Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Asymmetric Housing Co-adaptation Reshapes Neighborhood Recovery Trajectories after Hurricanes

Susu Xu

sxu83@jhu.edu

Johns Hopkins University https://orcid.org/0000-0001-7170-648X

Xuechun Li

Johns Hopkins University

Ruxiao Chen

Jonhs Hopkins University

Yanling Sang

Johns Hopkins University

Graham Hults

Johns Hopkins University

Jingxiao Liu

UC Berkeley

Kenneth Harrison

National Institute of Standards and Technology

Juan Fung

National Institute of Standards and Technology

Yongjun Zhang

Stony Brook University

Cristina Dragomir

New York University

Maurizio Porfiri

New York University Tandon School of Engineering https://orcid.org/0000-0002-1480-3539

Luis Ceferino

University of California Berkeley https://orcid.org/0000-0003-0322-7510

Article

Keywords:

Posted Date: November 2nd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10887426/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: There is **NO** Competing Interest.

Asymmetric Housing Co-adaptation Reshapes Neighborhood Recovery Trajectories after Hurricanes

Xuechun Li[1], Ruxiao Chen[1], Yanling Sang[1], Graham Hults[1], Jingxiao Liu[2], Kenneth Harrison[3], Juan F. Fung[3], Yongjun Zhang[4,5], Cristina Dragomir[6], Maurizio Porfiri[7], Luis Ceferino[2], and Susu Xu[1,8]

[1]Department of Civil and Systems Engineering, The Johns Hopkins University, Baltimore, MD 21218, USA

[2]Department of Civil and Environmental Engineering, UC Berkeley, Berkeley, CA 94720, USA

[3]National Institute of Standards and Technology, Gaithersburg, MD 20899, USA

[4]Institute for Advanced Computational Science, Stony Brook University, Stony Brook, NY 11790, USA

[5]Department of Sociology, Stony Brook, NY 11790, USA

[6]School of Liberal Studies, New York University, New York, NY 10012, USA

[7]Center for Urban Science and Progress, New York University, New York, NY 10012, USA

[8]Data Science and AI Institute, The Johns Hopkins University, Baltimore, MD 21218, USA

Abstract

Post-hurricane housing co-adaptation is a collective process in which household socioeconomic conditions, housing damage, and past activities interact with household-to-household influences to shape repair and/or sell decisions and further change community recovery trajectories. Existing research provides limited process-level understanding, constraining when and where recovery resources can be effectively targeted. We develop the Household-level Decision Diffusion (HDD) framework to reconstruct and forecast dynamic housing co-adaptation processes for the first time. We evaluate HDD by constructing a first-of-its-kind longitudinal dataset linking damage, housing and household conditions, time-series repair permits and time-series property transactions for more than 1.3 million households in nearly 2,800 communities after Hurricanes Harvey, Ida and Ian. HDD recovered latent neighborhood structure with higher edge precision than distance-based or classical Hawkes reconstructions, and forecasted both future repair and sale with AUC value of upto 0.94 and 0.82, respectively, significantly improving both decision prediction accuracy compared to state-of-the-art models and providing a process-level basis for effective recovery policy design. The revealed dynamics show that same-type neighborhood associations were positive for both decisions in all three events, but repair reinforcement was bounded and early while sale reinforcement accumulated across the observation window. Cross-type coupling was asymmetric in both magnitude and sign: the sale-

to-repair pathway was more strongly coupled than the reverse in 69–82% of communities, yet its realized effect was negative, with a greater sale history among connected neighbors associated with a lower probability of repair in almost every community after Ian and in most after Ida. Neighborhood influence is therefore better described as a two-channel process than as a single social multiplier: within-channel histories reinforce continued activity, whereas the exit-to-investment channel inhibits it. Neighborhood influences extended to selling among undamaged households as well. Recovery policy should therefore concentrate repair support in the early window, sustain exit-related monitoring long after it closes, avoid delivering buyouts and repair assistance into the same community at once, and assess eligibility at community rather than parcel scale.

1 Introduction

Tropical cyclones, also known as hurricanes or typhoons, are the costliest natural hazard in the United States. Since 1980, they have caused more than $1.5 trillion in damage, averaging approximately $23 billion per event [1]. As of 2025, more than 60 million people lived in U.S. Atlantic and Gulf coastline counties [2], and the hurricane hazards those communities face keeps intensifying [3, 4]. During the past decade, three of the five costliest hurricanes in continental U.S. history struck the Gulf Coast, including Harvey (2017), Ida (2021), and Ian (2022) [5]. Their strong winds, intense rainfall, and storm surge caused extensive building damage and displaced more than 3.5 million people [6–9].

After hurricanes, housing adaptation is a critical problem of enormous societal scale. Given the large population living in hurricane-prone areas, whether damaged houses are restored affects residents' ability to return and remain after a hurricane and further impact the long-term recovery of local businesses and economic activity. In particular, homeowners face a sequence of housing adaptation decisions about whether and when to repair or rebuild a damaged house, sell it, or defer action. These decisions are not necessarily one-time events, for example, a property may be repaired and later sold, or sold before a new owner initiates repairs. Moreover, households can influence one another because the neighborhood recovery environment is shaped partly by the actions of all households within it. One hypothesis that part of existing studies employed is that nearby repairs can encourage homeowners to remain and reinvest, whereas property sales or abandonment can discourage further investment and even owners of undamaged properties may consider selling and leaving [10–12]. These household-to-household interactions and

influences are often substantial, for example, after Hurricane Katrina, neighborhood influences were estimated to amplify rebuilding induced by the Road Home assistance program by 27% [10]. On the other hand, distressed exits have also been estimated to trigger 0.3–0.6 additional nearby foreclosures within a 0.10-mile radius [13]. Such neighborhood influences are deeply coupled with households' own decision histories and operate through various complex combination of distinct channels, including social ties that transmit repair-related information and facilitate mutual assistance and coordination [12, 14], and housing-market externalities through which nearby sales change neighborhood conditions and property values, thereby altering households' expected returns to reinvestment even without direct social contact [10, 11, 15]. Housing co-adaptation is therefore a recurrent, stochastic collective process: household decisions respond to the evolving actions of others while also changing the environment in which subsequent decisions are made. However, ***do neighborhood influences always reinforce community recovery? How do neighboring repairs impact a household's sale versus sale impact repair? How are the influences of various neighborhood activity sequences coupled with household's own socio-demographic background and activity history to affect future housing decisions, and how do they accumulate into community-level recovery trajectories?*** There still remain many critical and fundamental questions about housing co-adaptation unresolved as we lack a systematical tool to look into the detailed dynamics of the housing co-adaptation process.

Understanding and predicting such spatio-temporal dynamics of parcel-level housing co-adaptation process across communities is important for effective housing recovery and strengthening long-term coastal resilience. Since 2017, the U.S. government has appropriated roughly $35 billion through the Community Development Block Grant–Disaster Recovery program [16]. However, post-hurricane community recovery remains largely uneven. Across 2,195 census tracts affected by 12 extreme-weather events, lower-income neighborhoods were substantially less likely to rebuild and often did not return to their pre-disaster state, even with existing aid and insurance [17]. One major reason is that recovery resources are currently allocated largely using static household-level indicators, including housing damage, financial loss, and socioeconomic vulnerability. These indicators do not capture how repair and sale decisions interact and propagate among households over time. Households with similar damage and vulnerability can consequently be embedded in different recovery environments, and the same intervention may have different effects depending on when and where it is implemented. Therefore, process-level

understanding and prediction are critical and fundamental to estimate whether the post-hurricane recovery trajectory of each community will stall or accelerate, and to determine for each affected community, when and where limited interventions can be most effective in reducing persistent housing abandonment, population loss, and community resilience decline.

However, in the housing recovery research community, existing studies often ignore the dynamic and interactive nature of the housing co-adaptation processes. Most household-level models represent repair or sale as a single, terminal outcome rather than as a recurrent and co-evolving sequence of decisions (e.g. repair then sale) across interacting households. Based on this assumption, engineering and Markov-type models treat recovery as an individual process with no explicit neighbor influence [18], with statistical models attributing an observed repair outcome to damage, financial resources, socioeconomic characteristics, and local conditions [19]. A separate literature focuses on sale events alone and similarly treats the decision to leave in isolation from evolving neighboring decisions [20]. A few structural and agent-based models incorporate neighbor influence explicitly [10, 12]. But they mainly focus on neighboring repair alone and compress heterogeneous decision sequences into an aggregate or prespecified neighborhood effect stock. Such oversimplification overlook the dynamic nature of neighbor influences and represent them only through an aggregate recovery state. Consequently, variation generated by interdependent neighborhood decisions can instead be mis-attributed to factors of household, damage, or financial conditions, leading to inaccurate and misleading knowledge about the mechanism that actually drives recovery and poor generalization to future housing adaptation decision prediction.

To overcome these limitations, we develop a *first-of-its-kind framework*, Household-level Decision Diffusion modeling and inference (HDD), to reconstruct and predict post-disaster housing co-adaptation as an evolving process of interacting repair and sale decisions. HDD formulates household adaptation as a cross-type neural graph Hawkes process that jointly learns latent neighborhood dependence and the temporal dynamics of sparse repair and sale events, conditioned on household socioeconomic characteristics, housing and damage conditions, and community context. Rather than prescribing neighborhood connectivity or reducing surrounding decisions to an aggregate recovery state, HDD learns how preceding adaptation histories contribute to subsequent repair and sale propensities and recursively unfolds the learned process to forecast future household decision sequences. To evaluate this process-level inference at scale, we construct a *first-of-its-kind* integrated longitudinal, large-scale, parcel-level post-hurricane re-

covery dataset that links building damage and hazard exposure, housing characteristics, socioeconomic conditions, time-series repair permits, and time-series housing transactions for more than 1.3 million households across over 2,800 communities affected by Hurricanes Harvey, Ida, and Ian. Repair permits and property transactions are held in separate municipal and county systems with no common identifier, so post-disaster recovery is typically observed either as a one-time snapshot or as a small survey panel aggregated to census tracts. To our knowledge, no existing dataset provides household and housing conditions, building damage, and both sale and repair records for individual households through time and across multiple events.

Our results reveal that post-disaster housing co-adaptation is characterized by distinct and asymmetric diffusion processes that help explain divergent community recovery trajectories. Repair-related influence is strongest early and decays as recovery progresses whereas sale-related influence accumulates more persistently. These processes are strongly asymmetric: across the three hurricanes, sale-to-repair influence exceeds the reverse repair-to-sale pathway in approximately 69–82% of communities. Such sale-to-repair influence is predominantly inhibitory, in which neighboring sale activity with reduced repair. Exit therefore enters collective recovery dynamics not by propagating activity but by withdrawing it. Neighborhood influences also extends beyond directly damaged households, consistently increasing selling among undamaged households within affected communities, while its consequences vary markedly across hazard contexts—activating recovery after Harvey but suppressing repair and increasing prolonged inactivity after Ian. HDD further improves future decision prediction (AUC up to 0.94 for repair and upto 0.82 for sale, against 0.48–0.79 for the best alternative in each case) and reconstruction of independently referenced neighborhood dependence relative to alternative econometric, spatial, agent-based, and temporal point-process approaches. In summary, we found that community recovery is jointly shaped by initial damage and household vulnerability, as well as which adaptation decisions propagate, when their influence occurs, and how those interactions unfold across households, providing a process-level basis for four key improvements for recovery policy design and implementation. First, repair support has a deadline that exit monitoring does not: repair reinforcement fades within months while sale reinforcement keeps accumulating, so settlement speed, permitting throughput and contractor supply matter early, whereas displacement and distressed-sale monitoring must continue long after repair programs close. Second, early repair rates can be tracked during recovery to identify which communities are head-

ing toward majority repair, rather than diagnosed once recovery has ended. Third, buyouts and repair assistance should not be concentrated in the same community at the same time, because neighboring sales suppress repair, so how the two are sequenced may matter as much as how much is spent. Finally, eligibility should be assessed at community rather than parcel scale, because a disaster reshapes decisions in undamaged households as well as damaged ones.

