Leveraging the Microgeography of Neighborhoods to Estimate the Causal Effects of Flood Damage: Evidence from Hurricane Harvey

Wesley Austin Miller*

February 7, 2025

Abstract

The combination of climate change and economic growth in disaster-prone regions increasingly exposes the population to extreme weather events. The responses and resilience of households are key determinants of the social cost of climate change. By exploiting the microgeography of neighborhoods and exposure to flooding, I estimate the average treatment effects of household-level damage using a fuzzy regression discontinuity design. I leverage the relationship between a home's elevation and flooding during Hurricane Harvey, exploiting a discontinuous increase in damage from $0 to approximately $48,000 once water reaches the first floor of the structure. I provide novel evidence that disaster damage increases homeowner relocation while delaying property sales. The impacts on residential mobility attenuate over time, but I document a persistent divergence in the location and type of housing selected by damage-induced movers. My results indicate that flood damage makes people more likely to move shorter distances, and damage-induced movers are more likely to transition out of homeownership. Despite the combined shock to shelter and wealth, I find that flooded households are more likely to sort into higher-income census tracts. By evaluating the impact of direct household damage instead of neighborhood-level measures, this study provides new insights into how disaster damage affects residential choices.

* University of Southern California - Lusk Center for Real Estate, millerwe@usc.edu

1 Introduction

The combination of climate change and economic growth in disaster-prone regions increasingly exposes the United States population to extreme weather events, with the number of inflation-adjusted billion-dollar natural disasters nearly doubling during the 2010s (NOAA, 2023). The immediate cost of natural disasters is driven by economic disruption and property damage (Smith and Katz, 2013). Housing is the primary asset for a majority of US households, making residential property damage a mixture of shelter and wealth shocks. The individual responses to this blended shock are key determinants of the social cost of climate change. By exploiting the microgeography of housing markets and exposure to flooding, I leverage a natural experiment to identify the causal effect of disaster damage on residential choices for those who live at “ground zero” but who differ in terms of realized property damage.

Housing choices are central to individual well-being. The decision to remain in a home or move to a new location could impact access to certain labor markets, health care systems, school districts, and other amenities. Movers must decide where to relocate, and the characteristics of their neighborhood may affect themselves or other members of their household through environmental or peer effects (Chyn and Katz, 2021). They must also decide whether to rent or purchase their new residence, an investment decision that may impact wealth creation and intergenerational mobility.

Theoretically, disasters affect housing choices through both demand and supply channels. The destruction of local amenities as well as increased risk salience (or “new news”) pulls down the demand to reside in impacted areas. Simultaneously, disasters decrease the supply of inhabitable housing. Residents of these damaged units are forced to find alternative shelter, prompting various degrees of relocation. For homeowners, relocation decisions may coincide with repairs and homeselling considerations. These supply and demand dynamics evolve according to a multitude of factors including the magnitude of disaster aid and relief. Empirical evidence of the general equilibrium effects of disasters indicate higher residential mobility coupled with stagnation in housing market transactions (Boustan et al., 2020; Zivin et al., 2023).

It is difficult, however, to distinguish the causal effect of disaster damage from the regional economic shock due to the lack of administrative datasets detailing household-level exposure. Moreover, damage is endogenous to a variety of factors including individual risk preferences and mitigation measures as well as local investment in infrastructure and emergency preparations. These correlates may introduce selection bias into estimates of average treatment effects on individual outcomes. The literature addresses these challenges by comparing outcomes of all individuals in a disaster-struck area with outcomes of observably similar individuals in unaffected places (Vigdor, 2007; Groen and Polivka, 2008; Deryugina et al., 2018; Boustan et al., 2020). Another strand of the literature restricts analyses *within* a disaster area, relying on local variation in damage intensity to identify causal effects (Gallagher and Hartley, 2017; McCoy and Walsh, 2018; Bleemer and van der Klauww, 2019; Billings et al., 2022; Gallagher et al., 2023). These techniques, however, estimate the average treatment effect of living in damaged areas rather than the average treatment effect of damage itself.

In this paper, I estimate the average treatment effects of household-level damage using a fuzzy regression discontinuity design in the context of Hurricane Harvey, which damaged more than 200,000 homes in Houston, Texas in 2017. I leverage the relationship between flooding and a home’s elevation, exploiting a discontinuous increase in damage from $0 to approximately $48,000 once water reaches the first floor. I follow Dong and Kolesar (2023) and implement a donut design, allowing for valid inference despite measurement error in the running variable. To my knowledge, I offer the first causal estimates of household-level damage on individual outcomes.

I provide evidence of the hyper-local nature of disaster damage and response. In particular, I construct a household-level damage variable and exploit elevation differences between homes located in the same subdivisions, finding substantial variation in damage that causes a divergence in neighborhood and housing choices. My household-level results bolster the findings in the literature that leverages aggregate damage as the level of treatment.

While residential mobility spiked after Hurricane Harvey, Billings et al. (2022) and Gallagher et al. (2023) document little difference in out-of-Houston migration rates for those living in flooded and non-flooded census blocks. Similarly, I find parallel trends in housing transactions across damaged and non-damaged subdivisions, illustrating how local-area damage does not drive aggregate movement patterns. The lack of correlation with household movements is surprising given the established connection between local damage intensity and other individual outcomes in the literature.

My analysis, however, highlights the importance of *household-level* damage in individual decision making despite the lack of association at the local level. I find that exposure to property damage makes homeowners less likely to sell their homes relative to their non-flooded neighbors for at least two years. Homeowners on the margin of selling may delay that decision as they repair their homes. Moreover, post-disaster construction-labor shortages can prolong rebuilding efforts, and the full extent of disaster aid and insurance payments often takes months to materialize.

Despite decreasing housing transactions, I estimate that damage induces homeowners to move and relocate in the initial stages of the recovery process. This six-month short-run effect is dominated by local moves within the Houston metropolitan area. The estimated effect on move propensity attenuates over the recovery period, but I document a persistent divergence in the distance and location of post-storm residential choices. My results indicate that flood damage makes people less likely to move long distances, translating to more within-county moves and fewer out-of-county relocations. In addition to altering locational choices, I provide evidence that household-level flood shocks decrease homeownership.

The effects on homeownership are particularly striking given that home loans through the Small Business Administration (SBA) are the federal government's dominant form of individual-level assistance after natural disasters (Collier and Ellis, 2021). Billings et al. (2022) discuss the qualification hurdles for SBA loans and how their regressive allocation may limit their effectiveness to the marginal homeowner exposed to damage. This limitation may explain the damage-induced decrease in homeownership after Hurricane Harvey.

Despite the combined shock to shelter and wealth, I find that flooded households are more likely to sort into higher-income census tracts in the aftermath of Hurricane Harvey. Conditional on moving, flood damage is associated with the consumption of newer and more expensive homes. This relative improvement in physical and socioeconomic environments mirrors the long-run recovery patterns documented in the disaster literature (Sacerdote, 2012; Deryugina et al., 2018; Deryugina and Molitor, 2020). Given the importance of place-based spillover effects, the long-run welfare implications of suffering flood damage remain an open area of research.

The rest of my paper is structured as follows. Section 2 summarizes the natural disaster literature and how individuals respond to environmental shocks. Section 3 offers an overview of Hurricane Harvey and introduces the area of Houston that I analyze. Section 4 summarizes the sources of data and the analytical sample. Section 5 presents my empirical strategy, and Section 6 explores the validity of my approach. The main results are provided in Section 7, and I summarize the overall contribution of this project in Section 8.

2 Literature Review

Advances in quasi-experimental methods spurred a renaissance in the disaster literature that began with analyses of the victims of Hurricane Katrina. Hurricane Katrina survivors suffered in the short run, but negative average effects dissipated over time across a range of outcomes, including mortality (Deryugina and Molitor, 2020), employment, and earnings (Vigdor, 2007; Groen and Polivka, 2008; Deryugina et al., 2018). These papers all use some variation of a difference-in-differences framework, where the treatment group is a subset of individuals living in areas that were directly impacted by Hurricane Katrina and the control group consists of observably-similar individuals elsewhere in the United States. The estimands in these studies are generally interpreted as the average treatment effect of living in a disaster region, which reflects a weighted average across the spectrum of individual-level disaster exposure. For example, this weighted average includes the response of renter-occupied households as well as homeowners whose property may or may not have been damaged. This average treatment effect is particularly informative for evaluating community-wide impacts and recovery efforts across a region.

A parallel group of papers analyze the impacts of Hurricane Katrina by restricting comparisons of individuals living within the disaster area. Sacerdote (2012) documents a short-run decline and commensurate rebound in academic performance for Louisiana students who evacuated compared to their non-evacuating peers. Gallagher and Hartley (2017) leverage local flood variation in New Orleans and find that residents in the most flooded census blocks experience short-term spikes in financial distress relative to those in less exposed areas. Bleemer and van der Klaauw (2019) extend this analysis by examining how housing choices and household composition vary by census-block flood intensity. They find an immediate increase in move propensity that is persistently positive for more than a decade. Homeowners in New Orleans' inundated census blocks were 10 percentage points less likely to own a home by the end of the ten-year period.

More recently, a pair of papers has applied the census-block flood intensity approach to identify the impacts of Hurricane Harvey on individuals in Houston, Texas. Using hydrologic data from the Federal Emergency Management Agency (FEMA), Billings et al. (2022) calculate the weighted average flood depth of developed land for each census block to explore differential responses to damage based on access to insurance and credit. They find no signs of systematic financial distress for households living in the 100-year floodplain, an area where flood insurance takeup rates are higher. Households who live in flooded census blocks *outside* of the 100-year floodplain experience disproportionate increases in delinquent debt and bankruptcy. The lack of flood insurance in these areas makes residents more reliant on disaster assistance, but the allocation of these transfers is regressive and fails to counteract initial inequalities in financial health. Gallagher et al. (2023) apply the same empirical strategy to analyze investment in human capital, finding a reduction of student loans for college-aged individuals living in flooded census blocks relative to their peers in other parts of Houston. These effects are prevalent in areas with higher levels of homeownership, suggesting a relationship between disaster damage and consumption and investment decisions.

The emphasis on census-block-level flooding illustrates the intuitive connection between disaster damage and individual responses. While all individuals in a disaster area may be indirectly impacted by the aggregate economic shock, the first-order concern for most disaster recovery and relief efforts is to address the direct effects of these events. The identifying assumption of the flood-intensity approach relies on changes in outcomes of individuals in less-flooded census blocks providing a counterfactual to changes in outcomes for those in more-flooded census blocks (conditional on observable geospatial and socioeconomic factors). The literature supports this assumption by documenting the low explanatory power of pre-disaster census-block characteristics on flood intensity. The absence of correlation in aggregate data, however, does not imply individual-level exogeneity of flood exposure, and there may still be concerns of household selection across or within census blocks based on flood risk or unobservable factors.

While census blocks are the smallest geographic unit defined by the Census Bureau, they are typically delineated by physical features rather than by the characteristics of inhabitants. I contribute to this literature by restricting comparisons to people living in the same subdivisions, a locally-defined unit that is more homogenous across individual-level characteristics.[1] Moreover, I leverage the quasi-random peak water level reached in these subdivisions during Hurricane Harvey, allowing for the identification of the average treatment effects of household-level disaster damage.

3 Hurricane Harvey and Houston, Texas

Hurricane Harvey landed in Texas in August 2017, producing record-level rainfall that inundated more than 200,000 homes and caused $125 billion in direct damage. Figure 1 illustrates the widespread flooding across the Houston metropolitan area by mapping the location of emergency rescue requests throughout the region.

