

Social Vulnerability and Long-Run Tropical Cyclone Mortality: Evidence from Texas

Gizem Toplu Tutay

gizem@utexas.edu

Georgetown University <https://orcid.org/0009-0006-8628-9061>

John J. Hasenbein

The University of Texas at Austin

Erhan Kutanoglu

UT Austin: The University of Texas at Austin

Research Article

Keywords: Tropical cyclones, Long-run mortality, Social vulnerability, Disaster risk, Public health, Reduced-form panel model

Posted Date: September 10th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10831452/v1>

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

[Read Full License](#)

Social Vulnerability and Long-Run Tropical Cyclone Mortality: Evidence from Texas

Gizem Toplu-Tutay^{1,2*}, John J. Hasenbein¹ and Erhan Kutanoglu¹

¹*Operations Research and Industrial Engineering, The University of
Texas at Austin, Austin, TX, USA .

²The Earth Commons Institute, Georgetown University, Washington,
DC, USA .

*Corresponding author(s). E-mail(s): gizem@utexas.edu;

Abstract

Tropical cyclones can affect health long after landfall, but the persistence and magnitude of these effects may differ across communities. This study examines whether social vulnerability modifies the long-run mortality response to tropical cyclone exposure in Texas. We combine county-month all-cause mortality, population, and temperature data for 1968–2018 with reconstructed tropical cyclone wind exposure beginning in 1953 and county vulnerability classifications based on the Social Vulnerability Index. A reduced-form county-month panel model aggregates current and lagged wind exposure using a normalized long-run response profile and estimates separate effects for high- and lower-vulnerability counties. The specification accounts for county-specific seasonality, common month-year shocks, county trends, evolving seasonal patterns, and county-specific temperature effects. The estimated mortality response is positive and statistically significant among high-vulnerability counties, whereas the estimate for lower-vulnerability counties is smaller and statistically indistinguishable from zero. Future-exposure placebos and exposure-randomization tests support the interpretation that the high-vulnerability estimate is not generated mechanically by the model or exposure distribution. Translating the estimates into event-level and spatial burdens shows that long-run tropical-cyclone mortality is concentrated where repeated physical exposure overlaps with high social vulnerability. These findings indicate that disaster-mortality assessment and resilience planning should consider both hazard intensity and communities' capacity to absorb and recover from shocks.

Keywords: Tropical cyclones, Long-run mortality, Social vulnerability, Disaster risk, Public health, Reduced-form panel model

Introduction

Tropical cyclones (TCs) are typically evaluated through their immediate and visible consequences: wind damage, flooding, infrastructure failure, displacement, and deaths that occur during or shortly after landfall. Yet these direct impacts capture only part of the human burden of disasters. TCs can initiate longer chains of disruption that affect health through damaged housing, reduced access to medical care, financial strain, power outages, displacement, employment losses, psychological stress, and changes in local public resources (Young and Hsiang 2024a; Anttila-Hughes and Hsiang 2013; Kishore et al. 2018). These pathways may unfold gradually, making the resulting mortality difficult to attribute to the original hazard event using conventional disaster death counts.

Recent evidence suggests that this indirect mortality burden may be large. Young and Hsiang (2024a) estimate that TCs in the contiguous United States generate elevated all-cause mortality beyond 14 years after each geophysical event, with cumulative excess mortality far exceeding official counts of direct deaths. Their reduced-form approach treats TC wind exposure as an exogenous physical shock and estimates the long-run mortality impulse response using state-level monthly mortality data. Related health studies also link TC exposure to short-run increases in county-level mortality and hospitalization, including cause-specific mortality and hospitalizations among older adults (Parks et al. 2022, 2021; Burrows et al. 2023). This work establishes that TCs can have persistent population-level health consequences, but it leaves open an important question: whether these mortality responses are distributed evenly across populations exposed to similar physical hazards.

There are strong reasons to expect unequal mortality responses. Socially vulnerable populations may have fewer financial resources to absorb property damage, weaker access to insurance and credit, higher baseline health risks, greater dependence on disrupted public services, and more limited ability to relocate, evacuate, or recover after disasters (Flanagan et al. 2011; Hallegatte et al. 2020; Markhvida et al. 2020). These factors can convert the same physical TC exposure into different health consequences. A storm of similar wind intensity may therefore produce a larger mortality response in populations where pre-existing socioeconomic, demographic, housing, transportation, and institutional vulnerabilities constrain recovery. In this sense, social vulnerability may not only predict exposure to disaster risk; it may also mediate the severity and persistence of post-disaster health impacts.

A growing literature documents that disasters impose unequal burdens across households and communities. Studies of disaster recovery and well-being losses show that asset damage alone can understate the welfare consequences borne by poorer households, because households with fewer resources may experience slower recovery and larger losses in consumption and well-being for the same physical damage (Markhvida et al. 2020; Hallegatte et al. 2020). Related work on infrastructure disruptions shows that outages and recovery processes can impose larger burdens on vulnerable households and communities, including through unequal power-restoration times and disruptions to basic household capabilities (Clark et al. 2023; Peterson

et al. 2024; Best et al. 2023). However, quantitative evidence linking social vulnerability to long-run TC-induced mortality remains limited, especially at sub-state spatial resolution.

This paper examines whether socially vulnerable populations experience larger long-run mortality responses to TC exposure. We focus on Texas, a state with repeated TC exposure, substantial spatial variation in storm intensity, and marked differences in social vulnerability across exposed populations (Strader et al. 2026; Tate et al. 2021). We combine monthly county-level all-cause mortality, county population, reconstructed TC wind exposure, temperature, and CDC/ATSDR Social Vulnerability Index (SVI) data from 1968 through 2018 (Centers for Disease Control and Prevention 2024; Flanagan et al. 2011). Using county-month panel data, we specify a reduced-form econometric model that exploits within-county variation in TC exposure over time. The specification absorbs county-specific seasonality, common month-year shocks, county-specific trends, and evolving seasonal patterns, while controlling for county-specific temperature effects.

Our empirical design builds directly on the long-run TC mortality response documented by Young and Hsiang (2024a), but adapts it to the county level to study heterogeneity by social vulnerability. Because county-level mortality rates are noisier than state-level mortality rates and TC exposure is spatially sparse, we do not estimate a separate coefficient for every post-TC lag. Instead, we construct a compact exposure index that aggregates current and lagged TC wind exposure using the normalized quadratic lag-response profile implied by the state-level evidence. We then interact this exposure index with indicators for high- and lower-SVI populations. This approach preserves the long-run distributed-lag logic while allowing us to test whether the cumulative mortality response differs by social vulnerability.

