Floods and the Frictions of Price Adjustment in Housing Markets

Seonmin Will Heo*

January 8, 2026

Abstract

I investigate how housing markets adjust to realized flood events within the floodplain. Using nearly two decades of transaction-level data from Seoul, South Korea, I document that floods lead to delayed and gradual declines in housing prices: lease prices fall within two years of the event, while sales prices decline only after five years. Transaction volumes contract immediately in the lease market and after a lag in the sales market, indicating that market activity responds quickly even though prices adjust slowly. The empirical findings are inconsistent with pure informational frictions and are better explained by seller-side behavioral frictions: homeowners anchored to nominal purchase prices delay sales, thereby delaying price adjustment. The results show that reference-dependent behavior can shape the way how housing markets incorporate climate risk. (*JEL* Q54, R31, C23, D83)

---

[*] University of California, Santa Barbara, 2127 North Hall, Santa Barbara, CA 93106 (e-mail: sheo@ucsb.edu). I am deeply grateful to Olivier Deschênes, Kelsey Jack, Andrew Plantinga, and Heather Royer for their invaluable guidance and support. I would also like to thank Javier Birchenall, Kenneth Chan, Erik Eyster, Don Fullerton, Koichiro Ito, Alisa Tazhitdinova, Gonzalo Vazquez-Bare, and participants at the UCSB Environmental Reading Group for insightful comments.

1 Introduction

Natural disasters provide rare but powerful opportunities to study how markets process new information about risk. When a flood occurs, it transforms a previously latent hazard into a realized event, forcing households, investors, and policymakers to update their beliefs about the likelihood of future disasters. If an efficient market where the market internalizes risk once it is realized, then housing prices should adjust sharply and permanently upon the first post-disaster transactions, guiding households and investors toward optimal spatial allocation of capital. If, instead, information and behavioral frictions delay matching between home buyers and sellers, then observed price adjustment can be gradual.

This paper examines how floods alter housing prices, transaction volumes, and resale behavior within flood-risk zones using nearly two decades of transaction-level housing sales and lease data from Seoul, South Korea. I find that floods lead to gradual and persistent price declines as well as a relatively immediate decline in transaction volumes in the housing market. I then provide evidence of reference dependence and loss aversion that may result in the delay of price adjustment to avoid nominal loss. If households anchor to previous nominal transaction prices, then they would wait longer until they sell their properties, and transactions that are observed soon after the flood may be positively selected. I find that prices decline earlier for properties that were held for a longer period of time before the sale compared to properties that were recently purchased. To check for positive selection in early transactions, I predict what each property should sell for given its characteristics using a hedonic model and find that early transactions are positively selected relative to predicted fundamentals compared to later transactions.

Empirical research on natural disasters and housing markets documents heterogeneous impacts of flooding on property values. Most studies find that housing prices fall after flood events, but the duration of the effects differ. Some document rapid rebounds within a few years (Bin and Landry, 2013; Beltrán, Maddison and Elliott, 2019; Atreya, Ferreira and Kriesel, 2013), while others find persistent or even permanent discounts as markets gradually internalize long-term risk (Ortega and Taṣpınar, 2018; Ellen and Meltzer, 2024; Gibson and Mullins, 2020). In contrast, when floods severely contract local housing supply, temporary or lasting price increases can occur (Vigdor, 2008; Zivin, Liao and Panassie, 2023). The variation across studies highlights that the path of housing prices post flooding remains an open empirical question.

If housing prices adjust slowly to flooding, then households may continue to overinvest in high-risk locations, misjudge the expected costs of living in flood-prone areas, and underinsure against future losses. Such sluggish incorporation of flood discount generates both efficiency and equity concerns: it delays climate adaptation and shifts the burden of future damages toward less-informed, less-mobile households. For policymakers, understanding the speed of price adjustment is crucial for designing effective disclosure laws, insurance programs, and adaptation incentives that ensure market prices reflect true risk in a timely manner.

This paper contributes to this literature by first showing that floods lead to gradual and persistent price declines, rather than temporary price adjustment. Using a stacked event study design, I analyze the impact of flood events that occurred between 2010 and 2022. I compare properties within the same floodplain that were and were not flooded, examining the market response when a latent risk becomes salient through actual flooding. I find that lease prices begin to decline within two years after a flood, while sales prices fall only after about five years. The overall decline—7 percent for leases and 5 percent for sales—is both gradual and persistent, indicating that flood discount is slowly capitalized into the housing market.

The event-study evidence alone cannot fully explain why housing prices adjust with such delay. Several mechanisms could generate this temporal asymmetry. One possibility is that information about flood diffuses gradually, resulting in sluggish belief updating. Another is that flooded buildings physically deteriorate faster than non-flooded buildings. A third explanation involves behavioral frictions such as reference-dependent preferences, where homeowners anchored to their original purchase price resist selling at a nominal loss.

To examine mechanisms behind the gradual price decline, I analyze the impact of flooding on the number of transactions. I find that the number of transactions declines immediately following a flood in the lease market. In the sales market, the decline in transactions emerges with a two-year delay, which is shorter than the delay in price adjustment. These findings suggest that reductions in market activity precede price declines, serving not only as an early signal of adjustment in both tenure types but also as evidence that households respond quickly to the new information generated by the flood. This rapid response is inconsistent with gradual learning about flood risk on the demand side.

I then present the evidence of reference dependence and loss aversion that is consistent with the empirical results. A growing literature suggests that homeowners anchor their selling decisions to the nominal price at which they originally purchased the property, and are reluctant to realize losses relative to that reference point (Genesove and Mayer, 2001; Andersen et al., 2022). This form of loss aversion can induce sellers to delay listing or accepting offers when market prices fall below their original purchase price after a flood, thereby slowing price adjustment even in the presence of updated information about risk. By modeling this behavior and presenting empirical evidence consistent with bunching around zero nominal gains, I show that reference-dependent preferences can help explain why flood-affected housing markets exhibit slower adjustment. This behavioral rigidity in the asset market is also found in the rental market.

If a flood discount causes a drop in property values, then properties that were recently purchased are more likely to have their market values fallen below their previous nominal purchase prices than those purchased a long time ago. I further examine heterogeneity in price adjustment across duration of ownership. By distinguishing between properties that were purchased recently and those that have been held for longer periods, I show that the delayed decline in sales prices is concentrated among recently purchased properties. In contrast, long-held properties, for which nominal gains are more likely to remain positive, exhibit faster and more complete price adjustment.

If households are reference-dependent, then flood-affected properties that are transacted soon after flooding may have positive selection compared to those transacted in later periods. To assess whether the observed delayed response in average transaction prices reflects genuine slow adjustment or changing composition of transactions, I also conduct event-study regressions on residual prices, defined as the difference between observed transaction prices and model-implied predicted prices from a hedonic valuation model. This analysis provides a test for post-flood selection and the degree to which flood risk is capitalized in assessed property values. The results show that transacted properties have progressively negative residuals over time, indicating positive selection into earlier transactions. These findings suggest that selection may explain part of the early resilience in average prices, observationally appearing as a delay in price adjustment.

Finally, I find that flood-affected properties take longer to be resold or re-leased. While event study regressions show how prices and transaction volumes adjust over time after a shock, they are based on transacted properties and may miss how households self-select into market participation. Using survival analysis and accelerated failure time model, I find that flood-affected properties take 43-58% longer in duration to be re-transacted.

This study builds on and contributes to a few strands of the literature. First, this study contributes to a growing body of literature examining the effects of natural disasters on housing markets. A wide range of empirical studies have documented that natural disasters reduce property values, particularly in areas newly exposed to flood risk or where flood risk becomes more salient following a disaster (Kousky, 2010; Bin and Landry, 2013; Ortega and Taşpınar, 2018; Atreya, Ferreira and Kriesel, 2013; Atreya and Ferreira, 2015; Hirsch and Hahn, 2018; Bakkensen, Ding and Ma, 2019; McCoy and Walsh, 2018; Donovan, Champ and Butry, 2007). These studies have highlighted both the capitalization of physical damage and the informational effects of disasters on housing price dynamics. To my knowledge, no previous papers have analyzed the effect of a natural disaster on both sales and rent prices together. By using the universe of housing transactions that include both sales and lease in a market that is mostly comprised of multi-storied buildings, this study shows that the observed price adjustment post flooding can be gradual over time, highlighting the delayed market responses and behavioral frictions in the housing market.

Second, this paper contributes to the literature on housing price dynamics, particularly studies investigating the persistence, momentum, and heterogeneity in price responses across housing submarkets. Piazzesi and Schneider (2009) highlight the importance of expectations and belief formation in driving boom-bust cycles in housing markets, while Bakkensen and Barrage (2022) emphasize how evolving beliefs about climate risks influence long-run asset prices in exposed areas. The gradual and heterogeneous responses observed in this study suggest that price adjustments to flood shocks are not instantaneous, raising questions about market efficiency in the face of localized environmental risks. By documenting slow price adjustments even in high-information settings, I provide empirical evidence of frictions that may delay the incorporation of new risk information into housing prices.

Finally, this study contributes to the literature on household behavioral responses in housing markets, particularly research on reference dependence and loss aversion. Posited by (Kahneman and Tversky, 1979), reference dependence describes how individuals evaluate outcomes as gains or losses relative to a salient reference point—such as the original purchase price—rather than in absolute terms. Steffen Andersen, Cristian Badarinza, Lu Liu, Julie Marx and Tarun Ramadorai (2022) document that homeowners' price expectations and selling behavior are influenced by past purchase prices, consistent with the framework of reference-dependent preferences. The gradual decline in housing prices following flood events documented in this study is consistent with the view that sellers may resist price reductions when their reservation values are anchored to pre-flood purchase prices. This behavioral friction can help explain why market activity reduces first and then housing prices adjust with a delay in response to environmental shocks.

The remainder of the paper is organized as follows. Section 2 describes data used in the study and provides background information on the housing market and natural disasters in South Korea as well as reference dependence. Section 3 presents the empirical strategy of the study. Section 4 presents the main empirical findings. Section 5 discusses policy implications. Section 6 concludes.

2 Background and Data

2.1 Housing Market in South Korea

The South Korean housing market shares many characteristics with housing markets in other well-established urban economies, including high urban density, diverse housing types, and a mixture of owner-occupied and rental tenures. The predominant housing types include high-rise apartment complexes, officetels (mixed-use buildings combining residential and commercial functions), detached single-family homes, and low-rise rowhouses known as "villas."

Housing tenure is split between owner-occupiers and renters, with the national homeownership rate ranging between 55% and 60% over recent decades. A salient institutional feature of the Korean rental market is the prevalence of the *jeonse* (chônse) contract. In this arrangement, tenants provide a substantial lump-sum deposit—typically 50% to 80% of the property's market value—in lieu of monthly rental payments. The deposit is fully refundable at the end of the lease term, which usually spans two years. The jeonse system became popular since the 1960s when there was a shortage of mortgage supply. Jeonse contracts have functioned as a substitute for formal financial intermediation, with landlords using the deposit for investment or liquidity purposes. Jeonse contracts allowed renters to rent houses at a lower cost than the market price, which in turn allowed landlords to secure seed money for investment.

Because the deposit required for a jeonse tenant to put down up front is still substantial, jeonse tenants often borrow money from banks to pay the deposit, where the deposit is usually taken as collateral. As the interest rates dropped and mortgage loans became more available, the jeonse system has become less popular. Many owners transitioned to following the conventional lease system, requiring monthly rent with a small deposit instead.

However, even with recent development in housing finance and reliable banking system, the jeonse remains a popular lease option, which comprises about 40% of the lease contracts. This is because jeonse serves as a great tool for private financing for landlords and cheaper rent for tenants. The relative prevalence of jeonse and monthly contracts fluctuates with macroeconomic conditions, housing prices, and interest rates. In recent years, hybrid leases that combine moderate deposits with monthly payments have emerged, providing greater flexibility in contractual arrangements.

Currently, all lease types including jeonse, monthly rent, and a combination of the two are available in the housing market. Figure 1 shows proportion of households by contract status in South Korea. More than half of the South Korean population live as tenants on the lease. Real estate platforms now categorize various lease types into three: monthly rent, jeonse, and semi-jeonse. The lease type is semi-jeonse if there is a monthly rent and the required deposit is greater than 12 times the monthly rent. For example, a lease that requires a down payment of \$10,000 with the monthly rent of \$500 is considered a semi-jeonse lease, because the deposit is greater than a year's worth of monthly rent ($12 \times \$500 = \$6,000$).

[Figure 1 about here]

Prospective tenants broadly face three options to live in properties: they can buy a property, they can rent a property through jeonse, or they can rent a property by paying monthly rent. To analyze different lease contract regimes altogether, I use the borrowing interest rate for the jeonse loan (denoted $\tilde{r}$) to convert deposit into monthly rental rate. The combined monthly rental rate is calculated using the following formula:

$$\text{combined rental rate} = \left(\frac{\tilde{r}_t \times \text{deposit}}{12} + \text{monthly rent} \right) \times \frac{100}{\text{CPI}}.$$

This borrowing interest rate for the jeonse loan is lower than the general borrowing rate as well as the official conversion limit. The official conversion limit, set by the government to be 2% higher than the interest rate, is used when the landlord wishes to modify the lease during the contract renewal. For example, if the landlord would like to modify the current jeonse arrangement that involves the deposit of \$500,000 to a semi-jeonse lease that requires a down payment of \$300,000 instead, the highest monthly rent the landlord can ask with the official conversion rate of 6% is $\$200,000 \times 0.06/12 = \$1,000$. The official conversion rates are provided by the Korean Real Estate Board, which calculates the rates by region based on the previous transaction records. I use various conversion rates in my empirical analysis for robustness check.

2.1.1 Housing Liability

When a rental unit (including jeonse) is damaged by a flood, the landlord is responsible for restoring the unit to its original condition and providing temporary housing to the tenant.[1] Tenants are responsible only for minor repairs where each repair costs less than 100 dollars or less than 10% of the monthly rent. For example, when a rented unit needs a new light bulb, then the tenant bears the cost for the light bulb replacement. However, if the lighting system needs a repair, then the landlord bears the cost for the repair.

When a lower-floor unit is flooded, tenants on upper floors do not have to move out. The upper-floor tenants may get stranded if the ground floor is flooded, but upper-floor tenants are not forced to move out due to the flood on the ground floor.

2.2 Housing Transactions Data

The housing transaction data are provided by the Republic of Korea Ministry of Land, Infrastructure and Transport.[2] Real estate transactions are typically intermediated by li-

---

[1] Article 623 under Civil Law states that the lessor is obligated to deliver the property to the lessee and maintain the condition necessary for its use during the duration of the contract.

[2] The data are available in `https://rt.molit.go.kr`.

censed brokers and are systematically recorded through the national Real Transaction Price Disclosure System. They are required to report the details of the contract by law. This administrative infrastructure ensures a high degree of transparency and facilitates empirical analyses of transaction-level data. This universe of housing transactions includes transactions for all building types and land for the entire country. The dataset includes various hedonic housing variables such as size, area, year built, and floor. The dataset includes various housing types such as apartments, rowhouses, and detached houses.

To protect privacy, the dataset is masked at the individual unit level. Transactions within the same building that share identical observable attributes—such as address, floor, and unit size—are not distinguishable in the dataset. As a result, it is not possible to track repeated transactions for a specific unit over time.

To address this limitation, I construct a cell-level panel defined by the combination of building, floor, and area size. Each cell represents a homogeneous set of housing units with the same observable characteristics. I aggregate and merge transaction records at this level, recording average transacted prices, number of transactions, and timing for both sales and lease contracts. This approach preserves meaningful variation in market outcomes across space, floor, and dwelling size, while avoiding the identification issues that arise from masked unit identifiers. The resulting dataset enables comparison of price and volume dynamics within and across buildings, and between floors directly affected and unaffected by flooding events.

This cell-level construction is particularly advantageous in the context of Seoul's high-rise housing market, where many transactions occur within multi-story buildings. Using the floor levels, I can isolate the effects of flooding on properties whose buildings were inundated but whose specific units were not physically damaged. By leveraging this spatial and vertical structure, the analysis can separate the effect of restoration and rebuilding post flooding on prices.

Table 1 provides summary statistics of housing transactions used in the analysis. Although transactions data are available for the entire country, I use only the transactions on properties located in Seoul, because the accurate flood maps by year are available only for Seoul.

Figure 2 shows the number of lease transactions by lease type in Seoul. The number of lease transactions increases for all lease types, including jeonse. Those other forms of lease have become more popular in recent years, jeonse still remains as the most prominent form of lease as of 2023.

[Figure 2 about here]

2.3 Natural Disasters in South Korea

South Korea is subject to a range of natural hazards, among which floods and typhoons are the most frequent and economically damaging. The country's monsoon climate generates concentrated rainfall during summer months, resulting in periodic episodes of urban flooding, flash floods, and landslides. Historical flood events—particularly those affecting low-lying urban areas—have produced substantial economic losses and human casualties. According to the 2020 Natural Disaster Annual Report, heavy precipitation is responsible for more than 80% of the damages in South Korea. The second most damaging type of natural disasters is a typhoon (17%). Other natural disasters such as cold wave, heat wave, and wildfires do occur though their damages have been minor compared to heavy rain and typhoons.

Recent increases in flood frequency and severity have been attributed to rapid urbanization, insufficient drainage infrastructure, and climate change–induced shifts in precipitation patterns. Metropolitan areas such as Seoul and the surrounding Gyeonggi Province are particularly exposed due to dense development along river basins and flood-prone topographies.

Figure 3 shows the flood maps in recent years. During the sample period between 2006 and 2023, severe flood damages occurred in 2010, 2011, and 2022, affecting hundreds of thousands of people nationwide overall. Unlike most of the other studies that were discussed in this paper, almost all of flood damages were caused by intense rainfall that exceeded drainage capacity (pluvial flooding), which different in characteristics from coastal flooding or fluvial flooding, where the water-level rise in river channels exceeds the bank height.

This distinction is important for a few reasons. First, the risk of pluvial flooding is less salient than the risk of coastal flooding. This lower salience coupled with low frequency makes flooding in South Korea more random without the knowledge of a floodplain map. Second, pluvial flooding suffers less from the concern that flood risk is correlated with other amenities than coastal flooding or fluvial flooding. For example, the risk of coastal flooding or fluvial flooding can be mitigated if properties are distant from the ocean or the rivers. Meanwhile, major factors that increase the risk of pluvial flooding such as land gradient and land permeability, which are less salient if the flood did not occur in recent years.

[Figure 3 about here]

2.4 Flood Insurance

South Korea maintains a formal insurance program on natural disasters legislated under the Storm and Flood Insurance Act, designed to cover property losses caused by natural disasters such as typhoons, floods, and earthquakes. The program is overseen by the Ministry of the Interior and Safety (MOIS), with private insurers authorized to underwrite the policies. Storm and flood insurance policies are annual in duration and are sold either as stand-alone products or packaged with other non-life insurance coverage. For simplicity of exposition, I refer to these policies collectively as “flood insurance” throughout the paper. There are no insurance policies that cover only a specific natural disaster. All property owners are eligible but not required to purchase flood insurance. For full coverage, property owners must initiate and maintain the flood insurance contract that renews on an annual basis. To promote adoption, the government supports the program through premium subsidies covering between 55% and 100%.

As of 2024, out of 4.3 million households in Seoul, only 44,275 households purchased the flood insurance, amounting to the take-up rate of about 1%. The average take-up rate between 2017 and 2024 was 0.68%. When restricting to properties located within the "disaster-risk zone" officially designated by MOIS and making the conservative assumption that all insured households are located within the disaster-risk zone, the flood insurance adoption rate still remains below 10% (Table A.1).

Despite generous subsidies and public underwriting arrangements, the take-up of flood insurance in South Korea remains strikingly low. One reason for the low take-up is the role of post-disaster public relief, which can crowd out demand for private insurance. When areas are declared disaster zones, affected households are eligible for government relief payments covering roughly 30–35 percent of uninsured damages. Because public aid and insurance payouts are mutually exclusive, the marginal benefit of purchasing private coverage appears limited. For many households, the expected additional payout from insurance is small relative to the guaranteed baseline of government support, reducing incentives to enroll in the program. This mechanism echoes charity hazard documented in the disaster insurance literature, where generous ex-post transfers weaken ex-ante insurance demand.

A second reason is low public awareness and limited market visibility of flood insurance products, compounded by adverse selection. Surveys suggest that many households remain unaware of the availability of flood coverage, while private insurers have weak incentives to actively promote or market the product given its low profitability. These factors contribute to a thin market in which take-up is concentrated among those already at high risk, reinforcing adverse selection concerns. Although MOIS has attempted to raise awareness and expand coverage through outreach campaigns, empirical studies continue to find that flood insurance adoption is suppressed by insufficient salience of flood risk and the perception that existing government subsidies provide a sufficient safety net (Park and Yeo, 2013; Yeom et al., 2019).

2.5 Population and Housing Census

To understand the characteristics of South Korean households and to provide data required for policymaking, the South Korean government conducts the Population and Housing

Census surveys every five years. The data are available since 1975 and include housing and household variables such as type of living quarters, total floor area for residence, land area, year when the property was constructed, number of rooms, floor, current occupancy type, and household type.

2.6 Reference Dependence in Housing Markets

Kahneman and Tversky (1979) posit that agents evaluate outcomes relative to a reference point, typically the status quo or a salient past outcome. In the housing context, this implies that homeowners often assess prospective sale prices against their original purchase prices. Selling below this price is perceived as a loss, which can induce homeowners to set reservation prices above market-clearing levels and delay transactions. Suppose that the reference point of the seller is denoted by R and the sales price is denoted by P. Let $x = P - R$ be the outcome relative to the reference. Then sellers with reference-dependent preferences would derive utility from the following value function $v(x)$:

$$v(x) = \begin{cases} x^{\alpha} & \text{if } x \geq 0 \\ -\lambda(-x)^{\alpha} & \text{if } x < 0 \end{cases}$$

where $0 < \alpha < 1$ reflects diminishing sensitivity to gains/losses relative to the reference, and $\lambda > 1$ captures the degree of loss aversion. In this study, I use the previous nominal purchase price as the reference point, R.

A substantial literature documents the empirical relevance of this behavior. Genesove and Mayer (2001) provide early evidence of the reference-dependent behavior in the Boston condominium market, showing that loss-averse sellers kept properties on the market longer and accepted lower probability of sale rather than realizing a loss. More recent evidence from Andersen et al. (2022) corroborates these findings across multiple housing markets, emphasizing the role of reference-dependent preferences in generating asymmetric price adjustments.

In this study, I examine the reference-dependent responses to flooding in the housing market to provide more ways to understand why asset prices appear more rigid than rental prices within the same housing market. Although it may appear surprising that sales prices respond with a substantial delay, this pattern can be rationalized from the perspective of homeowners. Because sellers tend to care about their original nominal purchase price for various reasons, homeowners may choose to postpone selling and wait for market conditions to improve rather than realize a nominal loss. In this regard, the delayed adjustment of sales prices does not reflect irrationality, but rather the interaction of reference-dependent preferences and loss aversion with long investment horizons.