In the rest of this paper, we first introduce our HDD framework that models the complex, interacting dynamics of parcel-level housing co-adaptation process and reconstruct the underlying community influence networks. Then we introduce the evaluation of HDD framework on multiple costliest Gulf Coast hurricane events by applying to our novel large-scale parcel-level longitudinal datasets spanning three hurricanes and covering 1.3 million households. Based on the learned dynamics and reconstructed community influence networks, we further investigated whether neighborhood activities always have positive influences on a household's same-type decisions, and how neighborhood activities affect a household's cross-type decisions, how community influence networks are restructured before and after hurricanes, and how the neighborhood influences couple and accumulate to change the community recovery trajectories. We further provide discussions about the results, their societal and policy implications, and several limitations in Discussion section. We provide more detailed formulation and algorithm descriptions in Methods section. More figures and information can be found in supplementary file.

2 Results

Modeling Parcel-level Housing Co-adaptation Process

We develop a Household-level Decision Diffusion (HDD) framework to reconstruct and predict the parcel-level housing co-adaptation process from longitudinal repair and sale observations (Fig. 1; see Methods). We conceptualize each community as a coupled stochastic event system. In particular, conditioned on hazard damage, household and housing conditions, each household's probabilities to repair or sell jointly evolve with community context, their own decision history and the evolving decision histories of interdependent households. HDD decomposes these dynamics into a background adaptation propensity, self-history effects, and neighborhood-mediated effects, and jointly reconstructs the latent influence network structure through which the latter propagate. Given observed household and community con-

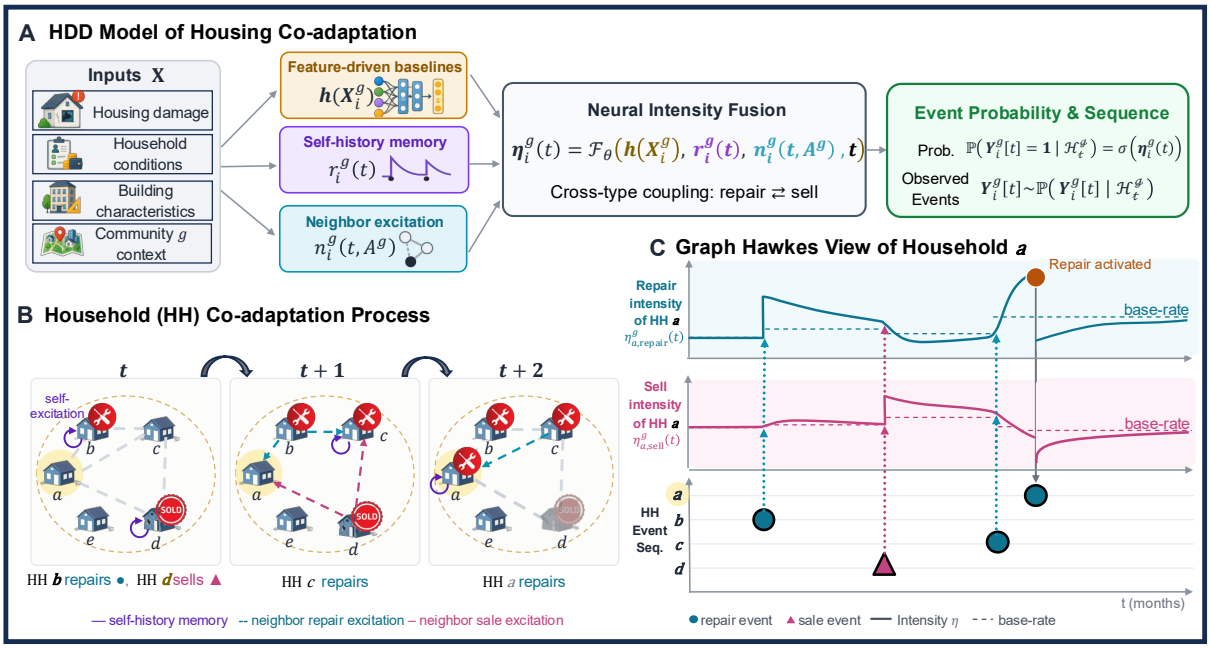


Figure 1: **Overview of the Household-level Decision Diffusion (HDD) framework and the modeled housing co-adaptation process.** **(A)** HDD represents each household's repair and sale propensities using three sources of information: feature-driven baselines derived from hazard damage and household, building, and community conditions; temporally decaying self-history states; and neighborhood histories propagated through latent influence networks. A neural intensity function integrates these signals, allowing repair and sale histories to interact through same-type and cross-type pathways. The resulting logits are transformed into monthly household-level event probabilities, from which repair and sale events are modeled as Bernoulli realizations. **(B)** Illustrative sequence of household co-adaptation within one community. At month t, household b repairs and household d sells. Their histories propagate through directed neighborhood pathways and contribute to subsequent household decisions: household c repairs at $t+1$, followed by household a at $t+2$. After selling, household d becomes inactive in subsequent network states, while signals already propagated to other households remain represented in their decaying histories. Purple loops denote self-history contributions; blue and magenta arrows denote neighborhood repair- and sale-history contributions, respectively. **(C)** History-dependent repair and sale propensities for household a corresponding to the sequence in panel B. Vertical dotted lines identify preceding repair events at households b and c and the sale event at household d; dashed horizontal lines denote feature-driven baseline propensities. These histories jointly update the subsequent repair and sale propensities of household a. After a's repair event is realized, both propensities decrease and gradually return toward their respective baseline levels. The curves are schematic representations of monthly model dynamics.

ditions and repair and sale histories, the framework jointly reconstructs this latent influence network structure and the temporal co-adaptation dynamics governing it, and estimates time-varying repair and sale probabilities for each household. This formulation shifts the unit of analysis from isolated household outcomes to the co-evolving process through which adaptation decisions interact across a community.

The main challenges associated with formulating the complex co-evolving process across a community are that (1) the neighborhood influence network is unobserved and distinct from conventional social network defined by distance, mobility, or family ties [21–23], for instances, households can respond to others' visible repairs or sale activity without directly interacting with them; (2) housing adaptation events are extremely sparse in time domain, but the latent process generating them is high-dimensional which jointly couples household conditions, community environments, neighboring decision histories and decision types, making the underlying co-adaptation dynamics fundamentally difficult to identify from data; (3) housing co-adaptation mechanisms (i.e. how neighbor influence propagates through network and couples with individual condition and self-history) are highly heterogeneous across communities, making mechanisms learned from single community difficult to generalize.

HDD addresses these challenges by formulating housing co-adaptation as a cross-type neural graph Hawkes process. First, rather than assuming a neighborhood network based on geographic or social proximity, HDD learns a latent community influence structure from observed adaptation histories conditioned on household characteristics and spatial context. Second, HDD jointly models sparse repair and sale events through neural intensity functions that integrate household-specific adaptation propensity, self-history, and neighbor-history effects. Unlike conventional Hawkes processes with predefined parametric excitation, or neural temporal point processes that model complex event histories without explicitly learning the underlying neighborhood structure, this formulation couples event dynamics with the latent network over which they propagate. It therefore allows different types of decisions to interact through distinct temporal pathways, rather than assuming a common or predefined diffusion mechanism. Third, HDD shares the parameterization of the neural functions governing these co-adaptation mechanisms across communities but still learn community-specific influence structures. Such design allows statistical strength to be pooled across sparse event sequences without assuming identical neighborhood dynamics.

To jointly reconstruct the latent influence network and co-adaptation dynamics, we first approximate

the likelihood of observed decision sequences of all households in a community by factorizing it over the preceding household decision history, latent neighborhood influence network, and background factors like household and housing conditions and community context (see Methods). Following the Hawkes-process formulation, preceding repair and sale events are encoded as temporally decaying history signals and propagated through the latent network. A shared neural intensity function then maps these multivariate self- and neighborhood histories to the conditional probabilities of subsequent (group) repair and sale events. We finally define the inference loss as the class-weighted negative log-likelihood of the observed event sequences, measuring how well the inferred network and co-adaptation dynamics explain the observed timing and type of household decisions, together with a regularization term that promotes a sparse influence structure. We jointly optimize the network reconstruction and neural intensity parameters by backpropagating this loss. To make this high-dimensional inference feasible under sparse household observations, we share the neural co-adaptation parameters across community sequences while still reconstructing a community-specific influence structure for each community. Mini-batch stochastic gradient optimization across communities further enables the framework to scale to millions of households. The resulting fitted model jointly reconstructs the latent neighborhood influence network structure and the co-adaptation mechanisms that generate the observed repair and sale sequences, and can subsequently be applied to held-out histories to predict future household decisions. Moreover, our HDD framework can flexibly incorporate additional household, hazard, or community observations through its background propensity representation and additional adaptation decisions through its multivariate event representation. The joint inference formulation is therefore generalizable to other settings in which sparse, heterogeneous decisions co-evolve through an unobserved interaction structure.

Evaluation of HDD across Multiple Gulf Coast Hurricanes

We evaluate HDD using dataset from three major Gulf Coast hurricanes: Harvey in Harris County, Texas (2017), Ida in New Orleans, Louisiana (2021), and Ian in Lee County, Florida (2022). These events span contrasting hazard and recovery contexts: Harvey was dominated by widespread rainfall-driven flooding, Ida by severe wind damage, and Ian by coastal wind and storm-surge impacts [7, 24, 25]. For each hurricane, we construct longitudinal household-level observations that jointly track structural damage, repair activity, and housing transactions together with upto 48 features that characterize household,

