Post-Flood Recovery and Socioeconomic Disparities: A Causal Analysis of Mobility-Based Recovery in Missouri

Songyun Zhang[a,1], Zekai Jin[a,1], Nan Lin[a,*], Kenan Li[b,*]

[a] *Washington University in St. Louis, St. Louis, MO, United States*
[b] *Saint Louis University, St. Louis, MO, United States*

Abstract

Flooding is one of the most frequent and disruptive natural hazards in the United States, yet post-disaster recovery trajectories remain highly uneven across communities. This study investigates how socioeconomic characteristics shape post-flood recovery patterns across Missouri counties using high-resolution mobility data. We integrate two complementary causal inference frameworks, Difference-in-Differences (DiD) and Propensity-Score Weighting (PSW) to estimate the impacts of flooding on multiple dimensions of mobility recovery, including activity volume, travel distance, dwell time, and movement diversity. By decomposing recovery into short-run shocks, recovery speed, time to recovery, and long-run persistence, we reveal systematic socioeconomic gradients in recovery capacity. Counties with higher poverty rates, larger shares of renter-occupied housing, and limited vehicle access exhibit slower or qualitatively different recovery trajectories, while higher educational attainment mitigates adverse impacts. The results highlight that recovery is not a simple return to baseline but a reorganization of mobility behavior shaped by structural inequalities. These findings provide actionable insights for equitable disaster recovery and resilience planning in flood-prone regions.

Keywords: Flood recovery, Mobility data, Difference-in-Differences, Propensity score weighting, Socioeconomic inequality, Urban resilience

*Corresponding author.
[1] These authors contributed equally to this work.

1. Introduction

Flooding is among the most frequent and economically damaging natural hazards worldwide, with impacts that extend well beyond immediate physical destruction to shape long-run social, economic, and spatial trajectories of affected communities. In the United States, floods account for more direct disaster losses than any other hazard category, with cumulative damages exceeding hundreds of billions of dollars over recent decades [11, 24]. Climate change is expected to further exacerbate these risks by increasing the frequency and intensity of extreme precipitation events, particularly in inland riverine regions across the Midwest and central United States [16, 31, 29]. As a result, understanding not only the immediate impacts of flooding but also the processes through which communities recover has become a central concern in urban systems research, environmental governance, and resilience planning.

A large interdisciplinary literature has emphasized that disasters are not natural in their social consequences. Instead, hazard impacts and recovery trajectories are fundamentally shaped by pre-existing socioeconomic conditions, institutional arrangements, and spatial structures [7, 30]. Communities characterized by higher poverty rates, housing insecurity, limited transportation access, and weaker social capital consistently experience slower and more incomplete recoveries, often leading to the persistence or amplification of inequality long after floodwaters recede [13, 15, 32]. These dynamics are especially salient in flood-prone regions where repeated exposure interacts with structural disadvantage, creating cumulative vulnerability over time [10].

Despite broad recognition of these patterns, empirically measuring post-flood recovery remains challenging. Traditional assessments of disaster impacts rely heavily on physical damage estimates, insurance claims, or aggregate economic indicators [21, 23]. While informative, these measures provide limited insight into how daily life resumes within affected communities. Recovery is not solely a matter of rebuilding infrastructure or restoring property values; it also involves the re-establishment of routine activities, access to employment and services, and the normalization of mobility patterns that underpin urban and regional functioning [1]. Capturing these functional dimensions of recovery requires data that reflect lived behavior rather than static outcomes.

Recent advances in passively collected mobility data, particularly anonymized smartphone-based location data, have opened new avenues for observing disaster impacts and behavioral adaptation at fine spatial and temporal scales. A growing body of research uses mobility data to study evacuation, displacement, and short-term disruption following natural hazards such as earthquakes, hurricanes, and floods [22, 27, 28]. These studies demonstrate that mobility patterns are sensitive indicators of disruption and adaptation within urban systems and often reveal impacts that are not visible in conventional damage-based metrics. Related work in other large-scale shocks, such as the COVID-19 pandemic, further shows that mobility responses are highly stratified by socioeconomic status, occupation, and access to resources [5, 20].

However, existing disaster–mobility research exhibits two important limitations. First, most studies focus on short-run behavioral responses, such as evacuation or immediate reductions in movement, rather than the longer-term recovery process. Recovery is often implicitly treated as a return to baseline activity levels, with limited attention to how different dimensions of mobility normalize at different speeds or settle into new equilibria. Second, mobility-based recovery is frequently summarized using a single outcome dimension, such as overall activity volume or average travel distance. Such scalar measures can obscure substantial heterogeneity in recovery dynamics. For example, aggregate activity may rebound quickly even as residents are forced to travel farther to reach essential services, spend less time at destinations, or exhibit persistent changes in spatial routines.

In parallel, causal identification of flood impacts on recovery outcomes is complicated by non-random exposure. Flood-prone areas differ systematically from unaffected areas in socioeconomic composition, infrastructure quality, and baseline mobility patterns. Difference-in-Differences (DiD) designs have become a standard tool for addressing these challenges by exploiting temporal variation to control for time-invariant unobserved heterogeneity [9, 17]. When combined with panel mobility data, DiD enables the estimation of short-run treatment effects and dynamic responses. Nonetheless, DiD is less well suited for scalar recovery outcomes such as recovery speed, time to recovery, or long-run persistence, which summarize post-flood behavior over extended periods.

Propensity-score methods provide a complementary framework by reweighting control units to resemble treated units along observed covariates [26, 14]. In disaster research, propensity-score weighting has been used to study housing recovery, migration, and economic outcomes, particularly when treated populations are relatively small compared to potential controls [3]. Yet few studies integrate propensity-score approaches with mobility-based recovery metrics, and fewer still combine them with panel-based designs to study multiple dimensions of recovery simultaneously.

This study addresses these gaps by providing a multi-dimensional, causally identified analysis of post-flood recovery using high-resolution mobility data from Missouri counties. Rather than treating recovery as a single outcome, we conceptualize it as a dynamic process characterized by multiple dimensions: the level of normalization, the speed of recovery, the time required to recover, and the persistence of long-run behavioral change. We operationalize these concepts using four complementary mobility indicators: device activity, travel distance, dwell time, and movement entropy, which capture distinct aspects of population presence, access, routine stability, and spatial diversity.

Methodologically, we integrate Difference-in-Differences and Propensity-Score Weighting within a unified empirical framework. DiD is used to estimate dynamic treatment effects on normalized recovery levels, while PSW is employed to estimate causal effects on scalar recovery metrics such as recovery slope, time to recovery, and long-run persistence. This combined approach leverages both temporal and cross-sectional variation, allowing us to isolate flood impacts while accounting for confounding from pre-existing socioeconomic differences.

Empirically, Missouri provides a compelling setting for this analysis. The state experiences recurrent riverine flooding and exhibits pronounced heterogeneity in poverty, housing tenure, vehicle access, educational attainment, age structure, and racial composition. These factors long hypothesized to shape disaster vulnerability and recovery capacity [8, 12]. By examining how these structural characteristics interact with flood exposure to shape mobility recovery trajectories, we contribute to a more nuanced understanding of resilience and inequality in flood-prone regions.

The contributions of this study are fourfold. First, we advance disaster recovery measurement by introducing a multi-dimensional mobility-based framework that moves beyond single-metric assessments. Second, we provide causal estimates of flood impacts by combining panel-based DiD with propensity-score weighting. Third, we document systematic socioeconomic gradients in recovery trajectories, showing that disadvantaged communities not only recover more slowly but often recover differently, with persistent reorganization of mobility behavior. Finally, our findings offer actionable insights for equitable disaster recovery and resilience planning by identifying which dimensions of mobility remain constrained and for whom.

The remainder of the paper proceeds as follows. Section 2 describes the data sources and research design. Section 3 presents the Difference-in-Differences analysis of normalized recovery levels. Section 4 reports Propensity-Score Weighting results for recovery slope, time to recovery, and persistence. Section 5 discusses implications for urban resilience and flood recovery policy, and Section 6 concludes.