If households are indeed reference-dependent when selling their properties, then one would expect to observe bunching at or just above the nominal gains of zero, accompanied by a sharp drop in frequency for transactions just below that reference point. This bunching pattern—particularly at zero nominal gain—is consistent with the notion that sellers are reluctant to realize nominal losses, even when doing so would be optimal under standard economic assumptions.

However, establishing the presence of reference-dependent behavior empirically requires a benchmark for what the distribution of gains would look like in the absence of such frictions. To this end, I estimate potential gains using a hedonic pricing model, which predicts sale prices based on observable housing attributes. A detailed description of the model specification and implementation is provided in Section 3. The nonparametric evidence of reference dependence and loss aversion is provided in Section 4.

2.7 Assessed Values of Housing Properties

A key input to the hedonic pricing model for constructing the counterfactual gains based on housing attributes is the official assessed property value. The assessed property and land values are published annually by the Ministry of Land, Infrastructure and Transport

(MoLIT) in South Korea.[3] These assessed values are used primarily for taxation purposes but are also used as baseline land and property values for new housing projects, health insurance contribution calculations, and eligibility for social welfare programs. Assessed values reflect estimated market prices as of January 1 of the same year.

The appraisal process begins with direct appraisal by certified real estate appraisers commissioned by MoLIT. These appraisers use a comparable sales approach to estimate a property's market value: they compare it to similar properties that have recently sold and account for unit-level characteristics including floor area, building age, floor level, and orientation, and locational attributes including transport access, commercial density, and school zones. Appraised values are subject to internal review by MoLIT and external review by regional appraisal review committees comprising professional appraisers and academic experts. A formal public inspection and objection period is provided prior to finalization. During this period, property owners may submit feedback or contest preliminary values. These objections are reviewed by administrative staff or appraisal committees, and values may be revised accordingly. Final confirmed prices are then published in relevant government websites.

Assessed values reflect public valuations of housing properties at the parcel level. Because they are systematically collected and standardized across the country, these values provide a consistent benchmark that captures spatial and temporal variation in property valuation, making them a useful component in predicting market prices in hedonic regressions.

---

[3] MoLIT is responsible for the public valuation of multi-unit housing, while local governments (specifically, mayors and district heads) assess detached housing, following detailed technical guidelines issued by MoLIT. This system is legally grounded in the Real Estate Price Disclosure Act, which mandates annual disclosure of assessed values for tax and administrative purposes.

3 Empirical Strategy

3.1 Estimating the Impact of Flood on Housing

To estimate the dynamic effect of flooding on housing prices, I employ a stacked event study framework that compares housing price changes before and after flood events. This strategy accommodates treatment effect heterogeneity over time and avoids biases inherent in traditional two-way fixed effects estimators, particularly when treatment timing is staggered. I also implement a doubly robust estimator to further mitigate potential biases due to observable differences between flooded and non-flooded properties. The estimator is doubly robust in the sense that the estimator is consistent for the average treatment effect on the treated if either the propensity score model or the outcome regression model is correctly specified. This method weights observations by the inverse probability of treatment conditional on pre-determined covariates such as the year the building was constructed, the unit's area size, and its floor level. This approach offers greater robustness against model misspecification and helps ensure more balanced comparisons between treated and control units. I also explore effect heterogeneity by floor level and contract type.

In the standard hedonic model, a residential unit i is a package with its J different attributes $(z_1, z_2, \ldots, z_J)$ where its attributes cannot be unbundled. The housing price is then denoted as $P_i = P(z_1, z_2, \ldots, z_J)$. In this model where both buyers and sellers of houses act as utility-maximizing agents and a sufficiently large number of differentiated products are available, the price function $P(\cdot)$ reflects the market equilibrium prices. Because the price function reflects points where the bid curves of buyers and the offer curves of sellers meet, I can use the hedonic price function to obtain willingness to pay for each attribute: the marginal price of an attribute at a given point along the hedonic price function is equal to the marginal willingness to pay for the set of buyers located on that point.

However, because environmental amenities may be correlated with unobserved attributes, estimates from standard hedonic models are vulnerable to omitted variable bias. To overcome the issue of endogeneity present in cross-sectional data, prior studies have applied difference-in-differences to the standard hedonic model. In this paper, I employ a stacked difference-in-differences hedonic price model restricting my sample to repeat transactions to address the potential concern for omitted variable bias.

I isolate the impact of flooding on housing prices using the following difference-in-differences regression equation:

$$\log P_{it} = \sum_{j \neq -1} \beta_j D_{i,t-j} + \alpha_i + \delta_t + \varepsilon_{it}, \tag{1}$$

where P_{it} indicates the sales or lease price of cell i at time t, α_i are cell-level fixed effects, δ_y are year fixed effects, and $D_{i,t-j}$ is an indicator that equals one if the cell i is affected at time $t - j$ relative to the flood year. I include the cell-level fixed effects to eliminate the effect driven by the change in composition of transacted properties. I set years to begin in July and end in June, because most floods occur between July and August. For the lease price, I use the combined rent rate where the deposit is converted to the monthly rate using the borrowing interest rate for the jeonse loan.

Recall from Section 2.2 that the dataset is at the cell level i, which is defined by the combination building b, floor f, and area size a. The housing transaction dataset is masked at the individual unit level to protect privacy. Each cell represents a homogeneous set of housing units with the same observable characteristics: building, floor, and area size. Although individual units within a cell cannot be distinguished, this data structure preserves the key treatment variation in my setting, which is at the building and floor level.

A growing body of literature has found that two-way fixed effects regressions may not provide consistent estimates under treatment effect heterogeneity or dynamic effects (De Chaisemartin and d'Haultfoeuille, 2020; Sun and Abraham, 2021; Callaway and Sant'Anna, 2021). I study the dynamics of treatment effects using the estimator proposed in Callaway and Sant'Anna (2021) to ensure that the treated group and the control group are well balanced in covariates with doubly robust inverse probability weighting. Figure 4 presents the event study plots of sales and lease price dynamics. Table 2 presents the average treatment effect coefficients from stacked difference-in-differences regressions.

To better understand the mechanisms driving housing price responses to flooding, I examine how floods affect transaction volumes. I estimate a set of stacked difference-in-differences event-study regressions using the number of transactions as the outcome variable as follows:

$$\log Q_{it} = \sum_{j \neq -1} \beta_j D_{i,t-j} + \alpha_i + \delta_t + \varepsilon_{it}, \tag{2}$$

where Q_{it} indicates the number of sales or lease transaction counts of housing cell i at time t, α_i are cell-level fixed effects, δ_y are year fixed effects, and $D_{i,t-j}$ is an indicator that equals one if the cell i is affected at time $t - j$ relative to the flood year. I construct a balanced panel of housing properties that were transacted at least once during the sample period. I include the cell-level fixed effects to eliminate the effect driven by the change in composition of transacted properties. I set years to begin in July and end in June, because most floods occur between July and August. Aside from the panel structure, the specification is structurally identical to the regressions on housing price, including the use of cell-level and year fixed effects and the application of doubly robust inverse probability weighting based on observed covariates. Figure 5 presents the event study plots of sales and lease transaction volumes. Tables 2 and 3 present the aggregated values of all post treatment effects from stacked event-study regressions.

Previous studies find large reduction on prices for properties located within the floodplain compared to properties located outside the floodplain. To make the comparison between flooded and non-flooded properties more credible, I restrict the control group to only the properties located within the floodplain that are never or not yet flooded. This restriction ensures that treated and control properties have similar baseline flood risk exposure, even if only a subset were directly affected during the flood event.

The main identifying assumption is that, conditional on the set of covariates including floor level, property size, and the year built, the average outcome for the treated and the control group would have evolved in parallel in the absence of treatment. Formally, setting up the potential outcomes framework, let $Y_{it}(0)$ denote the untreated potential outcome of cell i at time t if they were to not receive treatment across all time periods. Let X_i be a set of covariates on cell i, which are the year the building was constructed, the unit's area size, and its floor level. Then the identifying assumption can be written as

$$\mathbb{E}[Y_t(0) - Y_{t-1}(0) \mid D = 1, X] = \mathbb{E}[Y_t(0) - Y_{t-1}(0) \mid D = 0, X].$$

This assumption is known as the conditional parallel trends assumption. In addition, I assume that the treatment is irreversible; that is, $D_{t-1} = 1 \implies D_t = 1$. I also assume no anticipation of treatment. Formally, denote G as the year t in which a property first becomes treated. Let G_g be an indicator variable that is equal to 1 if a property is first treated in year g. Then the assumption of no anticipation can be expressed as

$$\mathbb{E}[Y_t(g)|G_g = 1, X] = \mathbb{E}[Y_t(0)|G_g = 1, X]$$

for all treated groups g and all pre-treatment years $t < g$. In this paper, this assumption means that whether the property will be flooded is not known to households prior to the realization of the flood event.

After I estimate stacked event study regressions for sales and leases, I compare the dynamic effects in the sales market with those in the lease market. Theoretically, as Poterba (1984) states, the value of a property is equal to the present value of the stream of future net service flow. Let R_t be the rent at time t and c_t be the the carrying cost at time t. Then the price P_t equals the present value of future net service flow $R_t - c_t$ discounted at the homeowner's real after-tax interest rate r_t:

$$P_t = \mathbb{E}_t \left[\sum_{s \geq 1} \frac{R_{t+s} - c_{t+s}}{\prod_{j=1}^{s}(1 + r_{t+j})} \right].$$

Under this frictionless benchmark, if a negative shock to R occurs at $t = k$, then the price would immediate drop at $t = k$ ($\Delta P_k - P_{k-1} < 0$), followed by a flat post-event path ($P_t - P_{t-1} = 0$ for $t > k$). In addition, in the absence of large, sudden movements in the carrying cost, sales prices and rents would move together for the same housing asset.

To formally test whether the post-flood effects are different between the sales and the lease markets, I exploit the influence function datasets of the stacked difference-in-differences event-study estimators for sales and leases. Let $\hat{\beta}^S_\tau$ and $\hat{\beta}^L_\tau$ denote the event-time coefficients (relative to $\tau = -1$) from the two event-study regression specifications. For each τ, I compute the difference $\hat{\Delta}_\tau = \hat{\beta}^S_\tau - \hat{\beta}^L_\tau$ and conduct clustered bootstrap at the unit level to obtain standard errors. I conduct pointwise tests and joint Wald tests to examine whether post-flood responses significantly differ in the early post-flood years.

I assume that buyers and sellers are fully informed. First, the information on all the available apartments for sale and lease has been available to the public through various online platforms and offline real estate agencies for more than a decade. Although the addresses of properties are masked at the floor level for multi-unit buildings and the street level for individual, detached houses, all the past transactions of residential properties are available to the public since 2006. Second, most transactions take place through real estate agencies, who provide tours for prospective buyers and renters. Finally, the housing market in Seoul is significantly large. The number of apartments listed for sale is around 50,000 on average, and the listings vary in size and transaction types that provide continuously varying amounts of attributes.

3.2 Estimating Potential Gains Using Hedonic Pricing Models

Establishing the presence of reference-dependent preferences empirically requires a benchmark for what the distribution of gains would look like in the absence of such preferences. Although it is impossible to observe such a counterfactual, I estimate the counterfactual distribution of gains using a hedonic pricing model based on assessed property values and observable housing characteristics. This preexisting empirical approach, commonly used to estimate the value of amenities such as environmental factors or other non-market attributes, allows us to generate a distribution of predicted gains that serves as a reference point against which observed sales can be compared. I estimate the hedonic price function using two complementary approaches: a traditional parametric OLS specification and a non-parametric gradient-boosted decision tree to capture complex non-linearities.

3.2.1 Baseline OLS Specification

The official assessed values already incorporate property-specific characteristics, but they may be subject to systematic biases such as lagged updates or political considerations, which can affect their accuracy as a measure of market value. By adding unit size, floor level, year built, and district and time fixed effects, I account for these characteristics directly in the regression, mitigating potential bias in the estimated gains. Specifically, I estimate the following hedonic pricing model:

$$\log\big(P_{idmy}\big) = \alpha \log\big(A_{idy}\big) + \theta_m + \theta_y + \theta_d + \gamma X_i, \tag{3}$$

where P_{idmy} is the sales price of property i in district d in month m and year y and A_{idy} is the assessed value of property i that is located in district d and assessed in year y. I include month, year, and district fixed effects as well as a set of time-invariant housing covariates X_i, which include unit size, floor level, and year built. I use only the properties within the floodplain. This specification produces model-implied hedonic prices for each observed transaction, which serve as estimates of the property's market value based on its observable characteristics.

I use predicted values to measure potential gains $\widehat{G} = \widehat{\log P} - \log R$, where the predicted property value $\widehat{P}$ is the housing price that a rational, unconstrained seller might expect under normal market conditions. I then use these potential gains to construct a counterfactual distribution of binned frequencies across nominal gains. By comparing the observed distribution of realized gains $G = \log P - \log R$ to this counterfactual, I can detect bunching patterns that are consistent with reference-dependent behavior.

3.2.2 Non-Parametric Estimation

While the baseline OLS specification provides interpretable coefficients, it relies on functional form assumptions that may not fully capture the complex pricing dynamics of Seoul's housing market. To address this, I employ a gradient-boosted decision tree (LightGBM) model as a robust alternative for constructing counterfactual prices. Unlike the parametric model, this machine learning approach approximates the hedonic price function $P(z)$ non-parametrically, allowing the data to determine the optimal functional form and interaction terms. I train the model on the same set of covariates used in Equation 3, including assessed values, housing attributes, and time and district fixed effects. Hyperparameters are tuned via 5-fold cross-validation to minimize out-of-sample Root Mean Squared Error (RMSE), ensuring that the generated counterfactual prices reflect fundamental market values rather than overfitting to noise.

4 Results

4.1 Post-flood Housing Price Dynamics

I begin by estimating the post-flood dynamics in housing prices using the event-study specification in Equation (1). Figure 4 presents event-study estimates of the impact of flooding on housing prices for both the sales and lease markets. The figure plots coefficients from the stacked event-study specification described in Section 3, separately for sales prices and lease values, with the year prior to the flood normalized to zero.

The results reveal a gradual and persistent decline in housing prices following flood exposure. Lease prices begin to decline soon after the flood event and continue to fall steadily over the subsequent years. This finding suggests that renters and landlords internalize the new flood information relatively quickly compared to the sales market. Sales prices, by contrast, exhibit a notably delayed response. For the first five years following the flood, sales prices in affected areas remain statistically indistinguishable from pre-flood levels. It is only after five years that a significant decline in transaction prices begins to emerge. This delayed response is quite surprising, even accounting for the fact that transactions of housing properties occur less frequently and thus respond to risk realization more slowly than the rental market.

To formally test whether the price adjustments in the sales and lease markets differ, I use the influence-function representation of the stacked difference-in-differences event-study estimators for sales and for leases to bootstrap standard errors. I estimate a pooled event-study specification that is equivalent to a pooled event-study regression with interactions by taking the difference in event-study coefficients and bootstrapping standard errors to account for the combined sampling variation. This approach allows me to test directly whether event-study coefficients differ between the two markets, rather than inferring from separate event-study regressions.

Figure A.6 shows the differences in event-study estimates between the sales and the lease markets. The results show that the timing of adjustment differs significantly across tenure types: lease prices decline quickly and gradually over time, while sales prices remain flat in the short run and only begin to fall several years after the flood. I additionally run a joint Wald test across early post-event periods to show that the decline in dynamic effects for the early post-event periods is slower in the sales market than in the lease market. Pointwise and joint Wald test results are provided in Table A.1.

In a frictionless asset market, prices would immediately adjust to absorb a shock. However, the observed gradual and delayed price response to a negative shock can be due to either gradual belief updating or search frictions delaying the match between sellers and buyers. To examine whether other margins in the housing market adjust to floods before prices adjust, I examine how flooding affects the number of housing transactions. Using the stacked event-study design with normalized transaction counts as the outcome in Equation (2), I find that sales transaction volumes decline by approximately 8 percent following a flood, while lease transaction volumes decline by 28 percent. The normalized transaction volume is obtained by dividing the coefficients by the control mean to be interpreted as a percentage of the control mean (Chen and Roth, 2024). These magnitudes show that floods meaningfully dampen market activity in both asset and rental markets. The results are shown in Table 3.

Figure 5 presents event-study estimates of the impact of flooding on housing transaction volumes for the sales and lease markets, from the stacked event-study specification described in Section 3. The event-study plots highlight important differences in timing between the two markets. First, lease transaction volumes decline immediately after a flood, whereas sales transactions volumes start to decline two years after the flood event, indicating a sluggish adjustment in the asset market. Second, the decline in housing transactions appears to anticipate the subsequent price decline in both markets. The sales transactions decline two years after the flood, whereas the sales prices drop five years after the flood. The lease transactions decline immediately after a flood, whereas lease prices begin to drop only about two years later. These findings suggest that (1) reduced market liquidity may serve as an early signal of the longer-run adjustment in prices, and (2) the housing market begins to respond to flood shocks fairly quickly, but the initial adjustment manifests through a slowdown in transactions rather than through prices.

Figures 4 and 5 show the contrasting dynamics of prices and transaction volumes between the sales and lease markets. These results motivate a closer examination of the institutional and behavioral mechanisms driving the observed differences.

First, the evidence of price declines in both markets indicates that flood discount is capitalized across tenure types. Prices ultimately fall in both the sales and lease markets, which suggest that households in both segments internalize flood discount into their valuations rather than perceiving floods as transitory shocks. This empirical evidence provides an important contribution to the existing literature, where findings have been mixed and it remains unclear whether the effects of flooding on housing prices are temporary or persistent.

Second, the delayed response in the sales market compared to the lease market implies that buyers and renters operate with different information sets. Renters, who engage in more frequent transactions and face lower search and moving costs, are more responsive to updated flood information. Leasing contracts can be renegotiated on an annual basis, making rental values more sensitive to contemporaneous shifts in risk salience through flooding. In contrast, the sales market is characterized by higher transaction costs, more complex financing arrangements, and substantially longer decision cycles. Purchases typically involve mortgage approvals and large upfront costs, all of which can slow the convergence of buyer and seller expectations and delay price adjustment.

Third, the observed lag between falling transaction volumes and subsequent price reductions points to frictions beyond gradual updating of belief on flood discount. Immediately after a flood, both buyers and sellers may face uncertainty about the appropriate price discount for affected properties, resulting in fewer transactions as the market gradually learns the new post-flood equilibrium. However, even as market participants update their expectations, prices remain slow to adjust.

In this context, behavioral frictions such as reference dependence offer a compelling explanation for the delayed price response. Homeowners who anchor to their original nominal purchase prices may be reluctant to realize a loss, which can cause many potential sellers to withdraw from the market rather than accept lower offers. This withdrawal can depress transaction volumes without immediately reducing prices. As nominal housing prices in the broader market rise over time, flood-affected property values that were below the previous nominal purchase prices also increase, allowing some owners to sell without incurring a nominal loss. The resumption of such sales can produce the observed delayed decline in prices. In the next section, I provide empirical evidence consistent with reference-dependent preferences and seller-side frictions.

4.2 Evidence of Reference Dependence

The preceding analyses on price and transaction volume show that flood-affected properties experience both a gradual price decline and a reduction in transaction volume. To explain the delay of price response to reduced transaction volume, I examine reference dependence behaviors among homeowners. In particular, I examine whether sellers are less likely to transact when the current market value of their property falls below its original purchase price. I also examine whether landlords are less likely to transact when the current rental rate of their property falls below its previous rental rate.

Panel A of Figure 6 plots binned frequencies of transactions across realized nominal gains $G = \log P - \log R$, where P is the transacted price and R is the previous nominal purchase price. Each dot represents the observed frequency of sales within each 1 percentage-point bin of realized gains. The plot also shows binned frequencies of transactions across counterfactual gains $\widehat{G} = \widehat{\log P} - \log R$ that would result if households were to sell their properties at their property values implied by the hedonic pricing model. In other words, the hedonic pricing model is used to predict the estimated property value $\widehat{\log P}$ based on the assessed value and other housing attributes. Panel B plots excess mass of observed sales relative to the level of the counterfactual distribution. Figure A.11 shows several counterfactual gains computed using different hedonic pricing models.

I create the same plots for the lease market, shown in Figure 7. I use the fixed deposit-to-rent conversion rate of 4.5% to combine deposit and annual rent. I do not use a variable conversion rate to capture nominal gains from the previous rental rate. The identical patterns emerge with different fixed conversion rates. Unfortunately, the assessed *rental* values of the properties are not available, and the assessed market values of the properties are used instead to obtain counterfactual gains. Figure A.12 shows several counterfactual gains computed using different hedonic pricing models, and these gains exhibit similar patterns.

Table 4 shows the regression results of hedonic price models on sales prices. The counterfactual gains shown in Figure 6 are computed using the predicted values from the hedonic pricing model (Column 1). I obtain an R^2 of 0.97 and similarly high fits for other specifications. Although the estimates cannot be interpreted causally, the results indicate that properties located within the floodplain sell at a discount of roughly 0.2 percent. By contrast, the experience of a single flooding event is associated with a more pronounced discount of about 0.9 to 1.3 percent relative to otherwise similar properties. Regression results without the assessed value are shown in Column 5. The results imply that much of flood risk is capitalized into the assessed market value.

Table 5 shows the regression results of hedonic price models on lease prices. The counterfactual gains shown in Figure 7 are computed using the predicted values from the hedonic pricing model (Column 1). I obtain an R^2 of 0.813 and similar fits for other specifications. The results are to be read with caution, as the assessed values are not assessed rental rates but assessed market values of the properties.

Figure 6 reveals several interesting features. First, the distribution of realized gains shows a pronounced jump at the reference point ($G = 0$), relative to the counterfactual gains. The jump at precisely $G = 0$ suggests that households anchor strongly to their original purchase price, using it as the reference point when selling their properties. Second, there is a noticeable dip immediately to the left of $G = 0$, implying that homeowners are reluctant to transact at even modest nominal losses. The combination of a spike at zero nominal gains and a deficit just to the left of $G = 0$ provides direct evidence of reference dependence and loss aversion in the post-disaster housing market. These patterns indicate that sellers anchor their decisions to the original *nominal* purchase price rather than to inflation-adjusted or real gains.

The features above can also be found in the lease market, shown in Figure 7. As expected, bunching at $G = 0$ is even more pronounced in the lease market than the sales market, most likely because lease contracts are frequently renewed (often annually). The combination of a big spike at zero nominal gains and a dip just to the left of $G = 0$ again shows that reference dependence and loss aversion are present in the lease market. These patterns indicate that landlords also anchor their decisions to the previous rental rate rather than to inflation-adjusted or real gains.