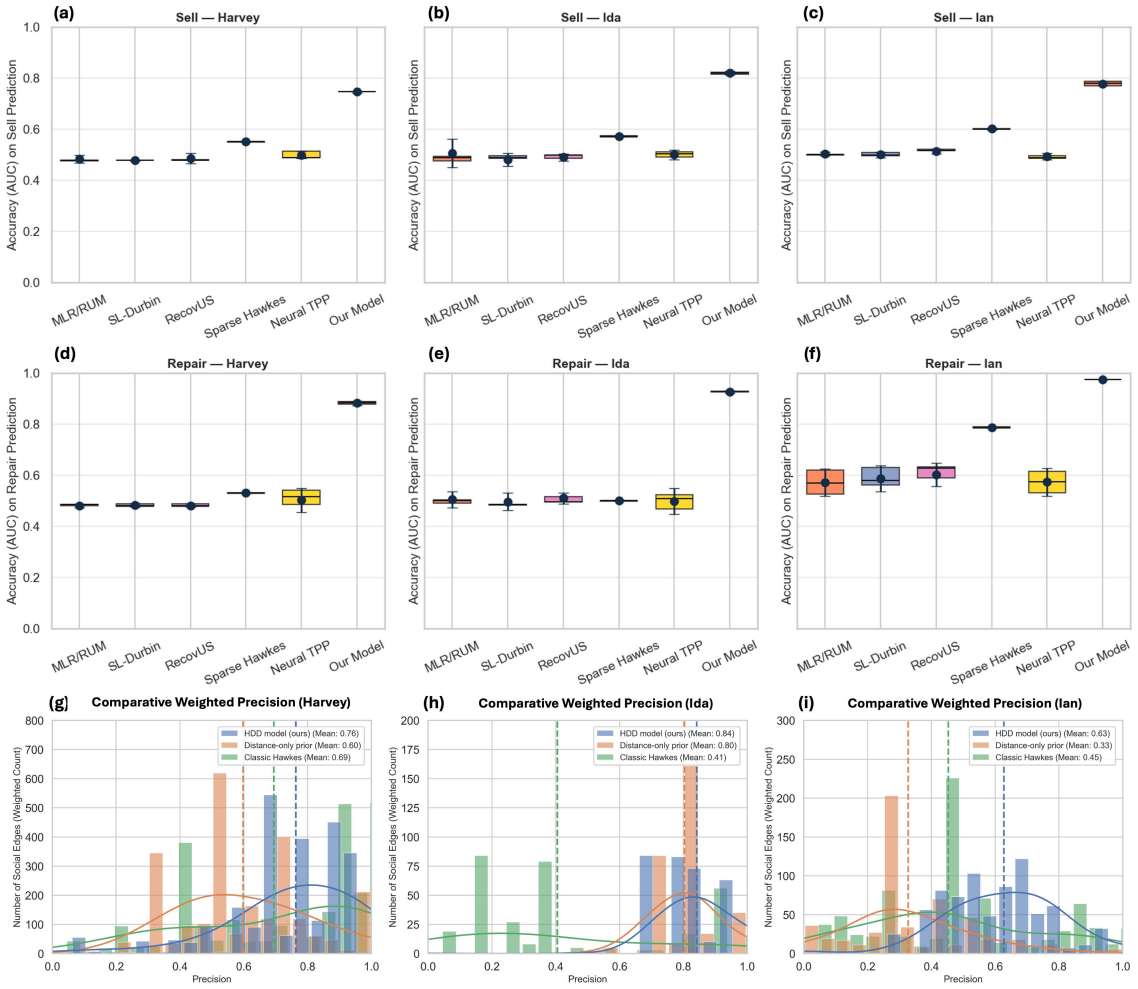


Figure 2: **HDD improves post-disaster decision forecasting and recovers externally plausible influence structure across three U.S. hurricanes.** (a–f) Event-state prediction performance for household sell and repair decisions following Hurricanes Harvey, Ida, and Ian. Across all events and both decision types, HDD achieves the highest AUC relative to benchmark approaches spanning covariate-driven discrete choice, spatial econometric spillovers, recovery simulation, and temporal point-process formulations, indicating that post-disaster adaptation is shaped by learned neighborhood interdependence beyond static covariates or distance-based proximity kernels. (g–i) Community-level weighted precision distributions for diffusion-network reconstruction, evaluated against mobility-derived proxy graphs as an independent behavioral reference. HDD consistently yields higher precision than both a distance-only prior and a classic Hawkes baseline, suggesting that adaptive, community-specific structure learning recovers more socially meaningful interaction pathways than geographic proximity or fixed temporal excitation alone. Dashed vertical lines denote mean precision across communities.

housing, and socioeconomic (including financial resource) conditions. The resulting dataset contains 1,017,324 households across approximately 2,000 communities for Harvey, 121,135 households across approximately 380 communities for Ida, and 181,659 households across 399 communities for Ian, providing more than 1.3 million households data with complete background, physical damage, and house adaptation histories across nearly 2,800 communities. The contrasting hazards, housing systems, and adaptation patterns across these events provide a heterogeneous empirical test bed for evaluating whether HDD can consistently reconstruct housing co-adaptation processes across communities.

We evaluate HDD along two complementary dimensions: (1) estimation of household decision dynamics and (2) reconstruction of the latent neighborhood influence network structure. For decision dynamics, We evaluate HDD's ability to predict held-out repair and sale events and compare its predictive performance against traditional models, including random-utility model, spatial econometric model, RecovUS-style agent-based model, conventional multivariate Hawkes process model, and neural temporal point process model (see Methods). Prediction performance is primarily evaluated using the area under the receiver operating characteristic curve (AUC), with average precision and Recall@K reported in Supplementary Information to assess rare-event performance under strong class imbalance. For influence-structure reconstruction, since direct empirical ground truth is unavailable, we use two complementary evaluations: (1) comparison with independently derived mobility-based connectivity in the same communities and (2) comparisons on simulation data with synthesized networks from real-world household characteristics.

Across all three hurricanes, HDD consistently provides the strongest held-out reconstruction of both repair and sale decisions (Fig. 2a–f). Specifically, HDD reached a AUC of 0.88 for repair and 0.75 for sale after Harvey, 0.93 and 0.82 after Ida, and 0.94 and 0.78 after Ian as shown in Fig. 2a–f. The margin is largest precisely where the benchmarks fail. For repair after Harvey and Ida, repair events are unpredictable from household covariates, spatial proximity, or flexible temporal dependence alone, and become predictable only once the community-specific relational structure over which histories interact is resolved. Compared to repair, sale events are more unpredictable was also the more predictable in all three hurricanes, indicating that sale timing carries more variation. The comparison with random-utility model shows that it is important to incorporate dynamic neighborhood influence. The improvement over spatial econometric and distance-based models indicates that neighborhood dependence cannot be adequately

represented by geographic proximity alone. HDD also outperforms the RecovUS-style simulation, suggesting that learning interaction structure from observed decision sequences provides information beyond prespecified neighborhood mechanisms. Finally, improvements over both conventional Hawkes and neural temporal point-process models show that flexible temporal dependence alone is insufficient: accurately reconstructing household adaptation also requires resolving the community-specific relational structure over which those histories interact. We next evaluate whether the latent structures supporting these improvements contain meaningful relational information. Because the true housing-adaptation influence network is unobservable, we construct mobility-based connectivity as an independent behavioral reference that is never used for HDD training. Across Harvey, Ida, and Ian, HDD achieves higher weighted edge precision against these reference networks than both geographic distance and conventional Hawkes reconstruction (Fig. 2g–i). Mobility connectivity is nevertheless not treated as ground truth for housing co-adaptation. Households can respond to place-based signals without direct mobility-based interaction. But the comparison can tests whether HDD recovers behaviorally plausible relational structure rather than whether it reproduces a conventional social network.

We further examine structural identifiability using synthetic co-adaptation processes in which both the event-generation mechanism and latent influence structure are known and household characteristics come from real communities. We systematically vary network density, the relative strength of self- and neighbor-mediated excitation, temporal decay, observation horizon, and training-data availability to reproduce different levels of difficulty encountered in sparse household adaptation data (Supplementary Figs. S5–S7 and Tables S4–S10). HDD maintains high edge precision across these regimes and recovers the underlying event dynamics. The recall decreases primarily when influence signals become too sparse to support exhaustive edge identification. This asymmetric behavior is desirable for the present application: when evidence for a household connection is weak, HDD preferentially removes the connection rather than introducing a spurious influence pathway. These evaluations establish complementary aspects of HDD. We use the reconstructed processes in the following sections to examine how repair and sale signals evolve, interact, and contribute to divergent community recovery trajectories.

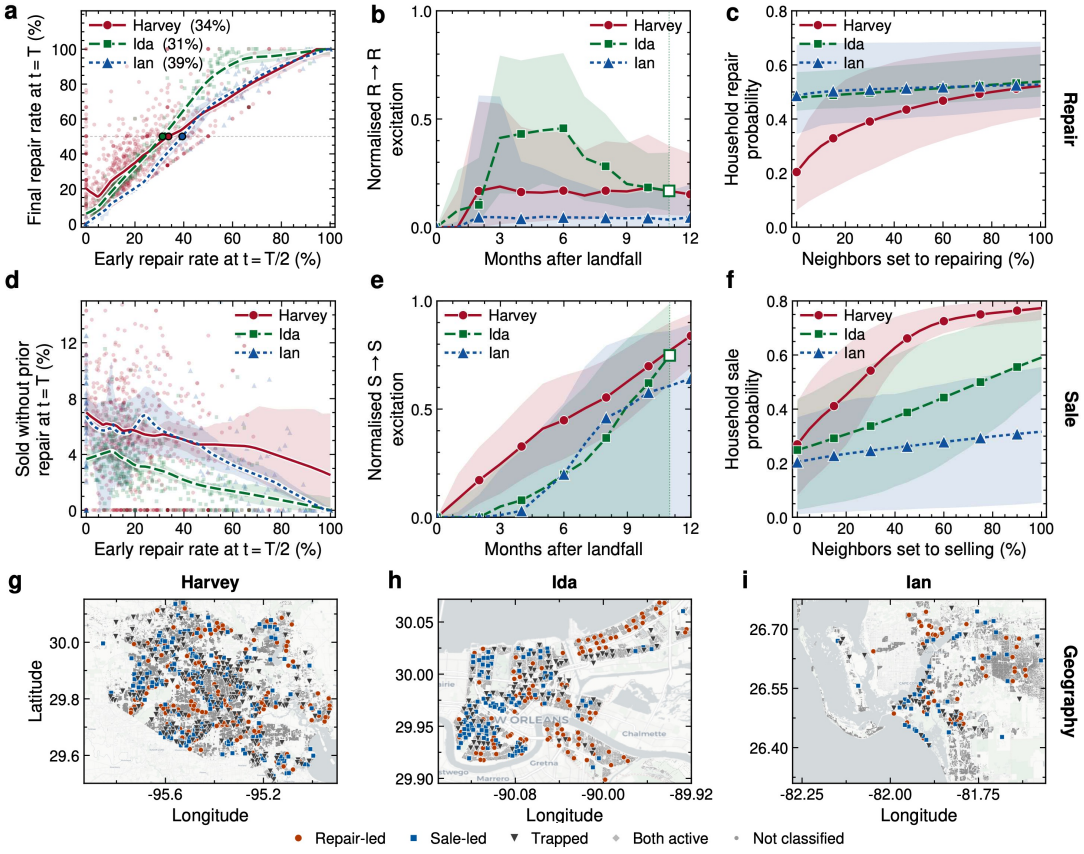


Figure 3: Community outcomes, neighborhood excitation, household responses, and the geospatial distribution of communities with different recovery trajectories after three hurricanes. (a) Final community repair rate at $t = T$ as a function of repair rate at $t = T/2$. Points denote communities, lines denote LOWESS fits, and shaded bands denote 95% bootstrap intervals. The horizontal dashed line marks 50% final repair; fitted crossing points were 34% for Harvey, 31% for Ida, and 39% for Ian. (b) Median normalized repair-to-repair (R→R) neighborhood excitation over time; shaded bands show the interquartile range across community–seed trajectories. (c) Mean fitted household repair probability when increasing fractions of neighbors were set to repairing; shaded bands show the interquartile range across households. (d) Fraction of households whose damaged properties were sold without prior repair at $t = T$, plotted against repair rate at $t = T/2$. Lines and shaded bands are LOWESS fits and 95% bootstrap intervals. (e) Median normalized sale-to-sale (S→S) neighborhood excitation over time; shaded bands show the interquartile range. (f) Mean fitted household sale probability when increasing fractions of neighbors were set to selling; shaded bands show the interquartile range. (g–i) Spatial distributions of community recovery trajectories after Harvey, Ida, and Ian. Classes are defined using within-hurricane median splits of final repair and sale rates: repair-led communities have above-median repair and below-median sale; sale-led communities have below-median repair and above-median sale; trapped communities are below the median on both; and both-active communities are above the median on both. Communities with insufficient number of damaged buildings are not classified. Basemap data © OpenStreetMap contributors.

