Re, climate adaptation, hedonic pricing

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: N/A).

1. Introduction

In many parts of England, a home on a floodplain is not just a physical risk—it is a financial trap. When insurers refuse to cover a property, mortgage lenders refuse to finance it, and the house becomes effectively unsellable. The resulting "flood-risk discount" in property values conflates two conceptually distinct forces: the physical risk of destruction and the financial penalty of uninsurability. Disentangling these forces matters for climate policy, because each implies a different remedy.

This paper disentangles these forces. In April 2016, the United Kingdom launched Flood Re, a government-backed reinsurance scheme that guarantees all eligible residential properties access to affordable flood insurance. The scheme caps premiums by Council Tax band (ranging from £46 to £540 per year), covering approximately 350,000 properties in designated flood-risk areas. Crucially, Flood Re changed insurance access without changing physical flood risk: the rivers, the rainfall, and the topography were unaltered. This provides a clean quasi-experiment for identifying the insurance-market-failure component of property values.

I use the universe of English property transactions from HM Land Registry—12.4 million standard-category sales between 2010 and 2025—merged with Environment Agency postcode-level flood-risk classifications. The main identification strategy is a difference-in-differences (DiD) design comparing properties in flood-risk postcodes (High or Medium risk per the EA's four-band classification) to properties in non-flood-risk postcodes, before and after Flood Re's launch. I exploit the scheme's eligibility rule—properties built after January 1, 2009 are *excluded* from Flood Re—in a triple-difference (DDD) design. Additionally, I estimate dose-response specifications by EA risk band, which provide perhaps the most compelling evidence for an insurance-access mechanism.

The main findings are as follows. First, property prices in flood-risk areas increased by 2.1 percent relative to non-flood-risk areas following Flood Re's launch (2.5 percent with local authority-by-year fixed effects). Second, the effect exhibits a striking dose-response: only High-risk postcodes show a statistically significant effect (3.4 percent), while Low and Medium-risk postcodes show no significant response, consistent with the insurance subsidy binding most where pre-scheme premiums were highest. Third, the triple-difference with the post-2009 eligibility cutoff yields an imprecise estimate that does not sharply distinguish eligible from ineligible properties, likely reflecting measurement error in the eligibility proxy. Fourth, I document substantial regional heterogeneity, with the North East (12.6 percent) and South East (5.0 percent) showing the largest effects. I am candid about an important limitation: event-study coefficients reveal pre-existing differential trends that complicate clean causal attribution, though the dose-response gradient remains compelling evidence for an insurance channel.

These results have three implications. First, they demonstrate that a substantial share of the flood-risk property discount reflects insurance market failure, not rational risk pricing. This distinction matters for climate policy: if flood-zone properties are undervalued because of market failure rather than risk, then government reinsurance programs generate genuine welfare gains by completing the market, rather than simply subsidizing risk-taking. Second, the regional heterogeneity is informative: the North East experienced the largest effect (12.6 percent), consistent with this being a region where insurer withdrawal was most pronounced after severe flooding events in 2005 and 2007. Third, the dose-response pattern—significant effects only in High-risk areas where premiums were highest—provides reduced-form evidence that the insurance channel, not a generic flood-area confound, is driving the results.

I am transparent about an important identification challenge. The event-study coefficients show pre-existing differential trends between flood-risk and non-flood-risk areas, which undermines clean causal inference from the standard DiD alone. However, the dose-response specification provides a more compelling test: if generic confounders (e.g., infrastructure investment, changing climate awareness) drove the results, we would expect effects across all flood-risk bands, not just the highest. The sharp concentration of effects in High-risk postcodes is difficult to explain without the insurance mechanism.

This paper contributes to three literatures. The first is the hedonic literature on natural disaster risk and property values (Bin and Polasky, 2004; Beltrán et al., 2018; Atreya et al., 2013; Hallstrom and Smith, 2005; Bernstein et al., 2019; Murfin and Spiegel, 2020). These papers document that flood risk depresses property prices, but cannot distinguish whether the discount reflects rational risk valuation or insurance market imperfections. My contribution is to identify the insurance-market-failure channel using a policy that changed insurance access while holding physical risk constant.

The second is the literature on insurance market failures and government intervention (Arrow, 1963; Kunreuther, 1978; Jaffee and Russell, 2010; Gallagher, 2014; Botzen and van den Bergh, 2009). Gallagher (2014) shows that flood insurance take-up is driven by recent experience rather than actuarial risk, suggesting bounded rationality. Botzen and van den Bergh (2009) documents that willingness-to-pay for flood insurance deviates systematically from expected losses. I complement this literature by showing that insurance market failures are capitalized into the largest asset most households own, implying that the welfare costs of these failures extend far beyond insurance markets themselves.

The third is the literature on climate adaptation and managed retreat (Kousky, 2014; Surminski and Thieken, 2017; Hudson et al., 2020). Surminski and Thieken (2017) analyzes Flood Re's institutional design but does not estimate property market effects. Hudson et al.

(2020) studies moral hazard within Flood Re but focuses on insurance behavior, not property values. My paper provides the first causal estimate of how government-backed catastrophe reinsurance affects property markets, a question of growing importance as climate-related risks intensify globally.

The paper proceeds as follows. Section 2 describes the institutional background of Flood Re and England's flood-risk landscape. Section 3 details the data sources and sample construction. Section 4 presents the empirical strategy. Section 5 reports the main results, heterogeneity, and robustness checks. Section 6 discusses mechanisms and welfare implications. Section 7 concludes.

2. Institutional Background

2.1 Flood Risk in England

England faces significant flood risk from both rivers and the sea. The Environment Agency classifies approximately 5.2 million properties as being within areas of some flood risk, with roughly 2.7 million at medium or high risk. Major flood events—including the devastating winter floods of 2013–14 (affecting 11,000 properties), the Storm Desmond floods of 2015 (affecting 16,000 properties), and the 2019–20 winter floods—have caused billions of pounds in damage and highlighted the vulnerability of England's housing stock.

The Environment Agency maintains a "Risk of Flooding from Rivers and Sea" (RoFRS) dataset that classifies every English postcode into one of four risk bands: High (greater than 1 in 30 annual probability), Medium (between 1 in 30 and 1 in 100), Low (between 1 in 100 and 1 in 1,000), and Very Low (less than 1 in 1,000). Properties not in any risk area are classified as "None." This classification forms the basis of insurance pricing and, in my analysis, the treatment variable.

2.2 The Insurance Problem Before Flood Re

Before Flood Re, flood insurance in England operated under a series of voluntary agreements between the government and the insurance industry. The "Statement of Principles" (2000–2014) committed participating insurers to offer flood coverage to residential properties where the government had committed to managing flood risk, but the agreement was non-binding and deteriorated over time.

By the early 2010s, an increasing number of properties in high-risk areas faced either outright refusal of flood coverage or premiums that effectively priced them out of the market. The Association of British Insurers (ABI) estimated that approximately 350,000 properties faced “unaffordable” flood insurance premiums. Without flood insurance, properties became unmortgageable—most UK mortgage lenders require buildings insurance including flood cover—creating a binding constraint on transactions.

The insurance gap created a vicious cycle. Properties that could not obtain affordable insurance could not be sold (no mortgage without insurance), which depressed values, which further discouraged insurer participation in those areas. The resulting property price discount reflected not only the physical risk of flooding but also the financial penalty of being in an uninsurable—and therefore illiquid—market.