Importantly, the contrast between realized and counterfactual gains demonstrates that these patterns cannot be explained purely by observable housing characteristics. The counterfactual distribution, derived from hedonic predictions, does not exhibit the same sharp discontinuity at zero, implying that the bunching is behavioral rather than mechanical. This echoes findings in Andersen et al. (2022), but in the context of flood-affected housing markets, it carries particular significance. It implies that reference dependence may slow the incorporation of a negative shock into prices, as homeowners delay selling their properties when market values fall below their previous purchase prices.

To ensure that the observed bunching is not an artifact of functional form misspecification in the standard hedonic regression model, I construct an alternative set of counterfactuals using a gradient-boosted decision tree (LightGBM) model. Unlike the OLS baseline, this non-parametric approach automatically learns complex interactions and non-linearities without manual restriction.

As detailed in Appendix B, the machine learning model outperforms the standard hedonic regression in out-of-sample predictive accuracy, increasing the R^2 by 2.3 percentage point. Despite this improvement in precision, the counterfactual distribution generated by the machine learning model exhibits the same smooth shape as the OLS counterfactual and does not replicate the sharp discontinuity at zero nominal gains. This robustness confirms that the bunching at the reference point is driven by behavioral frictions on the seller side rather than by unobserved structural characteristics or modeling errors.

To validate that the non-parametric machine learning model captures meaningful economic relationships rather than spurious correlations, I examine partial dependence plots (PDPs) for key housing attributes (see Figure 8). These plots visualize the marginal effect of each housing attribute on the predicted log transaction price, holding all other variables constant.

The results confirm that the model captures standard economic mechanisms without manual restriction. First, the relationship between logged assessed value and predicted price is approximately linear with a slope near one (a 45-degree line). However, the predicted market prices consistently lie slightly above the assessed values, indicating that while official assessments accurately rank properties by quality, they systematically undervalue market levels. Second, the floor area plot in Panel B exhibits a classic concave shape, reflecting diminishing marginal utility of space: each additional square meter adds value but at a decreasing rate. Third, the plot for year built shows no consistent trend, suggesting that once building age and location are accounted for, the specific vintage of construction carries little independent explanatory power.

Most notably, the plot for building age reveals a striking non-linear dynamic specific to the Seoul housing market. Normalized relative to a brand-new unit (age 0), the price premium increases with building age. Even though they depreciate physically, older units reflect the high option value of reconstruction inherent in aging complexes. Overall, the PDPs confirm that the gradient boosting model successfully identifies the standard economic mechanisms without requiring the manual specification of quadratic or interaction terms.

If households anchor on the previous nominal transaction price, they may hold their properties longer until the properties recover their pre-flood nominal prices after a flood. I examine whether the prices of long-held properties adjust differently from those of recently purchased properties. If the frictions stem solely from the buyer's side, then properties that were held for longer periods should be no different from those that were recently purchased. If the seller-side frictions such as the reference-dependent behavior delay the price response, then properties that were purchased a long time ago may respond more quickly to a negative shock than properties that were recently purchased.

Figure 9 shows the the dynamic price effects for properties that were previously transacted more than 3 years ago and for properties that were previously transacted in less than 3 years. The figure shows that prices of long-held properties decline gradually with a short delay of two years compared to those of recently purchased properties that do not decline for more than five years. This result implies that seller-side frictions delay the price adjustment in the housing market.

Finally, the delayed response in observed sales prices may partly reflect post-flood selection rather than slow price adjustment. If owners of properties with relatively favorable idiosyncratic characteristics are more likely to sell shortly after a flood, the average observed transaction price will initially overstate the true market-wide impact, and the subsequent decline will reflect compositional shifts as less desirable properties re-enter the market. This positive selection in earlier post-event periods may manifest as a delayed price response. To assess this possibility, I first compute the residual as the following:

$$\varepsilon_{it} = \log P_{it} - \widehat{\log P}_{it},$$

where $\widehat{\log P}_{it}$ is the potential value of the cell i at the time of transaction t based on the observed housing attributes of the cell, computed using the hedonic pricing model discussed in Section 3.2. I then run the stacked event-study regression in Equation 1 with the residual ε_{it} as the outcome variable. If the delayed price response is due to positive selection in early post-event periods, then the residuals in early post-event periods would be positive compared to the residuals in late post-event periods.

Figure 10 shows the event-study regression results on residual housing prices. The results show that the residual housing price declines over time after flooding, which is consistent with positive selection in early periods post flood. These findings suggest that selection may explain part of the short-term delay in price response and that the eventual price decline appears to be genuine revaluation of flood discount.

These findings from the stacked event-study regressions show how observed prices and transaction volumes adjust after a flood event. However, because these regression results are based on the transacted properties, they do not show how a negative shock affects market participation of the affected properties. To examine whether flood-affected properties do take longer time to be re-transacted, I employ survival analysis in the next section.

4.3 Probability of Resale After Flood

While event study regressions provide insight into the price dynamics of transacted properties, they do not capture the full range of behavioral responses to flooding. In particular, these regressions condition on transactions that have already occurred and thus may miss selection effects on market participation. Flooding may alter the perceived desirability of properties in affected areas and make them less attractive to prospective buyers or tenants. As a result, flood-affected properties could be less likely to sell. This reduced marketability could manifest as longer time-on-market, postponed listing decisions, or outright withdrawal from the market.

To examine these margins of adjustment, I next examine how flooding affects the timing of transactions. Specifically, I employ survival analysis using Kaplan-Meier estimators to visualize the probability of resale across flooded and non-flooded properties, and I test differences in survival time using an accelerated failure time (AFT) model. I define the sample based on properties that experienced at least one transaction prior to the flood. The Kaplan-Meier survival curves indicate the fraction of properties that have not been resold as a function of time since the flood. Let T denote the time (in days) until a property is transacted following the flood event, and let $R_i = 1$ if property i is resold and $R_i = 0$ if it is not resold by the end of the sample period. For each group—flooded ($F = 1$) and non-flooded ($F = 0$) properties— I estimate the survival function, defined as

$$S(t) = \Pr(T > t),$$

which gives the probability that a property remains unsold beyond time t. The Kaplan–Meier estimator is a nonparametric maximum likelihood estimate of $S(t)$, constructed as the product of observed survival probabilities over discrete time intervals.

Figure 11 shows the fraction of properties that have been re-transacted cumulatively as a function of time since the flood. The figure compares flood-affected properties to non-flooded properties after the 2010 flood for sales and the 2011 flood for leases, restricting the sample to those located within the floodplain to ensure comparability. The figure indicates that flood-affected properties exhibit a significantly lower hazard of transaction over time; that is, they are less likely to be transacted in the years following a flood, compared to non-flooded properties. In the sales market, this manifests as longer holding periods and a slower pace of turnover, suggesting that homeowners may postpone selling in the aftermath of a disaster. This delay could reflect both supply-side and demand-side constraints: homeowners may wait for prices to recover, while buyers may avoid recently flooded properties due to perceived risk or uncertainty about future flooding.

The pattern is even more pronounced in the lease market (Panel B). Flooded rental units take longer to be re-leased, and the cumulative probability of a lease transaction remains persistently lower than that of comparable non-flooded units over the entire post-flood period. This suggests heightened vacancy rates and reduced renter demand for previously flooded units, consistent with the earlier finding of declining lease prices.

The estimated coefficients from the AFT model indicate that flood-affected properties take 43.1% longer in duration to be resold compared to properties unaffected by flood. In the rental market, flood-affected properties take 57.9% longer in duration to be re-leased.

The results from the AFT regression are provided in Table 6. The larger coefficient in the lease market suggests that flooding has a more immediate and pronounced effect on the liquidity of rental units, likely reflecting greater mobility among tenants and a heightened sensitivity to recent flood exposure. In contrast, the more moderate delay observed in the sales market may reflect slower belief updating among homeowners and behavioral frictions such as reference dependence, as discussed in Section 5.

The observed patterns suggest heterogeneous timing and magnitude of liquidity shocks across the rental and sales markets. That the survival curves diverge more strongly in the rental market indicates that flooding has a larger immediate impact on lease transactions. This could be because renters, who are more mobile and often respond to short-term housing conditions, are quicker to avoid flood-affected properties—leading to slower turnover in lease contracts. Conversely, the sales market—characterized by higher transaction costs and longer holding periods—shows a more delayed divergence, peaking around five years after the flood.

This temporal asymmetry implies that owner-occupiers may initially resist selling longer than lessors, consistent with reference-dependent preferences and loss aversion, and only gradually adjust their behavior in response to flood risk. The early peak in the lease market likely reflects immediate risk salience and temporary displacement, whereas the later peak in the sales market may signal slower belief updating and structural market adjustments. Because the survival curves capture the timing of actual transactions rather than listing behavior, the larger immediate impact observed in the lease market compared to the sales market may also reflect a sharper post-flood decline in rental demand, which would prolong vacancy durations and delay lease turnovers.

These findings reinforce the notion that floods impose persistent frictions in housing markets. Slower transaction dynamics in both markets suggest that disasters not only affect property valuations but also disrupt the frequency and fluidity of market activity. The sales market may be constrained by financial and behavioral frictions, while the rental market reflects more immediate demand responses and reputational spillovers. Taken together, the patterns observed in housing price dynamics and transaction timing suggest that flood events trigger persistent shifts in market behavior within affected housing markets.

5 Policy Implications

The empirical evidence in the previous section suggests that housing prices adjust gradually after a flood event. The evidence also suggests that although the observed transacted prices adjust slowly, housing markets manifest absorption of a negative shock through reduction in transaction volumes and longer time off the market for properties with less favorable unobserved traits, which implies that properties remain mispriced for extended periods. These findings offer three distinct implications for climate adaptation policy. They suggest that relying on standard market signals may lead to under-adaptation, particularly during economic downturns, and point toward specific interventions in information disclosure and insurance design.

5.1 Liquidity as the Early Warning Signal

Standard cost-benefit analyses for flood protection infrastructure typically rely on hedonic price differences to estimate the avoided damages of flooding. My results suggest that this method could be misleading in the short run. Because sales prices remain statistically flat for five years following a flood in my study, a policymaker observing prices during this window would erroneously conclude that the market does not value flood risk.

Instead, the results indicate that liquidity is the early warning signal of climate distress. The “market freeze” that I document where transaction volumes contract while prices remain sticky demonstrates that the market has detected the risk but prevented it from being priced by seller-side frictions. Adaptation policy must therefore consider closely monitoring liquidity freezes instead of relying on short-term price adjustments right after a shock.

Furthermore, this friction is inherently counter-cyclical. One important implication of reference dependence is that the pace of price adjustment may be further slowed when nominal housing prices are relatively stagnant. This means that adaptation barriers are highest when the economy is weakest. If homeowners are anchored to their original purchase price in nominal terms and if prices of flood-affected properties fall below the original purchase price, then they may be unwilling to sell at a perceived loss in the absence of broader nominal appreciation. Consequently, in low-growth housing markets, the behavioral friction arising from reference dependence can prolong the delay in price responses to adverse shocks such as flooding. On the other hand, in periods where nominal housing prices rise rapidly, price adjustment in the sales market may occur at a faster rate, reducing the delay in response and potentially leading to dynamics that more closely resemble those observed in the lease market. This suggests that public interventions, such as buyouts or managed retreat programs, are most critical during economic downturns when private market adaptation grinds to a halt.

5.2 Information Reform: Breaking the Stigma

In addition, a distinguishing feature of this study is that most of the transacted properties are located in multi-story buildings, where only the lower floors are physically exposed to floodwater. The observed price declines are therefore not driven primarily by direct structural damage but rather by informational and reputational spillovers that affect entire buildings. Even upper-floor properties that remain undamaged and fully habitable experience persistent depreciation following a flood event. This pattern suggests that buyers and renters internalize the stigma of flooding at the building level, treating the entire structure as risky once any portion of the building has been inundated. Such stigma depresses asset values beyond the directly affected units.

To mitigate these reputational damages, policy should aim to provide ways to signal restoration and restore market confidence. One way is to enact a climate resilience certification program that publicly recognizes buildings that have implemented effective flood-mitigation measures. Similar to energy-efficiency ratings, resilience certification would provide credible information to buyers, renters, lenders, and insurers. By signaling that a building has addressed structural vulnerabilities, certification can help reduce the stigma attached to previously flooded buildings and facilitate faster price recovery for non-damaged units.

5.3 Re-designing Flood Insurance as a Price Signal

Finally, the slow adjustment of prices points to incomplete private and public risk-sharing arrangements. Take-up of the existing public natural disaster insurance remains low, even in disaster-risk zones. One contributing factor is charity hazard: the expectation of generous post-disaster relief crowds out demand for private insurance. Currently, uninsured households can receive up to 30-35% of damage compensation through public relief funds, which undermines the incentive to purchase coverage.

To address this, enrollment and compensation rules can be redesigned to preserve humanitarian assistance while reinforcing ex-ante incentives to insure. Shifting public subsidies from post-disaster relief to premium support or introducing automatic enrollment in high-risk zones would strengthen participation and improve fiscal sustainability. Floods can have impacts on upper-floor properties that are not directly damaged and thus ineligible for insurance payout. With mandatory enrollment at the building level, the insurance can be restructured to provide climate retrofitting subsidy that can mitigate reputational damages in the case of flooding.

6 Conclusion

This paper provides new evidence on how housing markets incorporate information about flood events and how behavioral and institutional frictions shape that process. Using nearly two decades of transaction-level data from Seoul, I show that floods induce permanent price declines that unfold gradually over time. Lease prices decline within two years, while sales prices adjust only after five years. In both tenure types, transaction volumes contract first, signaling that market activity responds faster than prices. These findings indicate that while households react quickly to the new information generated by floods, the capitalization of that information into prices occurs only slowly.

The analysis further reveals that behavioral frictions amplify the persistence of flood discounts. Properties affected by floods take substantially longer to be resold or re-leased, and transaction-level evidence shows that homeowners anchor to their original nominal purchase prices, delaying sales that would realize losses. Long-held properties adjust more quickly than recently purchased ones, consistent with reference-dependent preferences on the seller side. Additionally, early post-flood transactions exhibit positive selection, suggesting that composition effects may mask true immediate declines in market valuation during the early post-disaster period.

These results are consistent with the view that flood-induced discounts are not driven by slow learning, but by behavioral frictions that prolong the recovery of transaction activity and asset values. The results challenge the notion that housing markets efficiently internalize environmental risk once it is realized. Persistent price discounts, especially for undamaged upper-floor units, suggest that reputational damage and nominal loss aversion jointly depress market values in both the sales and the lease markets.

The persistence of flood discounts has important welfare and policy implications. Delayed risk capitalization can hinder climate adaptation, misallocate capital toward high-risk areas, and exacerbate inequities by exposing less-informed households to repeated losses. Policies that enhance transparency such as mandatory flood-risk disclosure, public risk mapping, and resilience certification can accelerate the speed with which markets incorporate risk information. Likewise, restructuring flood insurance and disaster relief programs to promote broader and earlier coverage can reduce reliance on behavioral responses and mitigate reputational spillovers.

By combining detailed empirical evidence with behavioral insights, this study shows that the speed and completeness with which housing markets adjust to flood discounts are central to understanding the economic consequences of climate shocks. Housing markets do not simply reflect the physical impacts of floods but also manifest how societies learn and adapt to an environmental risk. For policymakers, this implies that liquidity is a more reliable short-term gauge of climate distress than price. Recognizing the frictions that slow the adjustment to an environmental shock is critical for designing policies that help markets and households better adapt to an era of increasing climate risk.

References

Andersen, Steffen, Cristian Badarinza, Lu Liu, Julie Marx, and Tarun Ramadorai. 2022. "Reference dependence in the housing market." *American Economic Review*, 112(10): 3398–3440.

Atreya, Ajita, and Susana Ferreira. 2015. "Seeing is believing? Evidence from property prices in inundated areas." *Risk Analysis*, 35(5): 828–848.

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

Bakkensen, Laura A, and Lint Barrage. 2022. "Going underwater? Flood risk belief heterogeneity and coastal home price dynamics." *The Review of Financial Studies*, 35(8): 3666–3709.

Bakkensen, Laura A, Xiaozhou Ding, and Lala Ma. 2019. "Flood risk and salience: New evidence from the Sunshine State." *Southern Economic Journal*, 85(4): 1132–1158.

Beltrán, Allan, David Maddison, and Robert Elliott. 2019. "The impact of flooding on property prices: A repeat-sales approach." *Journal of Environmental Economics and Management*, 95: 62–86.

Bernoulli, Daniel. 1766. "Essai d'une nouvelle analyse de la mortalité causée par la petite vérole, et des avantages de l'inoculation pour la prévenir." *Histoire de l'Académie Royale des Sciences*, 1–45.

Bin, Okmyung, and Craig E Landry. 2013. "Changes in implicit flood risk premiums: Empirical evidence from the housing market." *Journal of Environmental Economics and management*, 65(3): 361–376.

Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo. 2016. "Understanding booms and busts in housing markets." *Journal of Political Economy*, 124(4): 1088–1147.

Callaway, Brantly, and Pedro HC Sant'Anna. 2021. "Difference-in-differences with multiple time periods." *Journal of econometrics*, 225(2): 200–230.

Chen, Jiafeng, and Jonathan Roth. 2024. "Logs with zeros? Some problems and solutions." *The Quarterly Journal of Economics*, 139(2): 891–936.

De Chaisemartin, Clément, and Xavier d'Haultfoeuille. 2020. "Two-way fixed effects estimators with heterogeneous treatment effects." *American economic review*, 110(9): 2964–2996.

Donovan, Geoffrey H, Patricia A Champ, and David T Butry. 2007. "Wildfire risk and housing prices: a case study from Colorado Springs." *Land Economics*, 83(2): 217–233.

Ellen, Ingrid Gould, and Rachel Meltzer. 2024. "Heterogeneity in the recovery of local real estate markets after extreme events: The case of Hurricane Sandy." *Real Estate Economics*, 52(3): 714–752.

Friedman, Jerome H. 2002. “Stochastic gradient boosting.” *Computational statistics & data analysis*, 38(4): 367–378.

Genesove, David, and Christopher Mayer. 2001. “Loss aversion and seller behavior: Evidence from the housing market.” *The quarterly journal of economics*, 116(4): 1233–1260.

Gibson, Matthew, and Jamie T Mullins. 2020. “Climate risk and beliefs in New York floodplains.” *Journal of the Association of Environmental and Resource Economists*, 7(6): 1069–1111.

Hirsch, Jens, and Jonas Hahn. 2018. “How flood risk impacts residential rents and property prices: Empirical analysis of a German property market.” *Journal of Property Investment & Finance*, 36(1): 50–67.

Kahneman, Daniel, and Amos Tversky. 1979. “Prospect theory: An Analysis of Decision under Risk.” *Econometrica*, 47: 278.

Kousky, Carolyn. 2010. “Learning from extreme events: Risk perceptions after the flood.” *Land Economics*, 86(3): 395–422.

McCoy, Shawn J, and Randall P Walsh. 2018. “Wildfire risk, salience & housing demand.” *Journal of Environmental Economics and Management*, 91: 203–228.

Ortega, Francesc, and Süleyman Taşpınar. 2018. “Rising sea levels and sinking property values: Hurricane Sandy and New York’s housing market.” *Journal of Urban Economics*, 106: 81–100.

Park, Chang-Gyun, and Eunjung Yeo. 2013. “Adverse selection in the Korean storm and flood insurance market.” *KDI Journal of Economic Policy*, 35(3): 39–63.

Piazzesi, Monika, and Martin Schneider. 2009. “Momentum traders in the housing market: Survey evidence and a search model.” *American Economic Review*, 99(2): 406–411.

Poterba, James M. 1984. “Tax subsidies to owner-occupied housing: an asset-market approach.” *The quarterly journal of economics*, 99(4): 729–752.

Sun, Liyang, and Sarah Abraham. 2021. “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects.” *Journal of econometrics*, 225(2): 175–199.

Vigdor, Jacob. 2008. “The economic aftermath of Hurricane Katrina.” *Journal of Economic Perspectives*, 22(4): 135–154.

Yeom, Jaeweon, Hyungjun Park, Dongoh Ha, Dalbyul Lee, and Juchul Jung. 2019. “The Impact of Risk Perception on Flood Insurance Adoption Intention in the Coastal Areas of South Korea.” *Journal of the Korean Society of Hazard Mitigation*, 19(5): 55–64.

Zivin, Joshua Graff, Yanjun Liao, and Yann Panassie. 2023. "How hurricanes sweep up housing markets: Evidence from florida." *Journal of Environmental Economics and Management*, 118: 102770.