Neighborhood Repair and Sale Histories Show Positive Same-type Associations with Household Decisions, but over Distinct Timescales

With the established the predictive performance of HDD, we examined how influences of a household's own and neighboring repair and sale dynamics reshape its subsequent activities across community and household scales (Fig. 3). A long-standing question we answered in this section is: *do neighborhood activities always have positive influences on a household's same-type activities?* Specifically, damaged house repairing is often treated as a strategic complement: each repaired structure can restore local amenity value, signal that the neighborhood remains viable, and increase the expected return to repairing nearby damaged properties or staying nearby undamaged properties. Recovery may thus produce a social multiplier. Structural equilibrium estimates following Hurricane Katrina support this interpretation, attributing to program-induced rebuilding a positive externality worth several thousand dollars for each neighboring household whose own rebuilding decision was unaffected [10]. However, neighboring repairs can also compete for locally scarce contractors, materials, and inspection capacity, thus increasing costs and delaying reconstruction which may possibly suppress nearby repairing. Distinguishing these mechanisms empirically is difficult because correlated exposure can produce the same spatial co-movement as household interaction. Households within a community often share flood depth, insurance coverage, income, and program eligibility, making clustered recovery difficult to distinguish from endogenous social influence. This is the identification problem formalized by Manski [26]. Within our learned HDD model, household-level decision timing and latent dependence structure are estimated jointly. This allows us to separate contributions from a household's own history and its neighbors' histories and to vary the neighborhood state while holding other model inputs fixed.

Our large-scale results confirmed that same-type associations were always positive for both decisions after all three hurricanes, but repair and sale differed in their duration as well as their strength. First, temporal dynamics of repair-to-repair influences and sell-to-sell influences exhibit different patterns. Repair-to-repair excitation increased during the first months after landfall and then declined or remained bounded, peaking within roughly six months in every event (Fig. 3b). Sell-to-sell excitation, by contrast, persistently increased across the entire observation window, indicating that sell-to-sell has longer influences across the community (Fig. 3e). Second, increasing of neighborhood repair and sale rate always leads to increasing of corresponding household repair probability and sale probability, which

confirms the positive association of same-type neighborhood influences, as shown in Fig. 3c,f. However, the magnitude of household response sensitivity to neighboring activities differs with adaptation decision type. Increasing the fraction of selling neighbors raised the household's own sale probability much more significantly than increasing the fraction of repairing neighbors raised the repair probability in every event: from 0.27 to 0.77 versus 0.20 to 0.52 after Harvey, from 0.25 to 0.59 versus 0.48 to 0.54 after Ida, and from 0.20 to 0.32 versus 0.49 to 0.53 after Ian (Fig. 3c,f). The sale response also remained larger on the log-odds scale, indicating that the difference was not solely attributable to different baseline probabilities. This contrast is consistent with structural differences between the two decisions. Repair is a coordination problem constrained by insurance settlements, contractor availability, and program deadlines, so complementarity may be strong but time-limited. Sale remains available over a longer horizon and may become more attractive as additional transactions provide information about neighborhood conditions and reduce the number of owners committed to remaining. The community-level relationship in Fig. 3a is consistent with this interpretation: final repair rate rises steeply with early repair rate and crosses 50% at early repair rates of 34%, 31%, and 39%. Repair reinforcement therefore appears to be concentrated early rather than accumulated throughout recovery. In contrast, the proportion of damaged properties sold without repair declined as early repair activity increased (Fig. 3d).

We next classified communities using within-hurricane median splits of final repair and sale rates (Fig. 3g–i). Repair-led communities combined above-median repair with below-median sale and sale-led communities showed the reverse pattern. Trapped communities were below the within-hurricane median on both outcomes, and both-active communities were above the median on both. All four trajectories occurred within each study area, although their spatial distributions differed among hurricanes. Because the classifications are relative to each hurricane, they distinguish within-event recovery trajectories rather than absolute levels of community recovery.

Asymmetric Cross-type Neighborhood Influences and Restructuring in Housing Co-adaptation Pathways after Hurricanes

Same-type reinforcement leaves a second question: how would a household's own and neighboring repair activities impact its future sale decision, and how would neighboring sale activities impact a household's future repair decisions? Neighborhood exits are known to propagate through housing markets [15].

Under irreversible-investment logic, repair is a sunk commitment whose return depends on an uncertain neighborhood trajectory. Observing neighboring sales may increase this uncertainty and the option value of waiting, thereby suppressing repair [27, 28]. Under a renewal mechanism, nearby sale may instead transfer damaged properties to owners with new capital and a greater willingness to invest and thus encourage nearby repair. Understanding which of these mechanisms dominate the cross-type influences (i.e. sale to repair and repair to sale) requires household-level observations of repair and sale for the same properties across multiple events, which remain rare in the recovery literature, and the spatially dispersed pattern of voluntary buyouts makes these relationships especially difficult to study using program records alone [20].

Our HDD framework answers this question at a large scale and through cross-event comparisons. As shown in Fig. 4, our results show an asymmetry in sign as well as magnitude of such cross-type influences. First, we found that S→R contributions dominated in most communities across all three hurricanes (Fig. 4a). Specifically, we calculate the absolute change in a household's decision intensity produced by one unit of cross-type history, evaluated from a common reference state in which no household has yet acted. For each community we set the sale history of a household and of its connected neighbors to one unit and record the resulting change in its pre-sigmoid repair intensity; we then repeat the procedure with repair history and record the change in sale intensity. Averaging each change over households and over months gives two values per community, and a community is classified as S→R dominant when the first exceeds the second in absolute value. Because both are read from the same zeroed history state and on the log-odds scale, this quantity measures the *strength of the cross-type coupling* given each community's learned housing co-adaptation mechanisms, i.e. how much weight the decision process places on each direction of transmission. This measurement is independent of the decision histories that community has accumulated, but the graph embedding, damage baseline and time baseline remain community-specific. It therefore describes coupling capacity rather than realized effect, and the absolute value is taken because the comparison concerns which direction is more strongly coupled, not which way it acts. The results show that the absolute S→R contribution exceeded R→S in 79.7% of Harvey communities, 69.4% of Ida communities, and 81.9% of Ian communities. Mean contribution magnitudes showed the same ordering. Each contribution is the marginal change in a household's decision probability produced by a unit perturbation of the corresponding neighbour-history state ($\epsilon = 0.01$, see Methods). The S→R

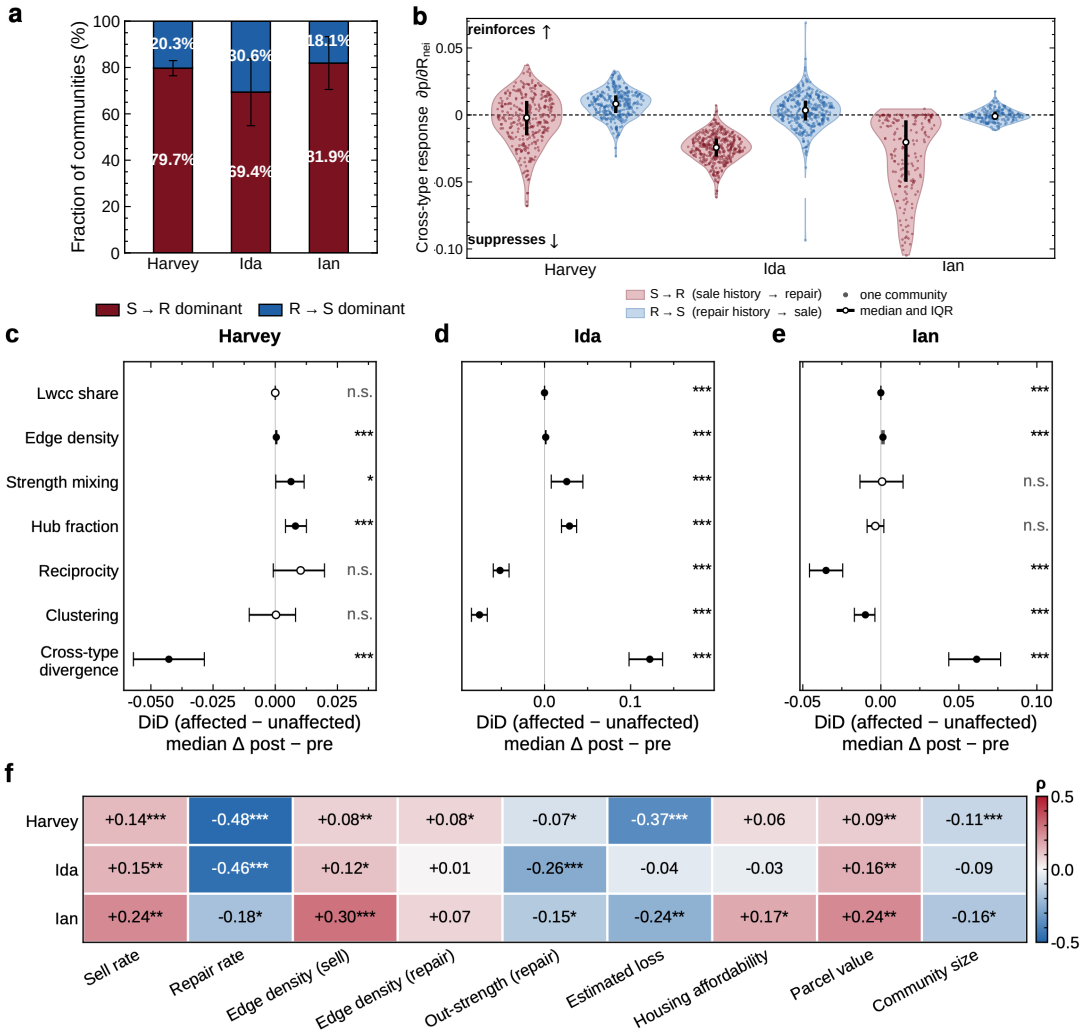


Figure 4: Asymmetric cross-type diffusion and post-hurricane shifts in learned community networks. (a) Fraction of communities in which total sell-to-repair (S→R) influence exceeds total repair-to-sell (R→S) influence. S→R dominates in most communities across all storms (80% Harvey, 69% Ida, 82% Ian). Error bars denote standard deviation across model seeds. (b) Neighborhood-mediated sell-to-repair and repair-to-sell, $\partial p/\partial R_{\text{nei}}$, for every community containing damaged buildings. Red distributions show the response of household repair probability to neighborhood sale history (S→R) and blue distributions show the response of household sale probability to neighborhood repair history (R→S). Each dot is one community. White circles and vertical bars denote medians and interquartile ranges. Values above the dashed line at zero indicate that neighborhood activity reinforces probability of the target decision and values below indicate suppressing. On average, neighborhood sale tends to suppress household's repair across three hurricanes. (c-e) Difference-in-differences analysis of learned graph metrics comparing affected and unaffected communities before and after each hurricane. Points show DiD estimates (affected − unaffected in median post-minus-pre change) and horizontal lines show bootstrap 95% confidence intervals from 5,000 resamples; filled points denote $p < 0.05$ and open points non-significant estimates. (f) Spearman correlations between community conditions and community cross-type asymmetry, defined as $|\text{S}\to\text{R}| - |\text{R}\to\text{S}|$, computed across communities within each hurricane. Positive values indicate that communities scoring higher on that condition tend to show greater S→R dominance. Asterisks denote two-sided significance: $*p < 0.05$, $**p < 0.01$ and $***p < 0.001$.

pathway exceeded R→S by 1.3X after Harvey, 1.7X after Ida, and 3.4X after Ian. The difference was largest after Ian, where the estimated S→R contribution was more than three times the reverse pathway. The results suggested that household repair probabilities were generally more sensitive to sale histories than sale probabilities were to repair histories.