We test two related hypotheses. The first hypothesis is that TC exposure increases monthly all-cause mortality among exposed populations. The second hypothesis is that this mortality response is larger among high-SVI populations, reflecting the possibility that physically similar TC exposure produces unequal health consequences across populations with different baseline socioeconomic, demographic, housing, transportation, and institutional vulnerabilities.

This paper contributes to three strands of literature. First, it extends the emerging literature on indirect and long-run disaster mortality by examining heterogeneity in TC-induced mortality at a finer spatial scale (Young and Hsiang 2024a; Parks et al. 2022; Kishore et al. 2018). Second, it contributes to the literature on social vulnerability and disaster impacts by estimating whether vulnerability modifies the mortality response to a physical hazard, rather than only describing where vulnerable populations are located (Flanagan et al. 2011; Hallegatte et al. 2020; Han et al. 2024). Third, it provides empirical evidence relevant for equity-oriented resilience planning: if socially vulnerable populations experience larger long-run mortality responses to similar TC exposure, then disaster mitigation, recovery, and public health interventions should account not only for physical hazard intensity, but also for the population-level capacity to absorb and recover from that hazard (Markhvida et al. 2020; Clark et al. 2023; Peterson et al. 2024).

Empirical Strategy

We estimate the relationship between tropical cyclone (TC) exposure and monthly all-cause mortality using a reduced-form panel design that exploits within-county variation in TC exposure over time. The design follows a distributed-lag impulse-response logic: TC exposure is treated as an exogenous physical shock, and the mortality response is estimated conditional on high-dimensional fixed effects and time-varying controls that absorb persistent cross-county differences, common temporal shocks, seasonality, county-specific trends, and temperature-related mortality patterns.

The outcome variable is monthly all-cause mortality per 100,000 population in county c and month t , denoted by Y_{ct} . Let W_{ct} denote county-level TC wind exposure. We construct a scalar TC exposure index, X_{ct} that aggregates current and lagged wind exposure using a normalized quadratic lag-response profile:

$$X_{ct} = \sum_{\ell=0}^{172} h_{\ell} W_{c,t-\ell}. \quad (1)$$

Here, ℓ indexes months since TC exposure, $W_{c,t-\ell}$ is the county-level wind exposure observed ℓ months before month t , and h_{ℓ} is the normalized lag weight assigned to month ℓ . The weights h_{ℓ} are derived from the quadratic approximation to the post-TC state-level mortality response reported by [Young and Hsiang \(2024a\)](#). Let

$$q_{\ell} = 0.0237 + 0.000535\ell - 0.0000039\ell^2, \quad \ell = 0, \dots, 172, \quad (2)$$

denote the unnormalized monthly lag-response profile. We normalize this profile as

$$h_{\ell} = \frac{q_{\ell}}{\sum_{k=0}^{172} q_k}, \quad \ell = 0, \dots, 172.$$

This normalization preserves the temporal shape of the estimated long-run response while scaling the exposure index to represent a weighted average of post-TC wind exposure. We restrict the post-event window to 0–172 months because the estimated mortality response remains positive and statistically meaningful over this window ([Young and Hsiang 2024a](#)).

The construction of X_{ct} serves as a dimension-reduction approach. Rather than estimating a separate coefficient for each of the 173 monthly lag terms within each SVI group, which is not well identified with 42 counties and would lead to a highly parameterized and unstable model, we impose a smooth lag-response profile motivated by the state-level evidence. This approach preserves the long-run distributed-lag structure of TC mortality effects while reducing the exposure history to a single scalar index.

To test whether socially vulnerable counties experience larger TC-related mortality responses, we allow the effect of TC exposure to vary by county-level social vulnerability. Let SVI_c denote county c 's average CDC/ATSDR Social Vulnerability Index (SVI) ([Flanagan et al. 2011](#); [Centers for Disease Control and Prevention 2024](#)). The SVI measure is based on national percentile rankings and is not renormalized within

Texas. We define the county group $g(c) \in \{L, H\}$ as

$$g(c) = \begin{cases} L, & \text{if } \text{SVI}_c \leq 0.75, \\ H, & \text{if } \text{SVI}_c > 0.75. \end{cases}$$

Thus, each county belongs to exactly one of the two groups: lower-SVI or high-SVI. We use the term “lower-SVI” rather than “low-SVI” because the reference group includes all counties with $\text{SVI}_c \leq 0.75$, including counties with moderate vulnerability. The estimating equation is

$$Y_{ct} = \beta_{g(c)} X_{ct} + \Gamma_{ct} + \varepsilon_{ct}, \quad (3)$$

where Γ_{ct} collects the fixed effects and time-varying controls, and ε_{ct} is the remaining county-month error term. Equivalently, this specification interacts X_{ct} with mutually exclusive indicators for lower-SVI and high-SVI counties:

$$\beta_{g(c)} = \begin{cases} \beta_L, & \text{if } g(c) = L, \\ \beta_H, & \text{if } g(c) = H. \end{cases}$$

Therefore, β_L measures the cumulative mortality-rate response to a one-unit increase in the TC exposure index among counties with $\text{SVI}_c \leq 0.75$, while β_H measures the corresponding response among counties with $\text{SVI}_c > 0.75$. The difference

$$\beta_\Delta = \beta_H - \beta_L$$

captures heterogeneity in the TC mortality response by social vulnerability.

Our empirical analysis tests two related hypotheses. First, TC exposure increases monthly all-cause mortality, implying a positive mortality response to X_{ct} . Second, this response is larger in high-SVI counties than in lower-SVI counties. The primary heterogeneity test is therefore whether $\beta_\Delta > 0$.

The fixed-effect and control component is

$$\Gamma_{ct} = \mu_{c,m(t)} + \theta_c t + \kappa_{c,m(t)} t + \delta_{1c} T_{ct} + \delta_{2c} T_{ct}^2 + \lambda_t. \quad (4)$$

Here, t denotes the centered month-year index, and $m(t)$ denotes the calendar month. The term $\mu_{c,m(t)}$ denotes county-by-calendar-month fixed effects, which absorb persistent county-specific seasonality. The term $\theta_c t$ denotes county-specific linear time trends, allowing each county to follow its own gradual mortality trajectory. The term $\kappa_{c,m(t)} t$ denotes county-by-calendar-month-specific linear trends, allowing seasonal mortality patterns to evolve differently across counties over time. The variables T_{ct} and T_{ct}^2 denote monthly average temperature and its square, with coefficients allowed to vary by county. Finally, λ_t denotes month-year fixed effects, which absorb shocks common to all counties in a given month-year.