2. Data and methodology

2.1. Study area and flood events

This study focuses on the state of Missouri, a region that experiences recurrent riverine flooding driven primarily by the Mississippi and Missouri river systems. Missouri presents a suitable empirical setting for examining post-flood recovery due to its pronounced spatial heterogeneity in socioeconomic conditions, settlement patterns, and transportation infrastructure. Flood impacts in the state span dense urban areas, small cities, and rural counties, allowing for a comparative analysis of recovery processes across diverse community contexts.

We identify major flood events during the study period using official flood records from federal and state agencies, including declarations of flood-related emergencies and documented high-water events. Counties are classified as *treated* if they experienced severe flooding during the focal event window, while counties without recorded flood impacts during the same period serve as controls. The timing of the flood event is standardized across units to enable panel-based causal inference.

2.2. Mobility data and outcome construction

To capture post-flood behavioral recovery, we employ anonymized and aggregated mobility data derived from smartphone location traces. These data provide repeated observations of population movement at fine spatial and temporal resolution while preserving individual privacy through aggregation and anonymization. Mobility indicators are aggregated to the census block group (CBG)–month level and subsequently summarized at the county level to align with the temporal scale at which flood impacts and recovery dynamics are observed and policy responses are typically implemented.

2.2.1. Socioeconomic and baseline risk covariates

Baseline socioeconomic characteristics are drawn from the American Community Survey (ACS) five-year estimates at the CBG level. These variables capture pre-flood structural conditions that are widely recognized as shaping disaster vulnerability, adaptive capacity, and recovery trajectories. All covariates are measured prior to the flood event to avoid post-treatment bias.

We organize ACS covariates into three conceptually distinct groups. First, *demographic structure* variables describe population composition and household characteristics, including age structure (e.g., shares of children and older adults), race and ethnicity proxies (e.g., non-Hispanic White, non-Hispanic Black, Hispanic, and other race shares), and household composition indicators such as average household size and the prevalence of single-parent households. These variables capture differential mobility needs, care responsibilities, and potential evacuation or relocation constraints.

Second, *socioeconomic status* variables measure economic resources and human capital, including median household income, poverty rate, educational attainment (high school completion and bachelor's degree shares), and employment status. These indicators proxy for access to financial buffers, job flexibility, and the capacity to absorb or adapt to post-disaster disruptions.

Third, *housing and urbanicity proxies* capture aspects of the built environment and residential stability that influence mobility behavior and recovery. These include housing tenure (owner-occupied versus renter-occupied shares), indicators of housing stock and density, and vehicle access (e.g., share of households without a vehicle). Together, these measures reflect baseline exposure, displacement risk, and dependence on local versus long-distance travel.

CBG-level covariates are aggregated to the county level using population-weighted averages to ensure that larger and more populous block groups contribute proportionally to county-level characteristics.

2.2.2. Behavioral recovery outcomes from mobility data

Behavioral recovery outcomes are constructed using SafeGraph mobility data aggregated to the CBG–month level. These data provide consistent measures of population presence, movement intensity, and activity patterns across space and time. To capture multiple dimensions of post-flood recovery, we construct four complementary mobility indicators that reflect distinct aspects of behavioral normalization.

- **Device activity (DEVICE)** measures the overall level of observed mobility activity within a CBG, proxied by the number of active devices detected. This indicator reflects population presence and the resumption of routine activity, capturing whether people are physically present and active in affected areas.

- **Travel distance (DISTANCE)** captures the median distance traveled from inferred home locations. Changes in travel distance reflect access constraints, service availability, and the need to travel farther to reach employment, healthcare, or essential goods when local infrastructure or amenities are disrupted.
- **Dwell time (DWELL)** measures the median time spent at destinations, serving as a proxy for routine stability and engagement in sustained activities such as work, shopping, or social interactions. Lower dwell times may indicate fragmented routines or temporary activity patterns.
- **Movement entropy (ENTROPY)** quantifies the spatial diversity of movement origins and destinations. Higher entropy reflects more dispersed and heterogeneous movement patterns, potentially indicating displacement, relocation, or exploratory behavior during recovery, whereas lower entropy reflects consolidation around stable activity spaces.

In addition to these core indicators, we derive supplementary mobility measures related to at-home intensity and local activity concentration, such as time-at-home proxies and neighborhood-level activity dispersion. These additional indicators are used to validate and contextualize the primary recovery measures and to ensure that observed normalization patterns are not driven by a single behavioral dimension.

2.2.3. Normalization and baseline adjustment

Each mobility series is normalized relative to a pre-flood baseline specific to each CBG. Normalization is implemented by expressing post-flood mobility levels as deviations from average pre-event behavior over a fixed baseline window. This transformation ensures comparability across counties and outcome dimensions and allows recovery to be interpreted as a return toward typical pre-event activity rather than as an absolute level of mobility.

By constructing normalized recovery measures across multiple behavioral dimensions, we capture not only whether activity resumes, but also how mobility structure, access, and routine stability evolve during the recovery process. This multi-dimensional approach enables a more nuanced assessment of post-flood recovery than single-metric analyses.

2.3. Recovery metrics

Recovery is conceptualized as a dynamic, multi-dimensional process rather than a single outcome, consistent with resilience and disaster recovery frameworks that emphasize functionality, speed, duration, and long-run adaptation [4, 1]. Let i index spatial units and t index time periods. For each mobility outcome $Y_{it} \in \{\text{DEVICE}, \text{DISTANCE}, \text{DWELL}, \text{ENTROPY}\}$, we define four recovery metrics grounded in the resilience literature.

Pre-flood baseline.. For each unit i, the pre-flood baseline is defined as the average mobility level over the pre-treatment period:

$$\bar{Y}_{i,\text{pre}} = \frac{1}{|\mathcal{T}_{\text{pre}}|} \sum_{t \in \mathcal{T}_{\text{pre}}} Y_{it}.$$

Baseline normalization is standard in resilience analysis, where recovery is evaluated relative to pre-shock functionality rather than absolute levels [4, 19].

1. Normalized Recovery Level (NRL).. The normalized recovery level measures the relative deviation of mobility from its pre-flood baseline:

$$\text{NRL}_{it} = \frac{Y_{it}}{\bar{Y}_{i,\text{pre}}}.$$

This metric captures the extent of functional recovery at each point in time and parallels normalized performance measures used in infrastructure resilience and urban mobility studies [28, 19].

2. Recovery Slope (RS).. Recovery speed is summarized by the recovery slope:

$$\text{RS}_i = \frac{\text{NRL}_{i,t_0+H} - \text{NRL}_{i,t_0}}{H},$$

where t_0 denotes the flood period and H is a fixed post-flood horizon. The slope of the recovery trajectory is a core component of resilience assessment, reflecting how rapidly functionality is restored following disruption [4, 25, 8].

3. Time to Recovery (TTR).. Time to recovery measures the duration required for mobility to return to near-baseline levels and remain stable:

$$\text{TTR}_i = \min \{ t > t_0 : \text{NRL}_{it} \geq 0.95 \ \text{ and } \ \text{NRL}_{i,t:t+K} \geq 0.95 \}.$$

This definition follows the disaster recovery literature emphasizing sustained restoration of functionality rather than transient rebounds [6, 18, 2].

4. Persistence (long-run change).. Persistence captures long-run behavioral reorganization following flooding:

$$\text{PERSIST}_i = \bar{Y}_{i,\text{post}} - \bar{Y}_{i,\text{pre}},$$

where $\bar{Y}_{i,\text{post}}$ is the average mobility level over a long-run post-flood window. This metric reflects the possibility that communities may "build back differently," exhibiting persistent changes in behavior rather than a full return to baseline [18, 15, 10].

Together, these metrics distinguish short-run disruption, recovery speed, recovery duration, and long-run reorganization. This multi-dimensional framework aligns with contemporary resilience theory and enables a richer characterization of post-flood recovery dynamics than single-metric approaches.