Second, as shown in Fig. 4b, the neighboring sale activities mostly have negative impacts on a household's repair decision. We calculate here a different quantity from that in Fig. 4a, and the two are complementary rather than alternative measurements of the same thing. The model is evaluated at its fitted operating point, with every household at its observed history and its neighbors at theirs, rather than from a zeroed reference state. Only the neighborhood-aggregated history is perturbed, leaving the household's own history fixed, rather than both together. The perturbation is small ($\varepsilon = 0.01$) and the response is expressed as a change in decision probability rather than in log-odds; and the sign is retained rather than its magnitude. Fig. 4a therefore reports how strongly the two cross-type pathways are wired, and Fig. 4b reports what that wiring does at the state communities actually occupy: whether a marginal increase in neighboring sale activity raises or lowers a household's probability of repair, and by how much on the scale at which decisions are made. A pathway can be the stronger of the two in Fig. 4a and act as a brake in Fig. 4b. As shown in the results, in Ida and Ian, most communities have negative impacts with a median response gradient of -0.024 and -0.020, respectively. In these communities, a greater history of sales among connected neighbors was associated with a lower probability of repair. The reverse pathway did not mirror this pattern: the R→S response was predominantly positive after Harvey and centered near zero after Ian. Sale histories therefore did more than propagate activity between decision types. Within the sale channel, they reinforced subsequent sales (Fig. 3e,f), and across channels, they were associated with reduced repair. Our results revealed that neighborhood influence should be better described as a two-channel process than as a single social multiplier: same-type histories reinforce continued activity, whereas the exit-to-investment pathway can inhibit repair.

Harvey provides an important contrast. Its S→R response straddled zero, with negative values in only 54.7% of communities and a median of -0.002, while same-type repair coupling was strongest after this event (Fig. 3c). One possible interpretation is that, in Harvey's deeper and more liquid housing market, neighboring sales were more readily absorbed as ordinary transactions and therefore conveyed less information about whether the neighborhood would recover. Under this interpretation, sales had

little association with repair decisions, while direct repair coordination remained strong. In the thinner markets affected by Ida and Ian, where expectations of repeated exposure may have been more salient, neighboring sales may have carried more information about neighborhood viability and increased the value of delaying investment.

The community correlates of cross-type asymmetry shift in the same order (Fig. 4f). We calculate the absolute cross-type contribution of each pathway for every community and define the asymmetry as their difference, $|\mathrm{S}\to\mathrm{R}| - |\mathrm{R}\to\mathrm{S}|$, using the per-community marginal probability changes. This statistic is the continuous form of the classification in Fig. 4a: it is positive when the sale-to-repair pathway is the stronger of the two and grows with the margin between them, so a community's position on this scale indicates how far its cross-type coupling is tilted toward the sale-to-repair direction. Spearman rank correlations between this asymmetry and community attributes then identify which community conditions accompany a stronger tilt, without assuming a linear relationship or a common scale across attributes. As shown in the figure, after Harvey, asymmetry was more closely associated with community activity, including repair rate ($\rho = -0.48$) and sale rate ($\rho = 0.14$), while the association with sale-network edge density was weaker ($\rho = 0.08$). After Ian, sale-network edge density became more strongly associated with asymmetry ($\rho = 0.30$), while the association with repair rate weakened to $\rho = -0.18$. Accordingly, the ratio of the mean absolute correlation for the two activity measures to that for the two structural measures declined monotonically across events, from 4.13 to 1.61 and then to 0.93. As the negative S→R association became more consistent across communities, the attributes associated with that asymmetry shifted from what households were doing toward how their latent influence networks were organized. A complementary neighborhood-ablation analysis showed that after Ian, the presence of reconstructed neighborhood connections was associated with 20.2% more damaged parcels in the no-action state(Fig. 5c).

Another key question is whether hurricane exposure was accompanied by restructuring latent neighborhood influence networks. For each hurricane, we compared pre- to post-hurricane changes in affected communities with contemporaneous changes in unaffected communities (Fig. 4c–e). We found that edge density increased after Harvey ($\Delta_{\mathrm{DiD}} = 0.00043$), Ida ($\Delta_{\mathrm{DiD}} = 0.00140$), and Ian ($\Delta_{\mathrm{DiD}} = 0.00135$). However, the organization of these connections differed across events. After Ida and Ian, reciprocity declined by 0.052 and 0.035, but clustering declined by 0.075 and 0.010, respectively. The denser networks

were more directional and less locally redundant. In contrast, Harvey's reciprocity and clustering changed little, but hub fraction and strength mixing increased. Influence after Harvey consequently became more concentrated around high-strength households without the corresponding loss of local cohesion observed after Ida and Ian.

Our results show that neighborhood influence on housing co-adaptation is asymmetric and structurally organized. Rather than determining recovery outcomes directly, neighbor interactions reshape community trajectories through a set of asymmetric diffusion mechanisms. First, they act on existing sale and repair signals, and this action is asymmetric in both magnitude and sign: sale histories are more strongly coupled to subsequent repair decisions than repair histories are to subsequent sale decisions, but that coupling is predominantly negative. This asymmetry is consistent with greater susceptibility to exit-related cascades than to coordinated recovery. Second, post-disaster network reorganization, such as increased connectivity, reduced reciprocity and greater influence concentration, may also change directional pathways that reinforce dominant signals. Finally, these processes operate on different temporal scales, with repair influence concentrated early and sell influence extending over longer horizons. Combined with the preceding temporal results, the findings distinguish an earlier repair regime from a more persistent sale regime whose histories extend more strongly into cross-type household decisions.

Neighborhood Influence Contributions Extend beyond Damaged Households

We established the understanding of the complex and dynamic coupling of same-type influences versus cross-type influences, including self influences and neighborhood influences. In this section, we further investigated how such influence coupling jointly reshape community recovery outcomes. Rather than reducing recovery to a single stabilizing or declining outcome, we decompose parcel-level responses into five action states: sell without repair, repair without sale, repair then sell, sell then repair, and no action. Using the trained HDD model, we remove all household-to-household connection edges and keep household covariates, damage, and self-history fixed, and compare the resulting counterfactual action distribution with the full model.

Figure 5(a–c) shows that neighbor influence reallocates damaged buildings across distinct post-disaster action pathways, but not in a uniform direction. In Harvey, neighbor influence primarily activates repair-linked response: repair-only increases by 11.9% and no action falls by 10.6%. Sell-without-repair

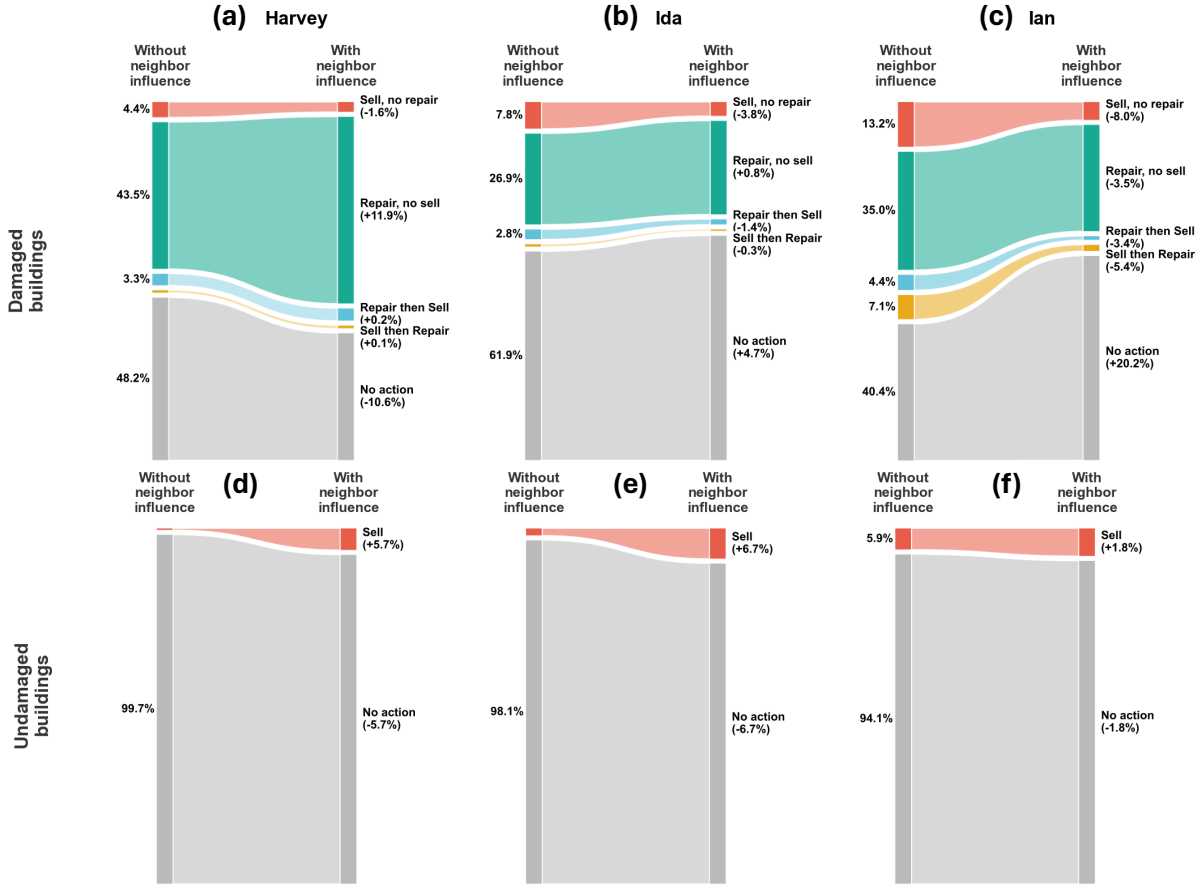


Figure 5: Neighborhood influence reshapes post-disaster housing decisions for both damaged and undamaged buildings. Alluvial diagrams compare the counterfactual distribution of parcel-level outcomes without neighbor influence (left bars) to the observed distribution with neighbor influence (right bars) across three hurricanes. (a–c) Damaged buildings. In Harvey (a), neighbors are associated with more repair and less inaction; unrepaired sale decreases slightly. In Ida (b), neighbors are associated with a small increase in repair-only and a decrease in unrepaired sale, alongside more inaction. In Ian (c), neighbors strongly suppress repair and push parcels toward inaction. (d–f) Undamaged buildings in the same damaged communities. Across all three storms, neighbors increase selling among undamaged parcels, indicating that neighbor activity spills over to buildings that were not directly damaged. Percentages are computed separately for damaged and undamaged buildings within communities containing at least one damaged parcel.

decreases slightly by 1.6%, indicating that neighbor influence diffusion in Harvey moves damaged parcels mainly out of inaction and into repair rather than unrepaired sale. Ida shows a more mixed response: no action increases by 4.7%, sell-without-repair decreases by 3.8%, repair-then-sell decreases by 1.4%, and repair-only increases only modestly (+0.8%). In Ian, neighbor influence increases no action by 20.2% but suppresses sell-without-repair, sell-then-repair, and repair-then-sell pathways. The results indicate that in Ian, neighbor influence appears less as immediate turnover but more as a shift away from active recovery pathways toward stalled action.