Figures

Figure 1: Proportion of households by contract status

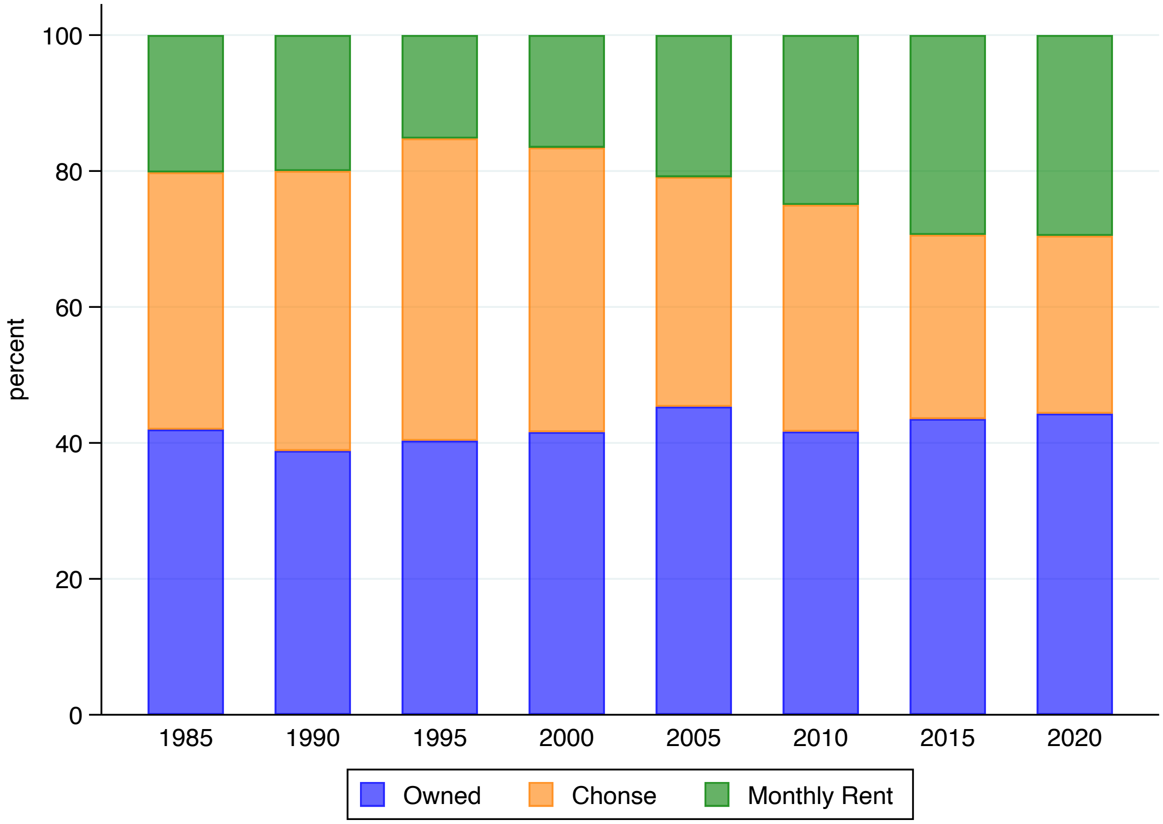

Figure 2: Number of lease transactions by lease type

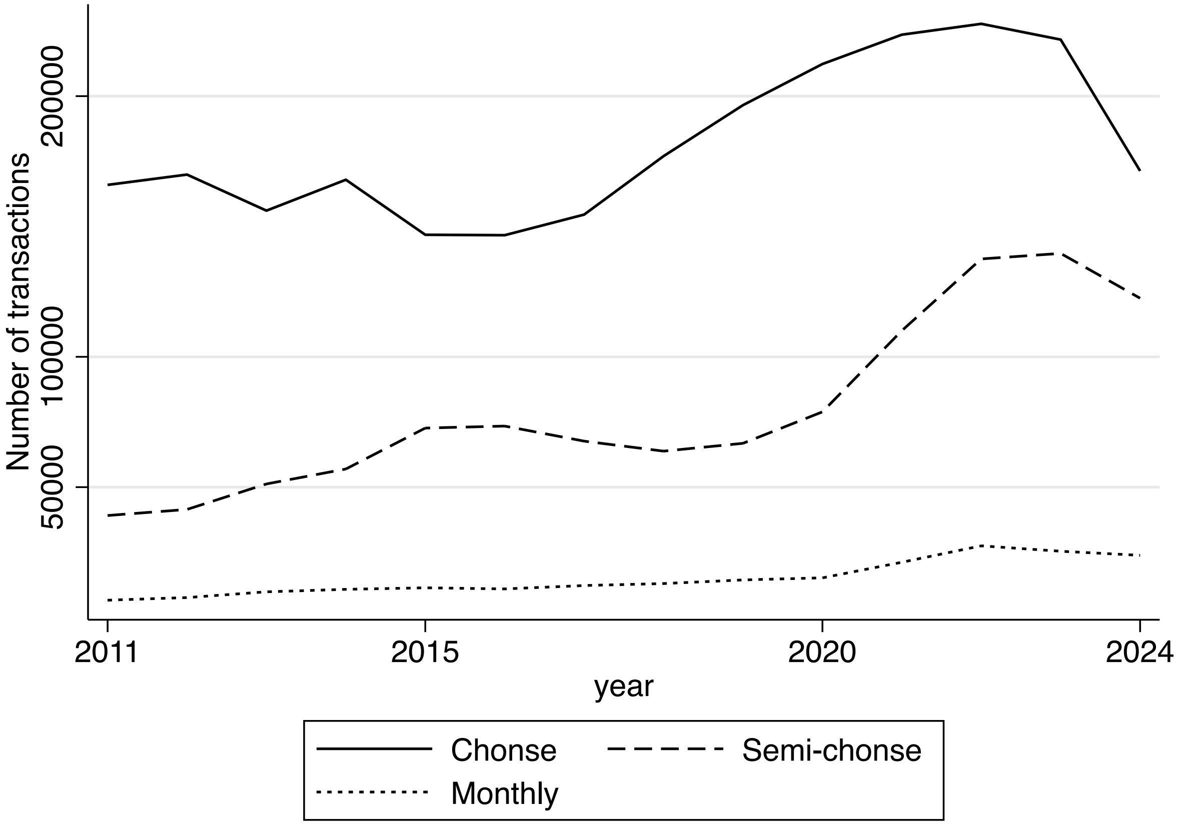

Figure 3: Selected Flood Maps of Seoul

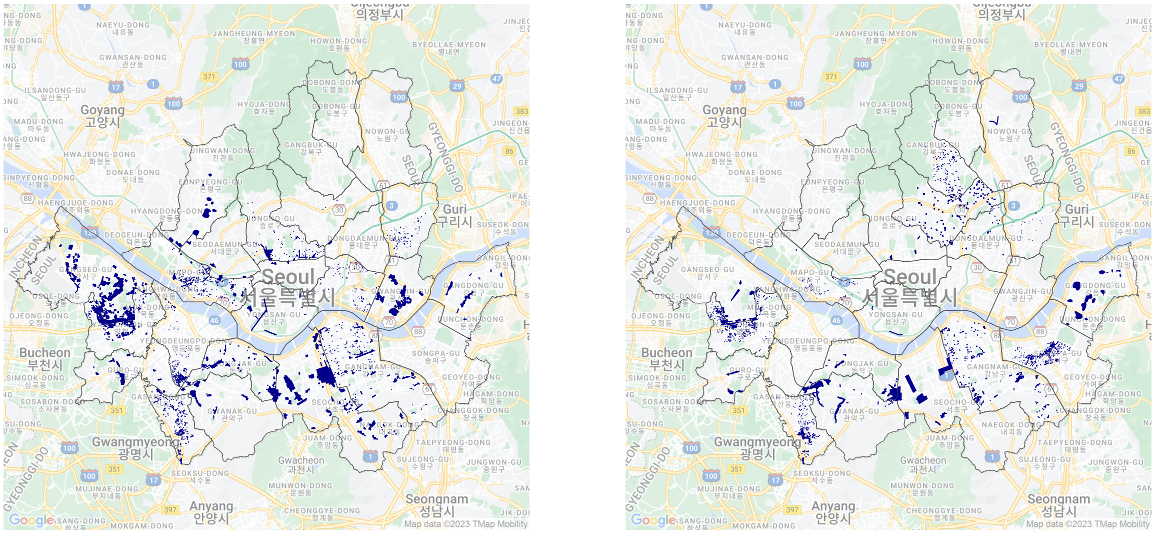
Flooded areas in 2010 (colored in blue)

Flooded areas in 2011 (colored in blue)

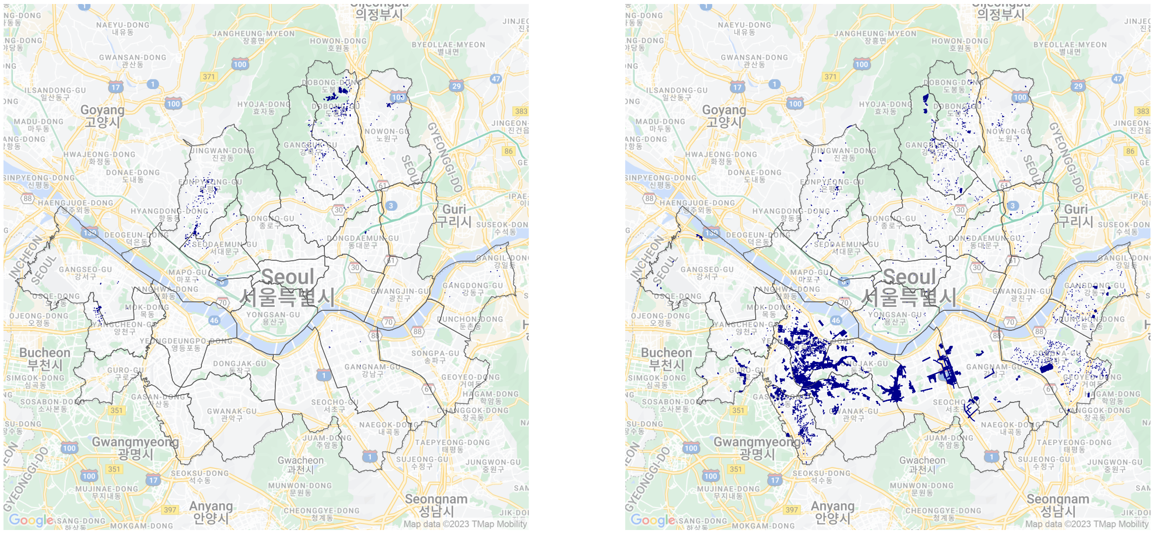
Flooded areas in 2018 (colored in blue)

Flooded areas in 2022 (colored in blue)

Figure 4: Event Study Estimates of Flood Effects on Real Sales Prices

Panel A: Sales

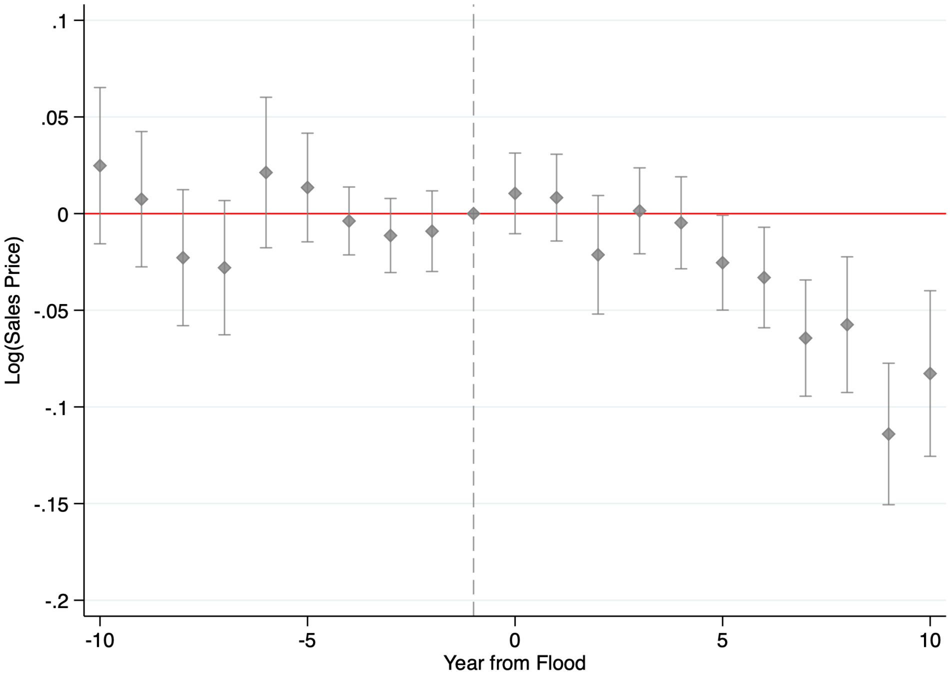

Panel B: Lease

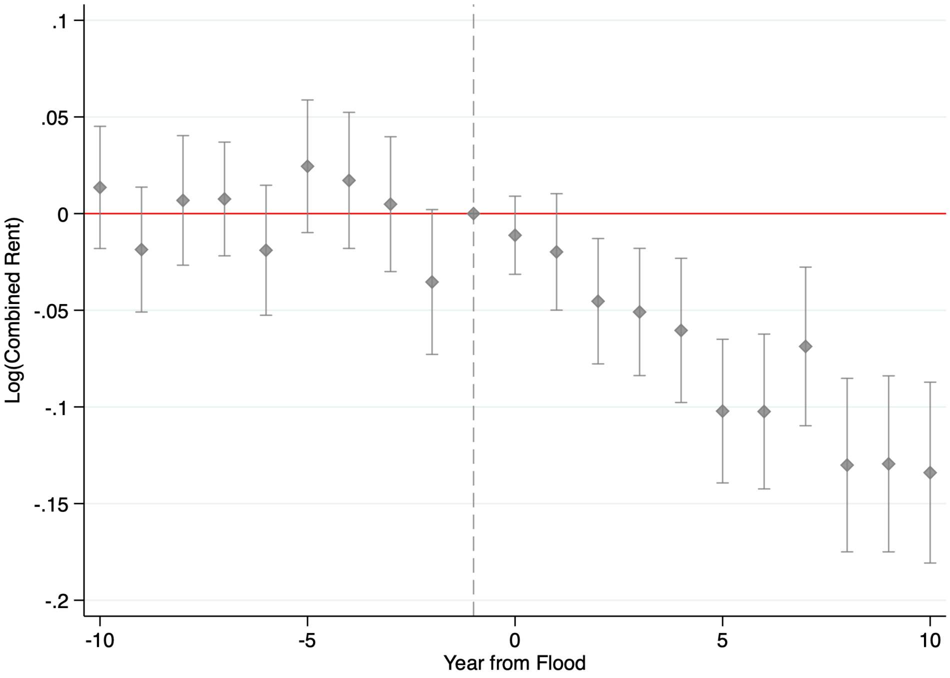

Note: This figure plots the coefficients and 95% confidence intervals from event-study regressions of flooding on logged housing price of the properties located within the floodplain using the Callaway and Sant'Anna (2021) doubly robust estimator with inverse probability weighting. Panel A plots the dynamic effects on logged sales price. Panel B plots the dynamic effects on logged rental price.

Figure 5: Event Study Estimates of Flood Effects on Housing Transaction Counts

Panel A: Sales

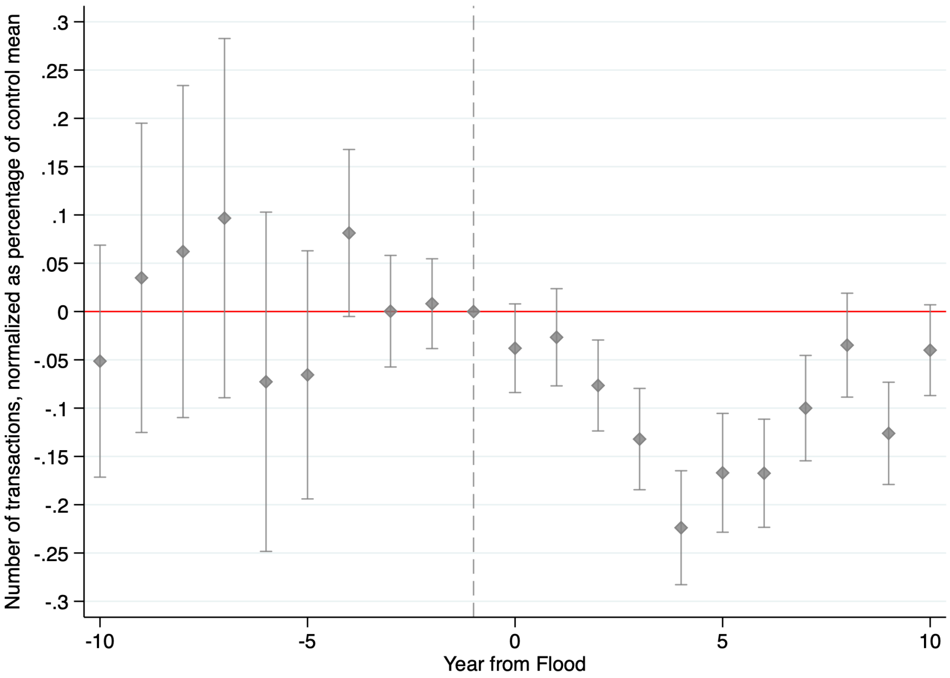

Panel B: Lease

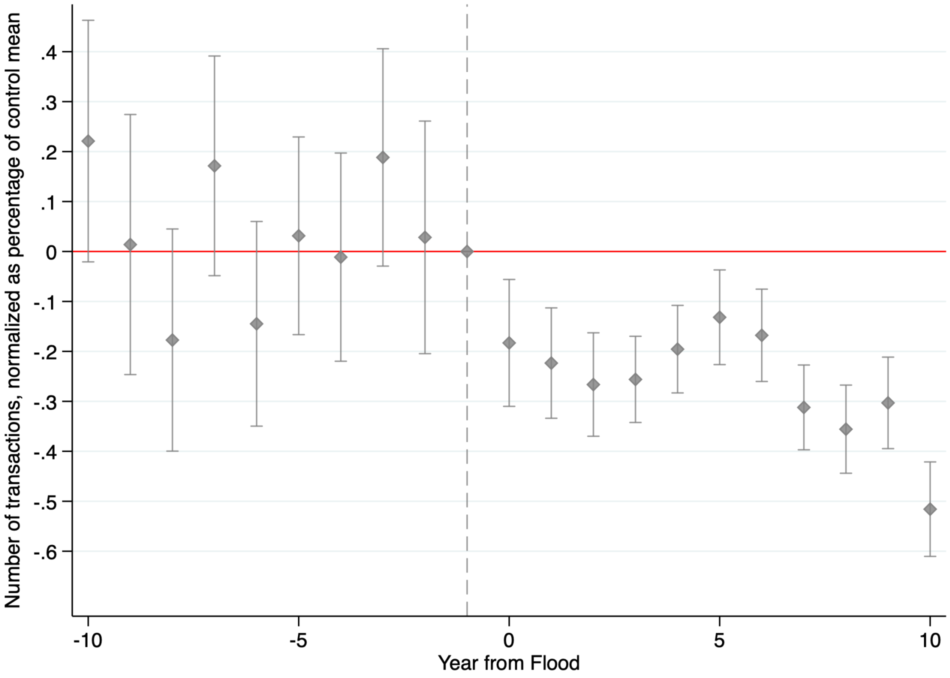

Note: This figure plots the coefficients and 95% confidence intervals from event-study regressions of flooding on logged transaction counts of the properties located within the floodplain using the Callaway and Sant'Anna (2021) doubly robust estimator with inverse probability weighting. Panel A plots the dynamic effects on logged transaction counts of housing sales. Panel B plots the dynamic effects on logged transaction counts of housing leases.

Figure 6: Reference Dependence and Loss Aversion in Sales Market

Panel A: Binned frequencies across nominal gains

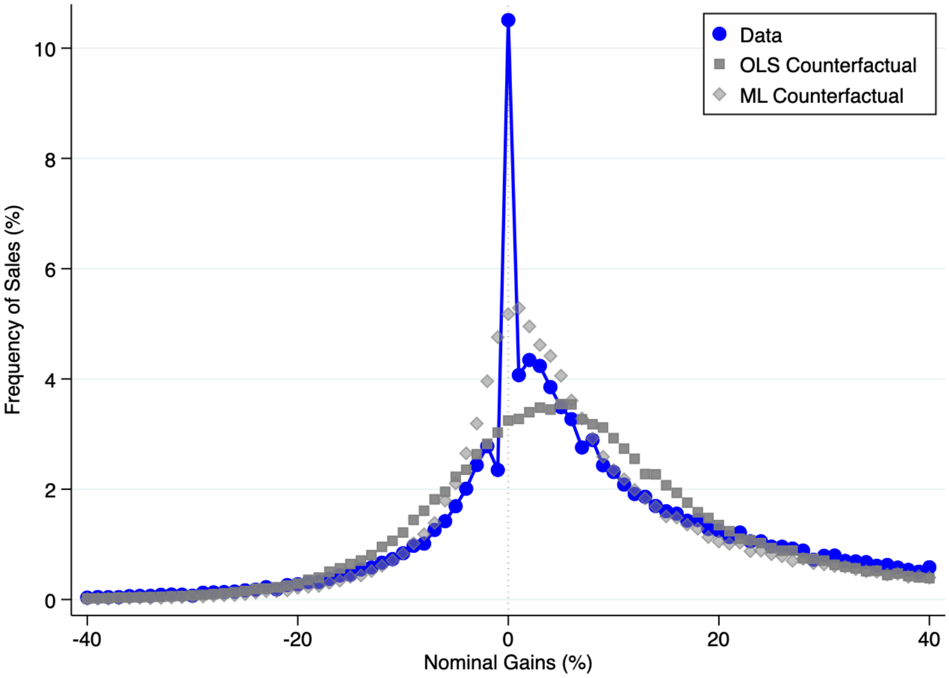

Panel B: Excess mass across gains

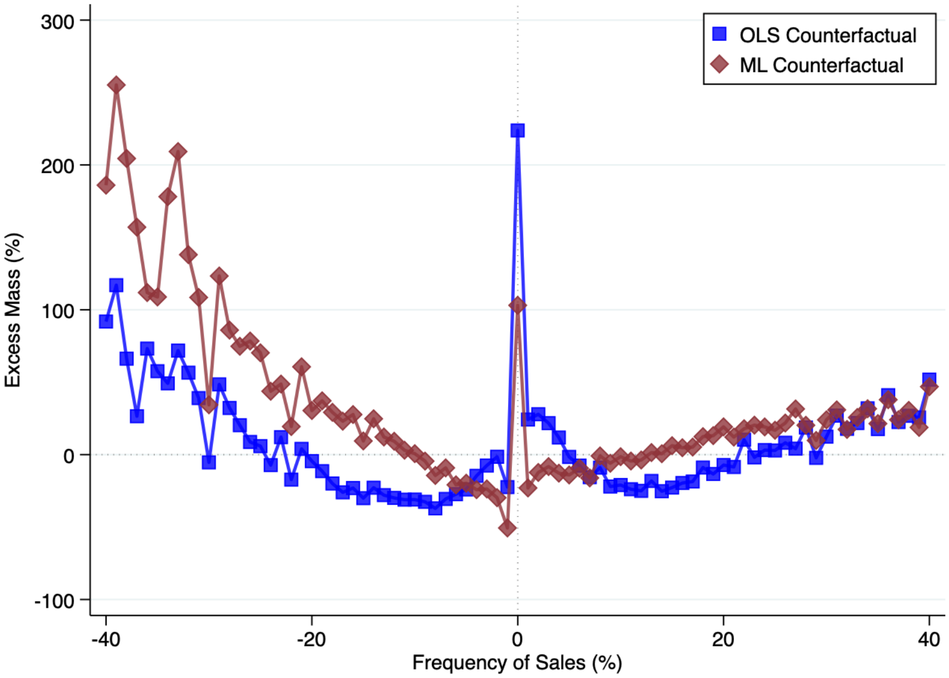

Note: Panel A plots binned frequencies of sales transactions in one percentage-point steps across realized nominal gains. The counterfactual gains are obtained using the hedonic pricing model. Panel B plots excess mass of transactions relative to the level of the counterfactual.