Figure 1: Hurricane Harvey Rescue Requests

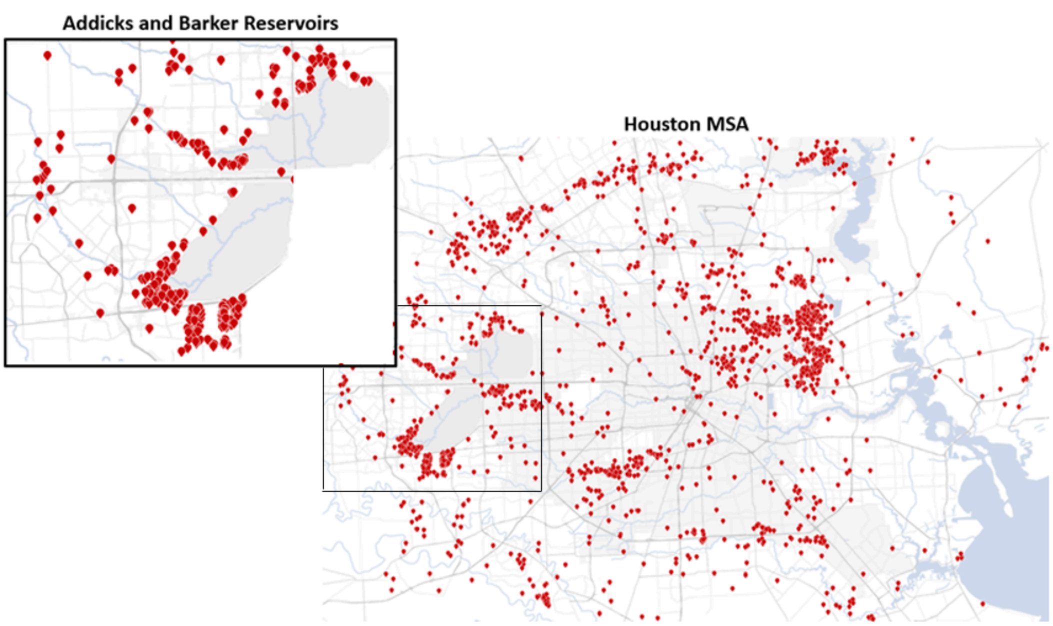

Notes: The zoomed-in box displays the area surrounding the Addicks and Barker Reservoirs (shaded in dark gray).
Source: New York Times (2017)

[1] See Table 2 for a comparison of within-group variation of housing characteristics for subdivisions and census blocks.

One of the major empirical challenges in the natural disaster literature is overcoming the endogeneity of damage. In the context of flooding, risk-averse individuals may select away from water sources or implement mitigation measures (e.g., the installation of flood vents). Others may value water as an amenity and prefer to live in close proximity to the natural resource, increasing their risk of flooding. Endogeneity concerns extend beyond individuals, as community-level infrastructure investment and maintenance are important determinants of local damage.

The Addicks and Barker Reservoirs highlighted in Figure 1 serve as Houston's primary flood infrastructure, as their earthen dams prevent upstream runoff from inundating the city center during storms. The Army Corps of Engineers developed the reservoirs in the 1940s by constructing the dams and acquiring 25,000 acres of abutting land. This government-owned land is designed to temporarily detain rainfall and allow for controlled drainage into the Gulf of Mexico (Furrh and Bedient, 2023).

Unlike lake-forming reservoirs, the government-owned land is perennially dry outside of extreme weather and is utilized as wooded parks, athletic fields, and other alternative uses, effectively masking its flood risk to surrounding suburbs. Moreover, the reservoirs are constructed as giant detention basins, creating a unique flooding mechanism that further distorts the salience of flood risk. While inland flood exposure typically depends on proximity to flowing water, Panel A of Figure 2 shows how the reservoirs create a pool based on a tub concept, where water rises uniformly within the basin irrespective of the location of water flow.

Panel B illustrates how Hurricane Harvey's unprecedented rainfall filled the Addicks and Barker Reservoirs, forcing water above the government-owned land for the first time and flooding thousands of homes. These homes lie outside of the 100-year floodplain, a classification that serves as the primary flood-risk signal in housing markets. I leverage the relatively unknown risk and unique flooding mechanism of the Addicks and Barker Reservoirs to circumvent selection issues that challenge the identification of causal effects of disaster damage on individual outcomes.

Figure 2: Flooding Mechanism

Panel A: General Tub Concept

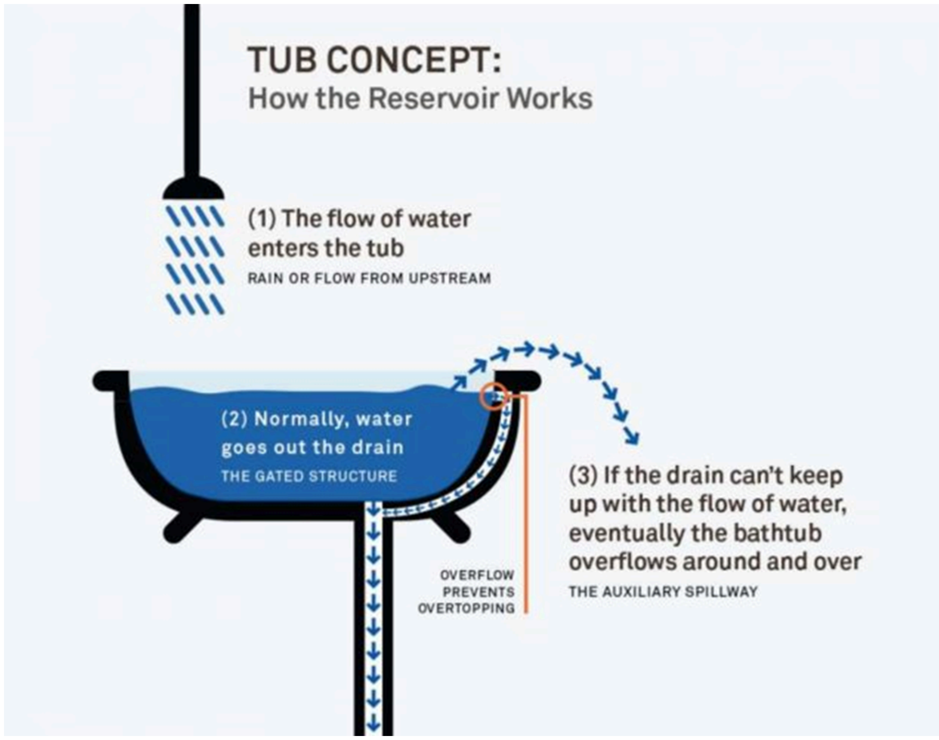

Panel B: Addicks and Barker Reservoir Overview

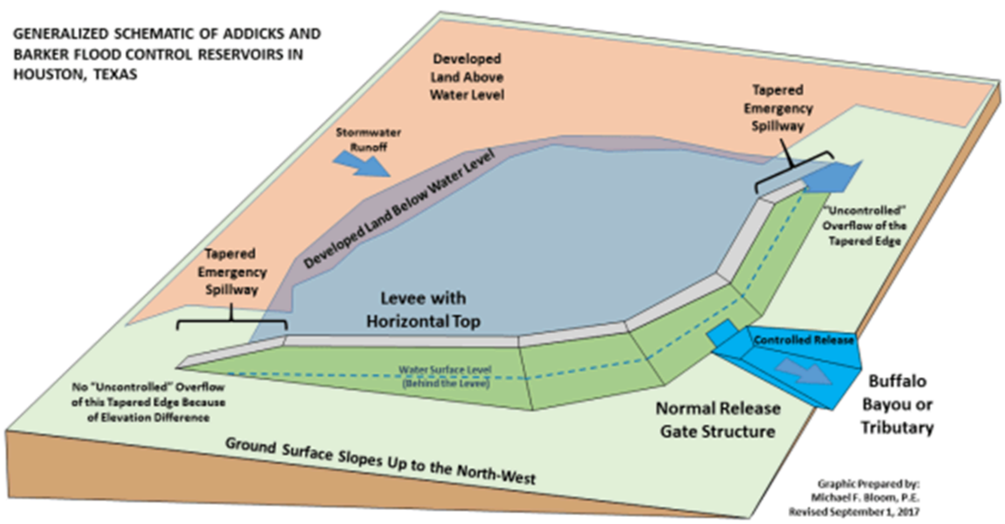

Source: O'Neil (2020) and Bloom (2017).

4 Data Sources and Sample Construction

I begin with the universe of owner-occupied housing units that were located in the Addicks and Barker Watersheds at the time of Hurricane Harvey.[2] Residential property data are obtained from the Harris and Fort Bend central appraisal districts and include the names of homeowners, occupancy status, property values, and physical housing characteristics (e.g., home size and year of construction).[3]

Table 1 summarizes these characteristics within the two watersheds, where there are roughly 100,000 homes distributed across 2,200 subdivisions. These homes were typically constructed in the mid-1990s, but homes in the Addicks Watershed tend to be smaller and lower-valued compared to homes in the Barker Watershed.

Table 1: Summary Statistics

Mean (Standard Deviation)	Addicks Watershed	Barker Watershed
Year of Construction	1995 (10)	1994 (11)
Square Feet	2236 (727)	2708 (892)
2017 Appraised Value	\$190,108 (77,431)	\$291,419 (129,775)
Estimated Damage	\$1,049 (12,251)	\$6,718 (30,554)
Maximum LiDAR Ground Elevation	1,522" (120)	1,432" (152)
Number of Subdivisions	1,634	844
Number of Owner-Occupied Homes	65,188	34,681

Notes: Summary statistics of owner occupied single-family residential properties located in the Fort Bend and Harris County portions of the Addicks and Barker Watersheds.

Measuring Disaster Damage

Since no administrative datasets detail household-level damage during Hurricane Harvey, I approximate this magnitude using pre- and post-storm property values from the Harris and Fort Bend central appraisal districts. In particular, I calculate

$$(1)\ Damage_j = MarketValue_{j,2017} - MarketValue_{j,2018},$$

where $MarketValue_{i,t}$ measures the market value of the structure of property j (not including land value) in year t. The 2017 values should not be influenced by Hurricane Harvey because those numbers were calculated and certified before the storm occurred.[4] On the other hand, $MarketValue_{i,2018}$ estimates the value of property as of January 1st, 2018, which was approximately four months after Hurricane Harvey.

[2] The Addicks and Barker Watersheds are contiguous watersheds in northwest Houston that both contain reservoirs with their respective names.

[3] Portions of Waller County are located in the Addicks and Barker Watersheds, but these are mostly nonresidential areas that lie several feet above the reservoirs.

[4] Some taxing jurisdictions allow for the reappraisal of property values after disasters. It is unclear if the central appraisal districts retroactively updated their certified tax rolls, and these updates would bias my damage estimates toward zero.

Consequently, Equation (1) reflects the change in property value due to disaster damage net of general home-price changes in 2017 and repairs that occurred before year-end. Summary statistics in Table 1 indicate an average damage of $1,049 and $6,718 in the Addicks and Barker Watersheds, respectively, but the size of the standard deviations reveal substantial variation. Appendix A1 details the strong, positive relationship between the damage estimates from Equation (1) and FEMA's estimates of flood depth, supporting the validity of this measure.

Measuring Property Elevation

My identification strategy outlined in Section 5 relies on using property-level elevation as the running variable in a regression discontinuity framework. In particular, properties are exposed to flood damage when water exceeds their first floor elevation (FFE), which I approximate using aerial light detection and ranging (LiDAR) data (TNRIS, 2022). Figure 3 illustrates the distribution of LiDAR points that measure the ground-surface elevation across two residential parcels in the sample.[5]

Figure 3: Property-Level Ground Elevation

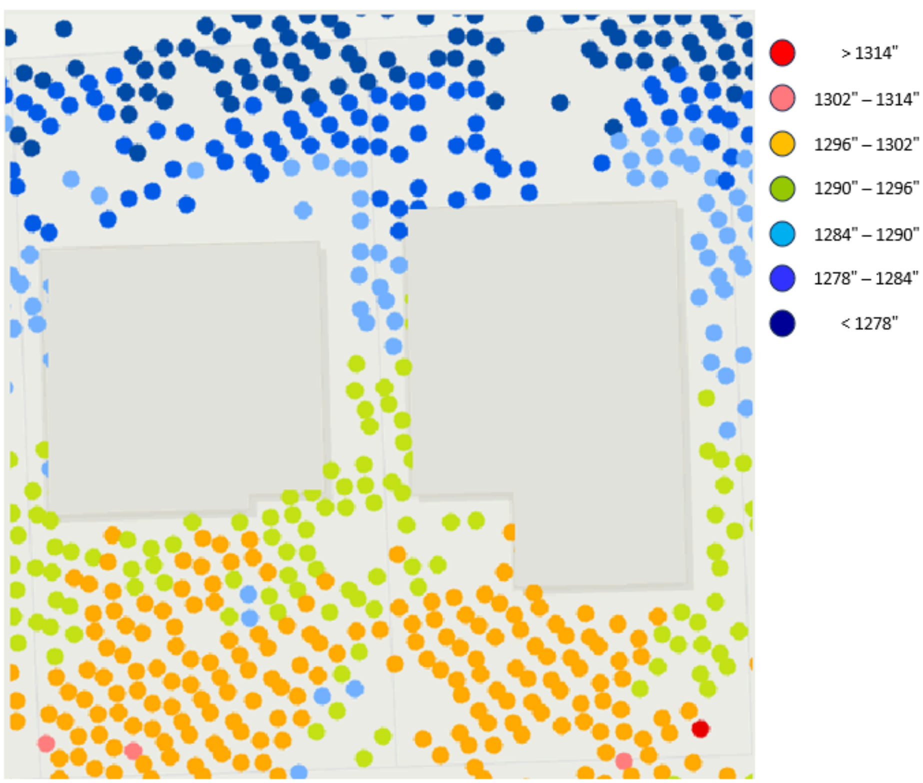

Notes: Aerial LiDAR ground elevation points (measured in inches) for two residential properties in Oak Park Trails Subdivision in the Barker Reservoir. Property parcel shapefiles are available at the Harris and Fort Bend Central Appraisal Districts. Aerial LiDAR data are accessed from TNRIS (2022).