Figure 5(d–f) further shows that neighbor effects extend beyond directly damaged buildings. Among undamaged buildings located in communities with damaged parcels, neighbor influence increases selling in all three storms: by 5.7% in Harvey, 6.7% in Ida, and 1.8% in Ian. This spillover to undamaged parcels is important because it shows that neighborhood diffusion is not simply a response to direct physical damage. Even households without recorded building damage respond to the surrounding recovery environment, suggesting that local market signals, perceived neighborhood viability, and nearby household actions can reshape decisions beyond the damaged building stock.

Supplementary Fig. S5 further links these pathway shifts to community characteristics. In Harvey, neighbor-induced repair is most strongly associated with damage exposure: communities with more damaged buildings, higher estimated loss, and greater damage fraction show stronger repair-without-sale shifts ($\rho = 0.42$–0.43, $p < 0.001$), while sell-network structure predicts sell-without-repair and reduced no action. In Ida, neighbor effects are more resource- and network-stratified: higher income and parcel value are associated with weaker movement into no action ($\rho = -0.40$ and -0.37, $p < 0.001$), while sell-network in-strength and inequality predict sell-without-repair shifts ($\rho = 0.47$–0.48, $p < 0.001$). In Ian, pathway shifts are most strongly tied to network structure and built value: estimated loss and repair-network inequality predict both suppression of repair-without-sale ($\rho = -0.70$ and -0.57, $p < 0.001$) and movement into no action ($\rho = 0.61$ and 0.58, $p < 0.001$), whereas improvement value is associated with weaker no-action shifts and stronger repair-then-sell pathways.

3 Discussion

In this study, we develop HDD, which, to our knowledge, provides the first framework for large-scale, parcel-level reconstruction of how housing adaptation decisions co-evolve across households to shape com-

munity recovery. Compared to traditional models focusing on individual decisions or terminal outcomes, HDD model post-disaster repair and sale as recurrent, interacting household decisions and jointly reconstruct the latent neighborhood influence structures and temporal dynamics underlying the co-adaptation process. We formulate this complex and interdependent process as a cross-type neural graph Hawkes process and develop a novel community-shared joint inference strategy to pool statistical strength across sparse household event sequences while retaining community-specific latent influence structures. HDD integrates household-specific background propensities, self-history, and neighborhood-history signals to estimate time-varying repair and sale probabilities. This formulation enables process-level analysis of when repair and sale histories contribute, how they interact across decision types, and how neighborhood dependence varies among communities. We further constructed a large benchmark dataset covering more than 1.3 million households across multiple Gulf Coast hurricanes to evaluate this framework. HDD consistently improved repair and sale prediction compared to existing state-of-the-art models, achieving a AUC of 0.88 for repair and 0.75 for sale for Harvey, 0.93 and 0.82 for Ida, and 0.94 and 0.78 for Ian. It also achieved the highest network reconstruction precision against mobility-based proxy networks. Controlled simulations further showed that HDD recovered event dynamics and network structure across diverse regimes.

We used the HDD outputs to answer long-standing questions about temporal and cross-type patterns of housing co-adaptation. First, does neighborhood activity always reinforce the same decision? It does, but not on a common timescale, and this qualifies rather than confirms the social-multiplier account of recovery. Same-type associations were positive for both decisions after all three hurricanes, consistent with structural estimates of positive rebuilding externalities after Katrina and against the competing prediction that locally scarce contractors, materials and inspection capacity would make neighboring repairs substitutes. The two channels are not, however, interchangeable. Repair-to-repair excitation rose and then declined or remained bounded, whereas sale-to-sale excitation accumulated across the full observation window; and the household response to selling neighbors exceeded the response to repairing neighbors in every event, on the log-odds scale as well as in probability. We read this as a structural difference between the two decisions rather than a difference in social salience. Repair is a coordination problem that involves settlements, contractor availability and program deadlines in a short time window, so complementarity is real but time-limited. Sale remains available indefinitely and becomes more

attractive as further transactions accumulate, each one both pricing the neighborhood and removing an owner committed to remaining. The community-scale counterpart is the steep dependence of final repair rate on early repair rate, which crosses 50% at early repair rates of 34%, 31% and 39%: the signature of a reinforcing process that resolves early rather than one that compounds throughout recovery.

Does activity in one type of decision carry into the other, and in which direction? Here the answer departs from what the existing literature would predict, and it is the finding we regard as most consequential. Cross-type coupling is asymmetric in magnitude, with the sale-to-repair pathway more strongly coupled than the reverse in 69–82% of communities. But the realized effect of that coupling is negative: evaluated at the state communities actually occupy, a greater sale history among connected neighbors was associated with a lower probability of repair in almost every community after Ian and in most after Ida. The reverse pathway was not its mirror image. Neighboring repair history was associated with slightly higher sale probability after Harvey and with no consistent direction after Ian. Existing work on neighborhood exits establishes that they transmit, for example, forced sales depress nearby transaction prices, and foreclosures raise the probability that neighbors default in turn, but that literature measures transmission in prices or within the same decision type. What it leaves open is whether an exit changes the behavior of households that do not leave. Our results indicate that it does, and in the direction predicted by irreversible-investment logic rather than by renewal: neighboring exits appear to raise the option value of waiting rather than to transfer property to owners more willing to invest. The practical consequence is that neighborhood influence in housing recovery is not a single social multiplier that scales all activity in proportion. It is a two-channel process in which within-channel histories reinforce continued activity while the exit-to-investment channel inhibits it, so the same relational structure that spreads repair also transmits withdrawal.

Neighborhood influence is neither stabilizing nor destabilizing in itself. Removing the reconstructed connections reallocated damaged parcels across action states in different directions by event: after Harvey the presence of neighborhood influence corresponded to more repair and less inaction, after Ian to substantially more inaction, and after Ida to a mixed pattern. The same intervention on the model therefore produces opposite outcomes depending on context, which argues against treating community connectedness as uniformly protective. Neighborhood influence also extended past directly damaged properties, corresponding to higher selling among undamaged households in affected communities after

all three hurricanes — evidence that the population responding to a disaster is wider than the population it physically damaged.

Implications for policy design and resource allocation. Four follow from the results rather than from general reasoning. First, timing is not symmetric across instruments: because repair reinforcement is bounded and early while sale reinforcement accumulates, repair-side capacity, such as settlement speed, permitting throughput, and contractor availability, has a window in which it can engage a reinforcing process, whereas exit-side monitoring must persist well beyond it. Second, the early-repair level at which communities cross to majority final repair is an observable operational quantity and could be tracked during recovery rather than diagnosed afterwards. Third, most consequential for allocation, repair assistance and exit management are well-coordinated instruments: because neighboring sales are associated with suppressed repair, concentrated exits and repair subsidies delivered into the same community at the same time work against one another, and sequencing them may matter as much as their level. Fourth, because neighborhood influence extends to undamaged households, eligibility rules keyed to parcel-level damage systematically under-cover the population whose decisions the disaster reshapes; community-scale assessment addresses a real gap rather than an administrative preference. The event contrasts also suggest what to condition on: where exits are absorbed by a liquid market, observed activity is a sufficient diagnostic, whereas where they carry information, the organization of the influence structure becomes the more informative signal.

Several limitations define the scope of these conclusions. Although the framework itself is generalizable to other hazards and events, the analysis results cover three U.S. hurricanes and may not generalize directly to other hazards or events. Permit and transaction records capture formally registered activity and may underrepresent informal repairs, rental displacement, and other adaptations that do not produce observable parcel-level events. Although the latent influence network can partially aligns with mobility-based proxies, its edges cannot be interpreted as verified interpersonal relationships. Future work could combine HDD with intervention-aware or quasi-experimental designs, richer mobility and communication data, and observations from additional hazards and institutional settings.

4 Methods

Data: We construct a large-scale parcel-level dataset from three major U.S. hurricanes: Ian (2022), Ida (2021), and Harvey (2017). Across events, we integrate and pair household attributes, structure-level damage assessments, permit-based repair activity, and property transaction records. Although the raw administrative sources differ across jurisdictions, all three datasets are processed through a common pipeline. We first define a parcel universe by retaining only geocoded single-family residential parcels and their corresponding valid Census Block Group (CBG) identifier. Repair events are inferred from building permit filings using the *filing date* to represent the timing at which a household is activated to initiate restoration, and similar approach applies to sell events. All event timestamps are binned to monthly resolution, producing a tensor $\mathbf{Y} \in \mathbb{Z}_{\geq 0}^{T \times N \times 2}$, where the two channels correspond to sell and repair events. Parcels in the same CBG form a community. Detailed data trends can be found in Supplementary Figure S1.

Hurricane Ian (2022), Lee County, Florida: Hurricane Ian made landfall in Lee County, Florida on September 28, 2022. It causes catastrophic wind and storm-surge damage estimated at $110 billion, and is the costliest Florida hurricane on record [7]. The National Flood Insurance Program received 28,852 claims in Lee County [29]. Building damage assessments used in this study are obtained by integrating data provided by FEMA, Lee County, and StEER reports [30–32], which classify damage level of 68,515 residential parcels. Lee County established an emergency permitting operation for storm-related building and well repairs, expedited permit processing, and temporarily waived permit fees. Repair activity is extracted from Lee County building permit filings disseminated through the Accela platform [33]. These processed datasets are cross-validated with aggregated records provided by FEMA [34], Redfin [35], and Shovels.ai [36]. Parcel and household characteristics are obtained from Regrid administrative data [37]. The final Ian dataset contains 181,659 residential parcels aggregated into 399 CBG communities, with 39 parcel-level features. The pre-hurricane window spans January 2021 to August 2022 (20 months) and the post-hurricane window spans September 2022 to October 2023 (14 months).

Hurricane Ida (2021), City of New Orleans: Hurricane Ida made landfall near Port Fourchon, Louisiana on August 29, 2021, inducing severe wind damage across the greater New Orleans metropolitan area and resulted in an estimated $75 billion in total losses [24]. FEMA disbursed $238 million in Individuals and Households Program (IHP) grants to Orleans Parish households, while the National Flood Insurance Pro-

gram received 2,027 claims in Orleans Parish [29], consistent with Ida's damages being predominantly wind-driven rather than flood-driven. The U.S. Army Corps of Engineers also deployed the Blue Roof program, installing approximately 33,600 temporary roof coverings across Louisiana [38]. We integrate City of New Orleans's building damage labels from FEMA Visual Damage Assessment products [39], which assessed 213,300 parcels into five damage classes. Because the City of New Orleans waived roofing permit requirements for storm-related repairs [40]. Repair activity is therefore reconstructed by integrating FEMA IHP habitability repair records with the City of New Orleans permit system. These processed datasets are further cross-validated with aggregated market and reconstruction records from Redfin [35] and Shovels.ai [36]. Parcel boundaries, parcel characteristics, and transaction records are obtained from Regrid [37], with closing dates converted to listing dates using Redfin median days-on-market. The final Ida dataset contains 121,135 residential parcels aggregated into approximately 380 CBG communities, with 46 parcel-level features. The pre-hurricane window spans September 2020 to July 2021 (11 months), and the post-hurricane window spans August 2021 to July 2022 (12 months).