We select this specification after comparing alternative fixed-effect structures, polynomial trend controls, temperature specifications, and exposure definitions. The current model provides the best balance between flexibility and parsimony: it absorbs

major sources of confounding variation while retaining sufficient within-county exposure variation to estimate heterogeneous TC mortality responses across SVI groups. The specification also passes the validation exercises, including placebo leads and exposure randomization tests.

Relative to state-level designs, county-level mortality rates are observed for smaller populations and TC exposure is more spatially sparse, producing a noisier signal-to-noise environment. Estimating a highly flexible distributed-lag model or very high-order county-specific time trends in this setting would substantially increase the number of parameters and risk absorbing the limited within-county variation in TC exposure that identifies the mortality response. We therefore use the state-level impulse-response evidence to discipline the lag structure and estimate a lower-dimensional county-level exposure index.

We estimate the model parameters using population-weighted least squares, with county population specified as analytic weights. This weighting gives greater influence to county-month observations representing larger exposed populations and yields coefficients interpretable as population-weighted mortality-rate responses rather than unweighted average county responses. We absorb the high-dimensional fixed effects in Equation (4) using `reghdfe` in Stata. Standard errors are clustered at the county level to allow arbitrary heteroskedasticity and serial correlation within counties over time.

Data and Sample Analysis

Our analysis combines county-month mortality, population, tropical cyclone (TC) wind exposure, temperature, and social vulnerability data for Texas counties. The mortality panel covers January 1968 through October 2018. Monthly all-cause mortality is used to construct the outcome variable, Y_{ct} , defined as deaths per 100,000 population in county c and month t . County population is used both to construct mortality rates and as analytic weights in the regression analysis.

Mortality and population data are assembled from two sources. For 1968–1998, county-level mortality records are constructed from the National Vital Statistics System multiple-cause mortality files, distributed through the [National Bureau of Economic Research \(2026\)](#). For 1999–2018, county-level mortality and population data are obtained from the CDC WONDER Multiple Cause of Death files ([Centers for Disease Control and Prevention 2025](#)). Both sources identify deaths by county of residence, which ensures consistency in the geographic definition of mortality throughout the sample. Population data are harmonized with the mortality records and used to convert death counts into mortality rates.

TC wind exposure is derived from the LICRICE tropical cyclone wind-field reconstruction used by Young and Hsiang ([Young and Hsiang 2024a,b](#)). LICRICE provides storm-specific maximum sustained wind speeds, measured in meters per second (m/s), on a $0.125^\circ \times 0.125^\circ$ grid for Atlantic tropical cyclones. We aggregate these pixel-level wind fields to the county-storm-month level by averaging wind speeds across pixels located within a 5 km buffer around each county boundary. This produces county-month measures of maximum TC wind exposure, denoted W_{ct} , measured in m/s. The TC exposure data begin in September 1953 and extend through October 2018.

The earlier start date relative to the mortality panel is necessary because the exposure index X_{ct} uses a 172-month distributed-lag window. Thus, for the first mortality observation in January 1968, the full lagged exposure history from September 1953 through January 1968 is available.

Temperature data are obtained from Berkeley Earth Surface Temperatures (BEST) ([Berkeley Earth n.d.](#)). We use monthly gridded land-surface air temperature data at the pixel level on a $1^\circ \times 1^\circ$ latitude-longitude grid. Temperature is measured in degrees Celsius ($^\circ\text{C}$) and reported as anomalies relative to the January 1951–December 1980 average. We use the climatological mean temperature dataset to convert anomalies to absolute temperatures. Pixel-level temperature data are then spatially aggregated to the county-month level by averaging absolute temperatures across pixels located within a 5 km buffer around each county boundary. Although the BEST temperature grid is coarser than the LICRICE wind grid, both variables are harmonized to the same county-month panel before estimation. Temperature enters the regression as a time-varying control to account for temperature-related mortality risk.

Social vulnerability is measured using the overall county-level national percentile rankings from the 2000, 2010, and 2018 CDC/ATSDR Social Vulnerability Index databases ([Centers for Disease Control and Prevention 2024](#)). The 2000 ranking is assigned to county-month observations from 1953 through 2000, the 2010 ranking to observations from 2001 through 2010, and the 2018 ranking to observations from 2011 through 2018. For each county, we average these assigned monthly rankings over the full 1953–2018 exposure panel to construct a time-invariant vulnerability measure. This procedure produces a duration-weighted average of the three index vintages. The resulting national percentile ranking is not renormalized within Texas. Counties with an average ranking above 0.75 are classified as high vulnerability, and the remaining counties are classified as lower vulnerability.

County boundaries are standardized to 2018 county boundaries before merging mortality, population, wind exposure, temperature, and SVI data. County boundary files are obtained from the U.S. Census Bureau TIGER/Line shapefiles ([U.S. Census Bureau 2018](#)) and the Atlas of Historical County Boundaries ([Newberry Library n.d.](#)). This harmonization ensures that all variables are measured using a consistent county geography.

The main estimation sample is restricted to counties with sufficient population, mortality counts, and TC exposure to support stable county-level inference. Specifically, we retain counties with population above 10,000 and exposure to at least 15 TC events for which county-average wind exposure exceeds 17.5 m/s, approximately the lower bound for tropical-storm-force winds. This filtering rule yields a sample of 42 out of 254 Texas counties. The restriction is motivated by the statistical limitations of county-level mortality analysis: counties with small populations and rare or no TC exposure generate highly noisy mortality rates and weak within-county exposure variation, which produce unstable estimates and large standard errors. Because the objective of this study is to evaluate whether more socially vulnerable populations exhibit stronger mortality responses to TC exposure, the relevant empirical

Tropical Cyclone Exposure Counts in Texas Counties
 Black outlines indicate the 42 counties used in the estimation sample




Fig. 1 Tropical cyclone exposure counts across Texas counties. County shading reports the number of TC events for which county-average wind exposure exceeds 17.5 m/s, approximately the lower bound for tropical-storm-force winds. Black outlines indicate the 42 counties included in the estimation sample.

sample is not all Texas counties, but counties with enough TC exposure and mortality information to identify such responses.

Figure 1 shows the spatial distribution of TC exposure across all Texas counties and highlights the 42 counties retained in the estimation sample. County shading reports the number of TC events for which county-average wind exposure exceeds 17.5 m/s, while black outlines identify the counties used in the regression analysis. The figure shows that the estimation sample is concentrated in the coastal parts of Texas, where repeated TC exposure is most prevalent, while still retaining variation in exposure intensity across the selected counties.