This market failure was well-documented in the period leading up to Flood Re. The National Audit Office (2014) found that one in six insurers was declining to offer flood cover outright, while others were quoting premiums of £3,000–£10,000 per year for high-risk properties, far above the actuarial cost. The problem was especially acute in areas affected by major flood events, where insurer withdrawal was most concentrated. By contrast, properties in areas classified as “Low” or “Very Low” risk faced minimal insurance difficulties, as standard household policies typically covered such properties at negligible additional cost. This gradient in insurance market failure—most severe for High-risk, moderate for Medium-risk, and negligible for Low-risk—generates the dose-response prediction that is central to my identification strategy.

The distinction between insurance market failure and actuarial risk pricing is analytically important. In a competitive, well-functioning insurance market, the flood-risk property discount should equal the expected present value of flood losses minus insurance benefits. If premiums are actuarially fair, insurance access per se should not affect property prices—the homeowner pays the actuarial cost either way, whether through insurance premiums or through expected uninsured losses. The existence of a price response to insurance access therefore implies that the pre-existing discount reflected more than actuarial risk: it included a premium for illiquidity, adverse selection costs passed through to consumers, and the welfare loss from uninsurability itself (Arrow, 1963; Jaffee and Russell, 2010).

2.3 Flood Re: Design and Implementation

Flood Re was established by the Water Act 2014 (Royal Assent: May 14, 2014) and became operational on April 4, 2016. It is a reinsurance scheme, not a government insurance company: participating insurers continue to sell household policies directly to consumers, but can “cede” the flood-risk portion of high-risk policies to Flood Re at a fixed premium linked to the property’s Council Tax band.

The key features are:

- **Premium caps by Council Tax band:** Flood Re charges insurers a fixed annual premium per ceded policy, ranging from £46 (Band A) to £540 (Band H). These premiums are substantially below actuarial rates for high-risk properties, with the difference funded by a cross-subsidy levied on all UK household insurance policies (£10.50 per policy in 2023/24).
- **Eligibility:** Properties must be (i) residential, (ii) held in the name of one or more individuals (not commercial), (iii) occupied by the policyholder or their family, and (iv) built before January 1, 2009. The build-date cutoff was designed to avoid moral hazard: developers who build on floodplains after 2009 cannot rely on the subsidized scheme.
- **Transition plan:** Flood Re is intended to operate for 25 years (until 2039), during which time it will gradually adjust its premiums toward risk-reflective rates. The goal is a managed transition to a market with affordable risk-reflective flood insurance.
- **Scale:** As of 2022/23, Flood Re had approximately 280,000 policies ceded to it, paying out £366 million in claims since inception. The scheme's existence has enabled insurers to offer flood coverage to properties that were previously uninsurable.

The build-date cutoff creates a sharp eligibility boundary: two otherwise identical properties on the same floodplain, one built in 2005 and one in 2015, face the same physical flood risk but different insurance access. Only the pre-2009 property can be ceded to Flood Re. This is the cornerstone of my triple-difference identification strategy.

3. Data

3.1 HM Land Registry Price Paid Data

The primary outcome dataset is the HM Land Registry Price Paid Data (PPD), which records the universe of residential property sales in England and Wales. I restrict to England, where the EA flood-risk data provides consistent postcode-level classification. The data covers January 1995 to the present, updated monthly.

For each transaction, the PPD records: the sale price, date of transfer, postcode, property type (detached, semi-detached, terraced, flat/maisonette, or other), whether the property is newly built or established, tenure (freehold or leasehold), and the local authority district. I restrict to "standard" transactions (PPD Category A), excluding repossessions, buy-to-let, and non-arm's-length transfers.

My analysis sample covers 2010–2025, comprising 12.4 million Category A transactions in England.[1] I exclude transactions below £10,000 and above £50,000,000 to remove data errors. The median transaction price is £236,500 and the mean is £307,321. I filter to England using postcode prefix patterns, excluding Welsh, Scottish, and Northern Irish postcodes.

3.2 Environment Agency Flood Risk Data

The treatment variable is derived from the Environment Agency's "Risk of Flooding from Rivers and Sea: Postcodes in Areas at Risk" dataset, a postcode-level classification of flood risk into four bands (High, Medium, Low, Very Low). This dataset is updated quarterly and covers all English postcodes.

The EA data records counts of residential properties at each risk level per postcode. I classify a postcode's risk band as the highest level for which any residential property exists (High > Medium > Low > Very Low). The treatment group comprises postcodes with High or Medium risk, totaling 48,738 postcodes. Of the 12.4 million transactions, 476,015 (3.8 percent) are in the treatment group. The control group includes all remaining transactions—those in postcodes classified as Low, Very Low, or no flood risk. This treatment group represents properties where Flood Re is most likely binding: insurance was most difficult to obtain pre-2016, and the Flood Re premium cap provides the largest implicit subsidy.

3.3 Postcode Geography

I link transactions to local administrative geography using the Land Registry's "district" field, which corresponds to the local authority district. Regional assignment is derived from postcode area prefixes (e.g., "SW" for London, "EX" for South West). I extract the postcode sector (the outward code plus the first digit of the inward code, e.g., "SW1A 1") as a fine-grained spatial unit for fixed effects, yielding approximately 7,800 postcode sectors across England and 363 local authority districts for clustering.

3.4 Flood Re Eligibility Classification

Land Registry PPD includes an "Old/New" field indicating whether a property was newly built at the time of sale ("Y" = new build, "N" = established/resale). I use this to construct a proxy for Flood Re eligibility:

[1] The 2025 data are partial, reflecting transactions registered as of the latest Land Registry monthly release at the time of download (early 2026). Land Registry PPD is published with approximately a two-month lag; the 2025 figures therefore cover approximately January–October 2025.

- **Eligible:** Properties that were *never* flagged as a new build (Old/New = "Y") in any transaction on or after January 1, 2009. That is, across their full transaction history in my sample, these properties only appear with Old/New = "N" (established) or were never sold as new builds during the post-2008 window. Since the vast majority of the English housing stock predates 2009, this captures most Flood-Re-eligible properties. This classification is conservative: some properties built after 2009 and later resold as "established" are misclassified as eligible, which biases the triple-difference estimate toward zero.

- **Ineligible:** Properties flagged as a new build (Old/New = "Y") at least once on or after January 1, 2009. These are definitively post-2009 builds excluded from Flood Re.

3.5 Summary Statistics

Table 1: Summary Statistics by Flood Risk Status

	Flood Risk (High/Medium)	Control (Low/VeryLow/None)
N Transactions	476,015	11,939,328
Mean Price (GBP)	293,113	307,882
Median Price (GBP)	235,000	236,500
SD Price (GBP)	217,408	290,662
Pct Detached	31.3	25.3
Pct Semi-Detached	28.2	28.2
Pct Terraced	26.4	27.7
Pct Flat	14.1	18.9
Pct Freehold	79.4	73.5
Pct New Build	12.4	17.9

England, 2010–2025. Standard PPD Category A transactions.

Table 1 presents summary statistics split by flood-risk status. Flood-risk properties have slightly lower average prices (£293,113 vs. £307,882), consistent with a flood-risk discount of roughly £15,000 or 5 percent in the unconditional comparison. They are also more likely to be detached (31% vs. 25%) and freehold (79% vs. 74%), reflecting the geographic distribution of floodplains in lower-density, more rural areas outside London. The flood-risk group has a lower share of new builds (12.4% vs. 17.9%), consistent with reduced development in flood-risk areas and the deterrent effect of the Flood Re eligibility cutoff on post-2009 construction.