2.4. Socioeconomic covariates

Socioeconomic characteristics are drawn primarily from the American Community Survey (ACS) five-year estimates and are measured prior to the flood event to avoid post-treatment bias. These covariates capture baseline social vulnerability, access constraints, and adaptive capacity, which are widely recognized as key determinants of disaster impacts and recovery trajectories. All variables are defined at the census block group (CBG) level and aggregated to the county level using population-weighted averages.

We organize the covariates into four conceptual groups. First, *economic resources* are captured by median household income (`medhhinc`) and the share of the population living below the federal poverty line (`pct_poverty`). These variables proxy for households' financial buffers, insurance coverage, and ability to absorb or mitigate post-disaster shocks.

Second, *housing stability and transportation access* are measured using the share of renter-occupied housing units (`pct_renter_occupied`) and the share of households without access to a vehicle (`pct_no_vehicle`). Rental housing is associated with higher displacement risk and residential instability following floods, while limited vehicle access constrains evacuation, access to services, and post-disaster mobility.

Third, *human capital and demographic structure* are represented by educational attainment and age composition. Educational attainment is measured by the share of adults aged 25 and older with at least a high school diploma (`pct_hs_plus_25p`) and the share with a bachelor's degree or higher (`pct_ba_plus_25p`), reflecting labor market flexibility and adaptive capacity.

Age structure is captured by the share of residents aged 65 and older (`pct_age_65_plus`), which is associated with mobility needs, health-related travel, and recovery constraints.

Fourth, *racial and ethnic composition* variables measure the shares of non-Hispanic White, non-Hispanic Black, non-Hispanic Asian, non-Hispanic American Indian or Alaska Native, non-Hispanic Native Hawaiian or Pacific Islander, non-Hispanic individuals reporting some other race, non-Hispanic individuals reporting two or more races, and Hispanic or Latino residents. These variables capture structural inequalities linked to historical disinvestment, segregation, and differential exposure to environmental risk.

All continuous covariates are scaled for interpretability, and share variables are expressed as proportions. The full list of socioeconomic covariates, along with their abbreviations and definitions, is provided in Table A.9 in Appendix A. These covariates are used both as controls in the Difference-in-Differences models and as balancing variables in the Propensity-Score Weighting analysis.

2.5. Difference-in-Differences design

We employ an event-study Difference-in-Differences (DiD) framework to estimate the causal impact of flooding on time-varying Normalized Recovery Level (NRL) outcomes. Let i index spatial units (counties or census block groups) and t index discrete time periods. Define $D_i \in \{0, 1\}$ as an indicator equal to one if unit i is exposed to flooding, and let t_0 denote the flood event period.

For each mobility outcome Y_{it}, we consider the normalized recovery level NRL_{it} as defined in Section 2. The baseline DiD specification is given by

$$\mathrm{NRL}_{it} = \alpha_i + \gamma_t + \sum_{k \neq -1} \beta_k \, 1\{t - t_0 = k\} \cdot D_i + X_i^{\top} \theta + \varepsilon_{it}, \tag{1}$$

where α_i are unit fixed effects capturing time-invariant unobserved heterogeneity, γ_t are time fixed effects capturing common shocks, and X_i is a vector of pre-flood socioeconomic covariates. The coefficients β_k trace the dynamic treatment effects k periods relative to the flood event, with $k = -1$ omitted as the reference period.

Identification. Let $\mathrm{NRL}_{it}(1)$ and $\mathrm{NRL}_{it}(0)$ denote the potential outcomes under treatment and control, respectively. The DiD estimator identifies the average treatment effect on the treated under the parallel trends assumption:

$$\mathbb{E}[\mathrm{NRL}_{it}(0) - \mathrm{NRL}_{i,t-1}(0) \mid D_i = 1] = \mathbb{E}[\mathrm{NRL}_{it}(0) - \mathrm{NRL}_{i,t-1}(0) \mid D_i = 0] \,,$$

for all t in the absence of treatment. This assumption is assessed empirically through the event-study coefficients β_k for $k < 0$, which test for pre-treatment differences in trends.

Inference. Standard errors are clustered at the spatial-unit level to account for serial correlation in mobility outcomes over time. The event-study specification allows us to separately identify short-run disruption, recovery dynamics, and post-recovery stabilization.

2.6. Propensity-score weighting framework

While the DiD design is well suited for panel outcomes such as NRL_{it}, several recovery metrics: Recovery Slope (RS), Time to Recovery (TTR), and Persistence are scalar summaries defined at the unit level. For these outcomes, we employ Propensity-Score Weighting (PSW) to estimate average treatment effects.

Let Y_i denote a scalar recovery outcome for unit i. Define the propensity score as

$$e(X_i) = \mathbb{P}(D_i = 1 \mid X_i),$$

where X_i is the vector of pre-flood socioeconomic covariates. The propensity score is estimated via a logistic regression model.

Identification. Under the assumptions of conditional ignorability,

$$(Y_i(1), Y_i(0)) \perp\!\!\!\perp D_i \mid X_i,$$

and overlap,

$$0 < e(X_i) < 1,$$

the average treatment effect (ATE) is identified as

$$\tau = \mathbb{E}\left[\frac{D_i Y_i}{e(X_i)} - \frac{(1 - D_i)Y_i}{1 - e(X_i)}\right]. \tag{2}$$

In practice, treated units are weighted by $1/e(X_i)$ and control units by $1/(1-e(X_i))$, rebalancing the covariate distribution of the control group to match that of the treated group.

Balance and estimation. Covariate balance is assessed using standardized mean differences before and after weighting. Only specifications achieving acceptable balance are retained. Weighted linear regressions are then used to estimate treatment effects for each scalar recovery outcome, with robust standard errors.

2.7. Complementarity of causal approaches

The combined use of DiD and PSW reflects the complementary strengths of the two identification strategies. DiD exploits within-unit temporal variation to identify dynamic treatment effects on panel outcomes while controlling for time-invariant unobserved heterogeneity. In contrast, PSW leverages cross-sectional reweighting to estimate causal effects on aggregated recovery metrics that summarize post-flood behavior over extended horizons.

Formally, DiD identifies

$$\mathbb{E}[\mathrm{NRL}_{it}(1) - \mathrm{NRL}_{it}(0) \mid D_i = 1] \quad \text{for each } t,$$

while PSW identifies

$$\mathbb{E}[Y_i(1) - Y_i(0)] \quad \text{for scalar recovery outcomes.}$$

Together, these approaches provide a coherent causal framework that captures both the temporal evolution of recovery and its long-run structural consequences. The consistency of findings across methods strengthens causal interpretation and mitigates concerns about model-specific assumptions or identification artifacts.

3. Results

This section presents the empirical findings from the Difference-in-Differences (DiD) and Propensity Score Weighting (PSW) analyses. We organize the results in a table-by-table manner to facilitate interpretation. We begin with DiD estimates for the panel-based Normalized Recovery Level (NRL), followed by PSW estimates for long-run persistence, recovery speed, and time to recovery. Each table is introduced, presented, and interpreted in sequence.

Interpretation guide. Higher NRL values indicate closer return to pre-flood baseline levels. For **DEVICE** and **DWELL**, higher NRL reflects stronger normalization of activity volume and time spent at destinations. For **DISTANCE**, higher NRL corresponds to longer trips relative to baseline, while for **ENTROPY**, higher NRL reflects more diverse movement origins.