Hurricane Harvey (2017), Harris County, Texas: Hurricane Harvey made landfall on August 25, 2017, producing unprecedented rainfall that flooded Harris County [41]. 165,076 single family households in Harris County reported home damage for FEMA Individual Assistance[34], and 110,107 of them were verified as damaged. Damage assessments used in this study are obtained by integrating FEMA HydroShare dataset and data sources from DesignSafe [42, 43], which identify 124,629 damaged single-family parcels. Parcel characteristics—including assessed and market values, building area, year of construction, and structural attributes—are drawn from the Harris County Appraisal District (HCAD) public data extracts [44]. Property transactions are extracted from HCAD deed records. Repair activity is extracted from HCAD building permit filings. The City of Houston expedited permitting for storm-related repairs and temporarily waived certain permit fees in the weeks following landfall, contributing to a rapid surge in permit filings. The final Harvey dataset contains 1,017,324 single-family residential parcels aggregated into approximately 2,000 census block group (CBG) communities, with 48 parcel-level features (42 numeric, 6 categorical). The pre-hurricane window spans May 2016 to July 2017 (15 months) and the post-hurricane window spans August 2017 to November 2018 (16 months).

Household-level Decision Diffusion (HDD) Framework

HDD framework aims to reconstruct post-disaster housing co-adaptation dynamics from sparse monthly

repair and sale histories. For each hurricane dataset, $g \in \{1, \ldots, G\}$ indexes census block group communities, and $i \in \{1, \ldots, N_g\}$ indexes households. Each household has covariates $\mathbf{x}_i^g$, coordinates $\mathbf{c}_i^g$, and a monthly event vector $\mathbf{Y}_i^g[t] \in \{0,1\}^K$, where $K = 2$ denotes sale and repair.

Household propensity and decision-history memory: HDD represents persistent heterogeneity using a household embedding, a decision-type-specific baseline, and a separate damage-related contribution:

$$\mathbf{h}_i^g = \phi_\theta(\mathrm{Norm}(\mathbf{x}_i^g)), \qquad \mathbf{b}_i^g = f_b(\mathbf{h}_i^g), \qquad \mathbf{d}_i^g = f_d(\mathbf{x}_{i,\mathrm{dmg}}^g), \tag{1}$$

where $\mathbf{x}_{i,\mathrm{dmg}}^g$ contains storm-specific damage variables and normalization uses training-data statistics. All three outputs are shared across communities through their parameterized functions. Because events are observed monthly, the household's previous decisions are represented by an exponentially decaying state:

$$\mathbf{r}_i^g(t) = \boldsymbol{\gamma} \odot \mathbf{r}_i^g(t-1) + \mathbf{Y}_i^g[t-1], \qquad \mathbf{r}_i^g(0) = \mathbf{0}. \tag{2}$$

The decision-type-specific persistence vector $\boldsymbol{\gamma} \in (0,1)^K$ is shared across communities. HDD also computes analogous histories over M prespecified decay rates, collected in $\mathbf{R}_{i,\mathrm{multi}}^g(t) \in \mathbb{R}^{K \times M}$, to represent both recent and persistent decision histories.

Latent influence network and neighborhood history: For each community g and source decision type k, HDD constructs a directed, weighted latent influence network

$$\mathbf{A}^{(g,k)} = \mathcal{G}_{\psi_k}(\mathbf{H}^g, \mathbf{X}^g, \mathbf{C}^g) \in \mathbb{R}_{\geq 0}^{N_g \times N_g}, \qquad A_{ii}^{(g,k)} = 0, \tag{3}$$

where $\mathbf{H}^g$, $\mathbf{X}^g$, and $\mathbf{C}^g$ collect household embeddings, covariates, and coordinates. An element $A_{ij}^{(g,k)}$ represents the influence of household j's history of type k on household i. The graph constructor combines directed embedding affinity, covariate similarity, and geographic distance. Candidate edges are restricted by a maximum distance and sparsified by retaining the m strongest incoming ties per household. Full graph-construction equations and parameter settings are provided in Supplementary Methods.

HDD distinguishes active network connections from previously transmitted influence. Let $a_i^g[t]$ indicate whether household i remains active at month t, and let $q_{ij,k}^g(t)$ denote the influence previously

transmitted from j to i. Neighborhood history is updated as

$$q^g_{ij,k}(t+1) = \gamma_k q^g_{ij,k}(t) + a^g_i[t]a^g_j[t]A^{(g,k)}_{ij}Y^g_{j,k}[t], \qquad n^g_{i,k}(t) = \sum_{j=1}^{N_g} q^g_{ij,k}(t). \tag{4}$$

A sale event is transmitted during the month in which it occurs. The selling household then exits the active network, preventing new interactions, while previously transmitted influence continues to decay.

Cross-type prediction and model estimation: A shared nonlinear prediction function combines household propensity, self-history, neighborhood history, multi-timescale memory, and temporal variation:

$$\boldsymbol{\eta}^g_i(t) = \mathcal{F}_\theta\Big(\mathbf{h}^g_i, \mathbf{b}^g_i, \mathbf{w}^g_i \odot \mathbf{r}^g_i(t), \mathbf{n}^g_i(t), \mathrm{vec}\Big[\mathbf{R}^g_{i,\mathrm{multi}}(t)\Big]\Big) + \boldsymbol{\tau}(t) + \mathbf{d}^g_i, \qquad p^g_{i,k}[t] = a^g_i[t]\sigma\Big(\eta^g_{i,k}(t)\Big), \tag{5}$$

where $\mathbf{w}^g_i$ contains non-negative self-history weights and $\boldsymbol{\tau}(t)$ captures temporal variation shared across households. Because the complete sale- and repair-history vectors enter both output logits, the model represents sale-to-sale, repair-to-repair, sale-to-repair, and repair-to-sale pathways. The graph index k denotes the source-history type, whereas the output index denotes the subsequent decision type.

HDD is trained end-to-end using a class-weighted Bernoulli loss with graph sparsity regularization:

$$\mathcal{L} = -\sum_{g,t,i,k} e^g_{i,k}[t]\left[\omega_{k,1}Y^g_{i,k}[t]\log p^g_{i,k}[t] + \omega_{k,0}(1-Y^g_{i,k}[t])\log(1-p^g_{i,k}[t])\right] + \lambda_{\mathrm{edge}}\sum_{g,k}\|\mathbf{A}^{(g,k)}\|_1, \tag{6}$$

where $e^g_{i,k}[t]$ indicates whether an event remains eligible and $\omega_{k,1}$ and $\omega_{k,0}$ address class imbalance. The top-m constraint and edge penalty discourage diffuse network solutions. The reconstructed networks are interpreted as latent representations of interdependent household behavior, rather than observed social or causal networks.

Evaluation Metrics: To evaluate rare-event performance relevant to disaster decision-making, we report both exact next-step prediction and windowed prediction over a horizon of H months, using AUC/AP and Recall@K computed strictly on held-out months. Specifically, we evaluate how well a method can predict future household events on held-out months only ($t \in [T_{\mathrm{train}}, T)$), using probabilistic outputs $\hat{p}^g_{i,k}[t] = \sigma(\hat{\eta}^g_{i,k}(t))$ for each household i in community g and decision type k (sell/repair). We report (i) ROC-AUC, which measures ranking quality independent of a fixed threshold, and (ii) Average Precision (AP), i.e., area under the precision-recall curve, which is more informative under strong class

imbalance typical of rare post-disaster decisions. Because decision support often targets a limited set of high-risk households, we additionally report Recall@K (or Recall@top-q%): among all positives in the test period, the fraction captured by the top-K highest predicted probabilities (e.g., top 1% and 5% of household-months), computed separately for each decision type. We evaluate in two complementary label modes: exact next-step prediction uses the binary label $Y^g_{i,k}[t]$ (event at month t), while windowed prediction uses $Y^{(H)} * i, k^g[t] = \mathbb{I}\{\exists, \tau \in (t, t+H] \text{ s.t. } Y * i, k^g[\tau] = 1\}$ to measure whether an event occurs within a horizon of H months, reflecting tail-relevant "will it happen soon?" operational queries. All metrics are computed by pooling household-month samples across selected communities.

We additionally evaluate social network reconstruction performance at the community level. For each community, we compute an edge-level precision, defined as the fraction of observed real-world (mobility-based) edges that also appear in the predicted network. Rather than reporting a single aggregate score, we report the distribution of community-level precision, and construct an edge-weighted precision distribution by weighting precision values according to the number of observed social edges when forming the distribution. Recall-based metrics are not considered due to the limited penetration rate of trajectory data (approximately 20%), which renders the observed real-world network incomplete.

To obtain a real-world reference network, we construct mobility-based proxy social networks using anonymized Cuebiq trajectory data. Cuebiq solely collects data from anonymized devices which opt-in to data collection for research purposes. To further preserve privacy, home areas are obfuscated to the census block group level, and a processing is performed within the Cuebiq secure research environment in accordance with platform privacy requirements; no raw mobility data are exported or stored locally. Following prior work on unsupervised mobility-based tie inference [45, 46], we quantify pairwise trajectory similarity using Dynamic Time Warping (DTW), which is robust to temporal misalignment and irregular sampling [47, 48]. Resident-level proximity/similarity signals are aggregated to construct a mobility-based social network. All analyses and reported outputs are privacy-preserving and aggregated (e.g., applying minimum group-size thresholds where appropriate), with no attempt to identify individuals or households.

Simulation Study for Diffusion Identifiability: To address the fundamental identifiability challenge that real-world diffusion structure is unobserved and proxy networks cannot establish causal validity, we conduct controlled simulations in which both the event-generation process and the underlying diffusion

graph are known. We generate multi-community sell and repair event sequences under graph-based excitation, with explicit control over network connectivity, excitation strength, temporal memory, and the balance between self-driven persistence and neighbor-driven propagation. The simulation study serves two complementary purposes. First, we evaluate predictive performance on held-out months to verify that HDD can accurately recover event dynamics generated by known diffusion mechanisms. Second, and more critically, we assess diffusion-graph recovery using edge-level precision and recall, enabling direct evaluation of whether the model introduces spurious influence ties when ground truth is available. We conduct systematic parameter sweeps over network density, excitation balance, temporal decay, and observation horizon (Supplementary Tables S4-S10). Results summarized in Supplementary figures and tables show that HDD consistently achieves high predictive accuracy while maintaining near-perfect precision across regimes, with recall varying predictably under sparse influence settings. This behavior reflects a deliberate design choice that prioritizes conservative structure recovery and guards against false social links in regimes where exhaustive edge identification is statistically ill-posed.

Additional Details about Benchmark Methods: We compare HDD against a set of representative baselines that span econometric choice modeling, spatial econometrics, agent-based recovery simulation, and temporal point-process learning. First, we include a random utility model (RUM; multinomial/binary logit) that predicts household sell/repair decisions from observed covariates only, serving as a feature-driven, non-diffusion reference [49]. Second, we consider spatial lag and spatial Durbin specifications that augment covariate effects with spatially weighted outcomes and/or covariates to capture local spillovers under a pre-specified proximity matrix [50, 51]. Third, we implement a RecovUS-style agent-based model in which heterogeneous households follow stochastic decision rules driven by damage, resources, and neighborhood context, allowing recovery trajectories to emerge from micro-level behavioral assumptions [12]. Fourth, we include a sparse multivariate Hawkes process baseline that estimates self- and cross-excitation via an ℓ_1-regularized influence matrix [52] to evaluate classic parametric diffusion recovery and graph reconstruction under limited events. Moreover, to isolate the effect of distance-only diffusion, we use a distance-decay prior that constructs candidate influence weights solely from geographic separation (e.g., exponential or kernel decay) [53], reflecting longstanding distance-decay formulations in spatial interaction modeling. Finally, we benchmark against a neural temporal point process (e.g., RMTPP/Neural Hawkes family) that uses recurrent neural representations of event history to model

non-linear excitation [54, 55], but does not explicitly reconstruct an interpretable, community-specific diffusion graph unless supplemented with an external candidate network. Across all baselines, we focus comparisons on held-out future-state prediction (and graph recovery where applicable), using identical train/test splits and the same household covariates and community subsets as HDD.