The panel contains 25,620 county-month observations from 42 counties over 610 months. Under this definition, the estimation sample includes 13 lower-SVI counties and 29 high-SVI counties. Table 1 reports a compact summary of the SVI composition and TC exposure characteristics of the estimation sample. We do not report all SVI quartile-level details in the main text. However, no county in the estimation sample falls in the lowest national SVI quartile. Therefore, the empirical contrast should be

Table 1 SVI composition and TC exposure characteristics of the 42-county estimation sample.

SVI group	Counties	SVI		Exposed events		Max wind	
		Mean	Range	Mean	Range	Mean	Range
Lower-SVI, $\text{SVI} \leq 0.75$	13	0.576	0.391–0.732	23.46	16–30	40.09	31.21–47.51
High-SVI, $\text{SVI} > 0.75$	29	0.887	0.757–0.990	20.97	16–30	42.73	31.70–59.28

Notes: SVI values are based on national CDC/ATSDR percentile rankings and are not renormalized within Texas. Exposed events denote the number of TC events for which county-average wind exposure exceeds 17.5 m/s. Max wind is each county’s maximum observed TC wind exposure during the sample period; the table reports the mean and range of this county-level maximum within each SVI group. Wind speed is measured in meters per second.

interpreted as a comparison between high-SVI counties and relatively lower-SVI counties within the TC-exposed estimation sample, rather than as a comparison between nationally low-vulnerability and high-vulnerability counties.

Results

Estimated Mortality Effects by Social Vulnerability

Table 2 reports the estimated TC mortality coefficients, $\hat{\beta}_g$, separately for lower- and high-SVI counties, along with standard errors, confidence intervals, and one- and two-sided p-values. Figure 2 translates these coefficient estimates into the implied monthly and cumulative mortality-rate responses. For group $g \in \{L, H\}$, the implied monthly response at lag ℓ is

$$\hat{\omega}_{g,\ell} = \hat{\beta}_g h_\ell, \quad \ell = 0, \dots, 172,$$

and the cumulative response through month τ is

$$\hat{\Omega}_{g,\tau} = \sum_{\ell=0}^{\tau} \hat{\beta}_g h_\ell = \hat{\beta}_g \sum_{\ell=0}^{\tau} h_\ell, \quad \tau = 0, \dots, 172.$$

Because the lag weights are normalized so that $\sum_{\ell=0}^{172} h_\ell = 1$, the cumulative response at month 172 equals the estimated group-specific coefficient, $\hat{\Omega}_{g,172} = \hat{\beta}_g$. We also report the heterogeneity contrast $\beta_\Delta = \beta_H - \beta_L$, which measures whether the TC mortality response is larger in high-SVI counties than in lower-SVI counties. For each coefficient, we test whether the estimate differs from zero and, given our directional hypothesis, whether it is positive.

The estimates indicate that the positive TC mortality response is concentrated among high-SVI counties. As shown in Table 2, the high-SVI coefficient is positive and statistically significant under both the two-sided and one-sided tests ($\hat{\beta}_H = 2.133$, $p_H^{2s} = 0.008$, $p_H^+ = 0.004$). In contrast, the lower-SVI coefficient is small and statistically indistinguishable from zero ($\hat{\beta}_L = 0.415$, $p_L^{2s} = 0.745$, $p_L^+ = 0.372$). The joint Wald test rejects the null hypothesis that both SVI-specific exposure coefficients are equal to zero ($H_0 : \beta_L = \beta_H = 0$; $F(2, 41) = 5.97$, $p = 0.0053$), indicating that TC

Table 2 Estimated TC mortality effects.

Group	Estimate $\hat{\beta}_g$	Std. error $SE(\hat{\beta}_g)$	95% CI	Two-sided p p_g^{2s}	One-sided p p_g^+
Lower-SVI, $g = L$	0.415	1.266	[-2.142, 2.971]	0.745	0.372
High-SVI, $g = H$	2.133	0.770	[0.577, 3.688]	0.008	0.004
High-lower difference, $g = \Delta$	1.718	0.941	[-0.183, 3.618]	0.075	0.038

Notes: The row $g = \Delta$ reports the heterogeneity contrast $\beta_\Delta = \beta_H - \beta_L$. Two-sided p -values test $H_0 : \beta_g = 0$ against $H_A : \beta_g \neq 0$, while one-sided p -values test $H_0 : \beta_g \leq 0$ against $H_A : \beta_g > 0$. Confidence intervals and standard errors are based on county-clustered standard errors.




Fig. 2 Estimated TC mortality-rate responses by SVI group. Panels (a) and (b) show the implied monthly responses, $\hat{\omega}_{g,\ell} = \hat{\beta}_g h_\ell$, for lower- and high-SVI counties, respectively; panel (c) shows the cumulative responses, $\hat{\Omega}_{g,\tau} = \sum_{\ell=0}^{\tau} \hat{\beta}_g h_\ell$; and panel (d) shows the cumulative heterogeneity contrast, $\hat{\Omega}_{\Delta,\tau} = \sum_{\ell=0}^{\tau} \hat{\beta}_\Delta h_\ell$, where $\hat{\beta}_\Delta = \hat{\beta}_H - \hat{\beta}_L$. Negative months report analogous future-placebo lead responses, used as negative exposure controls and discussed in Section 2. Shaded regions denote 95% confidence intervals.

exposure is associated with mortality in at least one SVI group. The group-specific estimates show that this rejection is driven by the high-SVI response.

Figure 2 shows how these coefficient estimates translate into monthly and cumulative mortality-rate responses. The lower-SVI response is plotted for comparison, but its 95% confidence interval includes zero throughout the cumulative response window, consistent with the statistically insignificant coefficient estimate in Table 2. We therefore interpret the lower-SVI estimate as statistically indistinguishable from zero, rather than as evidence of a precisely estimated null effect. In contrast, the high-SVI

Table 3 Model fit and residual diagnostic summary.

Diagnostic	Value
Adjusted R^2	0.836
Residual mean	0.127
Residual standard deviation	16.064
Share of $ \text{standardized residual} > 3$	1.56%

response is positive throughout the post-exposure window and accumulates gradually over time, reaching the cumulative coefficient estimate reported in the table by month 172.

The heterogeneity results provide directional evidence that socially vulnerable counties experience larger TC-related mortality effects. The estimated high–lower contrast is positive ($\hat{\beta}_\Delta = 1.718$), and the one-sided test rejects $H_0 : \beta_\Delta \leq 0$ in favor of $H_A : \beta_\Delta > 0$ ($p_\Delta^+ = 0.038$). The corresponding two-sided test of $H_0 : \beta_\Delta = 0$ is marginally significant at the 10% level ($p_\Delta^{2s} = 0.075$), so the heterogeneity evidence is stronger under the prespecified directional alternative than under a two-sided alternative. This pattern is also visible in Figure 2: the cumulative high-minus-lower contrast remains positive over the post-exposure window and increases gradually as the lag-response weights accumulate.