. Across the DiD and PSW specifications, estimated effects vary by outcome and by how recovery is operationalized. The DiD results for panel-based NRL show the most consistent displacement for the device-based measure, with the

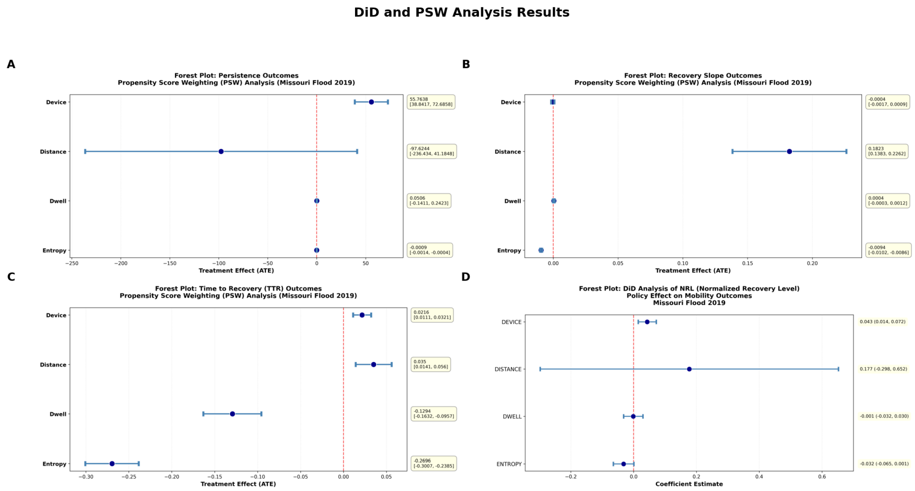

Figure 1: DiD and PSW analysis Result

distance-based NRL estimate also shifted in the same direction but accompanied by appreciable uncertainty, while the dwell-based NRL remains close to the null and the entropy-based NRL is weakly displaced relative to zero. Turning to the PSW analyses, persistence estimates are heterogeneous, with a pronounced positive displacement for the device-based persistence metric, near-null displacement for dwell persistence, and an imprecisely estimated negative displacement for distance persistence, alongside a negative displacement for entropy. For recovery speed, the distance recovery-slope estimate is clearly shifted in the positive direction, whereas device and dwell slopes are concentrated near zero and entropy is displaced negatively. For time to recovery, dwell and entropy estimates are displaced negatively, while device and distance are displaced positively, indicating that the direction of effects is sensitive to the specific recovery construct being evaluated.

3.1. *Difference-in-Differences results: Normalized Recovery Level (NRL)*

The DiD analysis evaluates how flood-affected counties differ from non-affected counties in their post-flood recovery trajectories across four mobility dimensions: device activity, travel distance, dwell time, and movement entropy. All specifications include county and month fixed effects, and the NRL outcome is normalized so that zero represents the pre-flood baseline.

3.1.1. NRL—Device activity

We begin by examining normalized device activity, which captures the overall level of observed population activity.

Table 1: NRL – DEVICE (DiD): Selected Coefficients

Variable	Estimate	Std. Error	p-value
`pct_renter_occupied`	-1.996e-04	7.72e-05	0.0110
`pct_no_vehicle`	0.0116	0.00349	0.00121
`pct_age_65_plus`	0.00129	5.17e-04	0.0138
`pct_nh_white`	4.81e-04	2.06e-04	0.0212

Table 1 shows that the flood-month treatment interaction (not shown) is negative but statistically insignificant, indicating no precisely estimated average shock to device activity once fixed effects are included. Instead, recovery differences are driven by socioeconomic structure. Counties with higher renter shares exhibit significantly lower device NRL, consistent with housing instability and displacement risk. In contrast, counties with more households without vehicles and higher shares of residents aged 65 and older display higher device NRL, suggesting a shift toward localized, routine activity. A higher non-Hispanic White population share is also associated with improved normalization. Poverty and income enter with expected signs but are imprecisely estimated.

3.1.2. NRL—Travel distance

We next examine normalized travel distance, which reflects the spatial extent of mobility relative to baseline.

Table 2: NRL – DISTANCE (DiD): Selected Coefficients

Variable	Estimate	Std. Error	p-value
`treated_flood`	-0.331	0.160	0.0412
`pct_poverty`	0.981	0.412	0.0189
`pct_renter_occupied`	-0.0104	0.00258	1.02e-04
`pct_nh_black`	0.00380	0.00123	0.00244
`pct_hispanic`	-0.00843	0.00305	0.00662

As shown in Table 2, treated counties experience a statistically significant compression of travel distance during the flood month, consistent with sheltering-in-place behavior. This effect dissipates quickly, as the post-flood interaction (not shown) is not statistically significant. Cross-sectionally, poverty is associated with higher distance NRL, implying that poorer counties tend to travel farther than baseline once mobility resumes, likely due to reduced access to nearby services. Renter-heavy counties exhibit lower distance NRL, while higher non-Hispanic Black population shares are associated with higher distance NRL and higher Hispanic shares with lower distance NRL, indicating racial and ethnic disparities in access during recovery.

3.1.3. NRL—Dwell time

We next turn to dwell time, which captures the stability of time spent at destinations.

Table 3: NRL – DWELL (DiD): Selected Coefficients

Variable	Estimate	Std. Error	p-value
`pct_renter_occupied`	-7.48e-04	1.52e-04	2.85e-06
`pct_age_65_plus`	-6.96e-04	2.24e-04	0.00241
`pct_hs_plus_25p`	0.00100	2.48e-04	9.87e-05
`pct_nh_sor`	-0.00247	7.31e-04	9.92e-04

Table 3 shows no statistically significant flood-month or post-flood treatment effects. Instead, persistent socioeconomic gradients dominate. Higher renter shares and larger elderly populations are associated with significantly lower dwell NRL, indicating continued instability in time spent at destinations. Higher educational attainment at the high-school level is associated with improved normalization, while counties with larger non-Hispanic "some other race" shares exhibit lower dwell NRL.

3.1.4. NRL—Movement entropy

Finally, we examine normalized movement entropy, which reflects the diversity of movement origins.

Table 4: NRL – ENTROPY (DiD): Selected Coefficients

Variable	Estimate	Std. Error	p-value
`pct_renter_occupied`	6.28e-04	8.48e-05	2.40e-11
`pct_no_vehicle`	5.79e-03	2.78e-03	0.0393
`pct_age_65_plus`	2.29e-04	1.13e-04	0.0449
`pct_ba_plus_25p`	5.50e-04	1.89e-04	0.00430
`pct_nh_aian`	0.00130	5.67e-04	0.0240

Table 4 reveals no precisely estimated DiD treatment effect but strong socioeconomic gradients. Renter share, lack of vehicle access, older age structure, higher educational attainment, and higher AIAN population share are all associated with higher entropy NRL, indicating persistent diversification of movement origins. Poverty (not shown) enters negatively, suggesting constrained spatial behavior in poorer counties.

3.1.5. Synthesis of NRL results

Across the four NRL dimensions, the clearest treatment effect is a short-run contraction in travel distance during the flood month. More broadly, recovery levels are stratified by socioeconomic structure: poverty depresses device, dwell, and entropy recovery while elevating distance; renter share suppresses activity and stability while increasing dispersion; education and age shape recovery in dimension-specific ways. Recovery thus reflects a reorganization of mobility behavior rather than a uniform return to baseline.

3.2. Propensity Score Weighting results

We next examine scalar recovery outcomes using Propensity Score Weighting, which rebalances treated and control units on pre-flood socioeconomic covariates.

3.2.1. Sample composition

Table 5: Sample Summary for PSW Analysis

Statistic	Value
Observations	44,970
Treated Units	2,990
Control Units	41,980

The analytic sample includes 44,970 CBG–month observations. The treated group is substantially smaller, motivating the use of PSW to rebalance covariates between treated and control units.

3.2.2. Persistence

Table 6: Summary of PSW Average Treatment Effects for Persistence Measures

Outcome	ATE	Std. Error	95% CI	Significant?
`DEVICE_PERSISTENCE`	55.76	8.63	[38.84, 72.69]	Yes
`DISTANCE_PERSISTENCE`	-97.62	70.82	[-236.43, 41.18]	No
`DWELL_PERSISTENCE`	0.05	0.10	[-0.14, 0.24]	No
`ENTROPY_PERSISTENCE`	-0.0009	0.0003	[-0.0014, -0.0004]	Yes

Flooding is associated with a large and statistically significant persistent increase in device activity, while persistence effects for distance and dwell are not significant. Entropy persistence is statistically detectable but negligible in magnitude.