Figure 7: Reference Dependence and Loss Aversion in Lease Market

Panel A: Binned frequencies across nominal gains

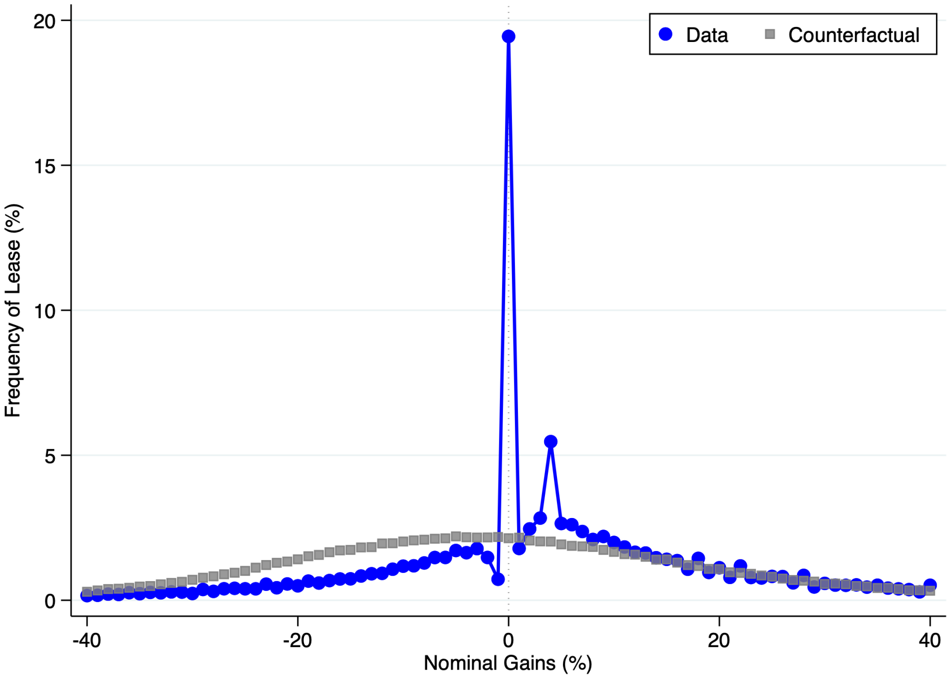

Panel B: Excess mass across gains

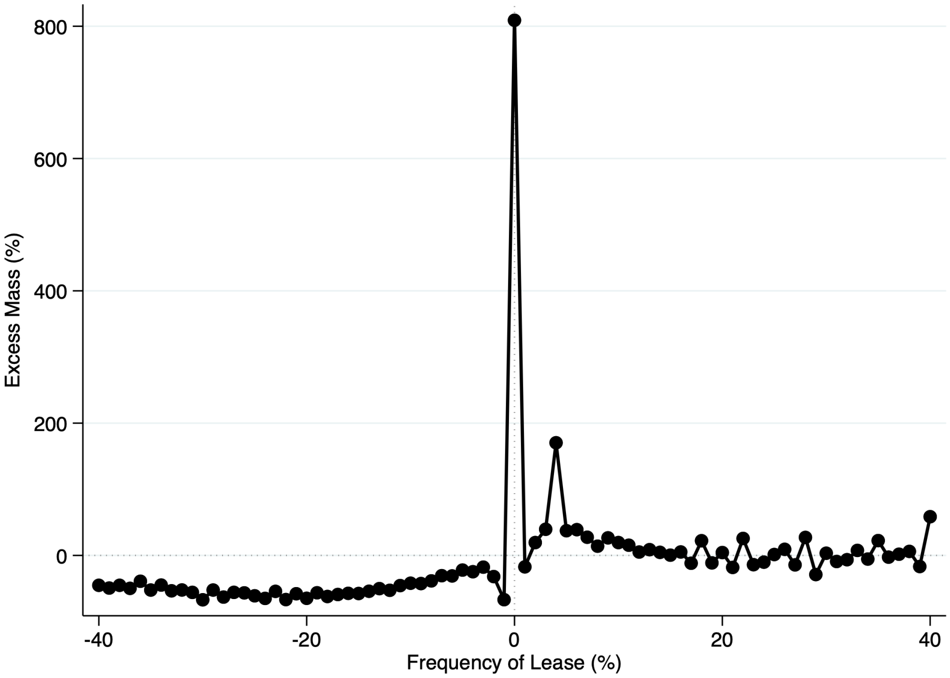

Note: Panel A plots binned frequencies of lease transactions in one percentage-point steps across realized nominal gains. The counterfactual gains are obtained using the hedonic pricing model. Panel B plots excess mass of transactions relative to the level of the counterfactual.

Figure 8: Partial Dependence Plots from ML Counterfactual Prices

(a) Panel A: Marginal effect of log assessed value (b) Panel B: Marginal effect of area size

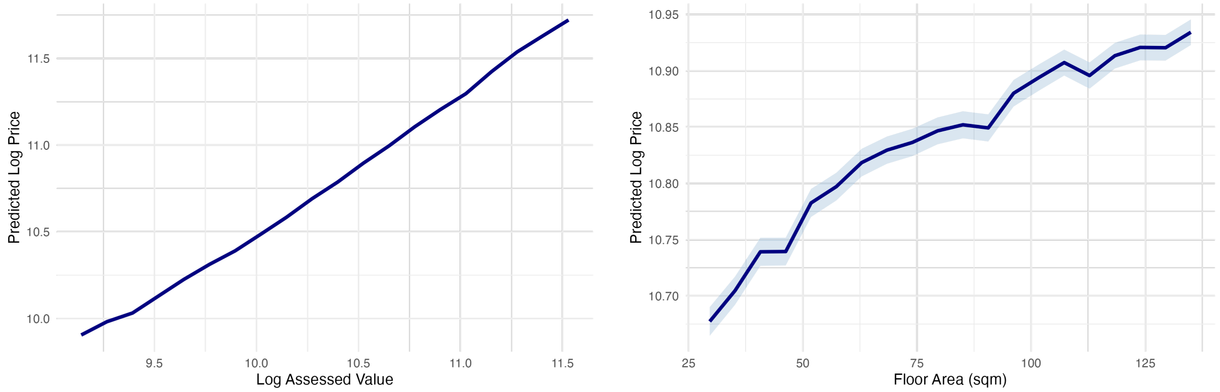

(c) Panel C: Marginal effect of being flooded (d) Panel D: Marginal effect of floor category

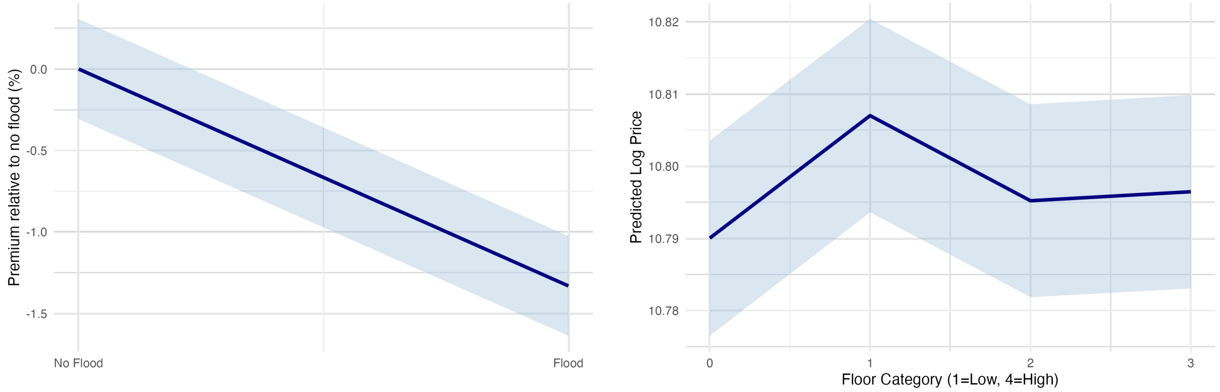

(e) Panel E: Marginal effect of year built (f) Panel F: Marginal effect of building age

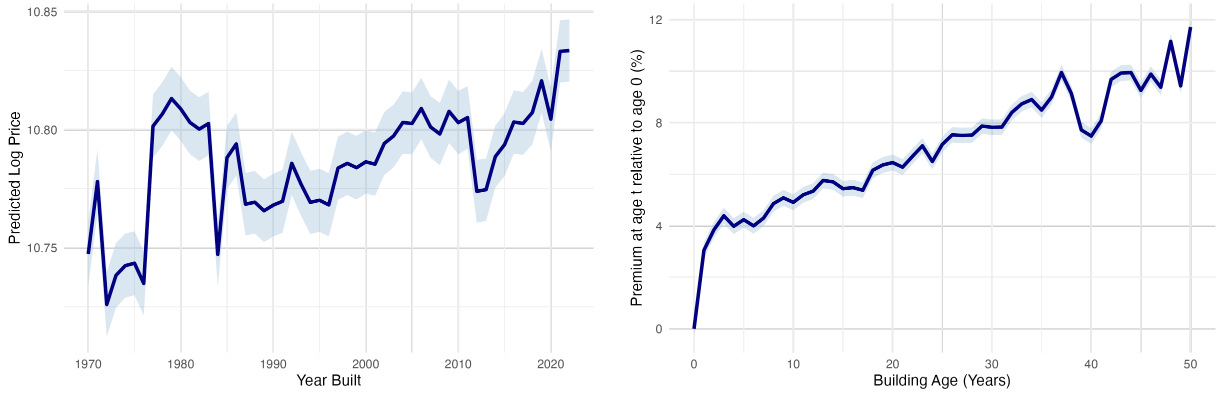

Note: The figures display Partial Dependence Plots (PDPs) generated from the gradient-boosted decision tree (LightGBM) model. The solid lines represent the average marginal effect of each feature on the predicted log transaction price. Panels A-F show the marginal effect of each housing attribute. The lower and upper 95% confidence levels are indicated with the light blue area.

Figure 9: Event Study Estimates of Price Effects by Time Since Previous Sale

Panel A: Long-Held Properties

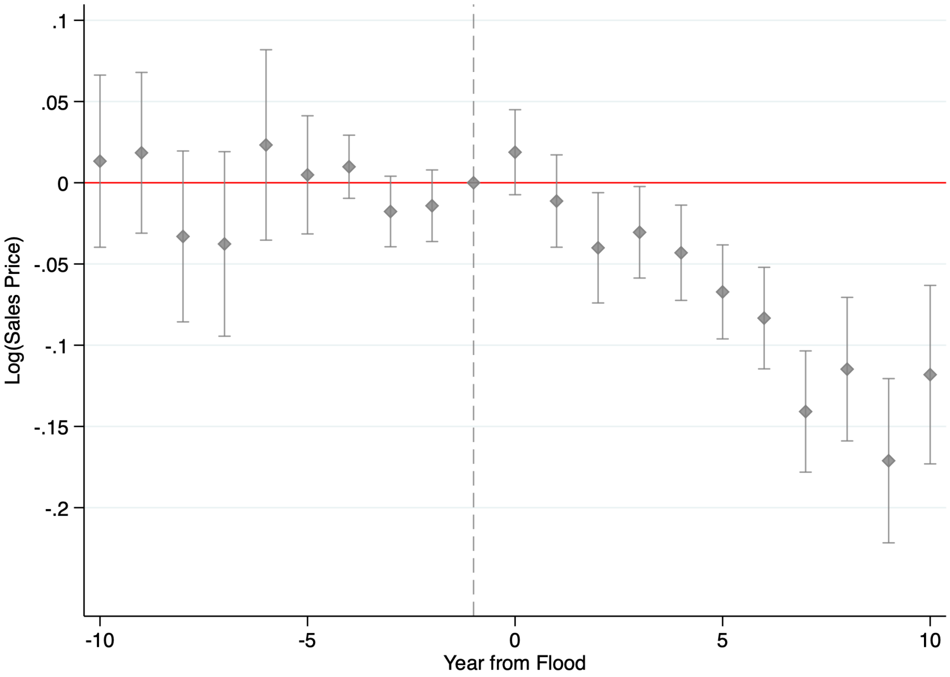

Panel B: Recently Purchased Properties

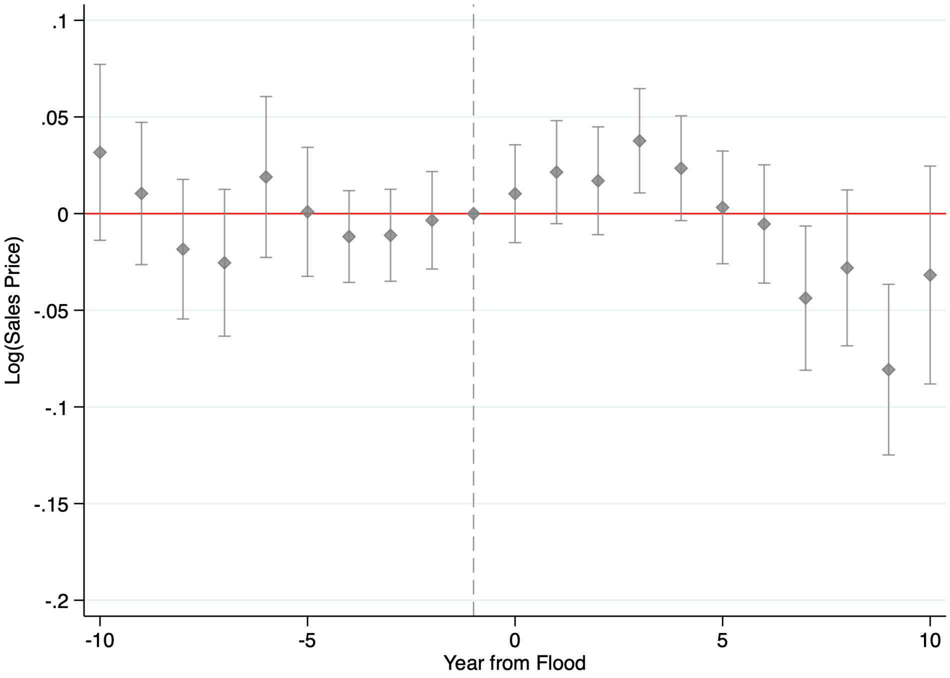

Note: This figure plots the coefficients and 95% confidence intervals from event-study regressions of flooding on logged housing price of the properties located within the floodplain using the Callaway and Sant'Anna (2021) doubly robust estimator with inverse probability weighting. Panel A plots the dynamic effects on logged sales price for properties that were previously purchased more than 3 years ago. Panel B plots the dynamic effects on logged sales price for properties that were previously purchased in less than 3 years.

Figure 10: Event-Study Estimates of Flood Impacts on Residual Housing Prices

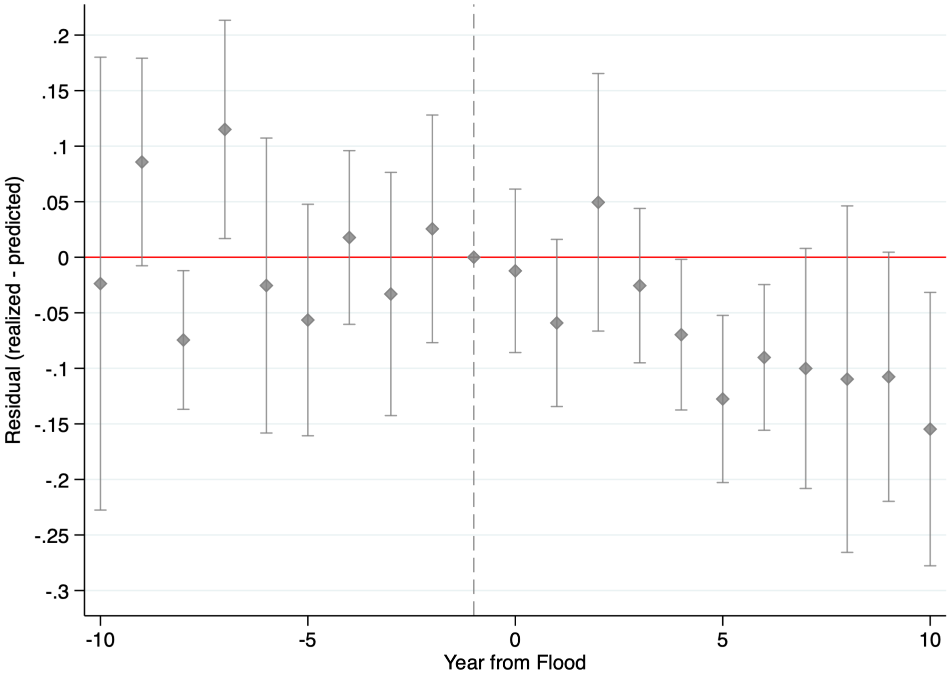

Note: This figure plots the coefficients and 95% confidence intervals from event-study regressions of flooding on residual housing price (actual price minus potential value from the hedonic pricing model) of the properties located within the floodplain using the Callaway and Sant'Anna (2021) doubly robust estimator with inverse probability weighting. The potential values of the properties are the predicted values from the hedonic pricing model in Equation 3.

Figure 11: Probability of Resale After Flood

Panel A: Sales

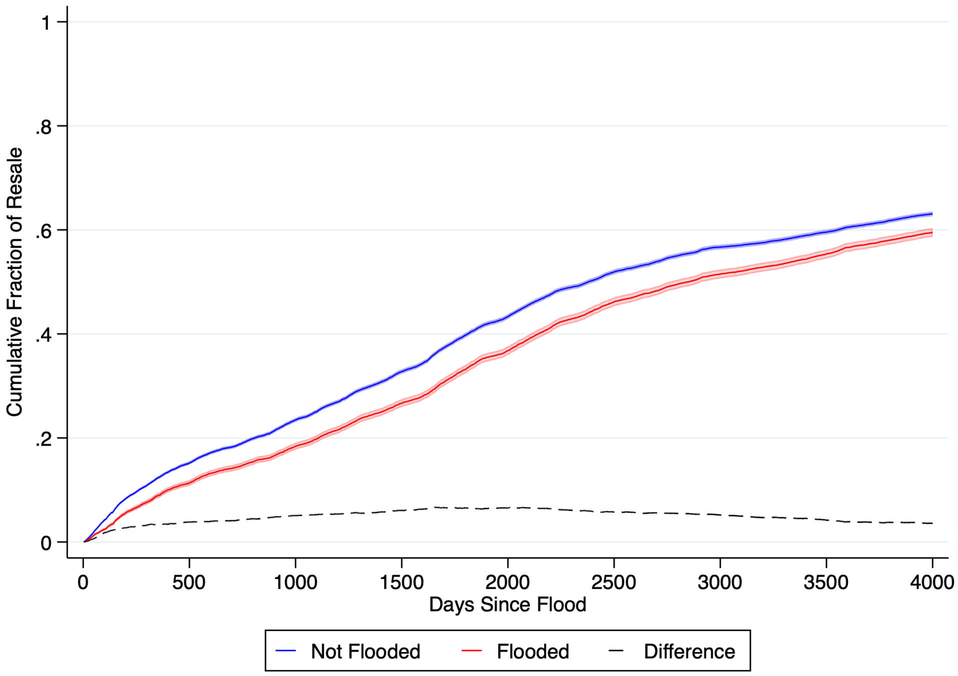

Panel B: Lease

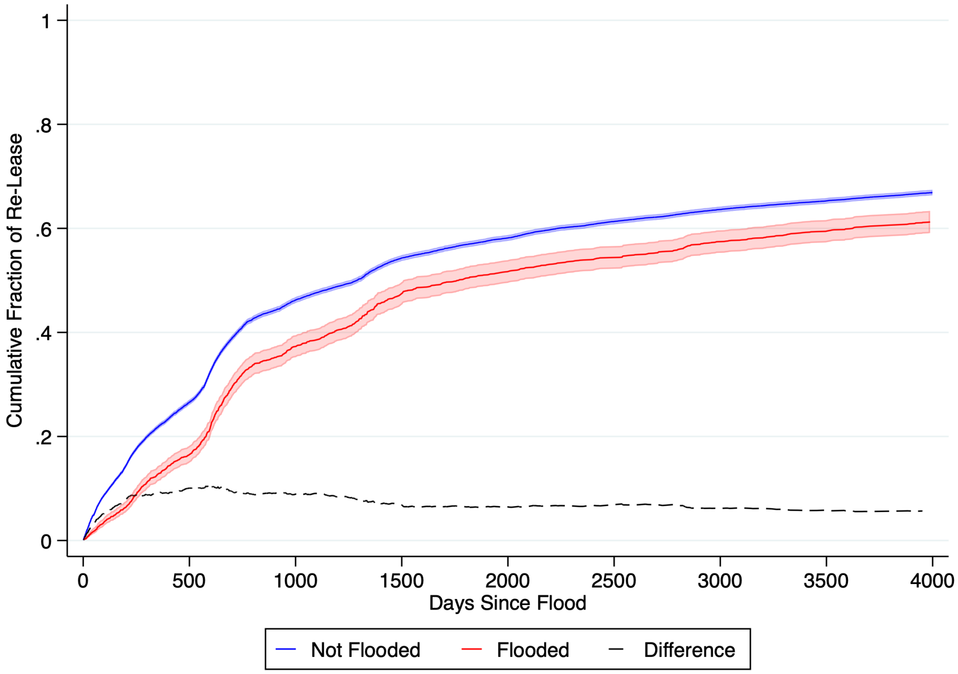

Note: Panel A plots the cumulative fraction of properties that have been resold as a function of time since the flood. Panel B plots the cumulative fraction of properties that have been re-leased as a function of time since the flood.

Tables

Table 1: Summary Statistics

Panel A: Sales vs. Lease

Variable	Sales				Lease			
	Mean	SD	Min	Max	Mean	SD	Min	Max
Sales Price	40791	29884	10000	200000				
Deposit					17060	18012	0	181000
Monthly Rent					22	34	0	415
Area Size (m^2)	71	26	15	317	59	28	10	294
Year Built	1998	8	1900	2009	1995	10	1006	2010
Floor	6	6	-1	64	6	6	-1	60
Observations	143326				341646			

Panel B: Flooded vs. Non-flooded

Variable	Non-flooded				Flooded			
	Mean	SD	Min	Max	Mean	SD	Min	Max
Sales Price	43447	30802	10000	200000	33490	25830	10000	200000
Deposit	17769	18565	0	181000	13169	13984	0	140000
Monthly Rent	22	35	0	415	20	29	0	400
Area size (m^2)	63	29	10	317	59	26	10	261
Year built	1996	9	1918	2010	1996	9	1939	2010
Floor	7	6	-1	64	5	5	-1	37
Observations	394090				90882			

Note: This table reports summary statistics of the housing transactions in Seoul, South Korea. Prices are reported in units of 10,000 South Korean won (KRW). Data are restricted to properties located within the floodplain that transact more than once over our time period. Dependent variables in regressions are log-transformed. For Panel B, “Flooded” means the building of the property is flooded, while “Non-flooded” means the building of the property is not flooded.

Table 2: The Impact of Flooding on Housing Prices

Dependent Variable:	Sales			Combined Rent		
(logged)	(1)	(2)	(3)	(4)	(5)	(6)
Post Flood	-0.0461***	-0.0359***	-0.0600***	-0.0711***	-0.0575***	-0.0939***
	(0.0103)	(0.0115)	(0.0244)	(0.0150)	(0.0163)	(0.0396)
Floor	All	> 1	≤ 1	All	> 1	≤ 1
Repeat Sales	Yes	Yes	Yes	Yes	Yes	Yes
Property FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	143,316	125,795	27,914	209,734	187,864	19,982
Mean of Dep. Var.	10.618	10.548	9.928	4.429	4.468	4.094

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The dependent variables are the logs of real sale price and combined rental price. Data are restricted to properties for which I observe more than one transaction during my sample period. I further restrict the data to properties that are located within the floodplain to ensure comparability. Standard errors in parentheses are clustered at the unit level.