[5] Based on conversations with housing developers, I assume that each property's FFE is equal to its maximum ground elevation. This approximation results in measurement error in my elevation running variable, and I explore the structure of this error in Appendix A2.

Outcomes
Given my analytical sample of owner-occupied households, I use deed transactions from CoreLogic's Owner Transfer dataset to determine if people sell their homes after Hurricane Harvey. I consider a home as sold if a deed associated with the property is recorded after August 25th, 2017, when the storm stalled over Houston. Approximately 15 percent of the sample recorded a deed within the 30 months leading up February 2020.

In order to examine homeowners' relocation decisions, I link the sample to a nationwide database of address histories from Infutor Data Solutions, an aggregator of address data that compiles voter files, property deeds, USPS address changes, etc.[6] The linked dataset allows me to analyze alternative measures of residential mobility such as out-of-county and out-of-state migration as well as the number of miles between pre- and post-storm residences. I also observe households' decisions regarding homeownership and housing characteristics by linking the dataset to CoreLogic's statewide property tax roll data. Census-tract socioeconomic factors of post-storm residences are obtained from the 2020 American Community Survey.

5 Empirical Framework

The idiosyncrasies of Houston's reservoirs offer a unique setting to identify the causal effects of disaster damage. My empirical strategy involves the comparison of outcomes for households in the same subdivision who live just above and just below the peak water level reached in the reservoirs during Hurricane Harvey. The US Army Corps of Engineers (2020) report that the Addicks and Barker Reservoirs reached peaks of 1309.2 and 1219.2 inches above mean sea level, respectively, which I use as thresholds in a fuzzy regression discontinuity design.

Consider a theoretical experiment where identical homes are randomly assigned different "doses" of flooding. The average treatment effect of dose d could be estimated by the difference between mean outcomes for those treated with d and non-flooded households.[7] Flood exposure, however, is a nonrandom event at a hyper-localized level due in part to individuals' mitigation measures (e.g., the installation of flood walls or flood vents), community development decisions (e.g., drainage infrastructure), and atmospheric and topographic variation. Table 2 documents heterogeneity in property characteristics across a range of geospatial units in my sample.[8] Properties are relatively homogenous within subdivisions, but there remains variation in the value and size of homes as well as in elevation. Elevation (and therefore flood risk) correlates with the value, size, age, and other (potentially unobservable) property characteristics that may also be related to outcomes, introducing omitted variable bias into the ordinary least squares estimator.

[6] The details of my linking process are provided in Appendix A3.
[7] The *average treatment effect* at dose d describes the level effect of the dose-response relationship. Alternatively, the slope effect captures *average causal response* to an incremental change in the dose at d (Callaway et al., 2024).
[8] For example, the standard deviation of home size is 729sqft and 895sqft in the Addicks and Barker Watersheds, respectively, and the average of 812sqft across the two watersheds. The averages are similar when weighted by the number of properties in each group.

Table 2: Average (Within-Group) Standard Deviation

Group Level (Number of Groups)	Market Value (2010-2016)	Home Size	Year Built	Maximum Ground Elevation (Inches)
Watershed (2)	\$97,172	810sqft	10.4	136”
Zip Code (9)	\$83,440	757sqft	8.7	67”
Census Block (2864)	\$21,056	377sqft	1.9	8”
Subdivisions (2095)	\$20,282	385sqft	1.3	8”

Notes: Each row contains averages of within-group standard deviations. For example, the standard deviation of home size is 727sqft and 892sqft in the Addicks and Barker Watersheds, respectively, and the simple average of the two standard deviations is 810sqft. Averages are similar when weighted by the number of properties in each group.

I address this concern by analyzing properties at elevations near the peak water level that was reached in the Addicks and Barker Reservoirs during Hurricane Harvey. Property and household characteristics, as well as potential outcomes, should be smooth as you rise in elevation through the peak water level because Hurricane Harvey’s precise magnitude was unknown *a priori*. Households could not predict the peak water level when making residential-sorting decisions months or years before the storm. The only difference between homes slightly above and below this level should be flood exposure. This setting approximates a local randomized experiment, where quasi-identical distributions of homes receive different damage doses.

My identification strategy exploits the fact that the intensity of flood damage decreases with a property’s elevation up to the peak water level. For example, National Flood Services LLC estimates approximately \$37,000 in structural damage for a 2,500sqft, one-story home that is exposed to six inches of water compared to \$24,000 in damage for the same home exposed to a single inch. No damage is expected for homes lying above water (FloodSmart, 2019). While the hydrology literature documents several types of flood damage functions, a common characterization is a discontinuous increase in damage at the first floor elevation (Theodosopoulou et al., 2022). Figure 4 illustrates a piecewise linear damage function with a discontinuity where the first floor elevation equals the peak water level.

Figure 4: Theoretical Damage Function

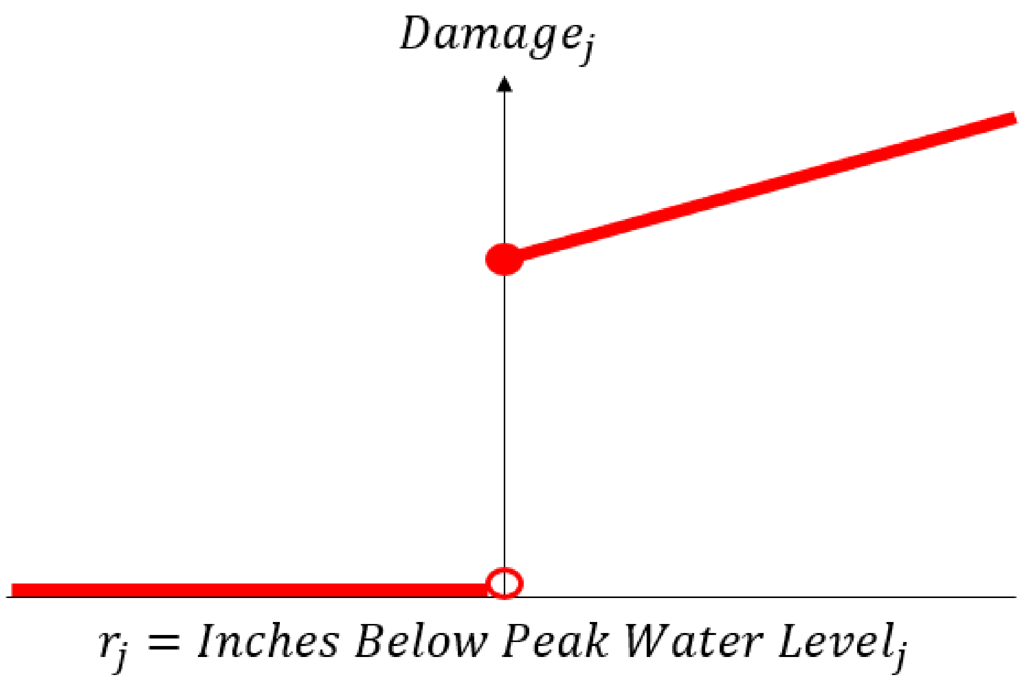

Notes: The running variable is the difference between the peak water level and property *j*'s first floor elevation. For example, a property whose first floor elevation is one foot below the peak water level has a running variable $r_j = 12$.

I formalize this approach to estimate the average treatment effects of disaster damage by using a fuzzy regression discontinuity design, where the expected damage function jumps discontinuously at the peak water level. I include subdivision fixed effects to restrict comparisons to individuals who live in observably-equivalent homes but on different sides of the peak water level.

I approximate each property's first floor elevation with its maximum ground elevation, admittedly using a running variable with measurement error.[9] Following Dong and Kolesar (2023), I implement a donut design to allow for valid inference despite mismeasured elevation. The donut solution requires two key assumptions. First, potential outcomes must be smooth in the mismeasured variable, which may hold mechanically as measurement error in the running variable tends to smooth conditional expectation functions. To illustrate this phenomenon in my context, I simulate a $20,000 damage discontinuity as a function of property elevation, comparing measurements with and without 12-inches of elevation error. Panel A of Figure 5 reveals how measurement error distorts the true discontinuity by smoothing the conditional expectation function.

The second assumption of the Dong and Kolesar (2023) solution requires that the mismeasured running variable correctly classifies treatment assignment. The simulated measurement error on the right side of Panel A illustrates how some undamaged (damaged) homes are misclassified below (above) the peak water level. Panel B illustrates how removing the misclassified observations resolves the

[9] Dong and Kolesar (2023) find that nearly a quarter of regression discontinuity designs that are published in top economics journals suffer from this threat to identification. Their proposed donut solution provides for valid inference for the local average treatment effect of units with values of the *mismeasured* running variable near the true threshold. In contrast, the canonical regression-discontinuity framework that estimates the local average treatment effect of units with values of the *true* running variable around the same threshold.

measurement-error issue. In particular, implementing a 12-inch donut in the presence of 12 inches of measurement error results in a point estimate that is much closer to the simulated discontinuity. Note that the standard challenges of donut trimming apply in this context (e.g., lost sample size and a modified local average treatment effect).

Figure 5: Simulated Damage Function

Panel A: Regression Discontinuity

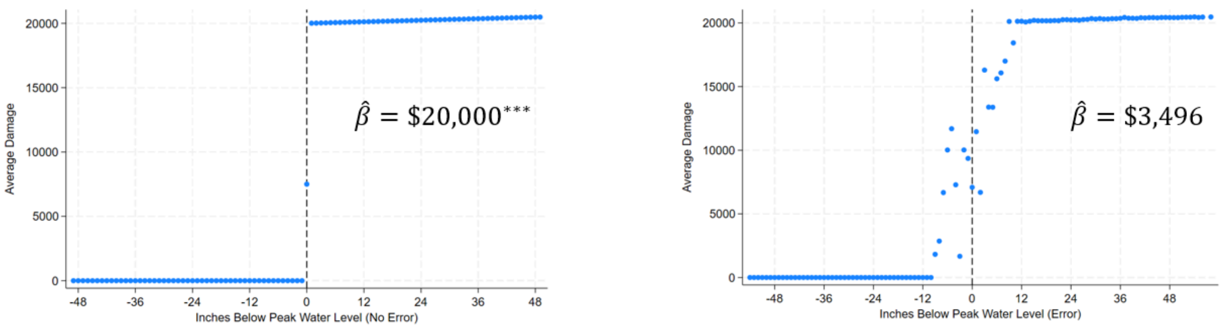

Panel B: Regression Discontinuity with 12-Inch Donut

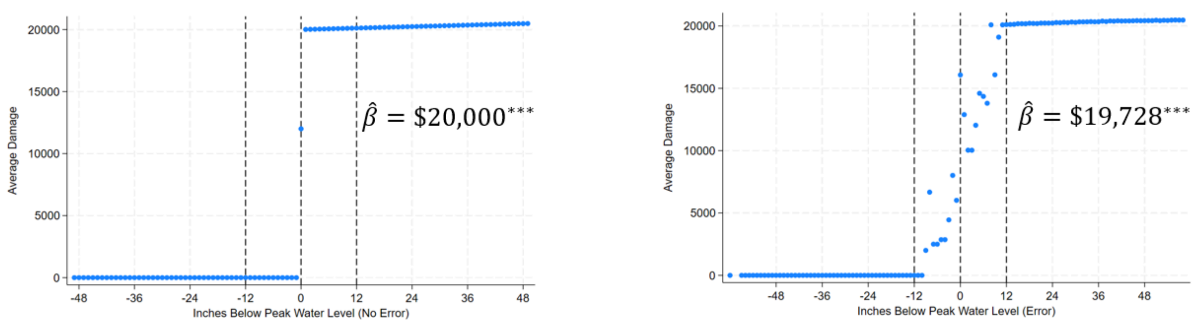

Notes: Figure 5 plots the average damage at each elevation inch for simulated data with a $20,000 discontinuity at $r = 0$. The left sides of Panels A and B illustrate the mean plots in the presence of accurate elevation data. The right sides of Panels A and B illustrate the mean plots in the presence of 24 inches of measurement error (12 inches above and below the true value). Panel B depicts dashed lines at ±12 inches, revealing how sample trimming can remove the measurement-error distortion of the conditional mean plots.