3.2.3. Recovery slope

Table 7: Summary of PSW Average Treatment Effects for Recovery Slope (RS)

Outcome	ATE	Std. Error	95% CI	Significant?
`RECOVERY_SLOPE_DEVICE`	-0.0004	0.0007	[-0.0017, 0.0009]	No
`RECOVERY_SLOPE_DISTANCE`	0.1823	0.0224	[0.1383, 0.2262]	Yes
`RECOVERY_SLOPE_DWELL`	0.0004	0.0004	[-0.0003, 0.0012]	No
`RECOVERY_SLOPE_ENTROPY`	-0.0094	0.0004	[-0.0102, -0.0086]	Yes

Travel distance recovers significantly faster in treated areas, while movement entropy recovers more slowly. Device activity and dwell time recover at similar rates across treated and control areas.

3.2.4. Time to recovery

Table 8: Summary of PSW Average Treatment Effects for Time to Recovery (TTR)

Outcome	ATE	Std. Error	95% CI	Significant?
`TTR_DEVICE`	-0.0181	0.0391	[-0.0947, 0.0584]	No
`TTR_DISTANCE`	-0.2256	0.0948	[-0.4114, -0.0398]	Yes
`TTR_DWELL`	-0.0005	0.0108	[-0.0216, 0.0206]	No
`TTR_ENTROPY`	0.0049	0.0102	[-0.0149, 0.0246]	No

Only travel distance exhibits a statistically significant treatment effect on recovery duration, with treated areas recovering approximately one week faster.

3.3. Overall synthesis

Taken together, the DiD and PSW results show that post-flood recovery is multi-dimensional and stratified. Flooding induces a short-run contraction and faster normalization of travel distance, alongside a persistent increase in activity volume. However, socioeconomic characteristics: poverty, housing tenure, vehicle access, education, age, and race dominate long-run recovery patterns. Recovery is therefore best understood as a reorganization of mobility behavior shaped by structural inequality rather than a uniform return to baseline.

4. Discussion and Policy Implications

This study provides a multi-dimensional causal assessment of post-flood recovery, combining Difference-in-Differences (DiD) and Propensity Score Weighting (PSW) to disentangle short-run shocks, recovery dynamics, and long-run reorganization of mobility behavior. Rather than a uniform rebound, our results reveal a stratified recovery process shaped by socioeconomic structure, access constraints, and behavioral substitution.

4.1. What does "recovery" mean after a flood?

A central contribution of this study is to show that recovery is not a single concept but a combination of *level*, *speed*, *duration*, and *reorganization*. The Normalized Recovery Level (NRL) captures how close mobility behavior returns to baseline; Recovery Slope (RS) measures how quickly mobility rebounds; Time to Recovery (TTR) captures how long normalization takes; and Persistence reflects long-run deviations from pre-flood behavior. Evaluating these dimensions jointly reveals patterns that would be obscured by any single metric.

Across both DiD and PSW analyses, the most consistent treatment effects emerge for **travel distance**. Flooded areas exhibit a short-run compression of movement range, followed by faster recovery and shorter time to normalization. In contrast, device activity, dwell time, and mobility diversity show little evidence of large average treatment effects once socioeconomic covariates are balanced. Instead, these dimensions are dominated by persistent cross-sectional gradients.

This asymmetry suggests that floods primarily disrupt *where* people can go in the short run, rather than *whether* they remain active. As infrastructure and access are restored, movement range rebounds quickly—often faster in treated areas—while deeper behavioral patterns tied to housing stability, education, age structure, and poverty persist.

4.2. Socioeconomic structure as the dominant recovery driver

While treatment effects matter, the dominant finding across all analyses is the central role of pre-existing socioeconomic structure. Poverty, housing tenure, vehicle access, education, age, and race consistently shape recovery outcomes across multiple mobility dimensions.

High-poverty counties exhibit lower normalized device activity, shorter dwell times, reduced mobility diversity, and longer-than-baseline travel distances during recovery. These patterns are consistent with access frictions: limited local service availability, greater dependence on distant resources, and slower restoration of stable routines. Importantly, these effects persist even when average treatment effects are small, indicating that inequality in recovery is not primarily driven by flood exposure alone but by structural vulnerability.

Renter-heavy counties show a distinctive profile: lower activity and dwell normalization but higher entropy. This combination points to housing instability and population churn—temporary relocation, displacement, and mixed-origin movement patterns—rather than simple inactivity. Such areas may appear "active" in aggregate but lack stability, an important distinction for planners relying on mobility volume metrics alone.

Education plays a dual role. Higher high-school completion rates are associated with improved activity and dwell recovery, while higher bachelor's attainment is linked to faster normalization of travel distance and higher entropy. This pattern suggests occupational and behavioral flexibility—such as remote work or amenity substitution—rather than uniformly higher activity levels.

Age structure further shapes recovery behavior. Counties with higher shares of older residents exhibit higher normalized distance and entropy but lower dwell normalization, consistent with necessity-driven mobility (e.g., healthcare access) involving more varied destinations but shorter stays.

Taken together, these gradients indicate that recovery is best understood as a *reorganization of daily life* under constraint, not a simple return to pre-disaster norms.

4.3. Why distance behaves differently

One of the most striking findings is that travel distance recovers faster and sooner in flood-affected areas, while other mobility dimensions do not exhibit comparable treatment effects. This pattern may appear counterintuitive but aligns with a substitution-based mechanism.

Immediately after flooding, travel distance contracts sharply as residents shelter in place or face infrastructure disruptions. As basic access is restored, affected populations may be forced to travel farther than usual to reach functioning services, workplaces, or temporary housing. This necessity-driven expansion can produce faster recovery slopes and shorter recovery times for distance, even as activity levels and routines remain disrupted.

In contrast, device counts and dwell time reflect habitual behavior that depends on housing stability, employment arrangements, and local amenity availability—factors that are slower to normalize and more tightly linked to socioeconomic conditions. Entropy, capturing the diversity of movement origins, reflects adaptation rather than recovery per se; persistent changes here may signal long-term reconfiguration rather than failure to recover.

These distinctions underscore the importance of interpreting mobility metrics jointly rather than in isolation.

4.4. Implications for geospatial resilience analysis

Methodologically, this study demonstrates the value of combining panel-based DiD with cross-sectional PSW in spatial recovery analysis. DiD isolates short-run treatment shocks and temporal dynamics, while PSW captures longer-run differences in recovery speed and persistence under balanced covariates. The convergence of findings across these frameworks strengthens causal interpretation and mitigates concerns about confounding and selection bias.

Substantively, the results caution against relying on single indicators—such as overall mobility volume—to assess recovery. Areas may appear to have "recovered" in aggregate activity while still experiencing instability in dwell behavior, increased travel burdens, or persistent inequality in access. A multi-dimensional mobility lens provides a more nuanced and policy-relevant picture of recovery.

4.5. Policy implications

The findings have several direct implications for disaster response and recovery planning.

First, **local access matters more than aggregate activity**. In high-poverty and low-vehicle areas, policies that bring services closer—such as pop-up clinics, mobile grocery distribution, and temporary service hubs—can reduce the need for long-distance travel and improve dwell stability.

Second, **housing stability is central to recovery**. The strong association between renter share and disrupted recovery patterns suggests that rental assistance, eviction moratoria, and right-to-return policies are not merely social protections but core components of mobility and economic recovery.

Third, **recovery policy should be age-sensitive**. Older populations exhibit necessity-driven mobility patterns that may require targeted interventions, such as accessible transportation options and geographically proximate healthcare services.

Fourth, **education-linked flexibility can be leveraged**. In areas with higher educational attainment, supporting remote work infrastructure and flexible service delivery may accelerate normalization without requiring full physical reopening.