5 Data Availability

The data collection and research project were preregistered at OSF (https://osf.io/jds9r).

6 Code Availability

The full replication analytical code is archived alongside the replication data and available via OSF (https://osf.io/jds9r).

References

1. NOAA National Centers for Environmental Information (NCEI). *U.S. Billion-Dollar Weather and Climate Disasters* 2025. `https://www.ncei.noaa.gov/access/billions/`.

2. U.S. Census Bureau. *County Population Totals and Components of Change: 2020–2025* `https://www.census.gov/data/tables/time-series/demo/popest/2020s-counties-total.html`. Mar. 2026.

3. Marsooli, R., Lin, N., Emanuel, K. & Feng, K. Climate change exacerbates hurricane flood hazards along US Atlantic and Gulf Coasts in spatially varying patterns. *Nature Communications* **10,** 3785. `https://doi.org/10.1038/s41467-019-11755-z` (2019).

4. Xi, D., Lin, N. & Gori, A. Increasing sequential tropical cyclone hazards along the US East and Gulf coasts. *Nature Climate Change* **13,** 258–265. `https://doi.org/10.1038/s41558-023-01595-7` (2023).

5. NOAA National Centers for Environmental Information (NCEI). *Costliest U.S. Tropical Cyclones (CPI-Adjusted)* excluding Maria. 2025. `https://www.ncei.noaa.gov/access/billions/dcmi.pdf`.

6. National Weather Service. *Service Assessment: Hurricane Harvey, August 25–September 1, 2017* tech. rep. (NOAA National Weather Service, 2018). https://www.weather.gov/media/publications/assessments/harvey6-18.pdf.

7. Bucci, L., Alaka, L., Hagen, A., Delgado, S. & Beven, J. *Tropical Cyclone Report: Hurricane Ian (AL092022)* tech. rep. AL092022 (National Hurricane Center, Miami, FL, Apr. 2023). https://www.nhc.noaa.gov/data/tcr/AL092022_Ian.pdf.

8. WWL-TV. *'Life-altering storm' — more than 250,000 residents ordered to evacuate* 2021. https://www.wwltv.com/article/weather/hurricane/hurricane-ida-emergency-evacuations-and-preparing-for-days-without-power/289-2b310a41-f2a0-418a-86f7-7985b42e7d2b.

9. Fox Weather. *2.5 million Floridians ordered to evacuate as Hurricane Ian barrels closer* 2022. https://www.foxweather.com/weather-news/hurricane-ian-florida-preparations-southeast-tampa.

10. Fu, C. & Gregory, J. Estimation of an Equilibrium Model With Externalities: Post-Disaster Neighborhood Rebuilding. *Econometrica* **87,** 387–421 (2019).

11. Harding, J. P., Rosenblatt, E. & Yao, V. W. The contagion effect of foreclosed properties. *Journal of Urban Economics* **66,** 164–178 (2009).

12. Moradi, S. & Nejat, A. RecovUS: An Agent-Based Model of Post-Disaster Household Recovery. *Journal of Artificial Societies and Social Simulation* **23,** 13 (2020).

13. Gupta, A. Foreclosure Contagion and the Neighborhood Spillover Effects of Mortgage Defaults. *The Journal of Finance* **74,** 2249–2301 (2019).

14. Aldrich, D. P. in *Building Resilience* (University of Chicago Press, 2012).

15. Campbell, J. Y., Giglio, S. & Pathak, P. Forced Sales and House Prices. *American Economic Review* **101,** 2108–2131 (2011).

16. Wharton Risk Management and Decision Processes Center. *Federal Disaster Rebuilding Spending: A Look at the Numbers* University of Pennsylvania, Wharton Climate Center. 2018. https://impact.wharton.upenn.edu/climate-center/federal-disaster-rebuilding-spending-a-look-at-the-numbers/.

765 17. Huang, T., Zanocco, C., Wang, Z., Hwang, J. & Rajagopal, R. Built environment disparities are
766 amplified during extreme weather recovery. *Nature* **648,** 349–356 (2025).

767 18. Sutley, E. J. & Hamideh, S. Postdisaster housing stages: A Markov chain approach to model se-
768 quences and duration based on social vulnerability. *Risk Analysis* **40** (2020).

769 19. Nejat, A. & Ghosh, S. LASSO Model of Postdisaster Housing Recovery: Case Study of Hurricane
770 Sandy. *Natural Hazards Review* **17,** 04016005 (2016).

771 20. Mach, K. J. *et al.* Managed retreat through voluntary buyouts of flood-prone properties. *Science*
772 *Advances* **5,** eaax8995 (2019).

773 21. Barthélemy, M. Spatial networks. *Physics Reports* **499,** 1–101 (2011).

774 22. González, M. C., Hidalgo, C. A. & Barabási, A.-L. Understanding individual human mobility pat-
775 terns. *Nature* **453,** 779–782 (2008).

776 23. Christakis, N. A. & Fowler, J. H. The spread of obesity in a large social network over 32 years. *New*
777 *England Journal of Medicine* **357,** 370–379 (2007).

778 24. National Hurricane Center. *Tropical Cyclone Report: Hurricane Ida (AL092021)* tech. rep. (Na-
779 tional Oceanic and Atmospheric Administration, 2022). `https://www.nhc.noaa.gov/data/tcr/`
780 `AL092021_Ida.pdf`.

781 25. National Hurricane Center. *Tropical Cyclone Report: Hurricane Harvey (AL092017)* tech. rep. (Na-
782 tional Oceanic and Atmospheric Administration, May 2018). `https://www.nhc.noaa.gov/data/`
783 `tcr/AL092017_Harvey.pdf`.

784 26. Manski, C. F. Identification of endogenous social effects: The reflection problem. *The review of*
785 *economic studies* **60,** 531–542 (1993).

786 27. Dixit, A. K. & Pindyck, R. S. *Investment under Uncertainty* ISBN: 9780691034102 (Princeton Uni-
787 versity Press, Princeton, NJ, 1994).

788 28. Bulan, L., Mayer, C. & Somerville, C. T. Irreversible investment, real options, and competition:
789 Evidence from real estate development. *Journal of Urban Economics* **65,** 237–251 (2009).

790 29. Federal Emergency Management Agency. *FEMA National Flood Insurance Program Claims Data*
791 `https://www.fema.gov/api/open/v2/FimaNfipClaims`.

30. Federal Emergency Management Agency, Response Geospatial Office. *FEMA Historical Geospatial Damage Assessments: Hurricane Ian* 2023. https://gis-fema.hub.arcgis.com/datasets/fema-historical-geospatial-damage-assessments.

31. Lee County Emergency Management. *Hurricane Ian Initial Damage Assessment* 2022. https://experience.arcgis.com/experience/27c1613132ac4bbf91c17f203b66ff0b.

32. Ceferino, L. *et al. StEER: Hurricane Ian Field Assessment Structural Team (FAST) Dataset* 2025. https://www.designsafe-ci.org/data/browser/public/designsafe.storage.published/PRJ-3709.

33. Lee County Government. *Administrative Building Permit Records for Lee County* https://aca-prod.accela.com/LEECO/Cap/CapHome.aspx. Accessed via the Accela Civic Platform. 2022.

34. Federal Emergency Management Agency. *OpenFEMA Dataset: Individual Assistance Housing Registrants — Large Disasters* 2023. https://www.fema.gov/api.

35. Redfin. *Redfin Housing Market Data Center* 2024. https://www.redfin.com/news/data-center/.

36. Shovels.ai. *Building Permit Data API* 2024. https://www.shovels.ai/.

37. Regrid. *Regrid Parcel Data Platform* https://app.regrid.com. 2022.

38. U.S. Army Corps of Engineers. *Operation Blue Roof: Hurricane Ida Response* 2021. https://www.usace.army.mil/Blue-Roof/.

39. Federal Emergency Management Agency. *Hurricane Ida Final Visual Damage Assessment* FEMA Hub (ArcGIS Feature Server). 2021. https://services.arcgis.com/XG15cJAlne2vxtgt/arcgis/rest/services/Hurricane_Ida_Final_Damage_Assessment_WFL1/FeatureServer/0.

40. City of New Orleans. *Department of Safety and Permits* 2025. https://www.nola.gov/safety-and-permits/.

41. Harris County Flood Control District. *Hurricane Harvey Flood Impact Summary* 2018. https://www.hcfcd.org/About/Harris-Countys-Flooding-History/Hurricane-Harvey.

42. Federal Emergency Management Agency. *Hurricane Harvey Building Damage Assessment* 2018.

43. Roueche, D. B., Lombardo, F. T., Krupar III, R. J. & Smith, D. J. *Collection of Perishable Data on Wind- and Surge-Induced Residential Building Damage During Hurricane Harvey (TX)* 2018. https://www.designsafe-ci.org/data/browser/public/designsafe.storage.published/PRJ-1668.

44. Harris County Appraisal District. *Harris County Appraisal District Public Data Extracts* 2023. https://hcad.org/pdata/.

45. Wang, D., Pedreschi, D., Song, C., Giannotti, F. & Barabasi, A.-L. *Human mobility, social ties, and link prediction* in *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Association for Computing Machinery, New York, NY, USA, 2011), 1100–1108. ISBN: 9781450308137. https://doi.org/10.1145/2020408.2020581.

46. Yang, D., Qu, B., Yang, J. & Cudré-Mauroux, P. LBSN2Vec++: Heterogeneous Hypergraph Embedding for Location-Based Social Networks. *IEEE Transactions on Knowledge and Data Engineering* **34,** 1843–1855 (2022).

47. McCool, D., Lugtig, P. & Schouten, B. *Dynamic Time Warping-based imputation of long gaps in human mobility trajectories* 2024. arXiv: 2410.16096 [stat.ME]. https://arxiv.org/abs/2410.16096.

48. Tao, Y. *et al.* A comparative analysis of trajectory similarity measures. *GIScience and Remote Sensing* **58,** 643–669. ISSN: 1548-1603 (2021).

49. McFadden, D. in *Frontiers in Econometrics* (ed Zarembka, P.) (Academic Press, 1974).

50. Anselin, L. *Spatial Econometrics: Methods and Models* (Springer, 1988).

51. LeSage, J. & Pace, R. K. *Introduction to spatial econometrics* (Chapman and Hall/CRC, 2009).

52. Xu, H., Farajtabar, M. & Zha, H. *Learning Granger Causality for Hawkes Processes* in *Proceedings of the 33rd International Conference on Machine Learning (ICML)* **48** (PMLR, 2016), 1717–1726. https://proceedings.mlr.press/v48/xuc16.html.

53. Fotheringham, A. S. Spatial Structure and Distance-Decay Parameters. *Annals of the Association of American Geographers* **71,** 425–436 (1981).

54. Du, N. *et al. Recurrent Marked Temporal Point Processes: Embedding Event History to Vector* in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)* (2016), 1555–1564.

55. Mei, H. & Eisner, J. *The Neural Hawkes Process: A Neurally Self-Modulating Multivariate Point Process* in *Advances in Neural Information Processing Systems (NeurIPS)* (2017). `https://arxiv.org/abs/1612.09328`.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- CLIMAHawkesNCsupplementary.pdf