Taken together, these results provide evidence that TC exposure is associated with elevated mortality in affected counties, and that the increase is concentrated among, and statistically larger for, high-SVI counties.

Model Fit and Residual Diagnostics

We assess the fit of the preferred specification using residual diagnostics that examine overall predictive fit, remaining temporal structure, exposure-related residual patterns, and event-time behavior around high-exposure months. The model explains a large share of the variation in monthly county mortality rates, with an adjusted R^2 of 0.84 based on 25,620 county-month observations from 42 counties.

Table 3 summarizes the main diagnostic statistics. The residual mean is close to zero relative to the residual standard deviation ($\bar{\varepsilon} = 0.127$, $SD = 16.064$), indicating that the model does not exhibit a large average residual bias. The standardized residuals have mean zero and unit variance by construction, although 399 observations, or 1.56% of the sample, have absolute standardized residuals greater than three. This indicates heavier tails than a Gaussian benchmark, which is not unexpected in county-month mortality data. We therefore do not rely on homoskedastic normal-error assumptions for inference; throughout the analysis, standard errors are clustered at the county level, and the main findings are further evaluated using future-exposure placebo and exposure-randomization tests.

Figure 3 provides graphical diagnostics. Panel (a) plots standardized residuals against fitted values. The residuals are centered around zero across the fitted-value range, with no strong systematic mean pattern, although the residual spread is not




Fig. 3 Residual diagnostics for the preferred model. Panel (a) plots standardized residuals against fitted values. Panel (b) shows the aggregate monthly mean residual across counties. Panel (c) reports the autocorrelation function of the aggregate monthly mean residual series. Panel (d) plots mean residuals in event time around high-exposure county-months, defined as months with county-level TC wind exposure of at least 25. Error bars denote approximate 95% confidence intervals.

constant. This heteroskedasticity reinforces the use of county-clustered standard errors rather than classical homoskedastic standard errors. Panel (b) shows the aggregate monthly mean residual across counties. The series fluctuates around zero without an obvious long-run drift, suggesting that the fixed effects and trend controls absorb the main aggregate mortality movements. We confirm this pattern using an augmented Dickey–Fuller test, which rejects the null hypothesis that the aggregate residual series contains a unit root (ADF statistic = -11.757 , $p \approx 0$, 1% critical value = -3.44). Since a unit root would imply nonstationary residual behavior and persistent drift over time, this result supports the interpretation that the aggregate residual series is stationary around zero.

Panel (c) reports the autocorrelation function of the aggregate monthly mean residuals. The autocorrelations are small at most lags and decay quickly, indicating limited remaining aggregate serial dependence. Within-county residual autocorrelation is somewhat larger at short lags but remains modest on average: the mean lag-1 autocorrelation is 0.105 and the median is 0.081. This suggests that some local persistence remains in county-level mortality shocks, but it is not large enough to indicate a major unmodeled time-series pattern.

Panel (d) examines residual behavior in event time around 291 high-exposure events, defined as county-months with county-level TC wind exposure of at least 25m/s after retaining only events within the same county that are separated by at least 12 months. If the model were systematically missing pre-existing mortality increases or post-event residual drift around high-exposure months, the event-time residual profile would show a clear pattern before or after exposure. Instead, the average residuals fluctuate around zero, and the uncertainty intervals generally overlap zero across the event-time window. This supports the interpretation that the estimated TC mortality response is not driven by an obvious residual pre-trend or omitted event-time pattern.

Additional diagnostics, not shown for brevity, are consistent with the main residual checks. Binned residual plots against the compact TC exposure index X_{ct} and calendar-month residual summaries show mean residuals close to zero, with no monotone exposure-related pattern or remaining seasonal pattern. However, the residual distribution is heavy-tailed, and the Q-Q plot shows clear departures from normality in the tails. Thus, the diagnostics support model adequacy in the sense of limited systematic residual structure, but not a Gaussian error process. Accordingly, inference does not rely on classical homoskedastic normal-error assumptions; it uses county-clustered standard errors, which are robust to heteroskedasticity and allow for within-county dependence, and is further evaluated using the future-exposure placebo and exposure-randomization tests in the following sections.

Future Exposure Placebo Test

As a falsification test, we examine whether future TC exposure predicts current mortality. Because future storms cannot causally affect current mortality, a statistically significant positive future-exposure coefficient would raise concerns about residual confounding, anticipatory trends, or model misspecification.

We construct the future-exposure placebo index as

$$X_{ct}^F = \sum_{f=1}^{12} h_f^F W_{c,t+f},$$

where $W_{c,t+f}$ denotes TC wind exposure in county c , f months after month t . The weights h_f^F use the same quadratic lag profile q_f defined in Equation (2), renormalized over the 12-month future window:

$$h_f^F = \frac{q_f}{\sum_{r=1}^{12} q_r}, \quad f = 1, \dots, 12.$$

Thus, X_{ct}^F is a normalized weighted average of TC wind exposure occurring in the 12 months after the mortality observation. We use this near-term window as a negative exposure control: future exposure should not affect current mortality, but it can reveal short-run pre-trends or residual confounding. The estimating equation is

$$Y_{ct} = \pi_{g(c)} X_{ct}^F + \beta_{g(c)} X_{ct} + \Gamma_{ct} + \varepsilon_{ct},$$

Table 4 Future-lead placebo tests.

SVI group	Estimate $\hat{\pi}_g$	95% CI	Two-sided p $p_g^{2s,f}$	One-sided p $p_g^{+,f}$
Lower-SVI, $g = L$	0.056	[-0.173, 0.285]	0.625	0.312
High-SVI, $g = H$	0.050	[-0.199, 0.299]	0.686	0.343

Notes: The coefficients π_g correspond to future TC exposure. Since future storms cannot causally affect current mortality, the placebo test supports the identifying interpretation when the future-exposure coefficients are small and statistically insignificant. Two-sided p -values test $H_0 : \pi_g = 0$ against $H_A : \pi_g \neq 0$, while one-sided p -values test $H_0 : \pi_g \leq 0$ against $H_A : \pi_g > 0$. Confidence intervals are based on county-clustered standard errors.

where X_{ct} is the actual post-TC exposure index defined in Equation (1), and Γ_{ct} denotes the same fixed effects and controls as in Equation (4). The coefficients π_g test whether future TC exposure is associated with current mortality for group $g \in \{L, H\}$.