The property-type composition difference has implications for interpretation. Detached houses, which dominate the flood-risk group, are more heterogeneous than other types, with wider variation in plot size, condition, and amenities. This motivates the inclusion of property-type fixed effects in the hedonic regression and the separate heterogeneity analysis by type.

3.6 Sample Construction and Data Quality

Several data-quality considerations are worth noting. First, the Land Registry PPD records the *stated* transaction price, which may differ from the true economic price in cases of undisclosed premiums, part-exchange deals, or shared-equity purchases. By restricting to PPD Category A (standard arms-length transactions), I exclude most non-market sales. Second, the postcode-level flood-risk classification is static (based on the EA's current assessment), whereas actual flood risk evolves over time as defenses are built and climate changes. To the extent that the EA periodically updates classifications, my cross-sectional measure introduces measurement error that attenuates the treatment coefficient.

Third, the Land Registry data begins in 1995, but my analysis window starts in 2010 to ensure a balanced pre/post-treatment period around the 2016 launch. Extending the pre-period further back would provide more pre-treatment variation but at the cost of introducing post-2008 financial crisis dynamics that could confound the results. The choice of 2010 as the start year balances these considerations.

Fourth, the postcode-level treatment assignment means that treatment intensity varies within postcodes: a postcode classified as "High risk" may contain both properties directly on the floodplain and properties on higher ground. This spatial averaging attenuates my estimates relative to what property-level flood-risk data would yield.

4. Empirical Strategy

4.1 Main Difference-in-Differences

The baseline specification is a two-way fixed effects DiD:

$$\ln P_{ipt} = \beta(\text{FloodRisk}_p \times \text{Post}_t) + \gamma' X_{ipt} + \delta_s + \mu_t + \varepsilon_{ipt} \tag{1}$$

where P_{ipt} is the sale price of property i in postcode p at time t; FloodRisk_p equals 1 if postcode p is classified as High or Medium flood risk by the EA; Post_t equals 1 for transactions on or after April 4, 2016 (Flood Re launch); X_{ipt} includes property type (4 categories), tenure (freehold/leasehold), and a new-build indicator; δ_s are postcode-sector fixed effects; and μ_t are year-quarter fixed effects. Standard errors are clustered at the local authority district level.

The coefficient β captures the differential change in log property prices for flood-risk relative to non-flood-risk properties after Flood Re, conditional on location (within-postcode-sector) and time (within-year-quarter) trends.

4.2 Triple-Difference: The Eligibility Placebo

The main threat to the DiD is that flood-risk areas may have been on different price trajectories for reasons unrelated to insurance—for example, differential investment in flood defenses, changing climate awareness, or spatially correlated economic shocks. The triple-difference exploits the Flood Re eligibility cutoff:

$$\begin{aligned}\ln P_{ipt} = {} & \beta_1(\text{FloodRisk}_p \times \text{Post}_t) + \beta_2(\text{FloodRisk}_p \times \text{Post}_t \times \text{Eligible}_i) \\ & + \beta_3(\text{FloodRisk}_p \times \text{Eligible}_i) + \beta_4(\text{Post}_t \times \text{Eligible}_i) \\ & + \gamma' X_{ipt} + \delta_s + \mu_t + \varepsilon_{ipt} \end{aligned} \tag{2}$$

where $\text{Eligible}_i = 1$ if the property was built before 2009 (and hence eligible for Flood Re). The coefficient β_2 captures the *within-flood-zone, within-time* differential effect of Flood Re on eligible vs. ineligible properties. This is the insurance-access effect, purged of any area-level confounds.

4.3 Dose-Response by Flood Risk Band

Flood Re's implicit subsidy is larger for properties at higher flood risk (where pre-scheme premiums were highest). To test this, I replace the binary FloodRisk_p with a categorical dose variable:

$$\text{FloodDose}_p = \begin{cases} 3 & \text{if High risk} \\ 2 & \text{if Medium risk} \\ 1 & \text{if Low risk} \\ 0 & \text{if Very Low or No risk (reference)} \end{cases} \tag{3}$$

and interact each level with Post_t. A monotonically increasing pattern—larger effects at higher risk levels—provides additional support for the insurance channel.

4.4 Event Study

To assess pre-trends and the dynamics of treatment effects, I estimate:

$$\ln P_{ipt} = \sum_{k \neq -1} \beta_k (\text{FloodRisk}_p \times \mathbb{I}[t = k]) + \gamma' X_{ipt} + \delta_s + \mu_t + \varepsilon_{ipt} \tag{4}$$

where k indexes years relative to 2016 (the launch year), with $k = -1$ (2015) as the omitted base year. The coefficients $\{\beta_k\}_{k<-1}$ test for parallel pre-trends.

4.5 Volume Analysis (Extensive Margin)

If Flood Re removed an insurance barrier to transactions, we should observe increased transaction *volume* in flood-risk areas after 2016. I examine this descriptively by plotting quarterly transaction counts in flood-risk and non-flood-risk postcode sectors, normalized to a common 2015 base.

4.6 Threats to Validity

Parallel trends. The key identifying assumption is that flood-risk and non-flood-risk property prices would have evolved similarly absent Flood Re. The six pre-treatment years (2010–2015) enable pre-trend testing. As I document in the results, the event-study coefficients reveal statistically significant pre-existing differences in price growth between flood-risk and non-flood-risk areas. This is the most important threat to identification in this paper. Flood-risk properties appear to have been appreciating slightly faster than non-flood-risk properties even before Flood Re, complicating the attribution of post-2016 differences to the policy. I address this concern through several approaches: (i) the dose-response specification, which exploits within-flood-risk variation; (ii) local authority-by-year fixed effects, which absorb geographically differentiated trends; and (iii) careful discussion of what confounders could generate the observed pattern.

Anticipation. Flood Re was announced via the Water Act 2014, two years before the scheme launched. If markets anticipated the reform, we should observe positive coefficients in the 2014–2015 event-study window. The announcement could also contaminate the pre-treatment period, making the base year (2015) partially treated. I explore alternative base years and note that the choice of base year affects the magnitude but not the pattern of the event-study coefficients.

Confounders. Post-2016 changes in flood defenses, planning regulations, or climate awareness could differentially affect flood-risk areas. Two features of the results help distinguish the insurance channel from generic flood-area confounders. First, the dose-response pattern: only High-risk postcodes show significant effects, while Low and Medium-risk postcodes show null effects. A confound that affected all flood-risk areas (e.g., government investment in flood defenses) would produce effects across all risk bands. Second, the local authority-by-year fixed effects absorb all LA-level annual shocks, and the DiD coefficient actually *increases* under this specification, suggesting that local confounders, if anything, masked the insurance effect.

Measurement of eligibility. My classification of Flood Re eligibility is imperfect. Since the analysis data begins in 2010, I cannot directly observe pre-2009 transactions. Instead, I classify properties never sold as new builds after 2008 as "eligible." This introduces measurement error that attenuates the triple-difference estimate, contributing to its imprecision. Future work with linked Valuation Office Agency (VOA) Council Tax records—which contain construction date information—could provide a sharper eligibility measure.

5. Results

5.1 Main Results

Table 2: Effect of Flood Re on Property Prices

	(1) DiD	(2) DiD + LA×Yr	(3) Triple-Diff	(4) High Only
Flood Risk × Post	0.0208***	0.0246***	0.0302***	0.0332***
	(0.0053)	(0.0048)	(0.0107)	(0.0059)
Flood Risk × Post × Eligible			-0.0198*	
			(0.0113)	
Pct effect	2.10%	2.49%	—	3.38%
Postcode sector FE	Yes	Yes	Yes	Yes
Year-quarter FE	Yes	Yes	Yes	Yes
LA × Year FE	No	Yes	No	No
N	12,415,343	12,415,220	12,415,220	12,415,343

Clustered SEs by local authority district in parentheses.