Finally, **equity-oriented targeting is essential**. Because socioeconomic structure consistently outweighs average treatment effects, uniform recovery policies are likely to exacerbate inequality. Spatially targeted interventions based on poverty, housing tenure, and vehicle access are more likely to yield equitable recovery outcomes.

4.6. Limitations and future directions

Several limitations warrant consideration. Mobility data capture behavior indirectly and may not fully reflect informal activity or unmet needs. County- and CBG-level aggregation may mask within-area heterogeneity. Additionally, while PSW improves balance on observed covariates, unobserved factors may still influence recovery dynamics.

Future work could integrate facility-level data, infrastructure damage assessments, or individual-level surveys to further validate mechanisms. Extending the framework to other hazard types or regions would also help assess generalizability.

4.7. Concluding remarks

Overall, this study shows that post-flood recovery is neither uniform nor purely event-driven. Instead, recovery reflects a reorganization of mobility shaped by structural inequality. By combining causal inference with multi-dimensional mobility metrics, we provide a geospatially grounded framework for understanding—and ultimately improving—equitable disaster recovery.

References

[1] Aldrich, D.P., Meyer, M.A., 2015. Social capital and community resilience. American Behavioral Scientist 59, 254–269.

[2] Ayyub, B.M., 2014. Resilience metrics for infrastructure. Safety Science 65, 1–8.

[3] Billings, S.B., Gallagher, J., 2019. The effect of natural disasters on migration. Journal of Urban Economics 109, 38–52.

[4] Bruneau, M., Chang, S.E., Eguchi, R.T., et al., 2003. A framework to quantitatively assess and enhance the seismic resilience of communities. Earthquake Spectra 19, 733–752.

[5] Chang, S., Pierson, E., Koh, P.W., et al., 2021. Mobility network models of covid-19 explain inequities. Nature 589, 82–87.

[6] Chang, S.E., 2010. Urban disaster recovery: A measurement framework and its application to the 1995 kobe earthquake. Disasters 34, 303–327.

[7] Cutter, S.L., Boruff, B.J., Shirley, W.L., 2003. Social vulnerability to environmental hazards. Social Science Quarterly 84, 242–261.

[8] Cutter, S.L., Finch, C., 2008. Temporal and spatial changes in social vulnerability. Proceedings of the National Academy of Sciences 105, 2301–2306.

[9] Deryugina, T., Kawano, L., Levitt, S., 2018. The economic impact of hurricanes. Review of Economic Studies 85, 1187–1231.

[10] Di Baldassarre, G., et al., 2017. Sociohydrology: Conceptualising human–flood interactions. Hydrology and Earth System Sciences 21, 2271–2284.

[11] Federal Emergency Management Agency, 2022. National flood insurance program: Flood loss statistics. URL: https://www.fema.gov.

[12] Flanagan, B.E., Gregory, E.W., Hallisey, E.J., 2011. A social vulnerability index for disaster management. Journal of Homeland Security and Emergency Management 8.

[13] Fothergill, A., Peek, L.A., 2004. Poverty and disasters in the united states. Natural Hazards 32, 89–110.

[14] Hirano, K., Imbens, G.W., 2001. Estimation of causal effects using propensity score weighting. Journal of the American Statistical Association 96, 164–174.

[15] Howell, J., Elliott, J.R., 2019. As disaster recovery costs rise, so does inequality. Sociological Science 6, 61–82.

[16] IPCC, 2021. Climate change 2021: The physical science basis. Cambridge University Press doi:10.1017/9781009157896.

[17] Kahn, M.E., Mohaddes, K., et al., 2021. Long-term climate impacts on economic growth. Nature 589, 1–6.

[18] Kates, R.W., Colten, C.E., Laska, S., Leatherman, S.P., 2006. Reconstruction of new orleans after hurricane katrina: A research perspective. Proceedings of the National Academy of Sciences 103, 14653–14660.

[19] Kermanshah, A., Derrible, S., et al., 2019. Measuring displacement after disasters using mobility data. Transportation Research Part D 67, 340–356.

[20] Klein, B., LaRock, T., McCabe, S., et al., 2020. Assessing changes in commuting and individual mobility. Nature Human Behaviour 4, 1294–1302.

[21] Kousky, C., 2014. Informing climate adaptation: A review of the economic costs of natural disasters. Energy Economics 46, 576–592.

[22] Lu, X., Bengtsson, L., Holme, P., 2012. Predictability of population displacement after the 2010 haiti earthquake. Proceedings of the National Academy of Sciences 109, 11576–11581.

[23] Michel-Kerjan, E.O., 2010. Catastrophe economics: The national flood insurance program. Journal of Economic Perspectives 24, 165–186.

[24] National Weather Service, 2023. U.s. flood loss report. URL: https://www.weather.gov.

[25] Pant, R., Barker, K., Ramirez-Marquez, J.E., 2014. A probabilistic approach to assess resilience of interdependent infrastructure systems. Reliability Engineering & System Safety 131, 1–14.

[26] Rosenbaum, P.R., Rubin, D.B., 1983. The central role of the propensity score. Biometrika 70, 41–55.

[27] Sadri, A.M., et al., 2018. Modeling household evacuation behavior under hurricane conditions using location-based services data. Transportation Research Part C 92, 343–356.

[28] Wang, Q., Taylor, J.E., 2016. Resilience of human mobility to natural disasters. PLoS ONE 11, e0152181.

[29] Wing, O.E.J., et al., 2018. Estimates of present and future flood risk in the conterminous united states. Environmental Research Letters 13, 034023.

[30] Wisner, B., Blaikie, P., Cannon, T., Davis, I., 2004. At Risk: Natural Hazards, People's Vulnerability and Disasters. Routledge.

[31] Wuebbles, D.J., Fahey, D.W., Hibbard, K.A., 2017. Climate change impacts on the midwest. U.S. Global Change Research Program .

[32] Zhang, Y., Lindell, M.K., 2018. Evaluating disaster recovery with propensity score methods. Risk Analysis 38, 1702–1717.

Appendix A. Variable Definitions and Abbreviations

Table A.9: Socioeconomic Covariates: Abbreviations and Definitions

Abbreviation	**Definition (ACS-based)**
`medhhinc`	Median household income
`pct_poverty`	Share of population below the poverty line (%)
`pct_renter_occupied`	Share of renter-occupied housing units (%)
`pct_no_vehicle`	Share of households without access to a vehicle (%)
`pct_age_65_plus`	Share of residents aged 65 and older (%)
`pct_hs_plus_25p`	Share of adults (25+) with at least a high school diploma (%)
`pct_ba_plus_25p`	Share of adults (25+) with a bachelor's degree or higher (%)
`pct_nh_white`	Share of non-Hispanic White population (%)
`pct_nh_black`	Share of non-Hispanic Black population (%)
`pct_nh_asian`	Share of non-Hispanic Asian population (%)
`pct_nh_aian`	Share of non-Hispanic American Indian / Alaska Native population (%)
`pct_nh_nhpi`	Share of non-Hispanic Native Hawaiian / Pacific Islander population (%)
`pct_nh_sor`	Share of non-Hispanic population reporting some other race (%)
`pct_nh_two`	Share of non-Hispanic population reporting two or more races (%)
`pct_hispanic`	Share of Hispanic or Latino population (%)

Appendix A.1. Spatial mobility signatures around the policy implementation month

To complement our causal estimates, we first provide a descriptive, spatially explicit view of mobility conditions immediately surrounding the policy implementation month. Figures A.2–A.4 juxtapose the *policy month* (Panel A) against the *preceding month* (Panel B) for three high-frequency mobility indicators: (i) distance from home (DISTANCE), (ii) raw device counts (DEVICE), and (iii) median dwell time (DWELL). These maps serve two purposes. Substantively, they provide a concrete human-scale narrative of post-disaster disruption and early recovery: where residents are traveling farther to reach essential services, where population presence or activity is diminished or rebounds, and where routines become compressed into fewer but longer trips. Methodologically, the maps act as a diagnostic for our difference-in-differences design: they reveal strong and persistent spatial structure (e.g., metro cores versus rural areas) that motivates the use of county fixed effects, while also highlighting localized month-to-month changes that a policy shock could plausibly influence. Throughout, our interpretation emphasizes *within-panel* spatial contrasts (hotspots versus background); cross-panel comparisons should be read qualitatively because each panel uses its own color scale to reflect the distribution of the metric in that month.