For each group, we test the two-sided null hypothesis $H_0 : \pi_g = 0$ against $H_A : \pi_g \neq 0$. We also report the one-sided test $H_0 : \pi_g \leq 0$ against $H_A : \pi_g > 0$, since a significantly positive future-exposure coefficient would be most problematic for the causal interpretation of a positive post-TC mortality response.

Table 4 shows that the future-exposure coefficients are small and statistically indistinguishable from zero for both groups. The lower-SVI placebo coefficient is $\hat{\pi}_L = 0.056$, with a two-sided p -value of 0.625 and a one-sided positive p -value of 0.312. The high-SVI placebo coefficient is similarly small, $\hat{\pi}_H = 0.050$, with a two-sided p -value of 0.686 and a one-sided positive p -value of 0.343. The joint placebo test also fails to reject the null that both future-exposure coefficients are zero ($F(2, 41) = 0.12$, $p = 0.885$).

The graphical evidence in Figure 2 is consistent with these tests, but it should be interpreted together with the statistical uncertainty. In panel (a), both the lower-SVI actual response and the lower-SVI future-placebo response are close to zero relative to their confidence intervals, and neither is statistically distinguishable from zero. Thus, the lower-SVI panel does not provide evidence of either a post-exposure mortality response or a problematic future-exposure placebo pattern. In panel (b), the contrast is clearer: the high-SVI actual response is positive over the post-exposure window, whereas the high-SVI future-placebo response at negative months remains centered near zero and is statistically insignificant. Panels (c) and (d) show the same pattern in cumulative form. The cumulative placebo responses remain close to zero, while the actual high-SVI response and the actual high-minus-lower heterogeneity contrast accumulate steadily over the post-exposure window.

Overall, the future-exposure placebo test provides no evidence that future TC exposure predicts current mortality. This supports the interpretation that the estimated post-exposure mortality response, particularly among high-SVI counties, is not driven by anticipatory patterns or residual short-run pre-trends correlated with subsequent TC exposure.

Exposure Randomization Tests

We next conduct exposure randomization tests to assess whether the estimated TC mortality response could arise mechanically from the model specification, the marginal distribution of wind exposure, county-level exposure propensity, or temporal clustering of storms. These tests use the same county-month panel, mortality outcome, population weights, SVI group assignments, and fixed-effect structure as the main specification. The only object randomized is county-level TC wind exposure.

For each randomized exposure panel, we reconstruct the compact lagged exposure index using the same lag-weighting function as in the main model. Specifically, for county c and month t , the observed exposure index is

$$X_{ct} = \sum_{\ell=0}^{172} h_{\ell} W_{c,t-\ell},$$

where $W_{c,t-\ell}$ is county-level TC wind exposure and h_{ℓ} is the normalized lag weight implied by the quadratic lag-response profile. In each placebo experiment, we replace W_{ct} with a randomized exposure series W_{ct}^r , reconstruct

$$X_{ct}^r = \sum_{\ell=0}^{172} h_{\ell} W_{c,t-\ell}^r,$$

and re-estimate the same weighted fixed-effects regression:

$$Y_{ct} = \beta_{g(c)}^r X_{ct}^r + \Gamma_{ct} + \varepsilon_{ct}. \quad (5)$$

Randomized TC exposure should not systematically predict mortality. If the observed coefficient is not unusual relative to the randomized coefficient distribution, the estimated TC mortality response may be consistent with noise or model-induced artifacts. Conversely, if the observed coefficient is extreme relative to the randomized distribution, this provides evidence that the estimated response is tied to the actual TC exposure history rather than to random exposure assignment.

We implement four county-level randomization designs, summarized in Table 5. Each design preserves a different feature of the exposure process while breaking the link between actual county-level TC exposure and mortality. For each design, we generate $R = 1,000$ randomized exposure panels, reconstruct the lagged exposure index, and re-estimate Equation (5). Empirical p-values are computed as

$$p_g^e = \frac{1 + \sum_{r=1}^R \mathbf{1}\{T_{g,r}^r \geq T_g^o\}}{R + 1}, \quad (6)$$

where T_g^o is the observed statistic from the actual exposure model and $T_{g,r}^r$ is the corresponding statistic from randomized replication r . For two-sided coefficient tests, we use $T_g^o = |\hat{\beta}_g|$ and $T_{g,r}^r = |\hat{\beta}_{g,r}^r|$, allowing randomized coefficients to be counted

Table 5 County-level exposure randomization designs.

Randomization design	Implementation	Structure preserved	Purpose of test
Global exposure shuffle	For each county-month, draw wind exposure from a randomly selected county-month in the full panel.	Preserves only the unconditional distribution of TC wind exposure.	Tests whether the marginal distribution of TC exposure alone can mechanically generate a mortality signal.
Within-county time shuffle	For each county-month, draw wind exposure from another month within the same county.	Preserves each county's overall exposure distribution but destroys exposure timing.	Tests whether county-level exposure propensity, rather than the timing of TC exposure, drives the estimated mortality response.
Within-month county shuffle	For each county-month, draw wind exposure from another county observed in the same month.	Preserves common storm-month timing but randomizes county-level exposure assignment within each month.	Tests whether the spatial allocation of exposure intensity within affected months drives the result.
Cross-county history reassignment	For each county, assign the complete wind-exposure time series of a randomly selected donor county.	Preserves the donor county's temporal storm sequence but assigns it to a different county.	Tests whether temporal clustering or serial structure in TC exposure generates spurious effects when disconnected from the correct county mortality series.

Notes: In all placebo designs, mortality outcomes, population weights, SVI group assignments, and fixed effects are held fixed; only county-level TC wind exposure is randomized. For each randomized exposure panel, we reconstruct the compact lagged exposure measure using the same lag weights and re-estimate the same regression specification.




Fig. 4 Implied cumulative mortality responses from the actual, future-placebo, and randomized exposure estimates. The dashed lines plot the cumulative responses implied by the mean randomized coefficients reported in Table 6, computed as $\hat{\Omega}_{g,\tau}^r = \hat{\beta}_g^r \sum_{\ell=0}^{\tau} h_{\ell}$, using the same normalized lag weights as in the main specification.

as extreme in either direction. For one-sided tests, we use $T_g^o = \hat{\beta}_g$ and $T_{g,r}^r = \hat{\beta}_{g,r}^r$, corresponding to the directional alternative that TC exposure increases mortality.