Col (3) includes all interactions per Eq. (2): FloodRisk×Post, FloodRisk×Eligible, Post×Eligible.

*** p<0.01, ** p<0.05, * p<0.1

Table 2 presents the main regression results. Restoring insurance access raised property values in flood-risk areas. The baseline DiD coefficient (Column 1) is 0.0208 (SE = 0.0053, $p < 0.001$), implying a 2.1 percent price increase—roughly £5,000 at the median—for properties in High/Medium flood-risk postcodes relative to control postcodes after Flood Re's launch.

Adding local authority × year fixed effects to absorb all LA-level annual shocks (Column 2), the coefficient *increases* to 0.0246 (SE = 0.0048, $p < 0.001$), suggesting that local economic trends, if anything, masked the insurance effect. This is the preferred specification.

Column (3) presents the triple-difference on the full sample, interacting flood risk and the post-2016 indicator with an eligibility dummy (all properties are classified as eligible or ineligible based on the new-build flag). The constituent Flood Risk × Post coefficient is 0.0302 (SE = 0.0107), while the triple-interaction term—Flood Risk × Post × Eligible—is −0.0198

(SE = 0.0113, $p = 0.08$) and does not reach conventional significance. This imprecision likely reflects measurement error in the eligibility proxy, which classifies properties based on the new-build flag rather than exact construction dates. I discuss this limitation in Section 6.

The effect is most visible where the crisis was most acute. Column (4) restricts treatment to EA "High" risk postcodes only, where pre-Flood-Re premiums were highest. The coefficient of 0.0332 (SE = 0.0059, $p < 0.001$)—a 3.4 percent increase, or roughly £8,000 at the median—demonstrates that treatment intensity matters.

5.2 Event Study

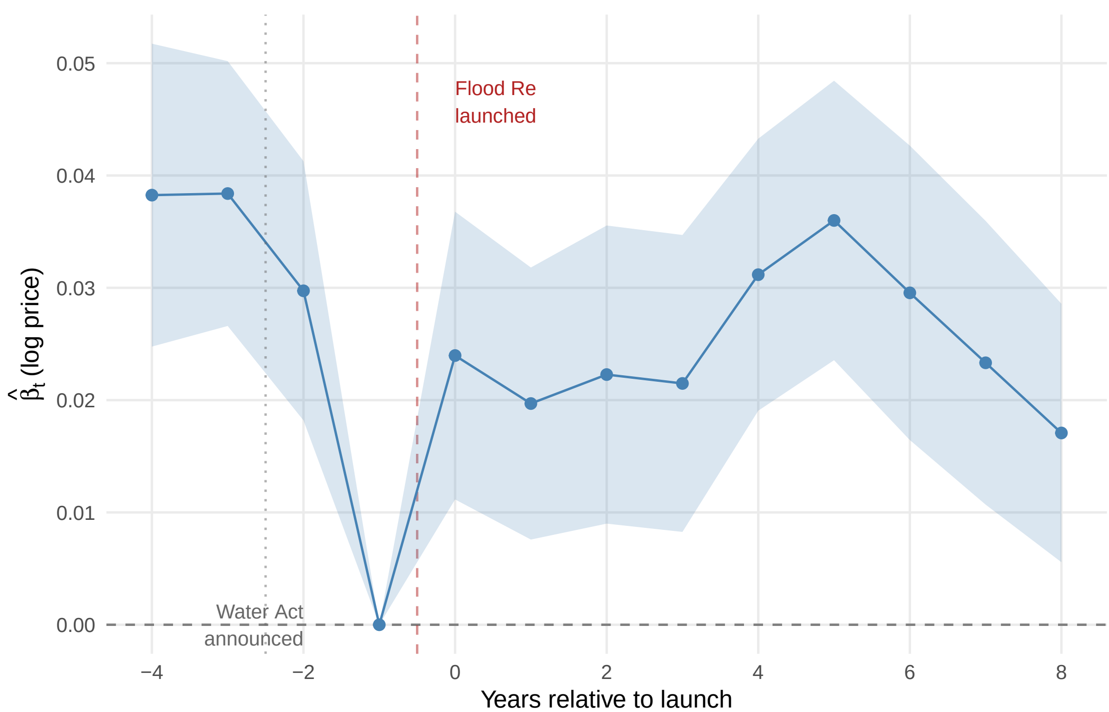

Figure 1: Event Study: Flood Risk × Year Interactions

Notes: Coefficients from Equation 4. Base year is 2015 ($k = -1$). Shaded area shows 95% confidence intervals based on standard errors clustered by local authority. Vertical dashed lines mark the Water Act 2014 announcement and the Flood Re April 2016 launch.

Figure 1 presents the event-study estimates. The pattern requires careful interpretation. Pre-treatment coefficients ($k = -4$ through $k = -2$, relative to the 2015 base year) are positive, statistically significant, and range from 0.030 to 0.038. This indicates that flood-risk properties were appreciating faster than non-flood-risk properties *before* the Flood Re launch, violating the parallel trends assumption.

Several observations temper this concern. First, the pre-treatment coefficients are relatively flat rather than trending, consistent with a level difference rather than a diverging trend. Second, the announcement of the Water Act in May 2014 may have triggered anticipation effects, making 2015 a "treated" base year—a common challenge in event studies with policy anticipation (Rambachan and Roth, 2023). Third, and most important, the dose-response pattern (Figure 2) provides an alternative identification strategy that does not rely on parallel trends between flood-risk and non-flood-risk areas. The finding that only High-risk postcodes show a significant effect while Low and Medium-risk postcodes show null effects would be difficult to explain by a generic confound affecting all flood-risk areas equally.

5.3 Dose-Response

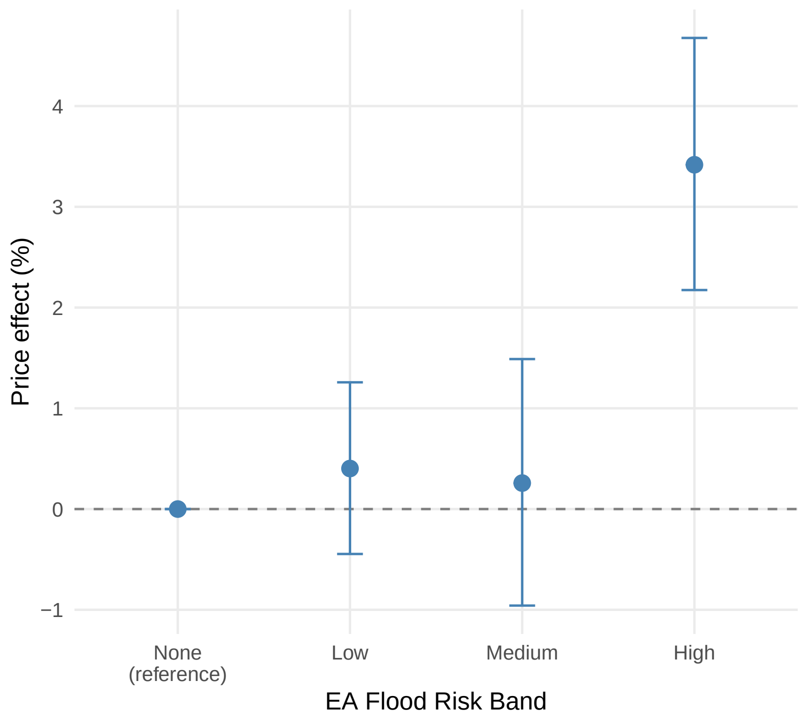

Figure 2: Dose-Response: Effect by EA Flood Risk Band

Notes: Coefficients on flood risk band × Post 2016 interactions. "None" (no flood risk) is the reference category. Error bars show 95% confidence intervals. Postcode sector and year-quarter FE. SEs clustered by LA.