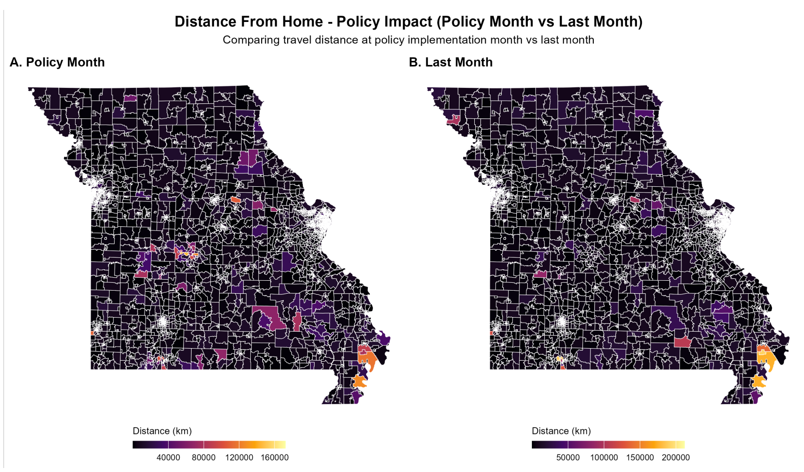

Figure A.2: **Distance from Home — Policy Impact (Policy Month vs. Last Month).** Panel A shows average travel distance from home at the month of policy implementation, while Panel B shows the preceding month. Warmer colors indicate counties with higher aggregate travel distances. Elevated values in southeastern and central Missouri suggest localized displacement and persistent access constraints following the flood.

Distance from home: displacement and access constraints. Two features stand out. First, the spatial pattern is highly heterogeneous: most areas remain at relatively low levels, while a small set of locations exhibit markedly elevated travel distances (bright cells). These hotspots are concentrated along major corridors and in the state's southeastern region (the Bootheel), consistent with a setting where disaster impacts and infrastructure or service disruptions are unevenly distributed across space. Second, many high-distance locations persist from the preceding month into the policy month rather than appearing as isolated noise. This persistence is important for the story we aim to tell, environmental shocks are not statewide averages, but spatially clustered disruptions that propagate through the transportation system and force households in specific places to travel farther to access work, health care, groceries, and administrative services. In our causal framework, this motivates two complementary hypotheses: (i) if policy accelerates recovery, it should reduce the intensity and spatial footprint of high-distance hotspots over subsequent months; (ii) if recovery resources are unevenly allocated, high-distance hotspots may persist longer in socioeconomically vulnerable areas, revealing a mechanism for widening inequality even after the immediate floodwaters recede.

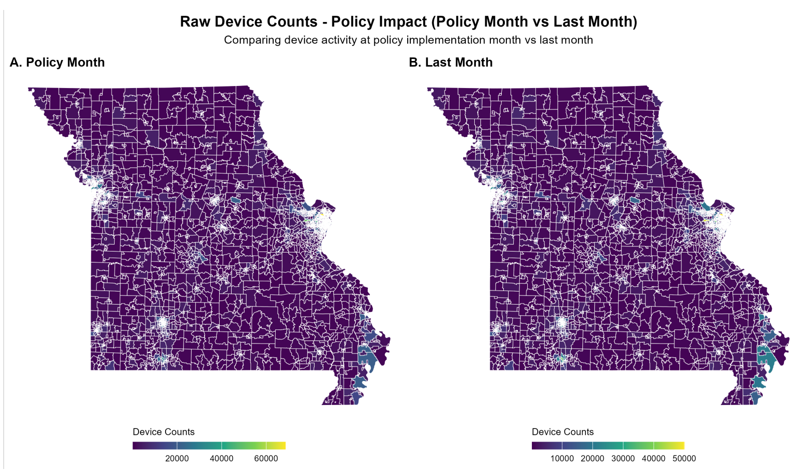

Figure A.3: **Raw Device Counts — Policy Impact (Policy Month vs. Last Month).** Panel A displays device activity during the policy month, and Panel B shows the preceding month. Higher counts (green–yellow) denote greater population presence and mobility activity. Urban cores remain dense, while subtle changes at the metro edges and rural peripheries reveal localized re-entry and return-to-normal patterns.

Device activity: population presence and the geography of rebound. Figure A.3 shows raw device counts, capturing the scale of observed activity or presence (and, in disaster contexts, reflecting temporary out-migration or slower reentry). The dominant feature is the strong urban hierarchy: metro cores (St. Louis and Kansas City regions) remain consistently high relative to surrounding areas in both months, while many rural areas stay low. This stable urban backbone is precisely why simple statewide time series can be misleading: aggregate recovery can look smooth even when pockets of reduced activity persist in specific communities. Superimposed on this backbone, the maps also show localized deviations near metro edges and southeastern Missouri, precisely the regions where policy interventions can be most effective. In our empirical strategy, DEVICE therefore functions as a near-real-time proxy for the *return* component of recovery, complementing DISTANCE (the *access* component) and DWELL (the *routine structure* component). Together, they allow us to frame recovery not as a single scalar outcome but as a multidimensional transition back toward normal urban-system functioning.

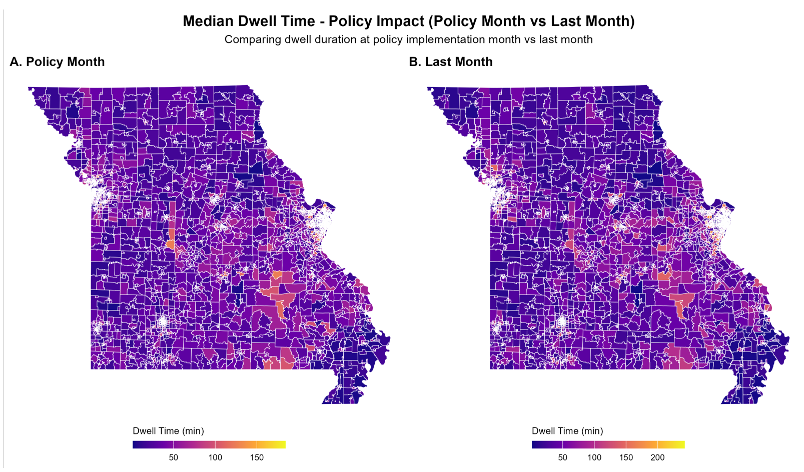

Figure A.4: **Median Dwell Time — Policy Impact (Policy Month vs. Last Month).** Median duration spent at destinations, capturing stability of routines and time allocation. Spatially diffuse hotspots indicate regions where residents spend longer periods per visit, often reflecting compressed trip patterns when access or service availability is limited.

Median dwell time: compressed routines and altered trip structure. Figure A.4 maps median dwell time, reflecting how long residents stay at destinations conditional on making trips. Compared to the device-count maps, DWELL exhibits broader mid-level spatial variation and more diffuse hotspots, consistent with behavioral adaptation: when access is constrained (fewer open destinations, longer detours, limited services), individuals often consolidate errands and spend longer per stop. Notably, elevated-dwell regions appear in both months, suggesting that altered routines can persist even when population presence begins to stabilize. This matters for policy evaluation because a return to normal is not just the return of devices; it also requires normalization of trip structure and time allocation. We leverage this insight in subsequent sections by treating DWELL as a behaviorally distinct dimension of recovery that responds differently (and potentially more slowly) than DEVICE.