Table 6 and Figure 4 show that the randomization tests are consistent with the main exposure model: the lower-SVI estimate is not unusual relative to randomized

Table 6 Randomization placebo tests.

Randomization design	SVI group	Observed estimate $\hat{\beta}_g$	Mean randomized estimate $\hat{\beta}_g^r$	Two-sided empirical p $p_g^{2s,e}$	One-sided empirical p $p_g^{+,e}$
Global exposure shuffle	Lower-SVI	0.415	-0.016	0.766	0.378
	High-SVI	2.133	-0.014	0.013	0.005
Within-county time shuffle	Lower-SVI	0.415	-0.013	0.738	0.369
	High-SVI	2.133	-0.020	0.008	0.005
Within-month county shuffle	Lower-SVI	0.415	-1.047	0.810	0.149
	High-SVI	2.133	0.456	0.121	0.102
Cross-county history reassignment	Lower-SVI	0.415	-0.904	0.765	0.106
	High-SVI	2.133	0.273	0.020	0.020
Pooled across all designs	Lower-SVI	0.415	-0.495	0.770	0.250
	High-SVI	2.133	0.174	0.040	0.032

Notes: The observed estimate $\hat{\beta}_g$ is estimated using the actual TC exposure series. The mean randomized estimate $\hat{\beta}_g^r$ is the average coefficient across randomized exposure replications. Each design-specific row is based on 1,000 randomized replications; the pooled row combines 4,000 replications across all four randomization designs. Empirical p -values are computed using Equation (6), with $\hat{\beta}_g$ as the observed statistic and $\hat{\beta}_{g,r}^r$ as its randomized counterpart in replication r . The pooled row is reported as a summary measure; the design-specific rows provide the primary placebo evidence.

exposure estimates, whereas the high-SVI estimate is large relative to most randomized exposure distributions. The figure provides the cumulative-response analogue of Table 6, showing that the mean randomized responses remain close to zero or below the observed responses once the mean randomized estimates are mapped through the normalized lag profile.

Since the actual lower-SVI coefficient is small and statistically indistinguishable from zero, the relevant question is whether the model nevertheless produces a spuriously large lower-SVI coefficient under randomized exposure. It does not: the lower-SVI coefficient is not extreme under any randomization design, with two-sided empirical p -values ranging from 0.738 to 0.810 and a pooled value of 0.770. This pattern is consistent with the main regression result that the lower-SVI response is statistically indistinguishable from zero.

In contrast, the high-SVI coefficient is unusually large relative to the randomized coefficient distributions in most designs. The one-sided empirical p -value is 0.020 under cross-county history reassignment, 0.005 under the global exposure shuffle, and 0.005 under the within-county time shuffle. The pooled one-sided empirical p -value is 0.032, indicating that the observed high-SVI coefficient is large relative to the combined placebo distribution across all four randomization designs. The two-sided empirical p -values for the high-SVI coefficient show a similar pattern.

The within-month county shuffle provides weaker evidence for the high-SVI coefficient ($p_H^{+,e} = 0.102$). This weaker result is expected in our single-state county-level setting. Unlike a national or multi-state randomization, this design preserves the timing of storm months within Texas and only reallocates wind intensity across counties observed in the same month. Because the estimation sample consists of repeatedly

exposed counties, many counties experience the same TC events in the same months, although with different wind intensities. For example, a given storm month may expose several counties simultaneously, with one county experiencing stronger winds than another. Thus, the within-month shuffle does not fully remove the common storm-month signal; it only tests whether the observed high-SVI response can be reproduced by randomly reallocating wind intensity across counties within those same TC-affected months. The weaker placebo result therefore suggests that the high-SVI estimate may be partly sensitive to common storm-month timing, but it does not imply that the estimate is generated mechanically by random county-level intensity assignment.

Overall, the exposure randomization tests support the interpretation that the observed mortality response among high-SVI counties reflects a TC-related signal rather than a spurious association arising solely from the marginal wind distribution, county-level exposure propensity, temporal clustering of storms, or model-induced noise. The within-month county shuffle is less sharp in this setting because it preserves common storm-month timing among repeatedly exposed Texas counties; therefore, we place more weight on the randomization designs that more fully break the link between actual county-level exposure histories and mortality.

Mortality Burden

Figure 5 translates the estimated SVI-group mortality responses into an event-level mortality burden for the 42-county Texas estimation sample. The calculation uses the event-county-month wind exposure W_{ect} , the normalized lag profile h_ℓ , and the estimated SVI-group coefficients from Equation (3). Over 1953–2018, 104 TCs affected Texas in the wind-exposure data. Because the 42 counties in the estimation sample account for approximately 23% of Texas’s population over the study period, the figure should be interpreted as the estimated burden for the exposed estimation sample rather than for all Texas counties.

Panel (a) shows monthly TC treatment incidence, measured as summed event-county wind exposure across the estimation sample. This panel captures the timing and intensity of TC exposure, not mortality. Panel (b) shows the fitted TC-induced mortality burden implied by the empirical model. For each storm, county-level wind exposure is distributed over the 172-month post-exposure window using the lag weights h_ℓ , multiplied by the relevant SVI-group coefficient, and converted from a mortality-rate response into excess deaths using county population.

The stacked responses in panel (b) allow us to read the estimated TC-induced mortality burden in any month and identify the storms contributing to that burden. Individual storms do not appear as one-month mortality spikes; instead, each event contributes gradually over the post-exposure window. Consequently, the monthly burden reflects the superposition of multiple active storm responses. The upper envelope rises after the early 2000s, when repeated TC exposure and persistent lagged responses generate a larger cumulative monthly burden. Because the TC exposure record in our study begins in 1953, the burden series accumulates responses only from TCs observed after 1953 and does not include any continuing mortality responses from pre-1953 TCs. Storm labels indicate the largest Texas county-level wind-exposure events.




Fig. 5 Estimated TC excess mortality burden in Texas for the 42-county estimation sample. Panel (a) shows monthly TC wind exposure summed across event-county observations. Colors indicate the decade in which each TC occurred. Panel (b) shows stacked fitted monthly excess-death responses by tropical cyclone. The black upper envelope gives the total estimated monthly TC-induced mortality across all storms with active lagged responses.

Figure 6 further summarizes the relationship between storm-level wind exposure and the fitted long-run mortality burden. Each point represents one TC. The horizontal axis reports total storm exposure across the estimation sample, while the vertical axis reports the full fitted 172-month TC-induced mortality burden. The positive association indicates that storms with larger summed county-level wind exposure generally produce larger fitted long-run mortality burdens, although the vertical dispersion also reflects differences in exposed population size and SVI composition across affected counties.