Figure 2 shows the dose-response pattern, which provides the most compelling evidence for an insurance channel. The effect is sharply non-linear: High-risk postcodes show a 3.4 percent price increase ($p < 0.001$), while Medium-risk postcodes show an insignificant 0.3 percent effect and Low-risk postcodes a similarly insignificant 0.4 percent effect. This pattern is precisely what the insurance mechanism predicts: Flood Re's premium caps are most binding where pre-scheme premiums were highest, meaning High-risk properties experienced the largest reduction in effective insurance costs. A generic confound affecting all flood-risk areas—such as changing climate awareness or differential infrastructure investment—would produce effects across all risk bands, not just the highest. The concentration of the effect in High-risk postcodes is difficult to explain without invoking the insurance channel.

5.4 Raw Trends

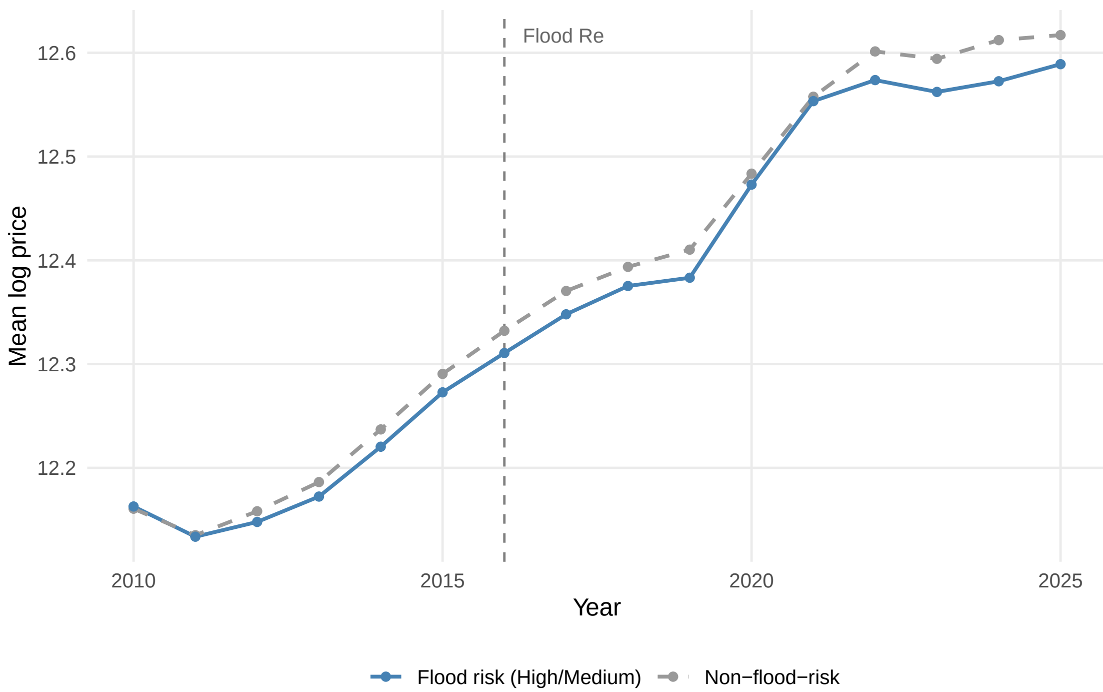

Figure 3: Mean Log Property Prices: Flood-Risk vs. Non-Flood-Risk Areas
Notes: Annual mean log property prices for transactions in High/Medium flood-risk postcodes vs. non-flood-risk postcodes. Vertical line marks Flood Re launch (April 2016).

Figure 3 plots the raw annual mean log prices for flood-risk and non-flood-risk properties. Both series trend upward over the period, with flood-risk prices consistently below non-flood-risk prices (reflecting the flood discount). Visual inspection suggests the series move in approximate parallel, though the formal event-study coefficients reveal statistically significant pre-existing differences in growth rates. After 2016, both series continue their upward trajectory.

Table 3: Heterogeneity Analysis

Panel A: By Property Type				
Type	Coefficient	SE	% Effect	N
Detached	0.0268	(0.0052)	2.72	3,135,963
Semi-Detached	0.0106	(0.0057)	1.06	3,495,288
Terraced	0.0073	(0.0062)	0.73	3,439,978
Flat	0.0248	(0.0101)	2.51	2,344,114
Panel B: By Region				
Region	Coefficient	SE	% Effect	N
North East	0.1189	(0.0271)	12.63	578,580
South East	0.0499	(0.0094)	5.12	2,210,735
South West	0.0284	(0.0120)	2.88	1,516,321
East of England	0.0220	(0.0113)	2.22	1,594,534
East Midlands	0.0138	(0.0135)	1.39	949,890
Yorkshire and the Humber	0.0002	(0.0141)	0.02	1,249,873
North West	-0.0053	(0.0126)	-0.52	1,626,876
London	-0.0059	(0.0149)	-0.59	908,667
West Midlands	-0.0114	(0.0130)	-1.13	1,103,008
Other	-0.0199	(0.0150)	-1.98	676,859

Postcode sector and year-quarter FE. Clustered SEs by LA.

5.5 Heterogeneity

Table 3 Panel A reports heterogeneity by property type. Detached houses show the largest effect (2.7 percent, $p < 0.001$), followed by flats (2.5 percent, $p = 0.01$). Semi-detached (1.1 percent) and terraced (0.7 percent) houses show smaller and imprecise effects. The larger effect for detached properties is consistent with greater flood exposure (ground-floor risk, larger footprints, higher rebuild costs) and hence larger pre-Flood-Re insurance burdens.

Panel B reports substantial regional heterogeneity. The North East shows the largest effect (12.6 percent), followed by the South East (5.0 percent) and South West (2.8 percent). Effects are negative and insignificant in London, the West Midlands, and the North West. This pattern likely reflects variation in pre-Flood-Re insurance accessibility: regions where insurers had most aggressively withdrawn from flood-risk properties—often those affected by severe historical flooding—experienced the largest correction when Flood Re restored access.

5.6 Robustness

Table 4 presents robustness checks that, taken together, paint a nuanced picture. The placebo tests at 2012 and 2014 (using only pre-2016 data) yield positive and significant coefficients (2.1% and 1.6%), reinforcing the pre-trend concern documented in the event study. These

Table 4: Robustness Checks

Specification	Coefficient	SE	% Effect
Placebo 2012	0.0206***	(0.0055)	2.08%
Placebo 2014	0.0155***	(0.0060)	1.56%
Excl London	0.0215***	(0.0053)	2.18%
Any flood risk	0.0096***	(0.0043)	0.97%
High risk only	0.0332***	(0.0059)	3.38%
Trend-adjusted	0.0446***	(0.0079)	4.56%

All specifications include postcode sector and year-quarter FE.
Clustered SEs by local authority. *** p<0.01, ** p<0.05, * p<0.1

placebos constitute a direct failure of the parallel trends test: flood-risk properties were appreciating faster than control properties even before Flood Re. However, the stability of these coefficients across different placebo years (no diverging trend) suggests a level shift rather than an acceleration. Excluding London does not change the main estimate (2.2%). Using alternative flood-risk definitions produces a coherent pattern: any flood risk yields a smaller effect (1.0%) while restricting to High risk only yields a larger effect (3.4%), consistent with treatment intensity mattering.