To implement the donut fuzzy regression discontinuity design, I start by estimating the discontinuity in damage around the peak water levels reached in the Addicks and Barker Reservoirs. I standardize the running variable by subtracting property j's maximum ground elevation from the peak water level of the reservoir in which j is located, resulting in a measure of the relative inches below the threshold. The estimating equation of the first stage is

$$(2)\ Damage_j = \alpha_n + \gamma r_j + \beta 1\{r_j \geq 0\} + \psi r_j 1\{r_j \geq 0\} + \xi_j \text{ for } r \in [h_l, h_l^{donut}] \cup [h_r^{donut}, h_r],$$

where *Damage* is determined by Equation (1), α_n represent subdivision-level fixed effects, *r* is the standardized running variable, $1\{r_i \geq 0\}$ is an indicator equaling 1 if *j* lies below the peak water level, h_l and h_r are the left- and right-sided bandwidths, respectively, and h_l^{donut} and h_r^{donut} are the left- and right-side lengths of the donut.

There are two primary motivations for the fuzzy component of the model. First, the magnitude of damage depends on the volume of water exposure to a home. The fuzzy design allows for causal identification of average treatment effects in the presence of different treatment intensities. Second, the probability of damage is non-binary across elevation because of other idiosyncrasies and determinants of flooding, i.e., there exists subsets of the population whose damage was not determined by their elevation relative to the peak water level. The fuzzy design allows for the identification of the local average treatment effect of the households who flooded because of this relative elevation (i.e., the "compliers").

The reduced-form regression equation illustrates how elevation is used to estimate the effect of flood damage on outcomes. Specifically, the reduced-form equation is

$$(3)\ Y_j = \alpha_n + \pi r_j + \lambda 1\{r_j \geq 0\} + \phi r_j 1\{r_j \geq 0\} + \eta_j \text{ for } r \in [h_l, h_l^{donut}] \cup [h_r^{donut}, h_r],$$

where *Y* is the outcome of interest. Note that the parameter λ is the local average treatment effect of living below the peak water level (or the intention-to-treat effect). The local average treatment effect of flood damage can be recovered by scaling λ by the parameter β from the first stage.

The causal interpretation of these parameters (and their ratio) relies on three key identifying assumptions in my setting. First, there must exist a relationship between a property's maximum ground elevation and flood damage that changes as you approach the edges of the donut surrounding the peak water level. Second, maximum ground elevation must accurately assign instances where flood damage occurred. Third, the exclusion restriction requires potential outcomes to be smooth through the peak water level and for the peak water level to affect outcomes *only through* its impact on flood damage. I evaluate these assumptions in Section 6.

6 Empirical Validation

I begin by exploring the damage-elevation relationship in the Addicks and Barker Watersheds that is specified in Equation (2). Figure 6 reveals a smooth curve similar to the simulation results with measurement error. There are fewer observations to the right of the threshold because the reservoirs' peak water levels only extended about three feet above the government-owned land and into neighborhoods.[10]

Figure 6: First Stage

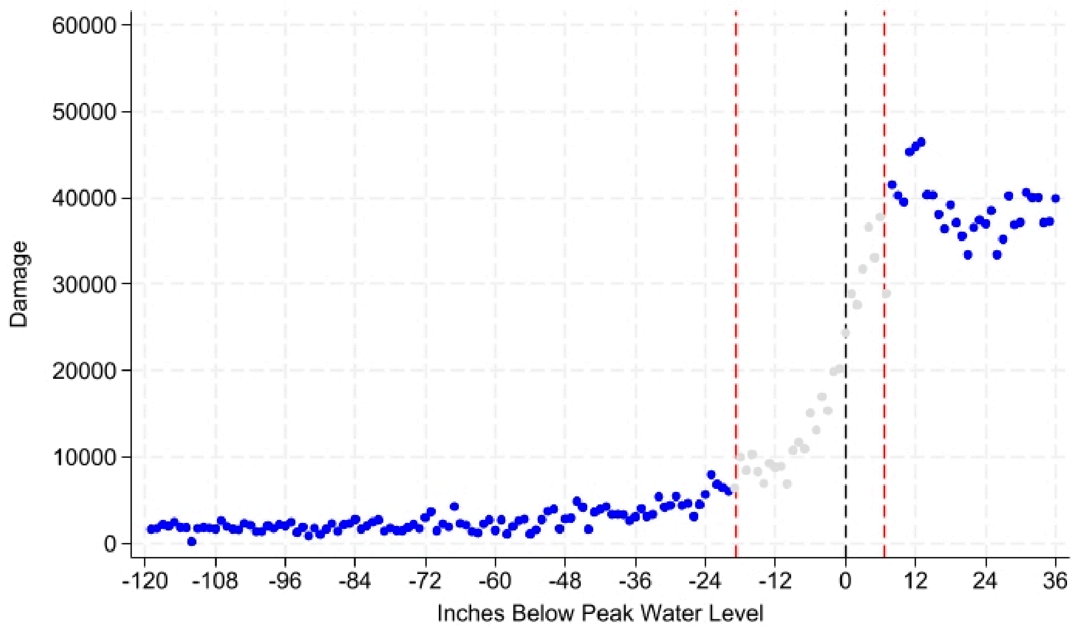

Notes: The coefficient plot is generated by rounding elevation to the nearest inch and estimating a subdivision fixed effects model saturated in 1-inch elevation dummy variables. The sample mean of the outcome variable is added back to the coefficient estimates for illustrative purposes. The red lines correspond to the preferred sample trimming at -18.7 and 6.6 inches below the peak water level.

Based on the measurement error structure explored in Appendix A2, I trim the sample from -18.7 to 6.6 inches to ensure that ground elevation reasonably assigns the occurrence of flood damage. Using this donut design, I estimate an average increase of $47,313 in flood damage for homes lying just below the peak water level. Robustness tests in Appendix 5 corroborate the stability of my estimates at different donut sizes and bandwidths.

The third identifying assumption requires potential outcomes to be smooth through the peak water level and for the peak water level to affect outcomes only through flood damage. While this is inherently untestable, I support this claim by documenting a relatively constant relationship between elevation and pre-storm property characteristics. Figure 7 illustrates the similarity in housing characteristics above and below the peak water level, which is expected given that Harvey's precise magnitude was unpredictable when households moved into these neighborhoods. The lack of pre-storm patterns supports the assumption that a property's distance from the peak water level only affected households through the impact of flood damage.

[10] An additional asymmetry appears around the threshold because of the right-skewness of elevation measurement error that is discussed in Appendix A2.

Figure 7: Pre-Storm Property Characteristics and Elevation

Panel A: Year of Construction

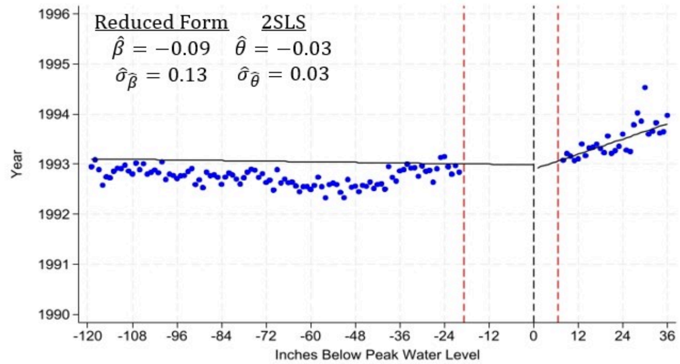

Panel B: Home Size

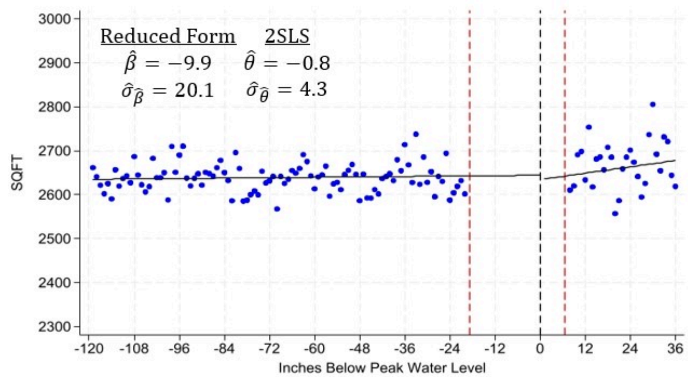

Notes: Panels A and B replace the outcome in Equation (2) with the property's year of construction and structure square footage, respectively. The regression estimates and lines are based on Equation (2), a local linear regression with uniform kernels and a preferred bandwidth of [-120.0,-18.7] and [6.6,36.0]. Standard errors are clustered at the neighborhood level. The coefficient plots are generated by rounding elevation to the nearest inch and estimating a subdivision-phase fixed effects model saturated in 1-inch elevation dummy variables. The sample mean of the outcome variable is added back to the coefficient estimates for illustrative purposes. The red lines correspond to the preferred sample trimming at -18.7 and 6.6 inches below the peak water level.

7 Results

Natural disasters are associated with increased residential mobility (Boustan et al., 2020). Billings et al. (2022) and Gallagher et al. (2023) indicate a similar pattern after Hurricane Harvey, when out-of-Houston migration spiked immediately after the storm. Migration rates, however, were roughly the same in flooded and non-flooded census blocks, suggesting that local level damage played little role in individuals' mobility decisions.

While individuals may not differentiate relocation decisions based on local damage intensity, Panel A of Figure 8 indicates the importance of property-level damage in influencing behavior. Households exposed to \$10,000 of flood damage are 0.4 percentage points (43 percent) more likely to move out of their pre-Harvey residence within six months of the storm compared to their non-flooded neighbors who lived above the peak water level. Panel B illustrates the estimated cumulative impact on residential mobility over the 10 quarter post-period. The spike in mobility is not captured in 3Q2017 data as Hurricane Harvey stalled over Houston during the last week of August. The initial wave of moves occurs in 4Q2017 and 1Q2018 before attenuating through 2019.

Figure 8 also reveals differential dynamics across move types. In Panel C, I estimate that damage causes an immediate increase for within-county moves that persists for nearly two years. The rapid shift toward short-distance moves may reflect the local nature of social networks as survivors pool resources across the community. Moreover, households with high placed-based attachment may minimize disaster distress by relocating and recovering within their established community.

On the other hand, Panel D suggests a delayed, negative impact on out-of-county moves that lasts several quarters before dissipating. Longer distance moves may require more planning and carry higher transaction costs that prevent an instantaneous response during disaster events. For example, households who have already secured future employment outside of a disaster-struck region may be more likely to proceed with their relocation plans despite suffering from property damage. The lagged, negative effect on out-of-county moves suggests that those who *would have* made longer-distance relocation plans in the two years after Hurricane Harvey instead substituted towards staying local.

The estimated residential mobility patterns are distinct from those documented in the context of Hurricane Katrina. In particular, Bleemer and van der Klaauw (2019) find that census-block level damage exposure in New Orleans stimulated migration out of Orleans Parish and Louisiana. These contrasting effects may be driven by differences in analytical samples (Houston homeowners vs. New Orleans residents), idiosyncrasies between Hurricanes Katrina and Harvey, the underlying economic health of the impacted regions, as well as the estimation of household- and census-block-level average treatment effects.