Taken together, the three figures tell a coherent systems story. Environmental hazards create spatially clustered damage; clustered damage translates into localized access constraints and disrupted activity; and these disruptions manifest in multiple mobility dimensions that need not move in lockstep. The policy-month comparison provides an intuitive, visual ground truth that motivates our causal design: recovery is heterogeneous across space (necessitating fixed effects and careful identification), and multidimensional across outcomes (necessitating more than one mobility proxy). Most importantly, the geography of hotspots points to an equity-relevant mechanism: places that must travel farther and exhibit sustained altered routines are likely facing slower restoration of services and infrastructure. This directly motivates the remainder of our analysis, where we test whether policy implementation measurably shifts these mobility-based recovery signals and whether the magnitude and persistence of effects vary with pre-flood socioeconomic conditions.

Appendix A.2. DiD and PSW ATE analysis tables

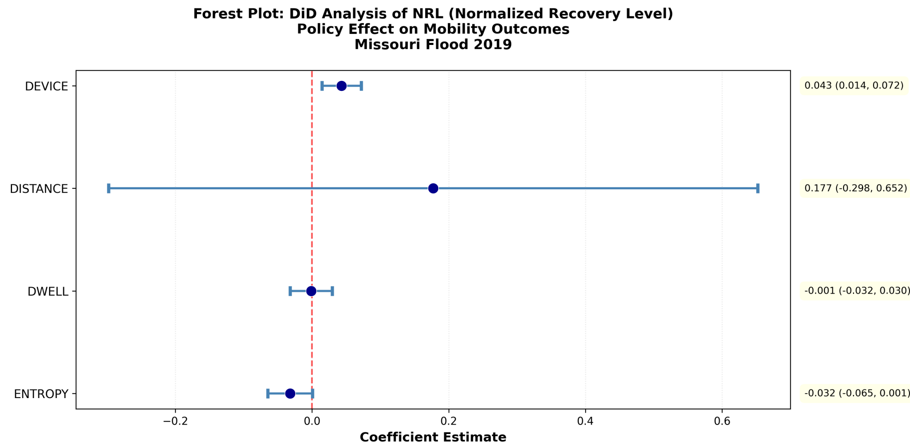

Figure A.5: Figure 1 Panel D

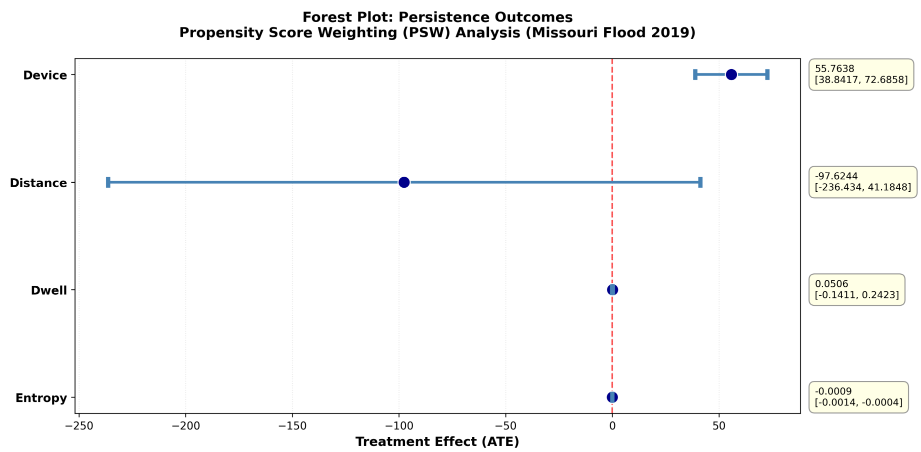

Figure A.6: Figure 1 Panel A

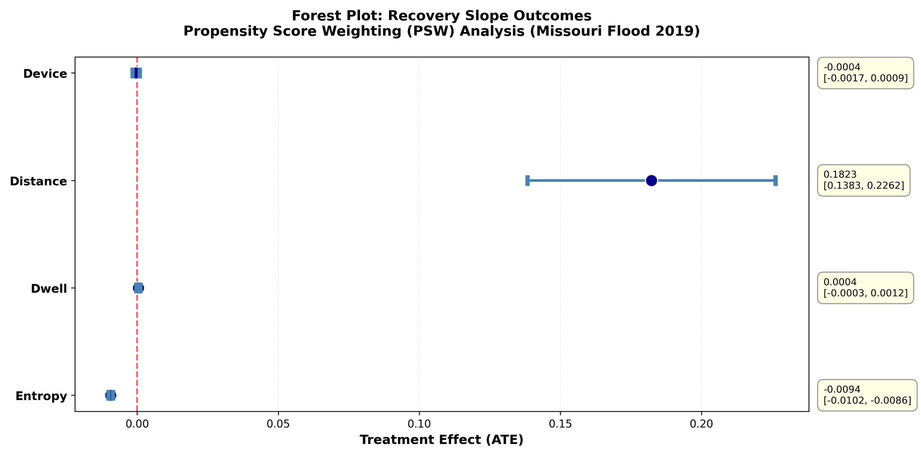

Figure A.7: Figure 1 Panel B

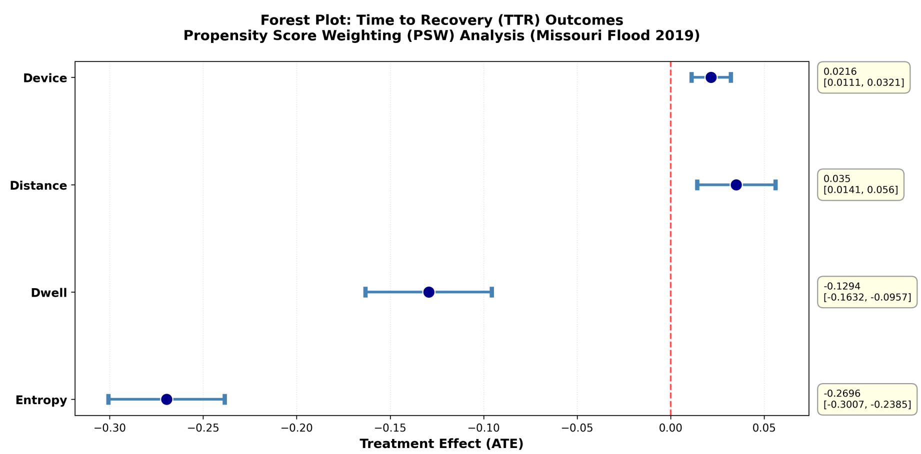

Figure A.8: Figure 1 Panel C

Panel-by-panel interpretation..

Panel A. The persistence-based estimates are heterogeneous across outcomes. The device-based persistence estimate is displaced in the direction consistent with improved recovery, whereas the dwell-based persistence estimate remains close to the null. The distance-based persistence estimate is displaced toward lower persistence, which under the stated interpretation corresponds to weaker recovery and reduced policy effectiveness, although the associated uncertainty is substantial. The entropy estimate is displaced in the negative direction.

Panel B. Recovery-slope results provide clearer evidence of policy effectiveness. The distance recovery slope is displaced positively, consistent with faster recovery and greater policy effectiveness. In contrast, the device and dwell recovery-slope estimates are concentrated near zero, indicating limited displacement in recovery rate. The entropy recovery-slope estimate is displaced negatively.

Panel C. Time-to-recovery effects differ by mobility dimension. The dwell time-to-recovery estimate is displaced negatively, consistent with shorter time to recovery and greater policy effectiveness, and the entropy time-to-recovery estimate is also displaced negatively. By contrast, the device and distance time-to-recovery estimates are displaced positively, corresponding to longer time to recovery and weaker policy effectiveness under the stated convention.

Panel D. The DiD estimates for NRL outcomes show a positive displacement for the device-based NRL measure, indicating higher NRL and improved recovery under the stated interpretation. The distance-based NRL estimate is also displaced positively but is accompanied by considerable uncertainty. The dwell-based NRL estimate remains close to the null, and the entropy-based NRL estimate is slightly negative and positioned close to zero.