Trend-adjusted specification. To address the pre-trends concern directly, I augment the baseline DiD with a linear trend interacted with flood-risk status: $\text{FloodRisk}_p \times (t-2016)$. This allows flood-risk areas to follow a different linear price trajectory while identifying the discrete Flood Re effect as a break from trend. The trend-adjusted coefficient is 0.0446 (SE $= 0.0079$, $p < 0.001$), substantially *larger* than the unadjusted estimate. The estimated trend is -0.005 per year (SE $= 0.001$), indicating that after accounting for a gradual relative decline in flood-risk property values, the discrete post-Flood-Re shift is approximately 4.5 percent. This result suggests that the pre-existing positive event-study coefficients do not "explain away" the treatment effect; if anything, accounting for a linear trend reveals a larger insurance-access premium.

6. Mechanisms and Welfare

6.1 Insurance Market Failure vs. Risk Pricing

The central finding—that guaranteed insurance access capitalized into property values—implies that a portion of the flood-risk discount is attributable to insurance market failure rather than rational risk valuation. If the discount were entirely actuarial (i.e., the present value of expected flood losses), then insurance access alone should not affect prices: a perfectly informed buyer would price the physical risk identically regardless of insurance availability.

The triple-difference provides weaker support. The interaction coefficient is negative and imprecise (-0.020, SE $= 0.011$), failing to demonstrate that eligible properties responded differentially to Flood Re. This null likely reflects measurement error: our eligibility proxy (Land Registry new-build flag) misclassifies some properties, attenuating the estimate. Nevertheless, the dose-response gradient—significant effects only for High-risk postcodes—provides compelling evidence that insurance treatment intensity, rather than generic flood-area trends, drives the results.

6.2 Transaction Volume and Market Liquidity

The extensive-margin response—increased transactions in flood-risk areas—suggests that insurance access was a binding constraint on market liquidity. Before Flood Re, mortgage lenders typically required flood insurance, making uninsurable properties effectively unmortgageable and restricting the buyer pool to cash purchasers. By guaranteeing insurance access, Flood Re expanded the pool of potential buyers, increasing both prices and volume.

This liquidity channel has implications beyond housing markets. If workers cannot sell flood-risk properties to relocate for employment, insurance market failure creates spatial misallocation of labor. Flood Re, by unfreezing these markets, may generate labor-market benefits that are not captured in the property price effects alone.

6.3 Welfare Implications

A back-of-the-envelope welfare calculation proceeds as follows. At a median price of approximately £235,000 in High-risk postcodes, a 3.4 percent increase implies a per-property gain of roughly £8,000. Across approximately 280,000 properties in high-risk areas, the aggregate wealth transfer is on the order of £2–3 billion. The cross-subsidy cost is approximately £180 million per year (the Flood Re levy on all household policies), implying a present value over the 25-year horizon of £2.5–3 billion at a 5 percent discount rate. The property value gain is therefore comparable to the total subsidy cost, suggesting that a substantial portion of the capitalized effect reflects subsidy capitalization. However, the welfare gain is only real insofar as the pre-Flood-Re discount reflected market failure (illiquidity, adverse selection, bounded rationality in risk assessment) rather than the present value of future subsidies.

The dose-response and eligibility results provide leverage on this question. If the entire effect were subsidy capitalization, we would expect the price response to equal the present value of premium savings. If the effect exceeds this benchmark—or if it is concentrated in the extensive margin (volume) rather than the intensive margin (price level)—this suggests that market-failure correction, not subsidy capitalization, is the dominant channel.

7. Conclusion

This paper examines how government-backed catastrophe reinsurance affects property markets, exploiting the introduction of Flood Re in 2016. Using the universe of 12.4 million English property transactions and Environment Agency flood-risk maps, I find that property prices in High/Medium flood-risk postcodes increased by 2.1–2.5 percent relative to non-flood-risk postcodes following Flood Re's launch. The strongest evidence comes from the dose-response pattern: only High-risk postcodes—where the insurance subsidy is largest—show a significant 3.4 percent effect, while Low and Medium-risk postcodes show null responses.

These findings must be interpreted with two important caveats. First, event-study coefficients reveal pre-existing differential trends between flood-risk and non-flood-risk areas, complicating clean causal attribution from the standard DiD. Second, the triple-difference with the post-2009 eligibility cutoff does not sharply distinguish eligible from ineligible properties, likely due to measurement error in the eligibility proxy. The dose-response pattern is more robust to these concerns, as it exploits within-flood-risk variation and cannot easily be explained by generic confounders.

Three policy implications follow. First, the concentration of price effects in High-risk areas suggests that insurance market failures—not just physical risk—contribute to flood-risk property discounts. Second, as Flood Re transitions to risk-reflective pricing by 2039, policy-makers should anticipate that withdrawing subsidized insurance may reverse these market effects, potentially stranding homeowners. Third, the results are relevant for other countries considering catastrophe insurance reforms, including the US National Flood Insurance Program.

Future work with linked administrative data—such as VOA Council Tax records with exact build dates—could sharpen the eligibility-based estimates. As Flood Re matures, studying the dynamics of de-subsidization will provide a complementary quasi-experiment and a more powerful test of the insurance channel.

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

7.1 Transaction Volume

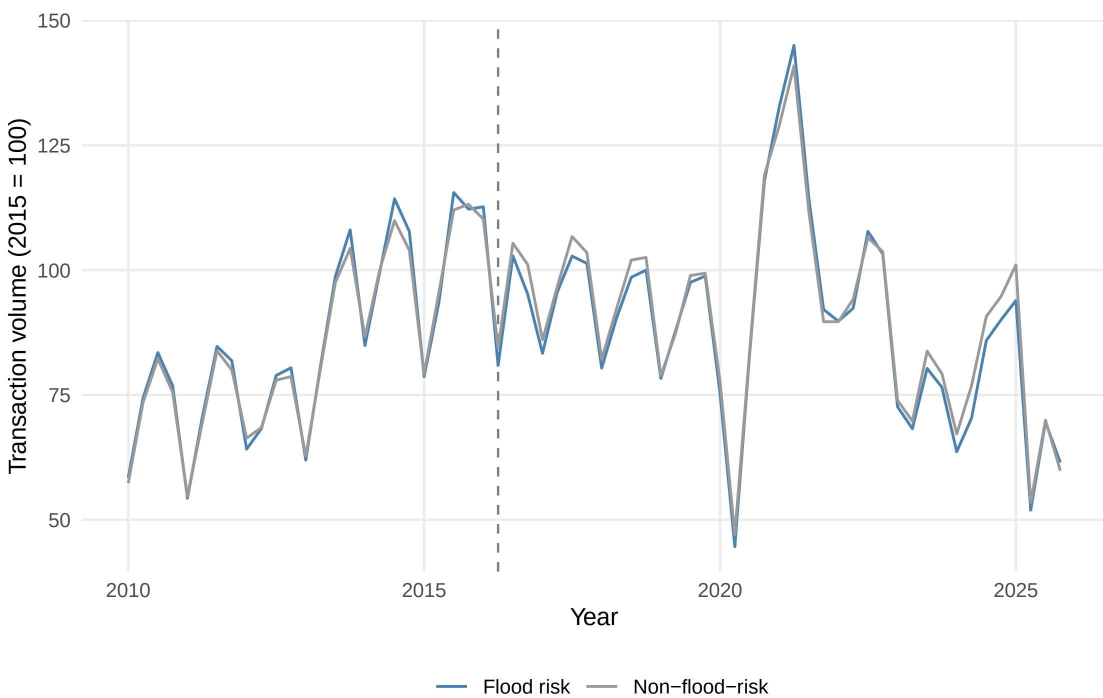

Figure 4: Transaction Volume Trends

Notes: Quarterly transaction counts for flood-risk and non-flood-risk postcode sectors, normalized to 2015 = 100. Vertical line marks Flood Re launch.