Figure 8: Residential Mobility

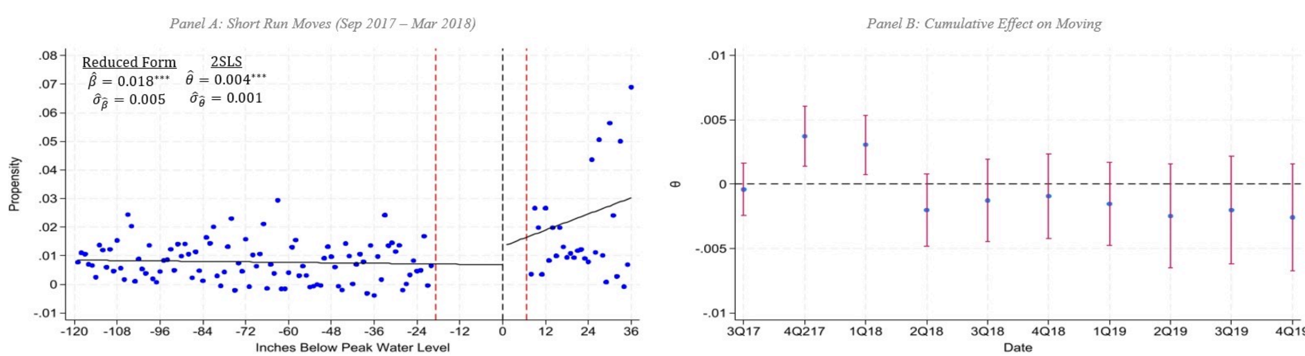

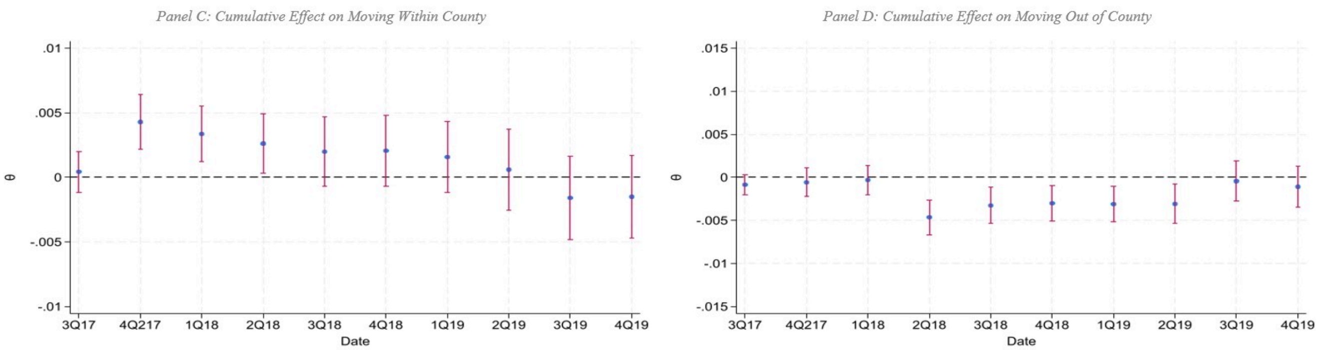

Notes: The point estimate $\hat{\theta}$ is the the estimated causal effect of \$10,000 of damage, reflecting the ratio of the reduced form effect $\hat{\lambda}$ from Equation (3) and the first stage effect $\hat{\beta}$ from Equation (2). Panel A is based on an indicator for Infutor observing at least one individual in household *j* moving between September 2017 and March 2018. The regression estimates and lines are based on Equation (3), a local linear regression with uniform kernels and a preferred bandwidth of [-120.0,-18.7] and [6.6,36.0]. Standard errors are clustered at the neighborhood level. The coefficient plots are generated by rounding elevation to the nearest inch and estimating a subdivision fixed effects model saturated in 1-inch elevation dummy variables. The sample mean of the outcome variable is added back to the coefficient estimates for illustrative purposes. The red lines correspond to the preferred sample trimming at -18.7 and 6.6 inches below the peak water level. Panel B plots the point estimates and 95% confidence intervals using a cumulative outcome variable for each post-storm quarter. For example, the estimates for 1Q2018 are based on whether at least one individual in a household moved between September 2017 and March 2018. Standard errors are clustered at the neighborhood level.

Standard economic theory predicts an aggregate decrease in housing market transactions after natural disasters that results from a combination of decreased housing stock and demand. Zivin et al. (2023) present evidence for this phenomenon from Florida's response to hurricanes between 2000 and 2016. McCoy and Walsh (2018) document a similar market-level decline in the context of Colorado wildfires. Similar to trends in residential mobility established in the literature, I display the lack of relationship between local-level damage and home sales in Figure 9. The number of home sales in high- and low-damage subdivisions track similarly before and after Hurricane Harvey. If anything, high-damage subdivisions suffer a stronger shock on impact and then rebound higher for longer.

Figure 9: Housing Sales in the Addicks and Barker Watersheds' Subdivisions

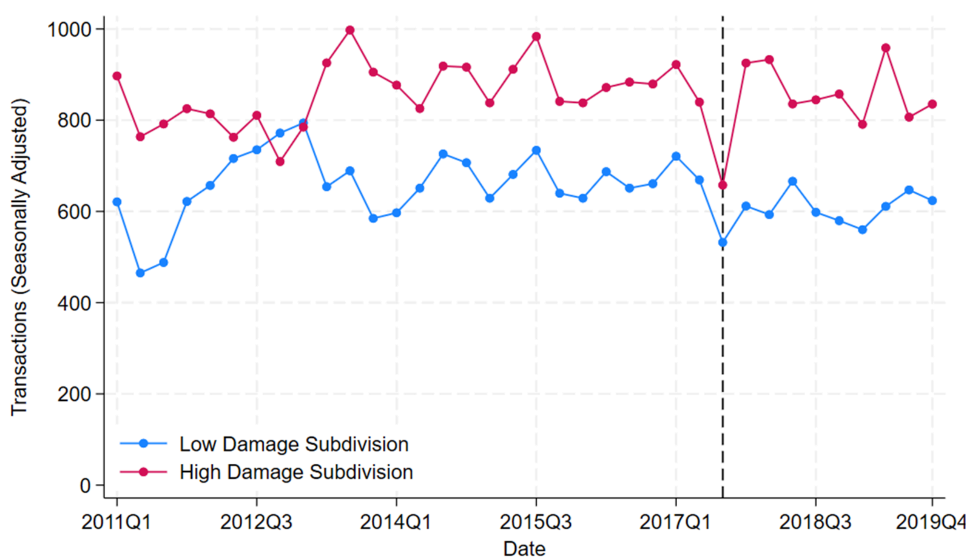

Notes: I collapse the count of CoreLogic deed transactions at the subdivision level. Subdivisions with above zero average property damage are classified as “high damage.” Note that damage is calculated from Equation (1) and can be less than zero due to general price appreciation and rebuilding efforts that occur prior to January 1st, 2018. The two series are residualized by quarter and rescaled by the constant term.

The results in Figure 10, however, confirm that property-level damage causes a decrease in housing sales in the short run. I estimate that \$10,000 of damage decreases homeowners’ propensity to sell within six months by 0.6 percentage points (23 percent) relative to their non-flooded neighbors who lived just above the peak water level. The cumulative impact grows for about a year and persists through 2019, suggesting that damage delays and may even prevent sales that would have occurred in the absence of flooding. Homeowners on the margin of moving may choose to repair their property before selling, and many flooded households were forced to wait months after Hurricane Harvey for full disbursement of disaster aid or insurance payments to help fund this investment.

Figure 10: Housing Sales

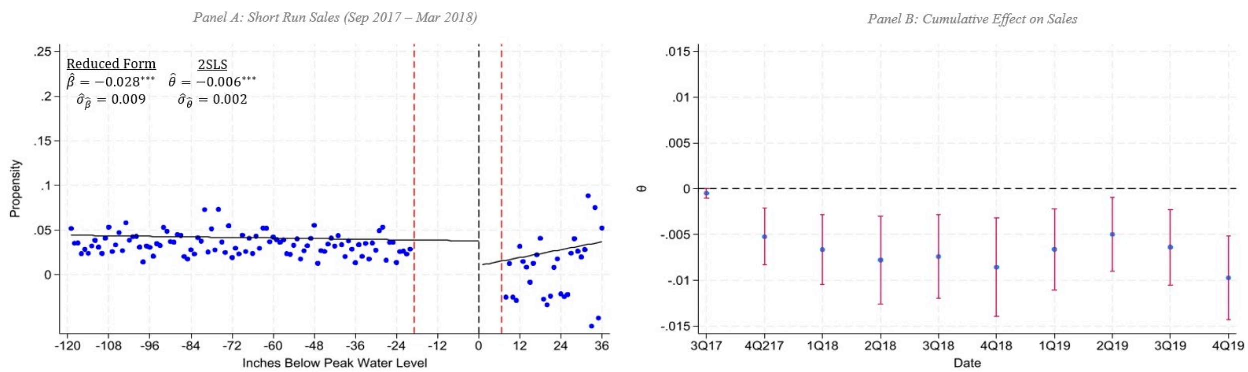

Notes: The point estimate $\hat{\theta}$ is the the estimated causal effect of \$10,000 of damage, reflecting the ratio of the reduced form effect $\hat{\lambda}$ from Equation (3) and the first stage effect $\hat{\beta}$ from Equation (2). Panel A is based on an indicator for property *j* having a deed recorded between September 2017 and March 2018. The regression estimates and lines are based on Equation (3), a local linear regression with uniform kernels and a preferred bandwidth of [-120.0,-18.7] and [6.6,36.0]. Standard errors are clustered at the neighborhood level. The coefficient plots are generated by rounding elevation to the nearest inch and estimating a subdivision fixed effects model saturated in 1-inch elevation dummy variables. The sample mean of the outcome variable is added back to the coefficient estimates for illustrative purposes. The red lines correspond to the preferred sample trimming at -18.7 and 6.6 inches below the peak water level. Panel B plots the point estimates and 95% confidence intervals using a cumulative outcome variable for each post-storm quarter. For example, the estimates for 1Q2018 are based on whether a property had a deed recorded between September 2017 and March 2018. Standard errors are clustered at the neighborhood level.

In addition to the persistent effects on home sales, I document impacts on housing consumption that last through at least 2019. For example, Figure 11 illustrates the effect of flood damage on the probability of households owning their post-Harvey residence. I estimate that \$10,000 of damage increases homeownership by 0.4 percentage points (1 percent) through 2019. This result is partially mechanical as the initial sample is comprised exclusively of homeowners, i.e., those who do not sell their homes are still considered homeowners even if they relocate to another property. Conditioning on the set of movers in Panel B, however, reveals a significant transition out of owner occupancy. In particular, households exposed to \$10,000 of damage are 2.7 percentage points (3 percent) less likely to own their next residence compared to their non-flooded peers. While smaller in magnitude, these results align with Bleemer and van der Klaauw (2019), who find persistently lower homeownership rates among people who lived in inundated parts of New Orleans during Hurricane Katrina.

Figure 11: Estimated Average Treatment Effect on Homeownership

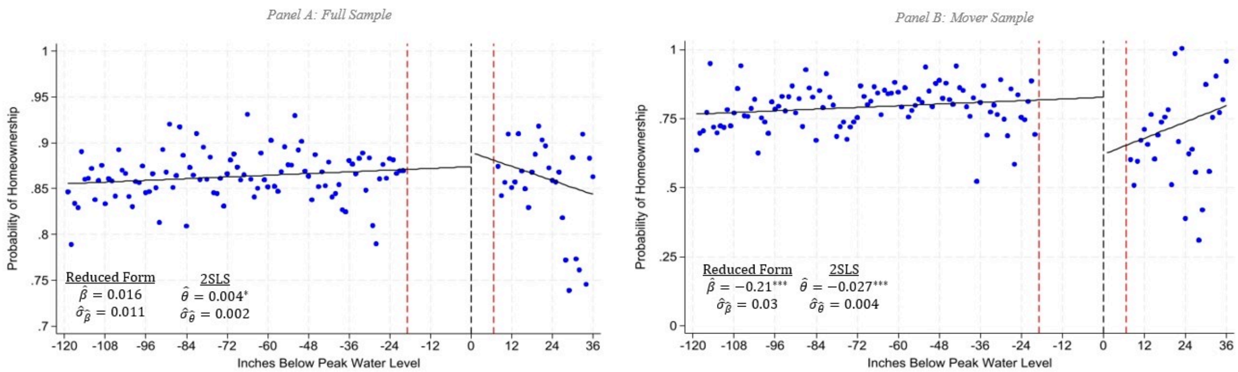

Notes: I consider an individual as a post-storm renter if they sold their home after Hurricane Harvey and did not appear in Texas appraisal-district data at their new address before 2020. The American Community Survey indicates that a majority of interstate movers rent their new residence, and due to data constraints I assume people who leave Texas are post-storm renters. The left side of Panel A uses the full sample including homeowners who do not sell after Hurricane Harvey. The right side of Panel B restricts the sample to movers. The point estimate $\hat{\theta}$ is the the estimated causal effect of \$10,000 of damage, reflecting the ratio of the reduced form effect $\hat{\lambda}$ from Equation (3) and the first stage effect $\hat{\beta}$ from Equation (2). The regression estimates and lines are based on Equation (3), a local linear regression with uniform kernels and a preferred bandwidth of [-120.0,-18.7] and [6.6,36.0]. Standard errors are clustered at the neighborhood level. The coefficient plots are generated by rounding elevation to the nearest inch and estimating a subdivision fixed effects model saturated in 1-inch elevation dummy variables. The sample mean of the outcome variable is added back to the coefficient estimates for illustrative purposes. The red lines correspond to the preferred sample trimming at -18.7 and 6.6 inches below the peak water level.

The impact of flood damage extends to the types of neighborhoods and homes where individuals choose to live. Results in Figure 12 indicate that impacted households tend to sort into higher-valued homes and higher-income neighborhoods. These changes in socioeconomic environment, combined with the estimated decrease in homeownership, highlight the potential tradeoffs in a post-disaster environment.