Table 3: The Impact of Flooding on Housing Transaction Counts

Dependent Variable:	Sales			Combined Rent		
(logged)	(1)	(2)	(3)	(4)	(5)	(6)
Post Flood	-0.082***	-0.072***	-0.152**	-0.282***	-0.160***	-0.167*
	(0.019)	(0.021)	(0.056)	(0.028)	(0.016)	(0.078)
Floor	All	> 1	≤ 1	All	> 1	≤ 1
Property FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,194,704	1,008,496	186,208	849,693	715,481	134,212
Mean of Dep. Var.	0.109	0.113	0.088	0.123	0.232	0.088

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The dependent variables are the number of transactions for sales and rents normalized by the mean of the control group. The balanced panel of housing properties includes only the properties for which I observe at least one transaction during my sample period. I further restrict the data to properties that are located within the floodplain to ensure comparability. Standard errors in parentheses are clustered at the cell level.

Table 4: Hedonic Pricing Models on Sales

Dependent Variable: Log Sales	(1)	(2)	(3)	(4)	(5)
Log assessed value	0.937^{***}	0.948^{***}	0.968^{***}	0.990^{***}	
	(0.001)	(0.000)	(0.000)	(0.000)	
Property size	0.000^{***}	0.000^{***}			0.005^{***}
	(0.000)	(0.000)			(0.000)
Year built	0.012^{***}	0.012^{***}			0.040^{***}
	(0.000)	(0.000)			(0.001)
Building age	0.014^{***}	0.013^{***}			0.042^{***}
	(0.000)	(0.000)			(0.001)
Within floodplain		-0.002^{***}			
		(0.000)			
Number of floods = 1	-0.009^{***}	-0.013^{***}			-0.150^{***}
	(0.001)	(0.001)			(0.002)
Number of floods = 2	-0.006^{***}	-0.008^{***}			-0.252^{***}
	(0.002)	(0.002)			(0.004)
Number of floods = 3	-0.039^{***}	-0.027^{***}			-0.061^{***}
	(0.006)	(0.005)			(0.011)
Properties	Floodplain	All	Floodplain	Floodplain	Floodplain
Year × Month FE	Yes	Yes	Yes	No	Yes
District FE	Yes	Yes	No	No	Yes
Observations	190,640	1,204,103	190,640	190,640	415,193
R^2	0.971	0.970	0.970	0.960	0.690
Adj. R^2	0.971	0.970	0.969	0.960	0.690

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The fitted values of the first regression model are used as the counterfactual sales price in Figure 6. Model (1) is chosen based on its highest adjusted R^2. The counterfactual distributions using the fitted values from other models are shown in the Appendix. The dependent variables are the logs of nominal sale price. Data are restricted to properties for which I observe more than one transaction during my sample period. Standard errors are in parentheses.

Table 5: Hedonic Pricing Models on Lease

Dependent Variable: Log Rent	(1)	(2)	(3)	(4)	(5)
Log assessed value	0.497^{***}	0.593^{***}	0.737^{***}	0.717^{***}	
	(0.001)	(0.001)	(0.001)	(0.001)	
Property size	0.006^{***}	0.005^{***}			0.014^{***}
	(0.000)	(0.000)			(0.000)
Year built	0.035^{***}	0.037^{***}			0.062^{***}
	(0.001)	(0.000)			(0.001)
Building age	0.024^{***}	0.027^{***}			0.051^{***}
	(0.001)	(0.000)			(0.001)
Within floodplain		0.038^{***}			
		(0.001)			
Number of floods = 1	0.006^{***}	0.025^{***}			-0.036^{***}
	(0.001)	(0.001)			(0.001)
Number of floods = 2	0.056^{***}	0.123^{***}			-0.040^{***}
	(0.003)	(0.003)			(0.002)
Number of floods = 3	-0.151^{***}	-0.156^{***}			0.004
	(0.007)	(0.008)			(0.004)
Properties	Floodplain	All	Floodplain	Floodplain	Floodplain
Year × Month FE	Yes	Yes	Yes	No	Yes
District FE	Yes	Yes	No	No	Yes
Observations	390,637	2,308,966	390,637	390,637	1,161,391
R^2	0.813	0.768	0.734	0.709	0.720
Adj. R^2	0.813	0.768	0.734	0.709	0.720

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The fitted values of the first regression model are used as the counterfactual lease price in Figure 7. Model (1) is chosen based on its highest adjusted R^2. The counterfactual distributions using the fitted values from other models are shown in the Appendix. The dependent variables are the logs of rental price, which combines deposit and monthly rent using the fixed conversion rate of 4.5%. Data are restricted to properties for which I observe more than one transaction during my sample period. Standard errors are in parentheses.

Table 6: Impact of Flooding on Time to Re-Lease: AFT Model Estimates

Panel A: Sales

	(1)	(2)	(3)	(4)	(5)	(6)
Flooded	1.374***	1.431***	1.457***	1.237***	1.305***	1.332***
	(0.027)	(0.031)	(0.034)	(0.019)	(0.022)	(0.024)
Distribution	Weibull	Log-logistic	Lognormal	Weibull	Log-logistic	Lognormal
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	No	No	No
Log-Likelihood	-87872.24	-87232.39	-87724.98	-88589.25	-88042.75	-88559.87
AIC	175810.48	174530.78	175515.95	177196.50	176103.50	177137.75
Observations	57321	57321	57321	57321	57321	57321

Panel B: Lease

	(1)	(2)	(3)	(4)	(5)	(6)
Flooded	1.384***	1.516***	1.579***	1.381***	1.519***	1.608***
	(0.072)	(0.081)	(0.088)	(0.066)	(0.075)	(0.083)
Distribution	Weibull	Log-logistic	Lognormal	Weibull	Log-logistic	Lognormal
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	No	No	No
Log-Likelihood	-55319.90	-54192.32	-54188.08	-55847.34	-54698.59	-54719.86
AIC	110697.80	108442.64	108434.16	111704.68	109407.19	109449.72
Observations	31227	31227	31227	31227	31227	31227

Note: The dependent variable for Panel A is the time until a property is resold, measured in days. The dependent variable for Panel B is the time until a property is re-leased, measured in days. The reported coefficients are exponentiated coefficients (time ratios). All specifications include controls for area and year built. I select Model 2 in Panel A and Model 3 in Panel B based on the Akaike Information Criterion (AIC), with lower values indicating better fit.

Appendix A Testing for the delay in price decline

To test whether there is a delay in price adjustment in the sales market compared to the lease market, I use the re-centered influence function (RIF) representation of the group–time average treatment effects, ATT(g, t). The influence function measures the sensitivity of an estimator to individual observations and forms the basis for variance calculations. The influence function is re-centered around the parameter of interest to obtain a random variable with two properties: (i) it has mean equal to the true treatment effect, and (ii) it preserves the influence-function expansion that characterizes the sampling distribution of the estimator.

In practice, the difference-in-differences regression model with multiple time periods by Callaway and Sant'Anna (2021) computes the RIFs for each unit and time period. Each RIF captures how an observation contributes to the estimation of ATT(g, t). Once constructed, these RIFs serve as the fundamental building blocks for aggregation and inference. For example, average treatment effects across groups and periods can be obtained by averaging the corresponding RIFs, and their sampling variability can be consistently estimated by exploiting the variance of the RIFs. Using RIFs avoids repeated re-estimation of nuisance parameters, while ensuring that standard errors remain valid under staggered treatment adoption and heterogeneous treatment effects.

This structure also enables flexible post-estimation analyses, such as joint Wald tests across periods, event-study plots, or comparisons of treatment effects across subgroups. In my application, I leverage the RIFs generated by the event-study regressions to conduct such analyses. To test for the delay in price decline in the sales market, I run the event-study regressions separately for the sales market and the lease market. I then use the RIFs generated by two regressions to conduct pointwise and joint Wald tests and draw plots showing the differences in event-study coefficients between sales and lease.

First, I conduct pointwise tests to examine whether each pair of event-study coefficients is different. I then conduct joint Wald tests to determine whether the first five post-event periods jointly differ. I conduct joint Wald tests on later post-event periods and all post-event periods to show that price adjustment paths of the sales market and the lease market differ. The results are summarized in Table A.1.

Second, to draw a plot showing differences in the event-study coefficients of the sales market and the lease market, I first append the RIFs from sales and lease. Since each RIF is centered at the corresponding treatment effect, the difference between two effects can be obtained by combining the two influence functions and taking their average. This is equivalent to appending the RIFs from the sales and lease estimations into a single dataset, while multiplying the lease-market RIFs by -1. I then generate 5,000 bootstrap samples of the combined RIFs to obtain clustered standard errors.

Formally, let $\hat{\beta}_\tau^m$ be the event-time coefficient for market $m \in \{S, L\}$ obtained from the stacked DiD event-study (property and time fixed effects, common controls), and let $\psi_{i\tau}^m$ denote the influence function of observation i to $\hat{\beta}_\tau^m$ such that

$$\sqrt{N}\big(\hat{\beta}_\tau^m - \beta_\tau^m\big) \ \xrightarrow{d}\ \mathcal{N}(0, \, \mathbb{V}[\psi_{i\tau}^m])).$$

Let us denote the difference between sales and lease $\hat{\Delta}_\tau \equiv \hat{\beta}_\tau^S - \hat{\beta}_\tau^L$. The influence-function representation for $\hat{\Delta}_\tau$ can be written as

$$\psi_{i\tau}^{\Delta} \equiv \psi_{i\tau}^S - \psi_{i\tau}^L,$$

which yields a consistent variance estimator

$$\widehat{\mathbb{V}}(\hat{\Delta}_\tau) = \frac{1}{N} \sum_i \big(\psi_{i\tau}^{\Delta} - \bar{\psi}_\tau^{\Delta}\big)\big(\psi_{i\tau}^{\Delta} - \bar{\psi}_\tau^{\Delta}\big)',$$

implemented via clustered bootstrap at the unit level, the same level used in the main analysis. Pointwise tests use

$$t_\tau = \frac{\hat{\Delta}_\tau}{\sqrt{\widehat{\mathbb{V}}(\hat{\Delta}_\tau)}},$$

and a joint test across all post-event $\tau \in \mathcal{T}_+$ uses the Wald statistic

$$W \;=\; \hat{\boldsymbol{\Delta}}' \, \widehat{\mathbf{V}}_{\boldsymbol{\Delta}}^{-1} \, \hat{\boldsymbol{\Delta}}, \quad \text{with } \hat{\boldsymbol{\Delta}} = (\hat{\Delta}_\tau)_{\tau \in \mathcal{T}_+},$$

with p-values obtained from the bootstrap distribution of W.

Appendix B Hedonic Pricing Using Machine Learning

In this section, I describe in detail the gradient-boosted decision tree model that I used to estimate the counterfactual gains based on housing attributes. Traditionally, the standard hedonic OLS specifications were used to estimate the values of housing attributes. However, a potential concern with the standard hedonic model is that it assumes a log-linear functional form, which may fail to capture non-linearities in housing valuation. This is critical especially if some housing attibutes such as building age have a non-monotonic relationship with housing price.

To ensure that the observed bunching is not due to misspecification of the functional form in the standard hedonic model, I complement the standard hedonic regression results with the results from the gradient-boosted decision tree models (LightGBM). LightGBM, which is short for Light Gradient-Boosting Machine, uses histogram-based algorithms to bucket continuous values into discrete bins. LightBGM selects the leaf that provides the maximum potential reduction in the loss function to split next, regardless of its depth. This approach focuses computational power where it matters most, leading to deeper, more complex, and often more accurate trees with fewer total splits.

I set the maximum number of boosting iterations to 10,000. This high upper-bound was chosen to accommodate a low learning rate of 0.05, which improves the model's ability to generalize to unseen data. To prevent overfitting, I employ early stopping: training is halted if the out-of-sample RMSE do not improve for 50 consecutive rounds. This ensures the model uses the optimal number of trees required to capture the signal without memorizing the noise. The optimal training round for my data sample was 8,526.

To prioritize out-of-sample prediction power over training fit, I utilize a 5-fold cross-validation strategy. A 20% hold-out set per fold provides a sufficiently large sample to estimate the generalization error with low variance, balancing statistical robustness with computational efficiency. I also employ Stochastic Gradient Boosting (Friedman, 2002) by subsampling 80% of the training data every 5 iterations. This injection of randomness reduces overfitting and improves out-of-sample generalization.

In addition, I restrict the tree complexity to a maximum of 31 leaves. This constraint acts as a regularization mechanism, limiting the model to capturing interactions of approximately fifth-order depth, thereby preventing the algorithm from fitting idiosyncratic noise in the housing data. I also impose a minimum leaf size of 20 observations. This prevents the model from generating price estimates based on outliers or thin data segments.

The resulting predicted log sales price $\widehat{\log P}_{ML}$ has an R^2 of 0.994. is used to create binned frequencies of transactions across counterfactual gains $\hat{G} = \widehat{\log P}_{ML} - \log R$, which can be interpreted as the value of the property if households were to sell their properties at their property values implied by the machine-learning hedonic pricing model.

Appendix C Accelerated Failure Time Model

To formally test whether the distribution of resale times differs between flooded properties and non-flooded properties, I run the accelerated failure time regression model, a parametric approach to survival analysis by modeling the logarithm of the resale time as a linear function of covariates. Specifically, let T_i denote the time it takes for property i to be resold, and $X \in$

$\mathbb{R}^k$ be a vector of covariates. The AFT model assumes the following log-linear relationship:

$$\log(T_i) = X_i^T \beta + \sigma \varepsilon_i,$$

where σ is a scale parameter and ε_i is a random error term following a known distribution.

For a binary treatment variable, $\exp(\hat{\beta})$ represents the multiplicative change in the expected time to event for the treated group relative to the control:

$$\frac{T^{Flood}}{T^{NoFlood}} = \exp(\hat{\beta}).$$

Appendix D Conceptual Model of Gradual Learning

The empirical evidence presented in Section 4 highlights a gradual adjustment of housing prices following flood events. In particular, I document that while both sales and lease prices gradually decline after flooding, sales prices respond with a significant delay, taking up to five years to show a measurable decrease. To illustrate how prices can slowly adjust through belief updating and social learning, I present a stylized dynamic model of the housing market in which agents gradually update their beliefs about flood risk based on observed events. The model is adopted from Burnside, Eichenbaum and Rebelo (2016). Although some households immediately update their beliefs on future flood risks, it is possible to generate gradual learning of future flood risk beliefs with belief heterogeneity and social interaction.

Following Burnside, Eichenbaum and Rebelo (2016), I model a population of households who form beliefs about the value of purchasing property in flood-prone areas. The economy consists of a continuum of agents, each with linear utility and beliefs over whether a recent flood signals a change in long-term flood risk. In this experiment, I consider only the properties within the floodplain. Agents can own one house or rent. There is a fixed stock of properties $k < 1$ available for sale, leaving the rental market consist of $1 - k$ homes. Each period, properties on the floodplain incur expected flood damage δ with probability π_t. I allow households to hold heterogeneous beliefs on future flood risk. To simplify, let us assume that the true flood risk is $\pi^* \in \{\pi_L, \pi_H\}$, where π_H indicates persistently higher flood risk and π_L indicates persistently lower flood risk.

Let us first consider the equilibrium. Each period t, households choose from the following options: (i) buy a property that was not flooded at price P_t^{NF} and receive the flow utility ε_h, (ii) buy a flood-affected property at price P_t^F and receive the flow utility $\varepsilon^h - \pi_t \delta$, or (iii) rent a property and receive the flow utility ε^r by paying rent w^{NF}. Then in equilibrium, prices of properties satisfy the following equation:

$$-P_t^{NF} + \beta\big(\varepsilon^h + \mathbb{E}_t\big[P_{t+1}^{NF}\big]\big) = -P_t^F + \beta\big(\varepsilon^h - \pi_t\delta + \mathbb{E}_t\big[P_{t+1}^F\big]\big) = \beta\big(\varepsilon^r - w^{NF}\big). \qquad \text{(A.1)}$$

Let $\Delta \equiv \varepsilon^h - (\varepsilon^r - w)$, which indicates the net utility flow of owning a property versus renting. Then the steady-state prices of properties are

$$\begin{aligned} P^{NF} &= \frac{\beta}{1-\beta}\Delta, \\ P^F &= \frac{\beta}{1-\beta}\left(\Delta - \pi\delta\right). \end{aligned}$$

In my study, I compare the price dynamics of flood-affected properties with properties that were not flooded but located within the floodplain. I can think of properties located within the floodplain having the expected flood damage of $\pi^L\delta$ and flood-affected properties having the expected flood damage of $\pi^H\delta$. However, to simplify the notations, I normalize Δ such that the low flood risk π^L is included in this net flow utility. Then the flood risk π simply captures the increase in flood risk after the flood: $\pi = \pi^H - \pi^L$. There is uncertainty about the true long-run flood risk.

Households are categorized into three types: optimistic (o), skeptical (s), and vulnerable (v). They have diffuse priors about future flood risks. Initially, all households have risk belief of π^L. Upon observing a flood event, optimistic and vulnerable households do not expect fundamentals of flood-affected properties to worsen and believe the expected costs of flood risk are still low after the flood:

$$\mathbb{E}^o(\pi^*) = \mathbb{E}^v(\pi^*) = \pi^L.$$

Skeptical households believe that the flood risk is higher than previously assessed, holding higher belief on flood risk: $\mathbb{E}^s(\pi^*) = \pi^H$. Almost all households are initially vulnerable, but they can be persuaded by optimists or skeptics through social interaction. Each type of household j has an entropy level:

$$e^j = -\sum_x f^j(\theta_x) \log f^j(\theta_x) \tag{A.2}$$

where $f^j(\theta_x)$ is the belief distribution over θ^*. When two households meet, the probability that household i adopts household j's belief is as follows:

$$g_{ij} = \max\left(1 - \frac{e^j}{e^i}, 0\right) \tag{A.3}$$

Suppose a household with higher entropy meets a household with lower entropy. This probability of belief adoption implies that the high-entropy household is more likely to adopt the belief of the low-entropy household if their entropy ratio is larger, whereas it is impossible for a low-entropy household to adopt the belief of a high-entropy household. I assume that the entropy of vulnerable households exceeds the entropy of skeptical and optimistic households. I further assume that in the case of flood events, the entropy of optimistic households is greater than the entropy of skeptical households:

$$e^s < e^o < e^v.$$

The basic intuition is that households can socially interact and "infect" each other with their own beliefs on future flood risk, the dynamics that are similar to models of infectious diseases proposed by Bernoulli (1766). In Burnside, Eichenbaum and Rebelo (2016), the population dynamics are used to generate booms and busts in housing markets. In this paper, I show how social interaction of households and belief heterogeneity can generate gradual learning of flood risk and price adjustment.

Initially, almost the entire population starts as vulnerable agents. As households interact with each other and adopt views of those with lower entropy, the shares of skeptical and optimistic households both grow as vulnerable households adopt their views via social interaction. Because the pdf of the entropy of skeptical households is the lowest, the share of skeptical households continues to grow and exceeds $1 - k$ at $t = t_1$, at which the marginal buyer becomes a skeptical household.

Burnside, Eichenbaum and Rebelo (2016) use this model to illustrate how housing booms and busts can appear even without observable indicators. There is uncertainty about which households hold the correct beliefs, and this uncertainty alone can generate a "fad" as certain types of households with stronger beliefs convince others with weaker beliefs. I adopt this model to demonstrate how social interaction and belief heterogeneity can generate gradual price adjustment in the housing market.

I simulate the economy with a simple numerical example to illustrate how belief heterogeneity and social interaction can generate gradual learning. I choose the normalization of $\Delta = 1$. I assume that $\pi\delta = 0.1$, $e^o/e^v = 0.820$, and $e^s/e^v = 0.800$. Each time period represents one month, and β is such that the implied annual discount rate is 6 percent. Following Burnside, Eichenbaum and Rebelo (2016), I start with almost all vulnerable households and a very small number of optimistic and skeptical households at time 0. The details of price paths and share flows are provided in Appendix D.1.

Panel A of Figure A.14 shows the evolution of shares of household types. As households interact, vulnerable households are initially persuaded by both optimistic and skeptical households. Because the pdf of the entropy of skeptical households is the lowest, both optimistic and vulnerable households eventually adopt the views of skeptical households. The shares of vulnerable and optimistic households slowly decline over time and converge to zero. Because skeptical households value the properties lower than the other household types, it is not until the share of skeptical households reaches $1 - k$ that the marginal buyer becomes a skeptical household. In other words, as long as the combined share of optimistic and vulnerable households exceeds k, the marginal buyer is an optimistic/vulnerable household.

Let t_1 be the time period when the share of skeptical households first exceeds $1 - k$. After t_1, the marginal buyer is a skeptical household, and the equilibrium price is the price of skeptical households. Between 0 and t_1, the marginal buyer is still an optimistic/vulnerable household. The market equilibrium price P_t is then equal to the price of optimistic/vulnerable households P^o minus the discounted value of the capital losses from selling the house to an optimistic/vulnerable household.

The simple experiment illustrates a scenario in which households with differing views about flood risk slowly update their flood risk beliefs via social interaction. Although the market price initially starts close to the original price after the flood event, the price gradually declines as more households become pessimistic about future flood risk.

One limitation to the conceptual model is that the model remains stylized and does not integrate reference dependence into the dynamic decision-making process of heterogeneous agents. The frictionless belief-updating model captures how perceptions of flood risk evolve over time through social interactions, but it abstracts away from the seller's optimal timing behavior under reference-dependent preferences. A more comprehensive model would explicitly incorporate reference-dependent utility into a seller's dynamic optimization problem, accounting for how anchoring to nominal purchase prices influences both the timing of sale and pricing decisions.

Appendix D.1 Share Flows and Price Paths

In this section, I describe in detail the flow of the share of households for each type and the path of the market price for the conceptual model discussed in Section Appendix D.

Appendix D.1.1 Share Flows

Recall that the pdf of the skeptical households has lower entropy than that of the optimistic households. Then in this experiment, the shares of optimistic, skeptical, and vulnerable households evolve according to the following laws of motion:

$$o_{t+1} = o_t + g_{vo}o_t v_t - g_{os}o_t s_t, \tag{A.4}$$

$$s_{t+1} = s_t + g_{vs}s_t v_t + g_{os}o_t s_t, \tag{A.5}$$

$$v_{t+1} = v_t - g_{vs}s_t v_t - g_{vo}o_t v_t, \tag{A.6}$$

where g_{ij} is the probability that household i adopts household j's belief given in Equation (A.3). Equation (A.4) states that the share of optimistic households in the next period is the share of optimists in the current period, plus the share of vulnerable households who interact with optimistic households and become convinced with optimistic views, minus the share of optimists who interact with skeptical households and become convinced with skeptical views. Equation (A.5) states that the share of skeptical households in the next period is the share of skeptics in the current period, plus the share of vulnerable households who interact with skeptical households and become convinced with skeptical views, plus the share of optimists who interact with skeptical households and become convinced with skeptical views. Finally, Equation (A.6) states that the share of vulnerable households in the next period is the share of vulnerable households in the current period, minus the share of vulnerable households who interact with optimistic households and become convinced with optimistic views, minus the share of vulnerable households who interact with skeptical households and become convinced with skeptical views.