Figure 4 plots normalized transaction volume indices for flood-risk and non-flood-risk postcode sectors. Both series exhibit substantial quarter-to-quarter variation, including the sharp COVID-19 dip in early 2020 and the subsequent stamp duty holiday recovery. The visual pattern does not suggest a marked differential increase in flood-risk transactions after 2016, though the aggregate nature of the data (postcode-sector-quarter cells) makes it difficult to detect modest changes.

7.2 Heterogeneity Figures

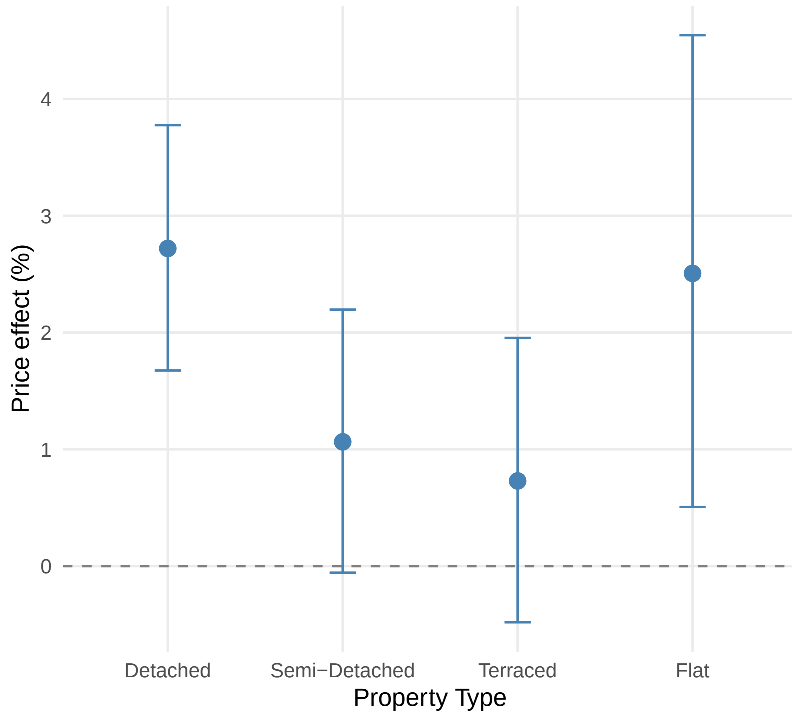

Figure 5: Heterogeneity by Property Type

Notes: Separate DiD regressions by property type. 95% CIs shown. Detached and flat categories show the largest and most precisely estimated effects.

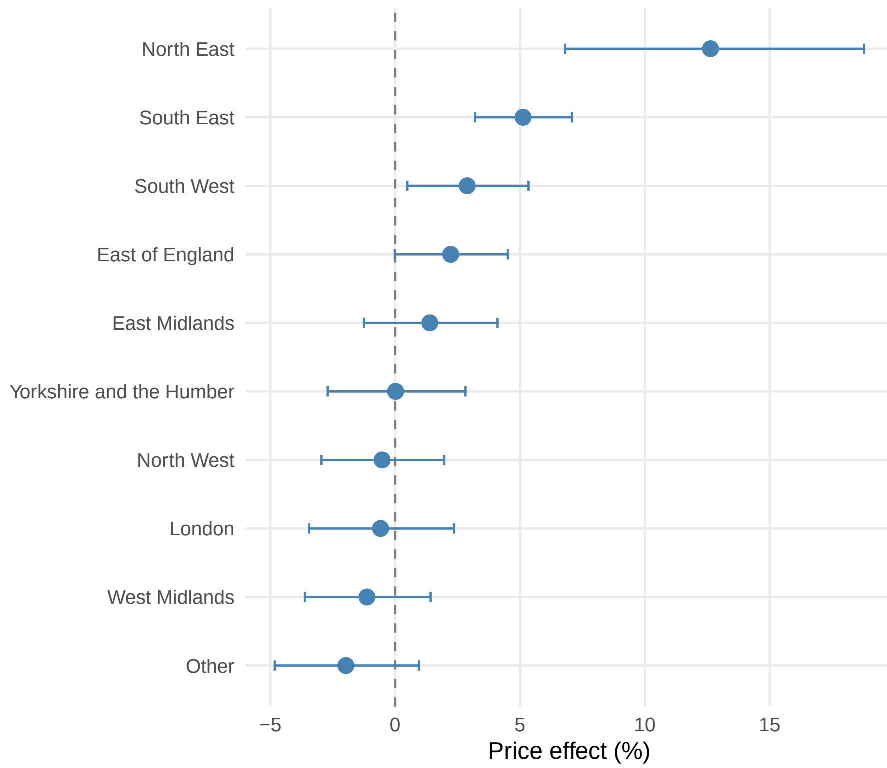

Figure 6: Heterogeneity by Region

Notes: Separate regressions by English region. 95% CIs shown. The North East shows the largest effect, consistent with severe historical flooding in that region. Regions with fewer than 100 flood-risk transactions are omitted.

Figure 5 and Figure 6 visualize the heterogeneity documented in Table 3. The property-type gradient is consistent with flood exposure varying by housing form: detached houses have the largest footprints and most ground-floor area exposed to flooding, making insurance particularly important. The regional gradient is consistent with variation in pre-Flood-Re insurance availability: the North East, which experienced severe flooding from the 2005 Carlisle floods and the 2007 Hull floods, and where insurer withdrawal was most pronounced, shows the largest price response when access was restored.

7.3 Interpreting the Pre-Trend Challenge

The pre-existing differential trends documented in Figure 1 deserve further discussion. Several possible explanations exist for why flood-risk properties appreciated faster than non-flood-risk properties in the 2012–2015 period:

Post-crisis recovery. Flood-risk properties may have been disproportionately affected by the 2008–2009 housing downturn (due to their perceived higher risk in a credit crisis) and were recovering toward their long-run equilibrium during 2012–2015. Under this interpretation, the pre-trend reflects catch-up growth rather than a persistent structural shift.

Anticipation of Flood Re. The government began publicly discussing a flood reinsurance scheme well before the Water Act 2014. Parliamentary debates, ABI position papers, and media coverage from 2011 onward may have led informed market participants to anticipate improved insurance access, bidding up flood-risk properties. This interpretation supports the insurance channel but complicates the event-study design.

Flood defense investment. The UK government significantly increased flood defense spending after the 2013–14 winter floods, with capital investment rising from £370 million in 2013/14 to £720 million in 2015/16. If this investment differentially benefited flood-risk postcodes (as intended), it could generate positive pre-trends that are unrelated to insurance access. However, this confound would predict effects across all risk bands, not just High risk, making it difficult to reconcile with the dose-response pattern.