Figure 12: Estimated Average Treatment Effect on Neighborhood and Home Choice for Movers

Panel A: Post-Storm Home Values (2020)

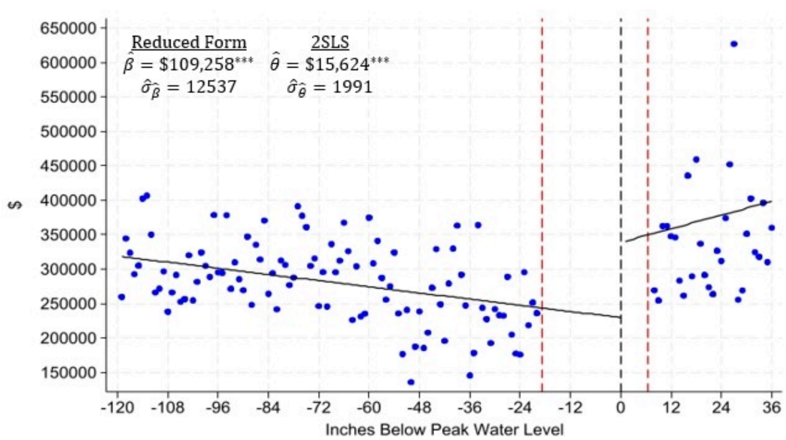

Panel B: Post-Storm Census Tract Average Income (2020)

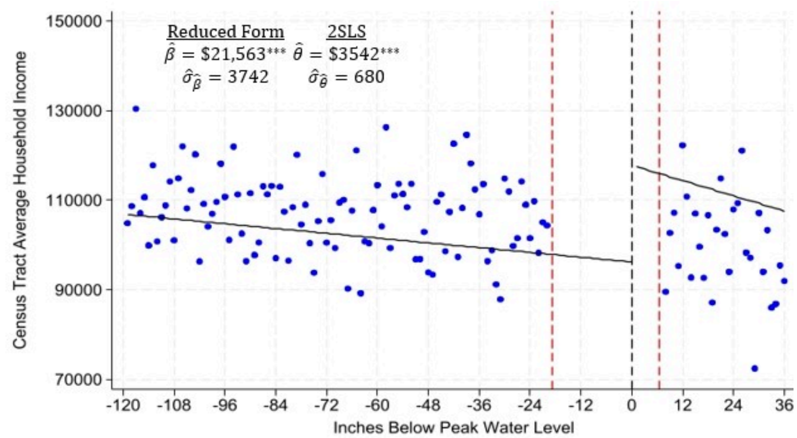

Notes: Figure 12 restricts the sample to individuals who sold their home after Hurricane Harvey and who have a post-storm address in the Infutor data. Panel A considers only those who relocated within Texas, as I am only able to match post-storm movers to appraisal district data within the state, while Panel B includes movers who relocated across the United States. The point estimate $\hat{\theta}$ is the the estimated causal effect of $10,000 of damage, reflecting the ratio of the reduced form effect $\hat{\lambda}$ from Equation (3) and the first stage effect $\hat{\beta}$ from Equation (2). The regression estimates and lines are based on Equation (3), a local linear regression with uniform kernels and a preferred bandwidth of [-120.0,-18.7] and [6.6,36.0]. Standard errors are clustered at the neighborhood level. The coefficient plots are generated by rounding elevation to the nearest inch and estimating a subdivision fixed effects model saturated in 1-inch elevation dummy variables. The sample mean of the outcome variable is added back to the coefficient estimates for illustrative purposes. The red lines correspond to the preferred sample trimming at -18.7 and 6.6 inches below the peak water level.

8 Conclusion

The natural disaster literature has documented a dynamic recovery process across a variety of outcomes for individuals living in disaster-struck areas. I contribute to this literature by using property-level data in a quasi-experimental design to estimate the average treatment effect of household damage exposure on residential mobility, home sales, and other housing choices after Hurricane Harvey. I leverage the relatively unknown risk and unique flooding mechanism of Houston's Addicks and Barker Reservoirs to circumvent selection issues that challenge the identification of causal effects of disaster damage. I overcome measurement error in the elevation running variable by sample trimming, resulting in a donut regression discontinuity design. I examine the relationship between flooding and a home's elevation, exploiting a discontinuous increase in damage from $0 to approximately $48,000 once water reaches the first floor. The results from this first stage are used to rescale reduced-form estimates to examine the impact of the $10,000 of flood damage on housing choices.

Although local-level exposure does not necessarily correlate with household movement patterns, I reject the notion that home sales and relocation decisions are unaffected by disaster damage. I estimate that flood damage increases residential mobility in the short run, particularly for shorter-distance moves. In contrast, flood damage delays home sales for about a year after Hurricane Harvey. These opposite effects align with theoretical predictions of disaster damage making a portion of the housing stock uninhabitable.

In addition to impacting locational choices, I provide evidence that household-level damage pushes movers out of homeownership. In particular, I estimate that $10,000 of damage causes a 3 percent decrease in homeownership. Despite the combined shock to shelter and wealth, I find that flooded households are more likely to sort into higher-income census tracts in the aftermath of Hurricane Harvey. This relative improvement in physical and socioeconomic environments mirrors the long-run recovery patterns documented in the disaster literature (Sacerdote, 2012; Deryugina et al., 2018; Deryugina and Molitor, 2020). Since disaster damage pushes people out of their neighborhoods and into new economic environments, the impacts of extreme weather may extend into other aspects of life. These neighborhood effects may augment or offset the transition into different types of housing or housing tenure.

The totality of my results raise important questions about the effectiveness of disaster aid. For example, the delayed effect on housing transactions may indicate that SBA loans are protecting people from losing their homes after catastrophic events. On the other hand, disaster damage net of relief efforts led to a substantial transition out of homeownership and into renter occupancy. The normative implications of this transition are unclear, especially as flooded households tend to relocate into higher income neighborhoods that may offer improved economic opportunities.

9 References

Billings, Stephen B., Gallagher, Emily A., and Ricketts, Lowell R. (2022) Let the Rich Be Flooded: The Distribution of Financial Aid and Distress After Hurricane Harvey, *Journal of Financial Economics*, VOl. 146(2), pp. 797-819.

Bleemer, Zachary and van der Klaauw, Wilbert (2019) Long-Run Net Distributionary Effects of Federal Disaster Insurance: The Case of Hurricane Katrina, *Journal of Urban Economics*, Vol. 111, pp. 70-88.

Bloom, Michael F. (2017) The History of Addicks and Barker Reservoirs, *Riparian Houston*, accessed online at https://riparianhouston.com/2017/09/03/the-history-of-addicks/ on 27 September 2023.

Boustan, Leah Platt, Kahn, Matthew E., Rhode, Paul W., and Yanguas Maria Lucia (2020) The Effect of Natural Disasters on Economic Activity in US Counties: A Century of Data, *Journal of Urban Economics*, Vol. 118.

Callaway, Brantly, Goodman-Bacon, Andrew, Sant'Anna, Pedro H. C. (2024) Difference-in-Differences with a Continuous Treatment, *NBER Working Paper* 32117, https://www.nber.org/system/files/working_papers/w32117/w32117.pdf.

Chyn, Eric and Katz, Lawrence F. (2021) Neighborhoods Matter: Assessing the Evidence for Place Effects, *Journal of Economic Perspectives*, Vol. 35(4).

City of Austin (2020) Appraisal District Property Value Updates, Special Report, City of Austin, Office of the City Auditor, April 2020.

Collier, Benjamin L., Ellis, Cameron, and Keys, Benjamin J. (2021) The Cost of Consumer Collateral:

Evidence from Bunching, *NBER Technical Report*.
Deryugina, Tatyana, Kawano, Laura, and Levitt, Steven (2018) The Economic Impact of Hurricane Katrina on Its Victims: Evidence from Individual Tax Returns, *American Economic Journal: Applied Economics*, Vol. 10(2), pp. 202-233.
Deryugina, Tatyana and Molitor, David (2020) Does When You Die Depend on Where You Live? Evidence from Hurricane Katrina, *American Economic Review*, Vol. 110(11), pp. 3602-3633.
Dong, Yingying and Kolesár, Michal (2023) When Can We Ignore Measurement Error in the Running Variable?, *Journal of Applied Econometrics*, Vol. 38(5), pp. 735-750.
Elliott, James R. and Wang, Zheye (2023) Managed Retreat: A Nationwide Study of the Local, Racially Segmented Resettlement of Homeowners from Rising Flood Risks, *Environmental Research Letters*, 2023.
FBCAD (2022) Data: Geographic Information System, *Fort Bend Central Appraisal District*, 2017 Certified Shapefile, accessed online at https://www.fbcad.org/gis-data/ on 08 June 2022.
FEMA (2017) National Flood Insurance Program, Flood Insurance Manual, U.S. Department of Homeland Security, Federal Emergency Management Agency (FEMA), effective 01 April 2017, accessed online at https://www.fema.gov/flood-insurance/work-with-nfip/manuals/archive on 21 April 2022.
FEMA (2018) Harvey Flood Depths Grid, Federal Emergency Management Administration, updated on 08 March 2023, accessed online at https://www.hydroshare.org/resource/e8768f4cb4d5478a96d2b1cbd00d9e85/ on 24 September 2023.
FEMA (2022) Recovery Support Function Leadership Group Spending Explorer, Federal Emergency Management Agency (FEMA), accessed online at https://recovery.fema.gov/spending-explorer on 27 April 2022.
First Street Foundation (29-06-2020) Flood Model Overview, *First Street Foundation*, 29 June 2020, accessed online at https://firststreet.org/research-lab/published-research/flood-model-methodology_overview/ on 02 July 2022.
Furrh, Jacob True and Bedient, Philip (2023) Upstream Addicks-Barker Reservoir Damages During Hurricane Harvey: A Case Study of Urban Hydrology and Policy Failure in Houston, TX, *Journal of the American Water Resources Association*, **https://doi.org/10.1111/1752-1688.13115**.
Gallagher, Emily A., Billings, Stephen B., and Ricketts, Lowell R. (2023) Human Capital Investment after the Storm, *Review of Financial Studies*, Vol. 36(7), pp. 2651-2684.
Gallagher, Justin and Hartley, Daniel (2017) Household Finance after a Natural Disaster: The Case of Hurricane Katrina, *American Economic Journal: Economic Policy*, Vol. 9(3), pp. 199-228.
Groen, Jeffrey A. and Polivka, Anne E. (2008) The Effect of Hurricane Katrina on the Labor Market Outcomes of Evacuees, *American Economic Review*, Vol. 98(2), pp. 43-48.
HCAD (2022) Public Data: Download GIS Data, *Harris Central Appraisal District*, accessed online at https://hcad.org/pdata/pdata-gis-downloads.html on 08 June 2022.
HCFCD (2020) Estimated Harris County First Floor Elevations, *Harris County Flood Control District*, derived from Cyclomedia LiDAR scans, shapefile accessed from ArcGIS Online Portal.
McCoy, Shawn J. and Walsh, Randall P. (2018) Wildfire Risk, Salience, & Housing Demand, *Journal of Environmental Economics and Management*, Vol. 91, pp. 203-228.
NCEI (2018) Costliest U.S. Tropical Cyclones Tables Updated, National Centers for Environmental

Information, National Oceanic and Atmospheric Administration, accessed online at https://www.nhc.noaa.gov/news/UpdatedCostliest.pdf on 13 April 2023.

New York Times (2017) Mapping the Devastation of Harvey in Houston, *New York Times*, 29 August 2017, accessed online at https://www.nytimes.com/interactive/2017/08/28/us/houston-maps-hurricane-harvey.html on 24 September 2023.

NOAA (2023) U.S. Billion-Dollar Weather and Climate Disasters, *NOAA National Centers for Environmental Information (NCEI)*, accessed online at https://www.ncei.noaa.gov/access/billions/.

O'Neil, Evan (2020) How the Barker and Addicks Dams Work, *The Houston Flood Museum*, accessed online at https://houstonfloodmuseum.org/how-the-barker-and-addicks-dams-work/ on 27 September 2023.

Sacerdote, Bruce (2012) When the Saints Go Marching Out: Long-Term Outcomes for Student Evacuees from Hurricanes Katrina and Rita, *American Economic Journal: Applied Economics*, Vol. 4(1).

Sheldon, Tamara L. and Zhan, Crystal (2019) The Impact of Natural Disasters on US Home Ownership, *Journal of the Association of Environmental and Resource Economists*, Vol. 6(6).

Smith, Adam and Katz, Richard W. (2013) U.S. Billion-dollar Weather and Climate Disasters: Data Sources, Trends, Accuracy and Biases, *Natural Hazards*, Vol. 67, pp. 387-410.

Texas Tribune (2014) Inside Intelligence: About Those Tax Appraisals, *Texas Tribune*, 01 May 2014, accessed online at https://www.texastribune.org/2014/05/01/inside-intelligence-about-those-tax-appraisals/ on 24 September 2023.

Theodosopoulou, Zafeiria, Kourtis, Ioannis M., Bellos, Vasilis, Apostolopoulos, Konstantinos, Potsiou, Chryssy, and Tsihrintzis, Vassilios A. (2022) A Fast Data-Driven Tool for Flood Risk Assessment in Urban Areas, *Hydrology*, Vol. 9(8).