Appendix D.1.2 Price Paths

The equilibrium price path is given by

$$P_t = \begin{cases} P^o - \big[\beta(1-\phi)\big]^{t_1-t}(P^o - P^s) & \text{if } t < t_1 \\ P^s & \text{if } t \geq t_1. \end{cases} \tag{A.7}$$

From period 0 until period t_1 when the marginal household becomes the skeptical household ($s_t >= 1 - k$), the equilibrium price is equal to the price of optimists minus the discounted expected value of the losses from selling the property to an optimistic or vulnerable household at time t_1. The equilibrium price becomes P^s after $t = t_1$ when the marginal household becomes the skeptical household.

The equilibrium price path resembles the path described in Proposition 2 of Burnside, Eichenbaum and Rebelo (2016). Two details are different. First, instead of optimistic households holding optimistic views on fundamentals of properties and the other types of households holding the same views on the fundamentals, this study examines the case in which skeptical households hold skeptical (pessimistic) views on fundamentals of properties, expecting higher flood risk in the future. Second, because skeptical households value the properties lower than the rest, it is not until the share of skeptical households reaches $1 - k$ that the marginal buyer becomes a skeptical household. In other words, as long as the combined share of optimistic and vulnerable households exceeds k, the marginal buyer is an optimistic/vulnerable household.

Appendix E Additional Figures and Tables

List of Figures in Appendix

List of Tables in Appendix

Figure A.1: Storm and flood insurance adoption

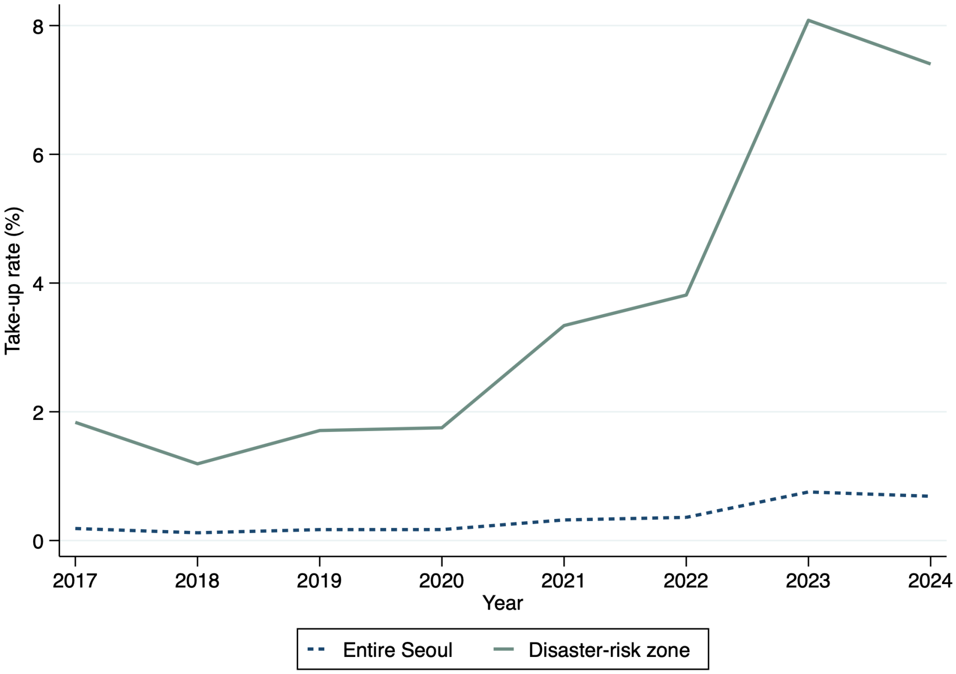

Note: This figure plots the adoption rates of flood insurance policies. The dashed line reports the fraction of all households in Seoul that hold flood insurance. The solid line reports the fraction of households with flood insurance relative to the number of households located within the officially designated "disaster-risk zone" in Seoul. For this calculation, I make the conservative assumption that all insured households are located within the disaster-risk zone.

Figure A.2: Event Study Estimates of Indirect Flood Effects on Real Sales Prices

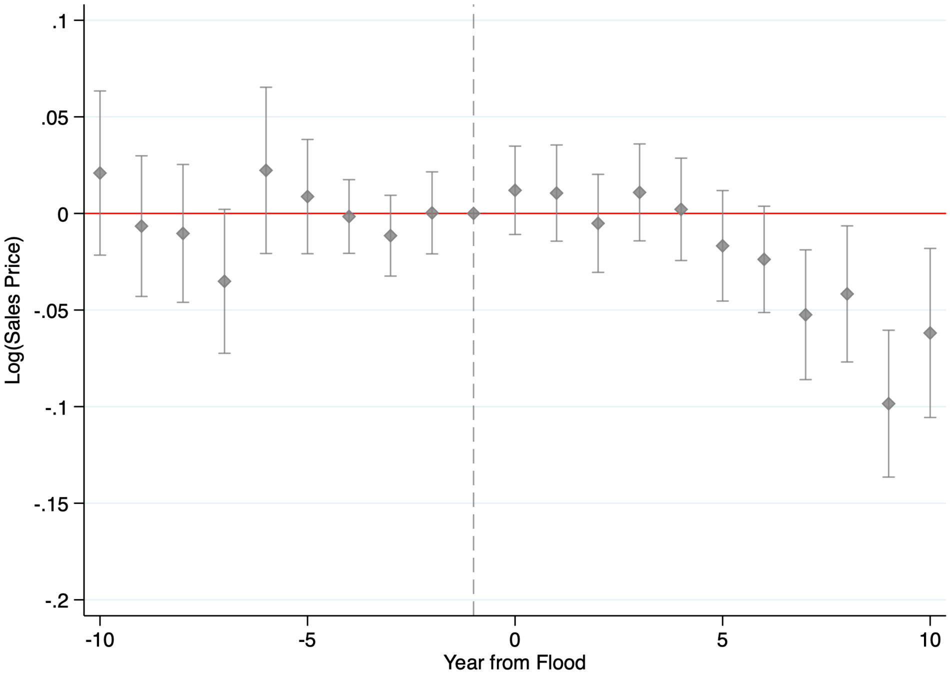

Note: This figure plots the coefficients and 95% confidence intervals from event-study regressions of flooding on logged sales price of the upper-floor properties located within the floodplain using the Callaway and Sant'Anna (2021) doubly robust estimator with inverse probability weighting.

Figure A.3: Event Study Estimates of Direct Flood Effects on Real Sales Prices

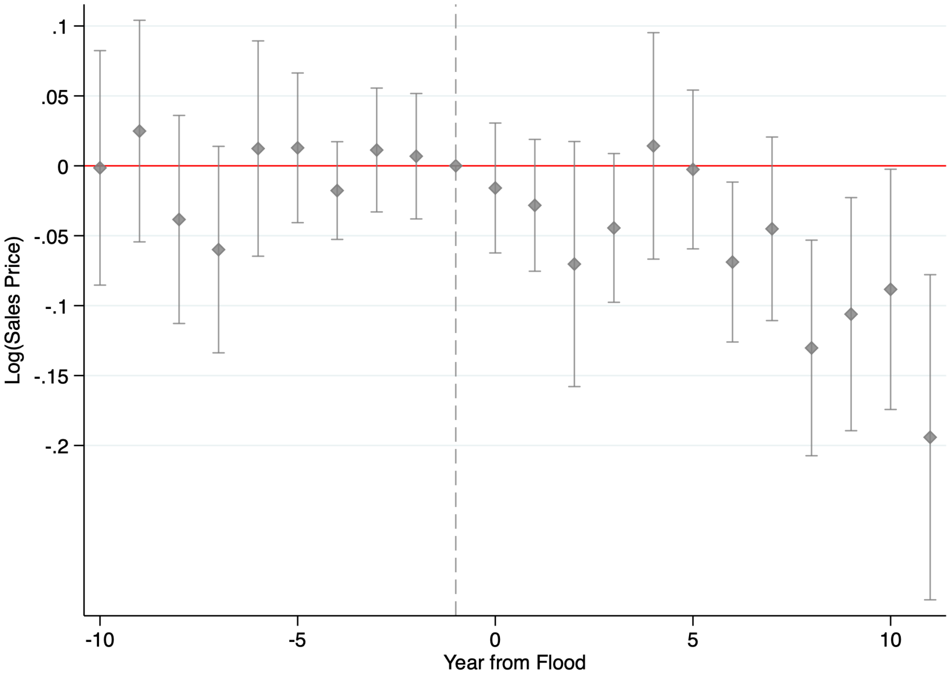

Note: This figure plots the coefficients and 95% confidence intervals from event-study regressions of flooding on logged sales price of the ground-floor properties located within the floodplain using the Callaway and Sant'Anna (2021) doubly robust estimator with inverse probability weighting.

Figure A.4: Event Study Estimates of Indirect Flood Effects on Combined Rent Prices

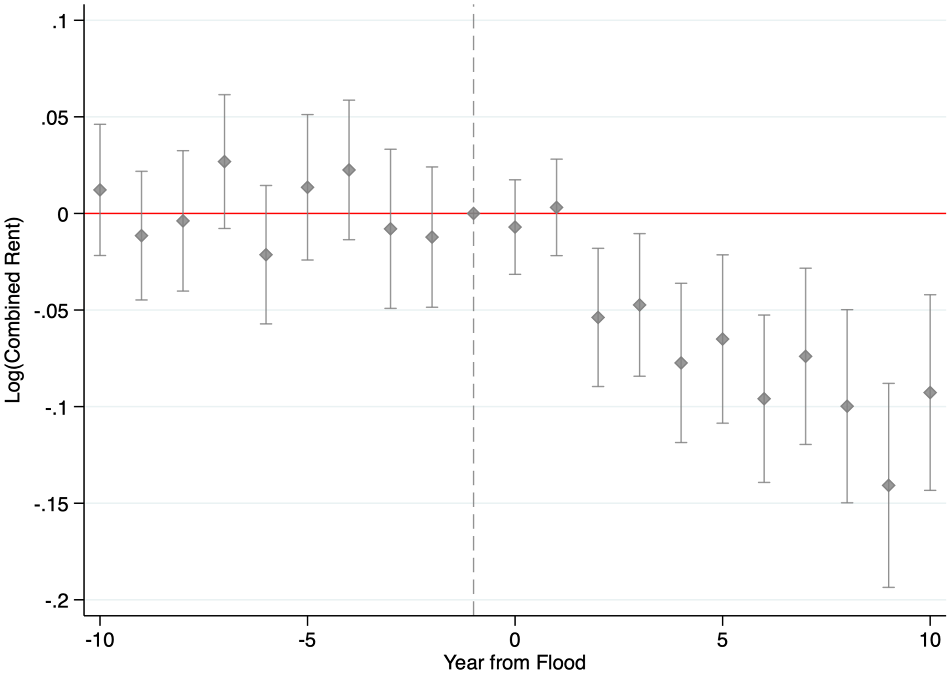

Note: This figure plots the coefficients and 95% confidence intervals from event-study regressions of flooding on the log of combined rental price of the upper-floor properties located within the floodplain using the Callaway and Sant'Anna (2021) doubly robust estimator with inverse probability weighting.

Figure A.5: Event Study Estimates of Direct Flood Effects on Combined Rent Prices

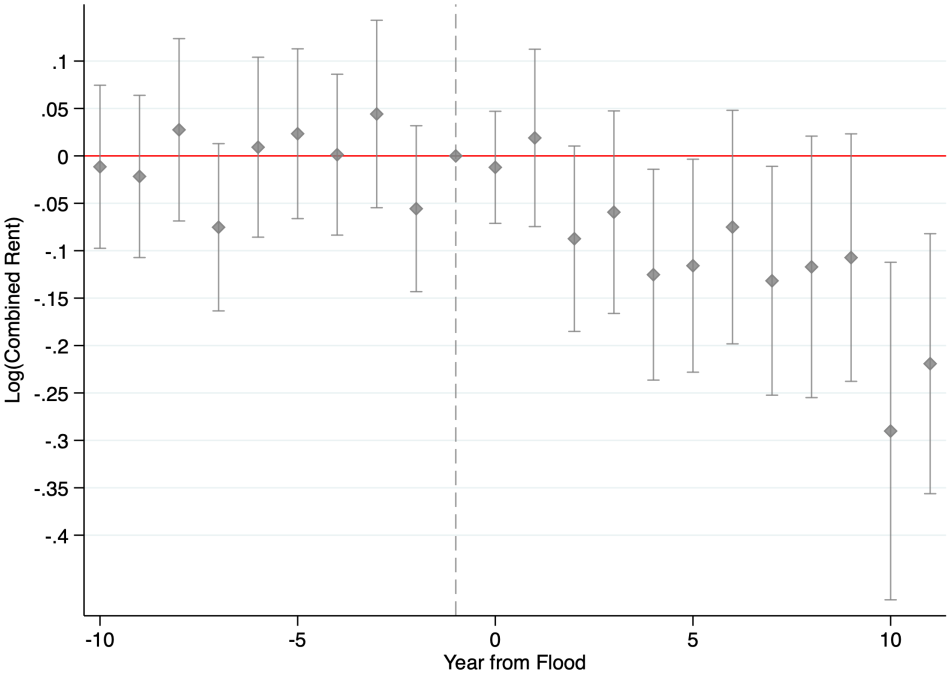

Note: This figure plots the coefficients and 95% confidence intervals from event-study regressions of flooding on the log of combined rental price of the ground-floor properties located within the floodplain using the Callaway and Sant'Anna (2021) doubly robust estimator with inverse probability weighting.

Figure A.6: Difference in Flood Effects between Sales and Lease

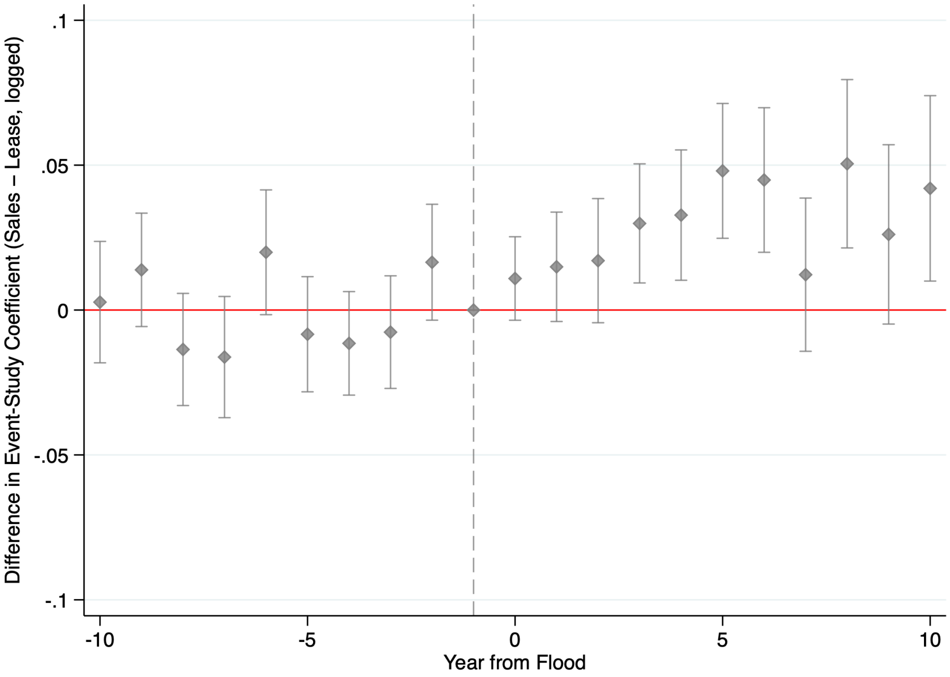

Note: This figure plots the difference in event-study coefficients between the sales market and the lease market. Standard errors are obtained by clustered bootstrapping at the property level. The event-study regressions for the sales market and the lease markets are pooled at the level of re-centered influence functions produced from the event-study estimations of Callaway and Sant'Anna (2021).

Figure A.7: Event Study Estimates of Indirect Flood Effects on Sales Transaction Counts

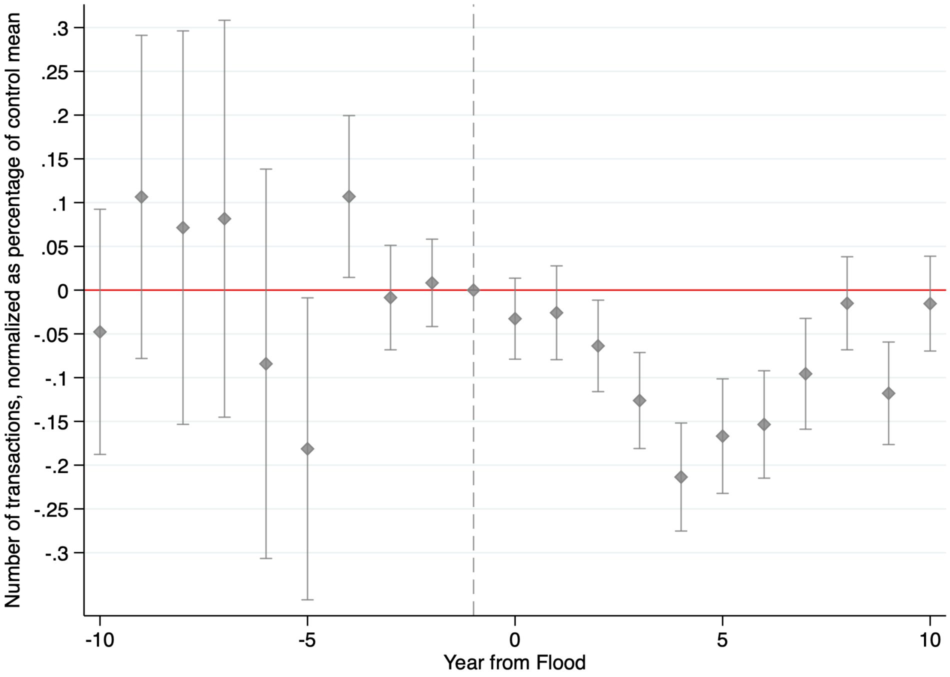

Note: This figure plots the coefficients and 95% confidence intervals from event-study regressions of flooding on logged sales transaction counts of the upper-floor properties located within the floodplain using the Callaway and Sant'Anna (2021) doubly robust estimator with inverse probability weighting.

Figure A.8: Event Study Estimates of Direct Flood Effects on Sales Transaction Counts

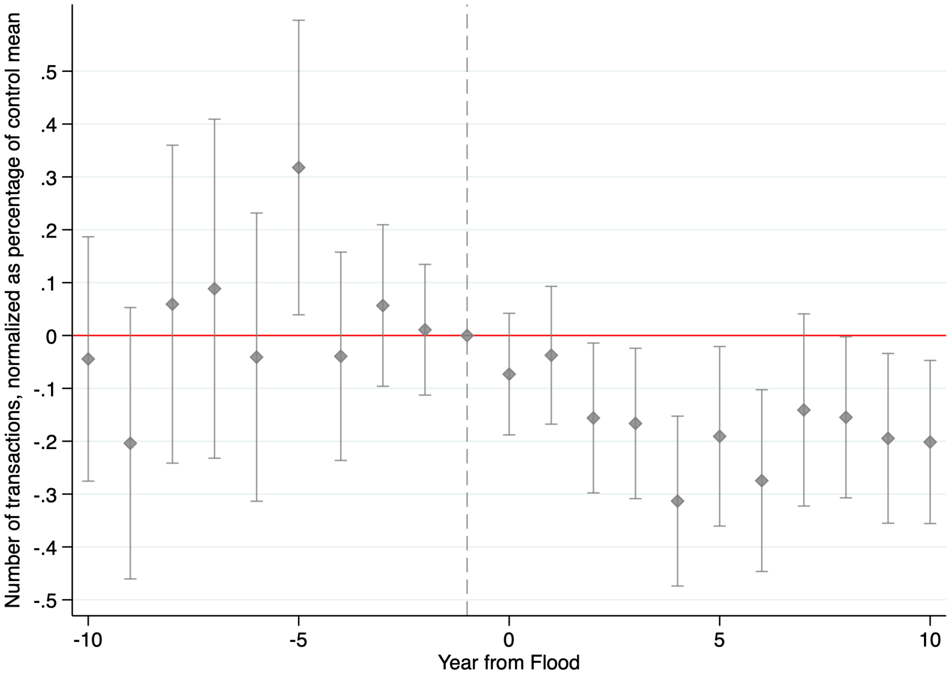

Note: This figure plots the coefficients and 95% confidence intervals from event-study regressions of flooding on logged sales transaction counts of the ground-floor properties located within the floodplain using the Callaway and Sant'Anna (2021) doubly robust estimator with inverse probability weighting.

Figure A.9: Event Study Estimates of Indirect Flood Effects on Rent Transaction Counts

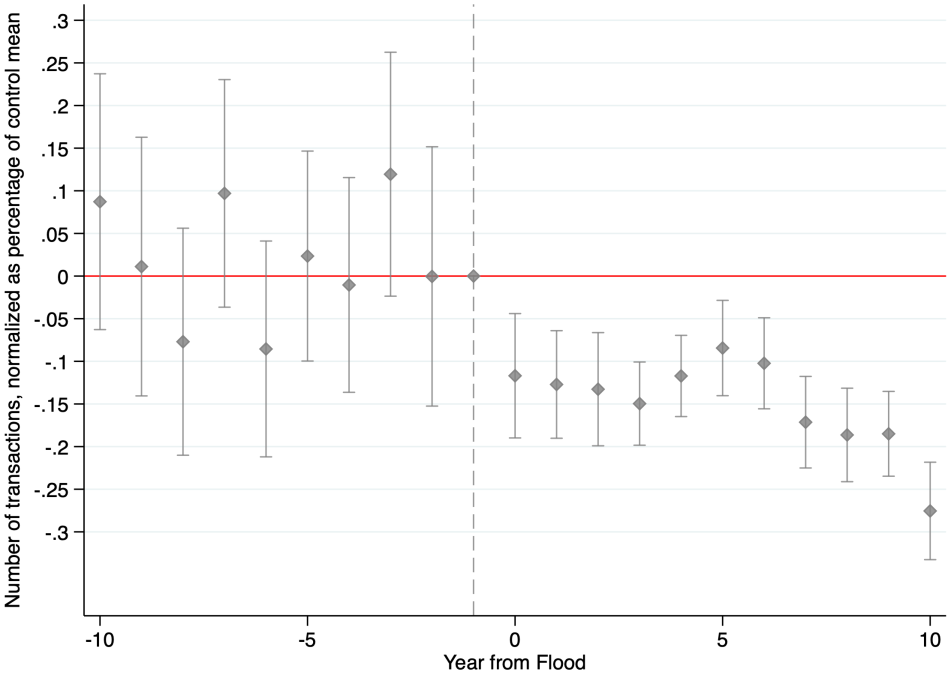

Note: This figure plots the coefficients and 95% confidence intervals from event-study regressions of flooding on the log of lease transaction counts of the upper-floor properties located within the floodplain using the Callaway and Sant'Anna (2021) doubly robust estimator with inverse probability weighting.