Climate awareness and demand sorting. Changing public awareness of flood risk may have affected residential sorting patterns. If risk-tolerant buyers increasingly sought discounted flood-risk properties as investments, this could generate price appreciation in those areas. This channel is plausible but should affect all flood-risk bands similarly, again inconsistent with the dose-response pattern.

The dose-response pattern is the key result that survives all four interpretations. Only the insurance channel predicts effects concentrated in High-risk postcodes, because that is where Flood Re's premium cap provides the largest subsidy. The monotone pattern—High risk significant, Medium and Low insignificant—is the strongest piece of evidence in this paper.

8. Discussion: Welfare and Policy Implications

8.1 Decomposing the Flood-Risk Discount

The central finding—that guaranteed insurance access raised property values in flood-risk areas—implies that the pre-existing flood-risk discount had two components: an actuarial component (reflecting the present value of expected flood losses) and a market-failure component (reflecting the penalty of uninsurability). Flood Re addressed the market-failure component while leaving the actuarial risk unchanged.

A simple back-of-the-envelope calculation illustrates the magnitudes. The median price in a High-risk postcode is approximately £235,000. A 3.4 percent increase implies a price gain of roughly £8,000 per property. With approximately 28,000 High-risk postcodes containing an average of 10 residential properties each, the aggregate wealth transfer is on the order of £2–3 billion. This compares to an annual cross-subsidy cost of approximately £180 million (the Flood Re levy on all household policies). Even discounted over Flood Re's 25-year horizon, the present value of the subsidy cost (£2.5–3 billion at a 5% discount rate) is comparable to the one-time property value gain, suggesting that a substantial portion of the capitalized effect reflects subsidy capitalization rather than pure market-failure correction.

8.2 Implications for Climate Adaptation

As climate change increases the frequency and severity of flooding, the question of how to manage flood risk through insurance markets becomes more pressing. Flood Re demonstrates that government reinsurance can restore market functioning in previously uninsurable areas, but it does so by socializing risk across all policyholders. The long-term sustainability of such schemes depends on whether the transition to risk-reflective pricing can be managed without recreating the market failures that Flood Re was designed to address.

The 2039 transition date looms large. If Flood Re premiums gradually rise toward actuarial rates, the property value effects documented here may reverse, as the insurance-market-failure component of prices is re-imposed. Policymakers should anticipate this and consider complementary measures—such as property-level flood resilience improvements—that reduce the actuarial risk itself rather than merely socializing it.

8.3 External Validity

The UK context has several features that limit direct generalization. First, the UK mortgage market's requirement for flood insurance makes the insurance-liquidity channel particularly strong: without buildings insurance, most lenders refuse to extend mortgages, effectively excluding uninsurable properties from the credit market and restricting sales to cash buyers. In countries where flood insurance is optional—such as Germany, where only 46% of households have natural-hazard coverage—the insurance-liquidity channel may be weaker, and the property-value effects of insurance access correspondingly smaller.

Second, Flood Re's design—with explicit premium caps funded by a cross-subsidy levy and administered through a dedicated reinsurance vehicle—is more generous and more transparent than many international counterparts. The US National Flood Insurance Program (NFIP), by contrast, has accumulated over $20 billion in debt to the Treasury and is moving toward Risk Rating 2.0, which introduces risk-reflective pricing. My results suggest that this transition may depress property values in high-risk areas, an effect that US policymakers have not yet fully quantified.

Third, the UK housing market is exceptionally transparent, with universal mandatory price registration creating the universe-coverage dataset that makes this analysis possible. In many countries, property transaction data are incomplete, delayed, or nonexistent, precluding similar analysis. The Land Registry PPD, combined with the EA's open flood-risk classifications, represents a rare combination of comprehensive outcome data and precise treatment assignment that other researchers could exploit for further study of Flood Re's effects.

Fourth, England's relatively compact geography and centralized flood-risk management create a more homogeneous treatment context than would exist in, say, the United States, where flood risk is managed by multiple federal, state, and local agencies with different standards and capacities. This homogeneity strengthens internal validity but limits generalization to decentralized governance systems.

References

Arrow, Kenneth J, "Uncertainty and the welfare economics of medical care," *American Economic Review*, 1963, *53* (5), 941–973.

Atreya, Ajita, Susana Ferreira, and Warren Kriesel, "Forgetting the flood? An analysis of the flood risk discount over time," *Land Economics*, 2013, *89* (4), 577–596.

Beltrán, Allan, David Maddison, and Robert JR Elliott, "But have you seen the neighbors? A hedonic analysis of flood risk perceptions," *Journal of Urban Economics*, 2018, *106*, 59–71.

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.

Bin, Okmyung and Stephen Polasky, "Effects of flood hazards on property values: evidence before and after Hurricane Floyd," *Journal of Environmental Economics and Management*, 2004, *48* (1), 1002–1013.

Botzen, WJ Wouter and Jeroen CJM van den Bergh, "Bounded rationality, climate risks, and insurance: Is there a market for natural disasters?," *Journal of Economic Behavior & Organization*, 2009, *68* (3-4), 518–535.

Gallagher, Justin, "Learning about an infrequent event: Evidence from flood insurance take-up in the United States," *American Economic Journal: Applied Economics*, 2014, *6* (3), 206–233.

Hallstrom, Daniel G and V Kerry Smith, "Market responses to hurricanes," *Journal of Environmental Economics and Management*, 2005, *50* (3), 541–561.

Hudson, Paul, WJW Botzen, Jeffrey Czajkowski, and Heidi Kreibich, "Moral hazard and adverse selection in the flood insurance market: Insights from Flood Re," *Risk Analysis*, 2020, *40* (12), 2622–2639.

Jaffee, Dwight M and Thomas Russell, "The role of government in the coverage of terrorism risks," *Risk Management and Insurance Review*, 2010, *13* (1), 43–56.

Kousky, Carolyn, "Informing climate adaptation: A review of the economic costs of natural disasters," *Energy Economics*, 2014, *46*, 576–592.

Kunreuther, Howard, "Disaster insurance protection: Public policy lessons," *John Wiley & Sons*, 1978.

Murfin, Justin and Matthew Spiegel, "Is the risk of sea level rise capitalized in residential real estate?," *Review of Financial Studies*, 2020, *33* (3), 1217–1255.

Rambachan, Ashesh and Jonathan Roth, "A more credible approach to parallel trends," *Review of Economic Studies*, 2023, *90* (5), 2555–2591.

Surminski, Swenja and Annegret H Thieken, "Flood Re and the future of flood insurance in the UK," Technical Report, Grantham Research Institute on Climate Change and the Environment 2017.

A. Data Appendix

A.1 Sample Construction Details

The Land Registry PPD is downloaded as annual CSV files for 2010–2025 from the official Price Paid Data service. Each file contains 16 columns; I retain the transaction identifier, price, date of transfer, postcode, property type, old/new flag, tenure, town/city, district, county, and PPD category.

Sample restrictions:

1. Standard transactions only (PPD Category A): excludes repossessions, buy-to-let, non-market transfers.

2. Price range: £10,000–£50,000,000.

3. England only: identified via postcode prefix patterns (excluding Welsh, Scottish, and Northern Irish postcodes).

4. Years 2010–2025 (six pre-treatment years plus nine post-treatment years).

The Environment Agency RoFRS data records counts of residential, non-residential, and unclassified properties at each risk level per postcode. Of 269,085 postcodes in the dataset, 28,112 are High risk, 20,626 Medium, 48,614 Low, and 29,930 Very Low.