TNRIS (2022) DataHub, Texas Natural Resources Information System, accessed online at data.tnris.org on 02 August 2022.

US Army Corps of Engineers (2020) Buffalo Bayou and Tributaries Resiliency Study, Texas, Interim Feasibility Report, October 2020, accessed online at https://www.swg.usace.army.mil/Portals/26/BBTnT_Interim_Report_202001001_Final_.pdf on 30 April 2022.

Vigdor, Jacob L. (2007) The Katrina Effect: Was There a Bright Side to the Evacuation of Greater New Orleans?, *The B.E. Journal of Economic Analysis*, Vol. 7(1).

Zivin, Joshua Graff, Liao, Yanjun, and Panassie, Yann (2023) How Hurricanes Sweep Up Housing Markets: Evidence from Florida, *Journal of Environmental Economics and Management*, Vol. 118.

Appendix A1: FEMA Flood Depth and Damage Estimates

Although Texas' central appraisal districts are required to value property at 100% of its market value as of January 1st of each year, the lack of mandatory sales-price disclosure challenges the accuracy of these estimates. Anecdotal evidence, however, suggests that public appraisers have historically accessed sales data, and the accuracy of market values for typical homes in large subdivisions are less of a concern (Texas Tribune, 2014; City of Austin, 2020).

In order to assess the validity of the flood damage estimates from Equation (1), I compare the results with water levels reached on each property during Hurricane Harvey. Immediately after the storm, FEMA began modeling flood depths (i.e., the difference between water elevation and ground elevation) across the disaster zone. FEMA's flood-depth data are based on a variety of sources including observed water levels at stream gauges, remote sensing, and other inspection data. These data are intended to be used for determining damage levels on specific structures (FEMA, 2018).

Figure A1 presents the relationship between three measures of property-level flood depth and my estimates of flood damage in the Addicks and Barker Watersheds. FEMA's model does not capture flood depths below 12 inches, resulting in left-censored data grouped at zero inches. Properties where the minimum flood depth is approximately 12 inches average almost $10,000 of damage. All three flood-depth metrics are positively related to my damage estimates.

Figure A1: Property-Level Flood Depth and Estimated Damage

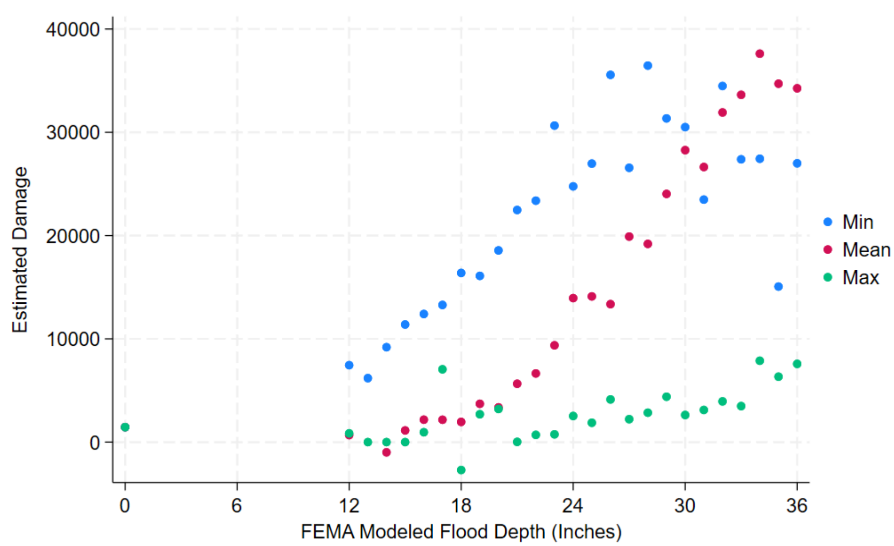

Appendix A2: Measurement Error in Elevation

One of the implications of the intra-parcel elevation variation is that I may not observe the true FFE. Based on conversations with housing developers, I assume that each property's FFE is equal to its maximum ground elevation. In order to explore the error structure of this approximation, I obtain validation data generated by mobile (rather than aerial) LiDAR technology that measures the elevation of the base of each home's front door (HCFCD, 2020). There is minimal overlap between the validation data and my analytical sample, but I am able to compare the maximum ground elevation and base-of-door elevations for 13,000 homes in other parts of Houston.

Figure A2 displays the distribution of deviations between the two measures. The maximum ground elevation tends to be higher than the base of the front door, resulting in positive skewness. Variation in landscaping, natural land gradients, and instrument imprecision likely explain this overestimation. The maximum ground elevation, however, is highly predictive of base-of-door elevation, with a 0.99 correlation coefficient and a majority of errors lying within seven inches. I use this error structure to refine my empirical strategy in Section 5.

Figure A2: Elevation Data Deviations

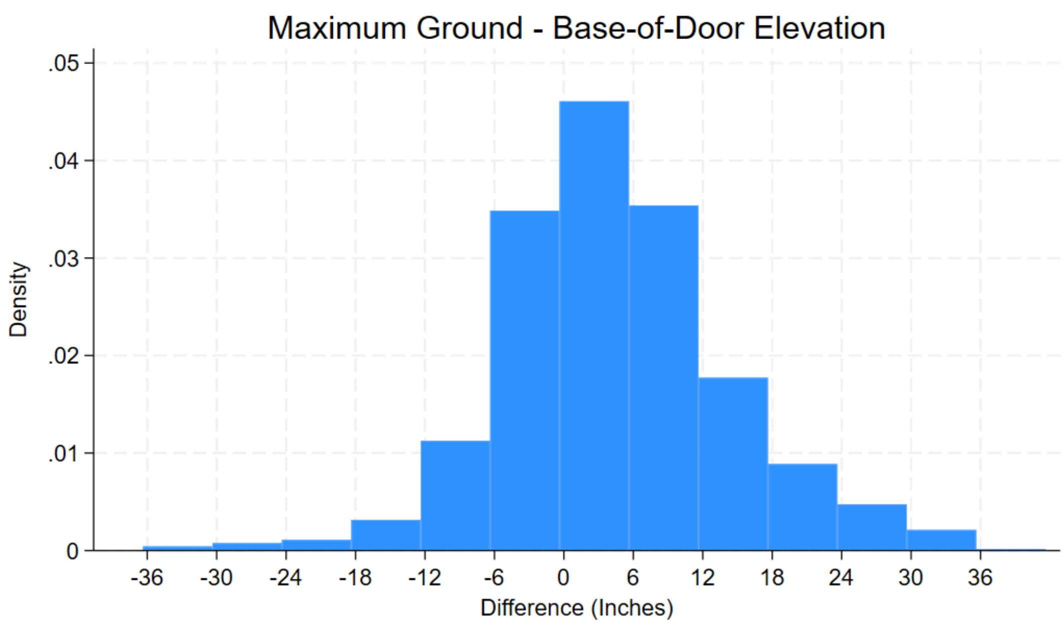

In order to satisfy the second identifying assumption outlined in Section 4, the mismeasured elevation must correctly classify when properties are damaged. In an ideal setting, I would delete observations whose maximum ground elevation lies within the support of measurement error estimated from the validation data. This support, however, extends from -30 to 48 inches, thereby removing my entire sample right of the threshold. I opt for a tighter trimming that spans 80 percent of the error support but maintains two-thirds of my sample right of the threshold.

Appendix A3: Data Linking Process

The linking process begins by cleaning the Harris and Fort Bend Central Appraisal Districts' property owner name fields. Plural referential phrases (e.g., "et uxor") and their abbreviations (e.g., "et ux.") are omitted as well as commonly used titles (e.g., "Mr.", "Mrs.", etc.). Suffixes (e.g., "Junior", "Senior", etc.) are standardized and abbreviated.

In Step 1 I link more than 17,000 movers (57 percent) to the Infutor data based on an exact match of their names and addresses at the time of Hurricane Harvey. The imperfect match rate is likely driven by lack of standardization of names and addresses across data sources. For example, the use of a middle initial rather than a middle name would result in a failure to match.

Figure A3: Linking Central Appraisal District Data to Infutor

Step 1: Identify Exact Matches

Movers	Jaro-Winkler Scores	Infutor
James D. Jones	1.0	James D. Jones
122 Main St	1.0	122 Main St
	2.0	
Jane Sue Doe	0.91	Jane Doe
456 W. Magill Ave	0.93	456 West Magill Avenue
	1.84	
Jerry Smith	0.39	Smith Jerry
789 Meadowland Blvd	1.0	789 Meadowland Blvd
	1.39	
Eric Lee	1.0	Eric Lee
321 Sunny Lane	1.0	321 Sunny Lane
	2.0	

Match List

Name	Address
James D. Jones	122 Main St
Eric Lee	321 Sunny Lane

Step 2: Identify High-Quality Unmatched Movers

Movers	Jaro-Winkler Scores	Infutor
Jane Sue Doe	0.91	Jane Doe
456 W. Magill Ave	0.93	456 West Magill Avenue
	1.84	
Jerry Smith	0.39	Smith Jerry
789 Meadowland Blvd	1.0	789 Meadowland Blvd
	1.39	

Match List

Name	Address
James D. Jones	122 Main St
Eric Lee	321 Sunny Lane
Jane Sue	456 West Magill Avenue

I allow for more flexible matching in Step 2 using the Jaro-Winkler string-distance algorithm, which scores the similarity of strings between 0 and 1 for no similarity and exact matches, respectively. Since I link movers based on both names and addresses, the combined Jaro-Winkler score ranges from 0 to 2. I consider combined scores above 1.8 to be extremely accurate, and the inclusion of these high-quality matches increases the number of movers whom I observe to 22,901 (77 percent).

I extend the analysis further by analyzing how movers' housing consumption differs between their pre- and post-storm addresses. Specifically, I take the movers who matched to Infutor, and I match them a second time based on their post-storm address to CAD data across Texas obtained from

CoreLogic.[11] This second phase of matching repeats the methodology used in the first phase, considering only matches with at least a 1.9 combined Jaro-Winkler score. I successfully match more than 3,000 movers across Texas, all of whom maintained owner occupancy despite selling their pre-Harvey home.

Note that if a mover transitions out of owner occupancy, they cannot be accurately matched in the second phase because they do not own their post-storm residence. Consequently, I perform an exact match based solely on post-storm addresses to the statewide CAD data to learn more about the housing consumption of households who transitioned into renter occupancy.

Appendix A4: Additional Outcomes

The Owner Transfer dataset also contains information on foreclosures, an outcome indicative of financial distress and hardship.[12] Foreclosures are a relatively rare event, with an average of 6 foreclosures per hundred households in Houston between 2010 and 2016.

Foreclosures offer an alternative (yet infrequent) mobility outcome that may be particularly sensitive to physical property damage. Homeowners faced with repair costs may default on mortgage payments if budget constraints are binding, and policymakers recognize this threat by often implementing foreclosure moratoria in the aftermath of a disaster. Appendix Table ? provides results for the impact of damage on foreclosures, where I find tightly estimated null effects through 2017 when the moratorium was effective. The 95 percent confidence intervals encompass zero throughout the post-Harvey period, but there are signs of a temporary spike in the summer of 2018, when the first Harvey-related foreclosure proceedings occurred. This potential uptick was short lived, and point estimates hover below zero through the remainder of the analytical period. Importantly, my outcome only accounts for foreclosures of a homeowner's pre-storm address, but damaged households may experience differential risk of foreclosing on their next residence.

[11] I do not have access to CoreLogic's national data, preventing me from linking out-of-state movers to their post-storm addresses.

[12] I use a broader definition of foreclosure that includes foreclosure deeds, deeds of trust that specify foreclosure, as well as deeds in lieu of foreclosure, the latter of which is colloquially known as a "friendly foreclosure."

Appendix A5: Robustness Tests

The general specification outlined in Equations (2) and (3) require an arbitrary bandwidth and donut-size selection. Based on sample-size constraints and the structure of elevation measurement error detailed in Section 5 and Appendix A2, I use a bandwidth of [-120.0,-18.7] and [6.6,36.0] to estimate the various discontinuity parameters. The following figures explore the robustness of this bandwidth selection.