Figure A.10: Event Study Estimates of Direct Flood Effects on Rent Transaction Counts

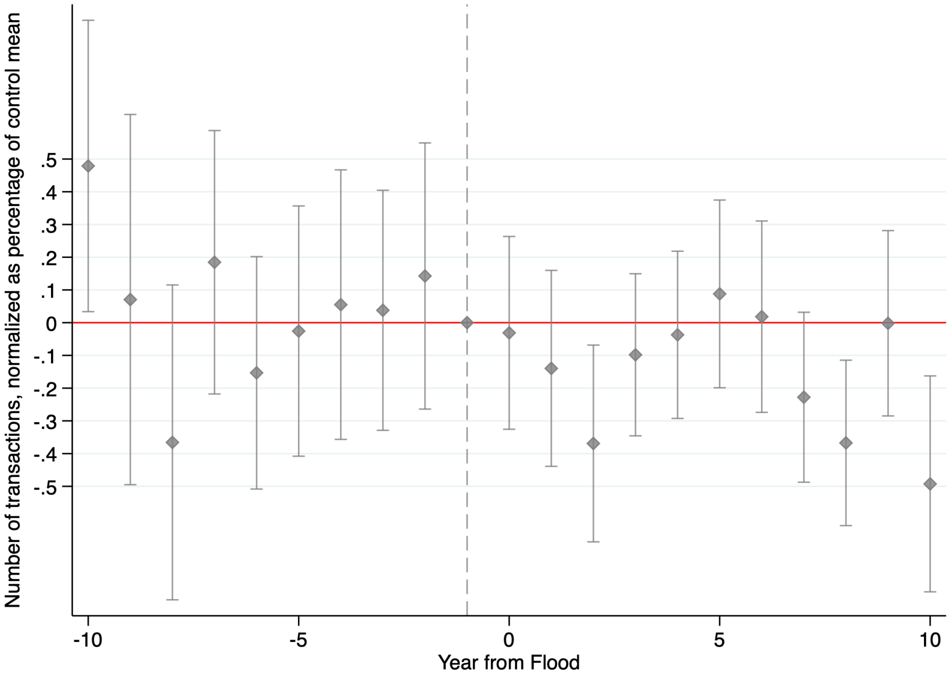

Note: This figure plots the coefficients and 95% confidence intervals from event-study regressions of flooding on the log of lease transaction counts of the ground-floor properties located within the floodplain using the Callaway and Sant'Anna (2021) doubly robust estimator with inverse probability weighting.

Figure A.11: Robustness Check of Counterfactual Gains in Figure 6

Panel A: Binned frequencies across nominal gains

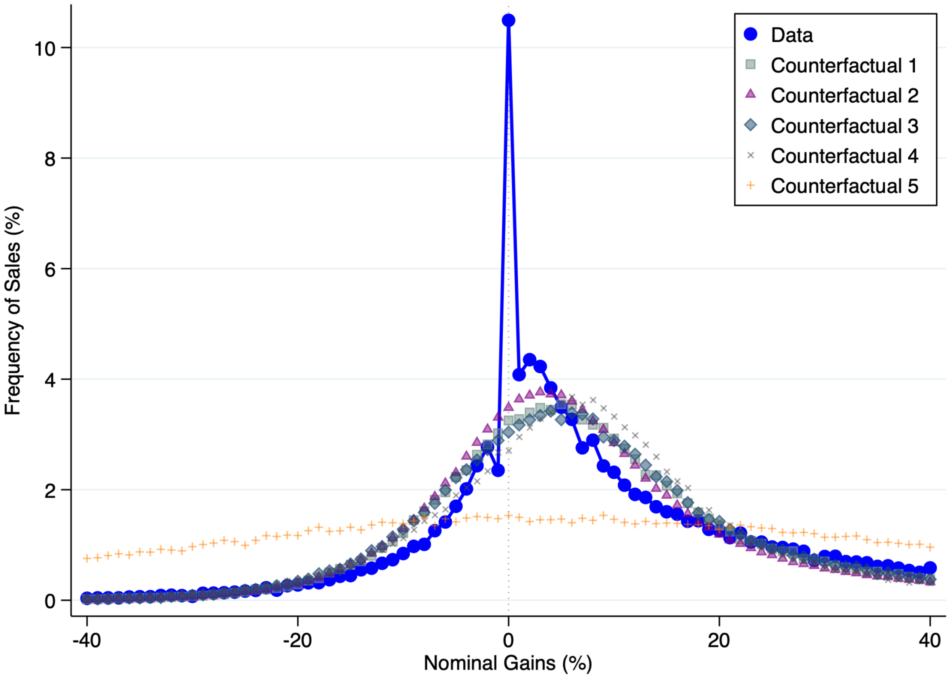

Panel B: Excess mass across gains

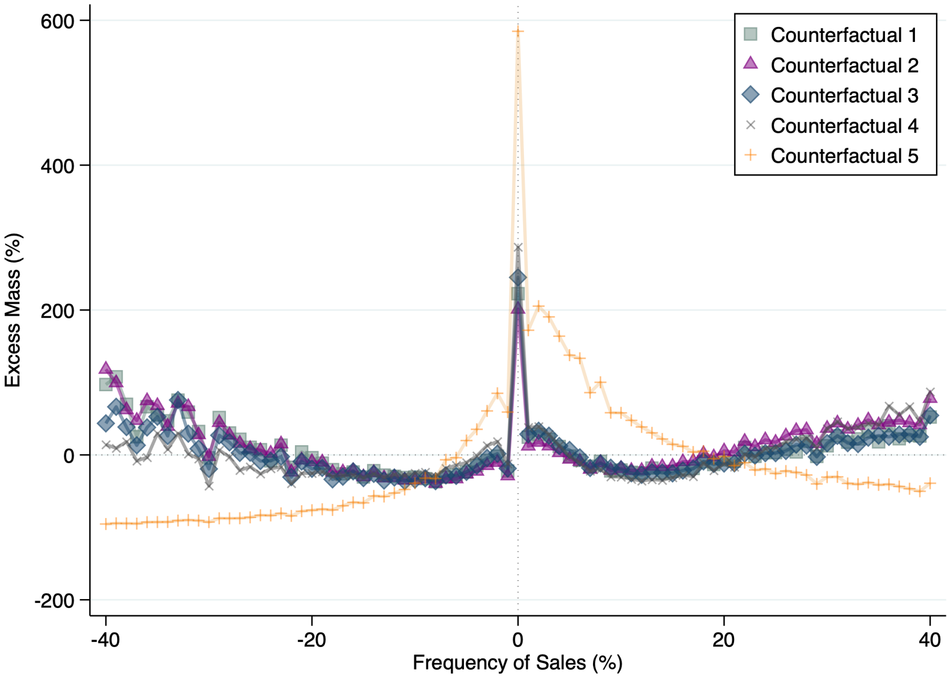

Note: Panel A plots binned frequencies of sales transactions in one percentage-point steps across realized nominal gains. The counterfactual gains are obtained using the hedonic pricing models. Model 1 is the regression model in Equation 3. Model 2 is the same regression model including all the properties in Seoul. Model 3 includes only the assessed value and year and month fixed effects. Model 4 includes only the assessed value. Model 5 is Model 1 without the assessed value. Panel B plots excess mass of transactions relative to the level of the counterfactual.

Figure A.12: Robustness Check of Counterfactual Gains in Figure 7

Panel A: Binned frequencies across nominal gains

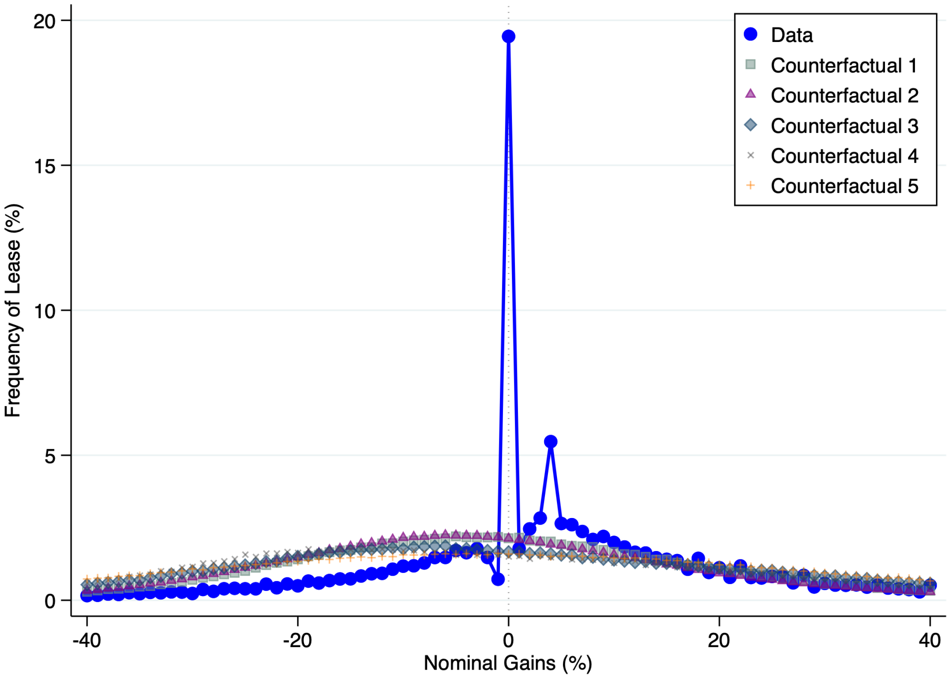

Panel B: Excess mass across gains

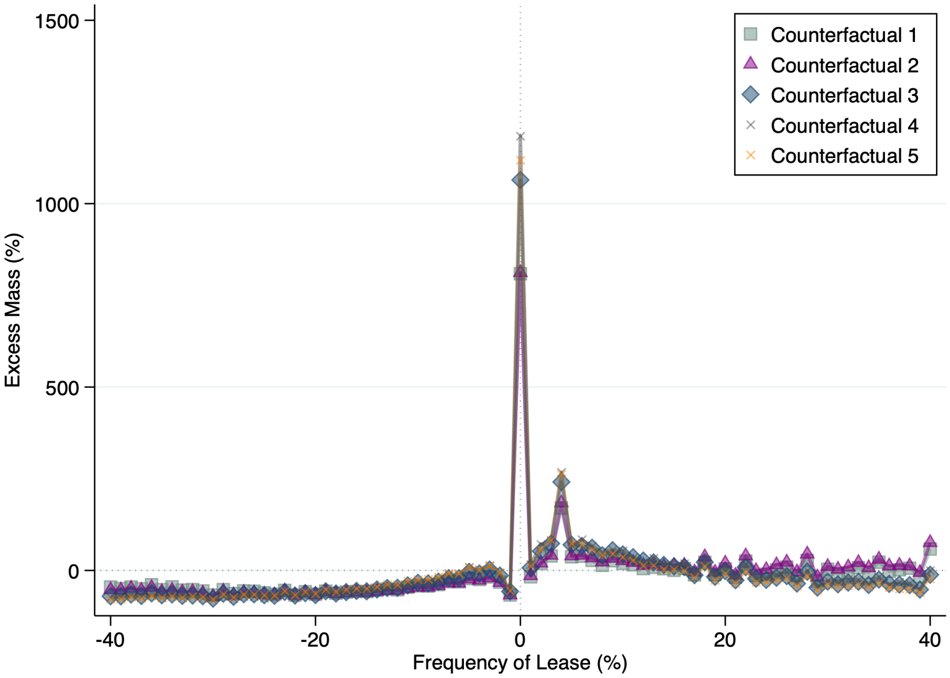

Note: Panel A plots binned frequencies of lease transactions in one percentage-point steps across realized nominal gains. The counterfactual gains are obtained using the hedonic pricing models. Model 1 is the regression model in Equation 3. Model 2 is the same regression model including all the properties in Seoul. Model 3 includes only the assessed value and year and month fixed effects. Model 4 includes only the assessed value. Model 5 is Model 1 without the assessed value. Panel B plots excess mass of transactions relative to the level of the counterfactual.

Figure A.13: Robustness Check: Probability of Resale After Flood

Panel A: Sales

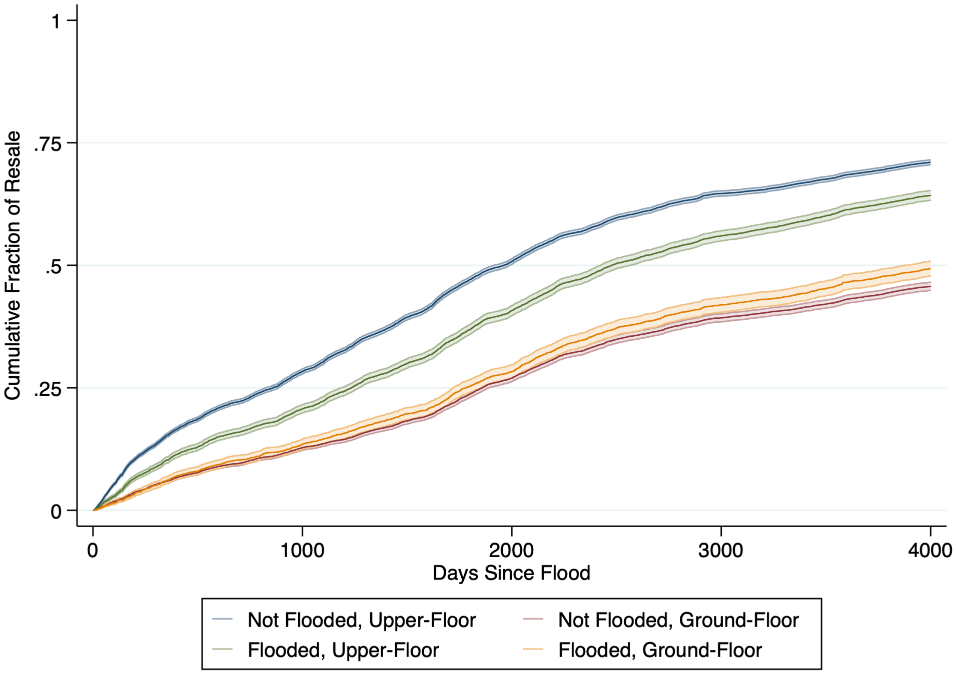

Panel B: Lease

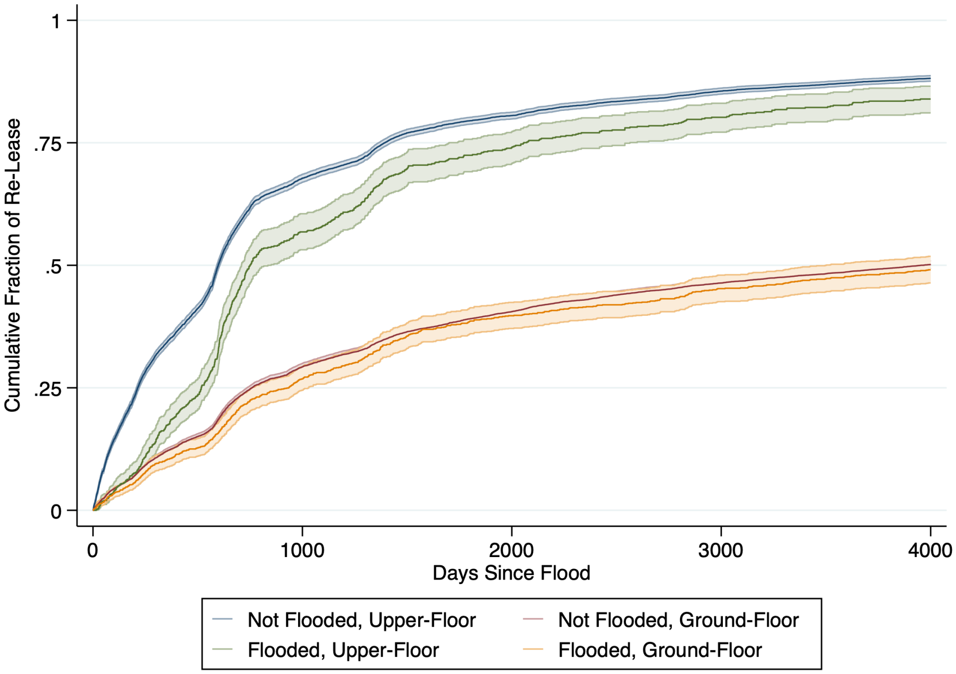

Note: The figure shows Kaplan–Meier survival curves that plot the probability that a property remains unsold over time since the flood event. The figure shows the curves separately for upper-floor properties and ground-floor properties. Panel A plots the cumulative fraction of properties that have been resold as a function of time since the flood. Panel B plots the cumulative fraction of properties that have been re-leased as a function of time since the flood.

Figure A.14: Share Flows and Price Path

Panel A: Share Flows

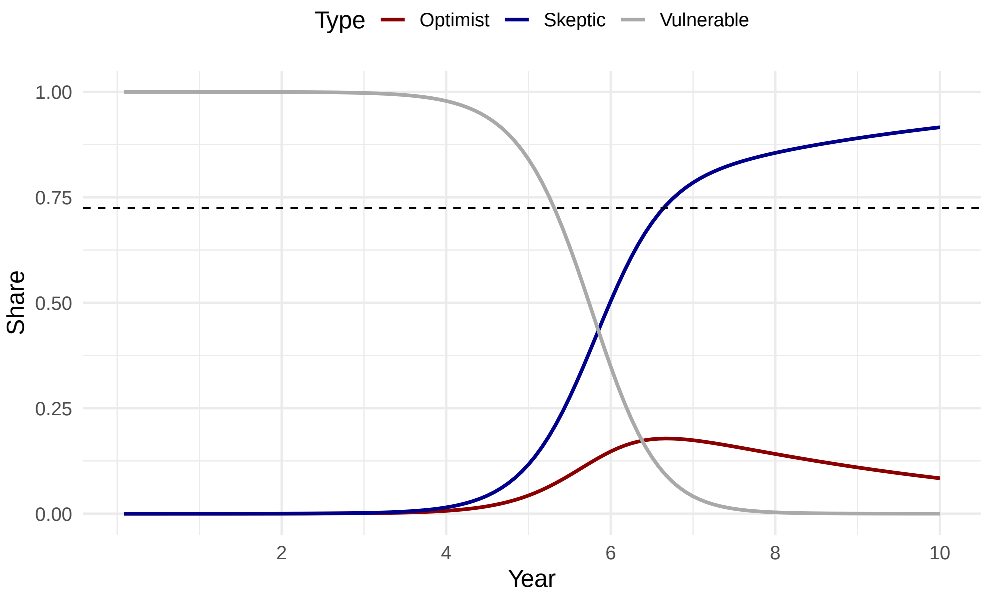

Panel B: Market Price

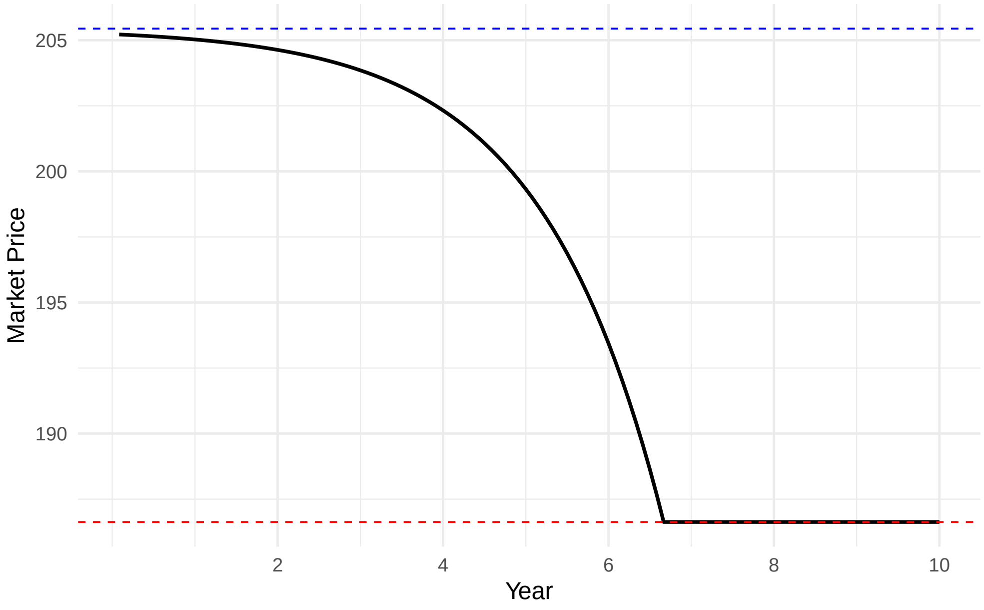

Note: The figure shows equilibrium paths of the conceptual model of social interaction and belief heterogeneity discussed in Section Appendix D. Panel A plots shares of household types. The dotted line indicates the share of renters, $1-k$. Panel B plots the path of the market price under the assumption that the uncertainty is not realized, which is computed using Equation (A.7). The blue and red dotted lines indicate the price for optimistic and skeptical households absent uncertainty, respectively.

Table A.1: Wald tests on dynamic effects between sales and lease

Post-event period τ	Wald χ^2	p-value
A. Individual Wald tests: $\beta_\tau^S = \beta_\tau^L$		
0	1.832378	.175847
1	1.968663	.1605898
2	1.101761	.2938799
3	5.919569	.0149736
4	5.989355	.0143925
5	10.63693	.0011085
6	7.673794	.0056029
7	.0261121	.8716268
8	6.228737	.0125693
9	.2585074	.6111475
10	2.541274	.1109053
B. Joint Wald tests: $\beta_\tau^S = \beta_\tau^L \;\; \forall \tau \in \mathcal{T}_{\underline{\tau},\bar{\tau}}$		
Early post-event periods (0–5)	13.09286	.0415848
Later post-event periods (6–10)	17.64569	.0034249
All post-event periods (0–10)	28.02767	.0032059

Note: The Wald tests are conducted on pooled datasets for the sales market and the lease market, at the level of re-centered influence functions. The first column shows the null hypothesis for the Wald test. For example, the first row shows the result of the Wald test with $H_0 : \beta_0^S = \beta_0^L$. The last row shows the result of the joint Wald test with $H_0 : \beta_\tau^S = \beta_\tau^L \;\; \forall \tau \in \{0, \ldots, 10\}$.