Figure A5.1: First Stage (Robustness)

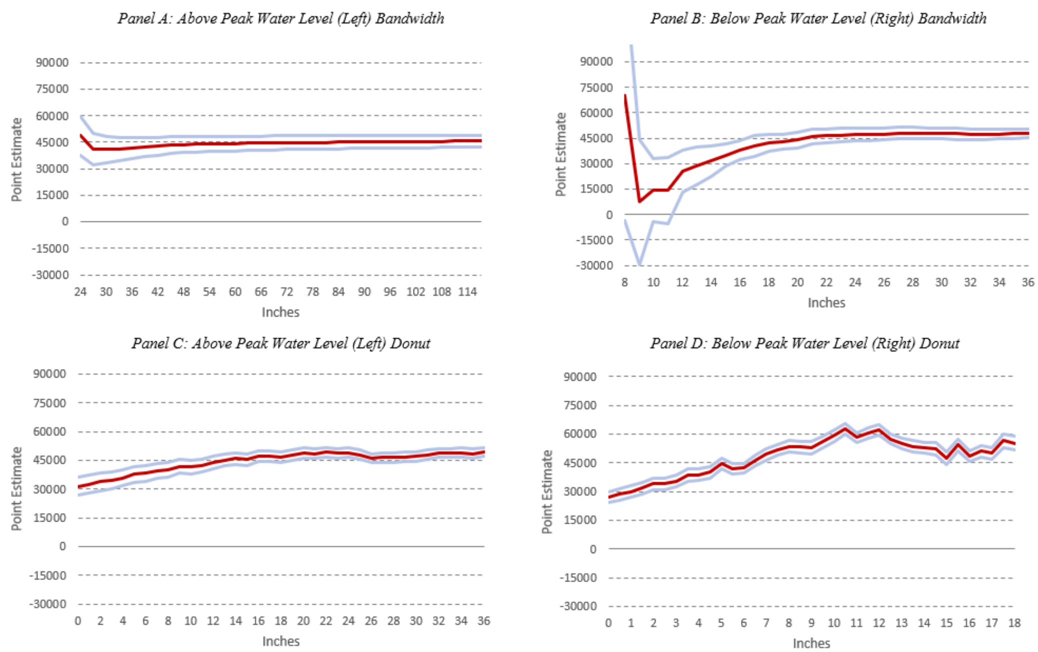

Notes: Panel A depicts how parameter estimates from Equation (2) evolve based on a bandwidth of $[a, -18.7] \cup [6.6, 36]$, where $a \in [-24, -120]$. Panel B depicts how parameter estimates from Equation (2) evolve based on a bandwidth of $[-120, -18.7] \cup [6.6, b]$, where $b \in [8, 36]$. Similarly, Panels C and D iterate through different donut sizes from 0 (no donut) to the -36 and 18 inches, respectively.

Figure A5.2: Year of Construction (Robustness)

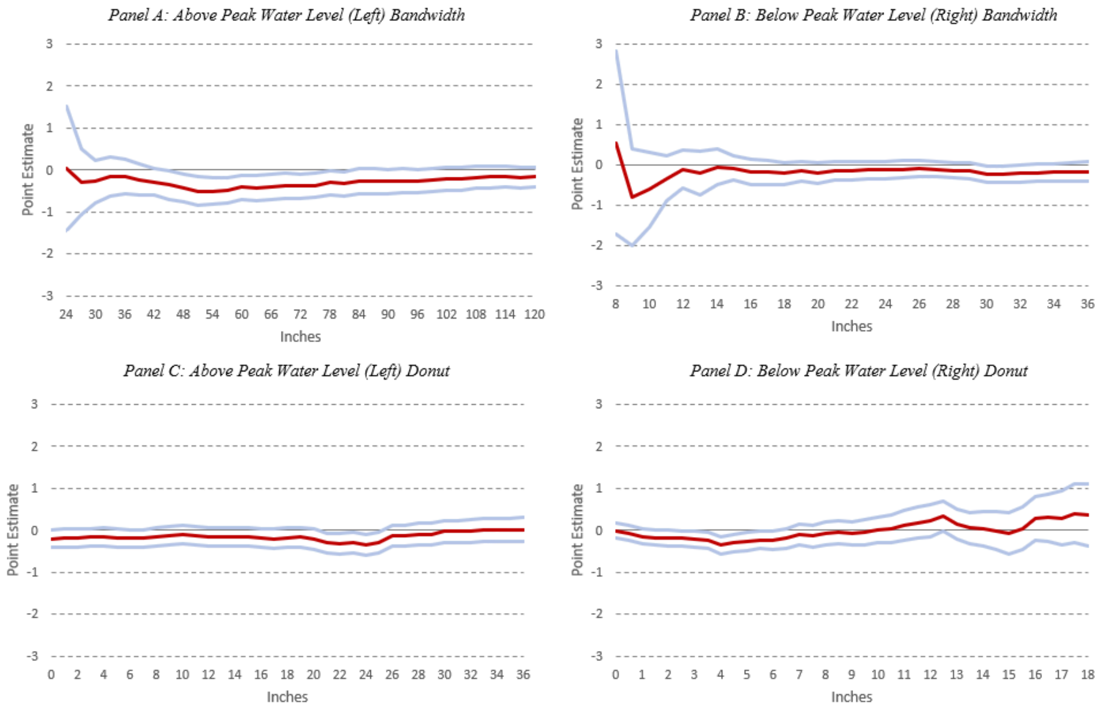

Notes: Panel A depicts how parameter estimates from Equation (3) evolve based on a bandwidth of $[a, -18.7] \cup [6.6, 36]$, where $a \in [-24, -120]$. Panel B depicts how parameter estimates from Equation (2) evolve based on a bandwidth of $[-120, -18.7] \cup [6.6, b]$, where $b \in [8, 36]$. Similarly, Panels C and D iterate through different donut sizes from 0 (no donut) to the -36 and 18 inches, respectively.

Figure A5.3: Home Size (Robustness)

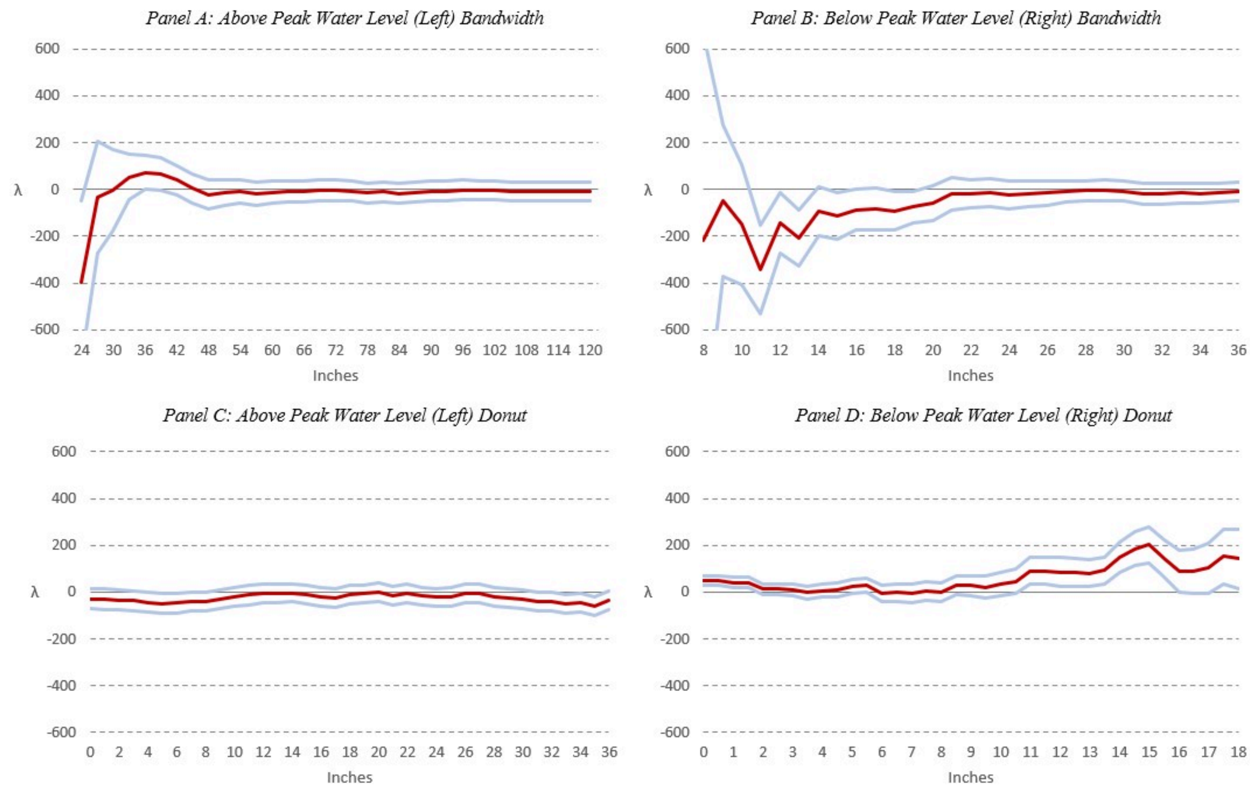

Notes: Panel A depicts how parameter estimates from Equation (3) evolve based on a bandwidth of $[a, -18.7] \cup [6.6, 36]$, where $a \in [-24, -120]$. Panel B depicts how parameter estimates from Equation (2) evolve based on a bandwidth of $[-120, -18.7] \cup [6.6, b]$, where $b \in [8, 36]$. Similarly, Panels C and D iterate through different donut sizes from 0 (no donut) to the -36 and 18 inches, respectively.

Figure A5.4: Short Run Moves (Robustness)

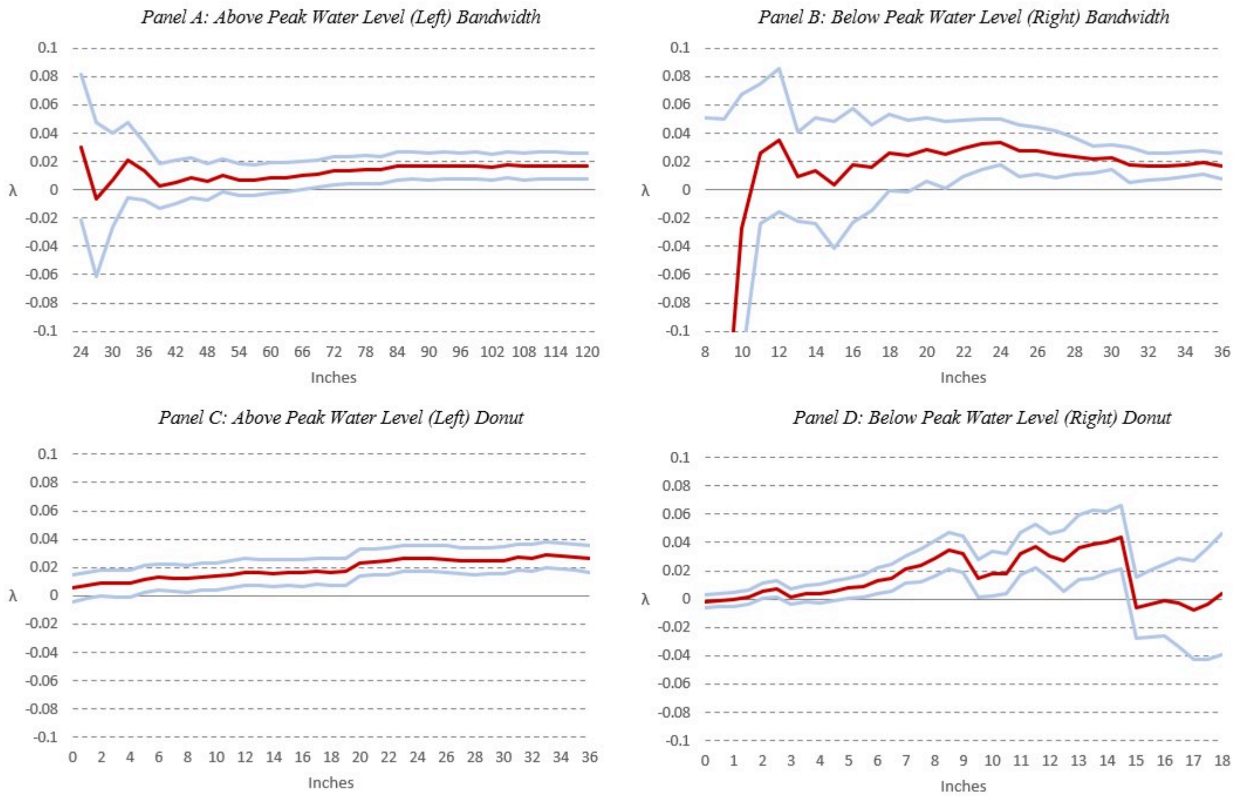

Notes: Figure A5.4 is based on an indicator for Infutor observing at least one individual in household j moving between September 2017 and March 2018. The regression estimates on Equation (3), a local linear regression with uniform kernels and standard errors are clustered at the neighborhood level. Panel A depicts how parameter estimates from Equation (3) evolve based on a bandwidth of $[a, -18.7] \cup [6.6, 36]$, where $a \in [-24, -120]$. Panel B depicts how parameter estimates from Equation (2) evolve based on a bandwidth of $[-120, -18.7] \cup [6.6, b]$, where $b \in [8, 36]$. Similarly, Panels C and D iterate through different donut sizes from 0 (no donut) to the -36 and 18 inches, respectively.