Table 7 reports official direct-death counts for billion-dollar tropical cyclone disasters affecting Texas and compares them with the fitted long-run excess mortality estimates from our model. These measures are conceptually distinct: official direct deaths measure deaths immediately attributed to the storm, while our estimates measure long-run all-cause excess deaths over 172 months in the 42-county estimation sample. Hurricane Harvey appears at the top of Figure 6, with the largest estimated long-run mortality burden in our study sample. Harvey also has one of the largest official direct-death counts among the listed Texas events. More broadly, the comparison illustrates that long-run excess mortality can substantially exceed the immediate direct-death counts used in conventional disaster reporting.

Figure 7 complements the temporal burden plot by showing the spatial distribution of the estimated TC mortality burden. All Texas counties are shown for geographic context, but only the 42 counties used in the estimation are colored. Panel (a) reports the cumulative fitted TC excess mortality burden per 100,000 population over




Fig. 6 TC wind exposure and estimated total TC-induced mortality in 42-county sample. Each point represents one tropical cyclone. Colors indicate the decade in which each TC occurred.

Table 7 Official direct deaths and estimated long-run TC-induced excess deaths for billion-dollar tropical cyclone disasters affecting Texas between 1980 and 2018 ([NOAA National Centers for Environmental Information \(NCEI\) 2025](#)).

Tropical cyclone	Official direct deaths	Estimated long-run excess deaths
Hurricane Allen (August 1980)	13	944
Hurricane Alicia (August 1983)	21	3,501
Tropical Storm Allison (June 1989)	11	2,177
Tropical Storm Frances (September 1998)	2	2,307
Tropical Storm Allison (June 2001)	43	2,120
Hurricane Rita (September 2005)	119	3,385
Hurricane Dolly (July 2008)	3	1,273
Hurricane Ike (September 2008)	112	4,480
Hurricane Harvey (August 2017)	89	4,782

Notes: Official direct deaths refer to deaths immediately attributed to each storm in disaster reporting. Estimated long-run excess deaths are fitted all-cause excess deaths over the 172-month post-exposure window for the 42-county estimation sample. These two columns are conceptually distinct and should not be interpreted as alternative estimates of the same mortality quantity.

Cumulative mortality burden across Texas counties, 1968–2018




Fig. 7 Spatial distribution of cumulative mortality burden in the 42-county Texas estimation sample. Panel (a) shows cumulative fitted TC excess mortality burden per 100,000 population over 1968–2018. Panel (b) shows cumulative observed all-cause mortality per 100,000 population over the same period. All Texas counties are shown for geographic context, but only the 42 estimation-sample counties are colored. Harris County is highlighted in blue.

1968–2018. Panel (b) reports cumulative observed all-cause mortality per 100,000 population over the same period. Thus, panel (a) isolates the TC-attributable component implied by the empirical model, whereas panel (b) reflects overall mortality conditions accumulated across the full mortality sample.

Across the 42 counties, the mean cumulative estimated TC excess mortality burden is 1,136.60 deaths per 100,000, compared with a mean cumulative observed all-cause mortality rate of 46,709.05 deaths per 100,000. Harris County has an estimated TC excess mortality burden of 1,711.49 deaths per 100,000 and cumulative observed all-cause mortality of 30,280.43 deaths per 100,000. The spatial pattern in panel (a) reflects both physical exposure and the estimated SVI-specific response: counties with greater wind exposure and high-SVI classification receive larger fitted TC mortality burdens.

Figure 8 summarizes how the estimated TC mortality burden varies with social vulnerability. Panel (a) plots county mean SVI against the estimated mortality-rate response per 1 m/s of TC wind exposure per 100,000 population. Because the empirical specification estimates separate treatment responses for lower- and high-SVI counties, the fitted relationship is a step function: lower-SVI counties are assigned $\hat{\beta}_L = 0.415$, while high-SVI counties are assigned $\hat{\beta}_H = 2.133$. This panel therefore isolates the vulnerability component of the estimated mortality response: holding physical exposure fixed, the model implies a substantially larger mortality response for high-SVI counties.

Panel (b) combines this SVI-specific response with realized county-event wind exposure. Each point represents one county-event pair and reports the full 172-month cumulative fitted TC mortality burden per 100,000 population. Thus, unlike panel (a), panel (b) reflects both differential vulnerability and differential exposure: high-SVI counties may experience larger fitted burdens because they have a larger estimated




Fig. 8 Social vulnerability and estimated TC-induced mortality burden. Panel (a) shows the SVI-group-specific mortality-rate response per 1 m/s of TC wind exposure per 100,000 population. Panel (b) shows county-event cumulative fitted TC mortality burden per 100,000 population. The dotted curve overlays a descriptive logistic fit to binned county-event means, as defined in Equation (7).

mortality response to a given wind exposure, because they are exposed to stronger TC winds, or both. The dotted curve is a descriptive logistic fit to binned county-event means,

$$y = 6.47 + \frac{28.20}{1 + \exp[-100(x - 0.75)]}, \quad (7)$$

with $R^2 = 0.98$. This curve should be interpreted only as a descriptive summary of the step-like pattern induced by the SVI-group specification, not as a separately estimated structural dose-response function.

This distinction is important for interpreting the mortality-burden maps. Because the lower-SVI coefficient is small and statistically imprecise, the absolute fitted burden for lower-SVI counties should not be overinterpreted. The stronger empirical finding is the relative difference between high- and lower-SVI mortality responses. For this reason, the burden estimates are especially useful for relative mapping: they identify where TC-related mortality risks are expected to be amplified by social vulnerability, rather than serving as exact county-level death counts.

These relative burden measures can inform equity-oriented emergency management. Disaster-preparedness models, such as those developed for emergency vehicle prepositioning (Kim et al. 2024), and infrastructure-resilience models (Toplu-Tutay et al. 2024, 2022) often require planners to account for the fact that a given physical hazard or disaster-induced infrastructure disruption can impose unequal health and well-being losses across communities. The estimated TC mortality-burden mapping provides one way to incorporate these disproportionate impacts into such planning models.

Overall, the mortality-burden results show that TC-induced mortality risk is not only a function of storm exposure. It also depends on the vulnerability of the exposed population. Panel (a) isolates this vulnerability channel by comparing the estimated response across SVI groups under the same unit of exposure, while panel (b) shows the

combined implication of vulnerability and realized wind exposure across county-event observations. Together with the temporal and spatial burden figures, these results suggest that long-run TC mortality burden is concentrated where repeated physical exposure overlaps with high social vulnerability.

Conclusion

This paper examines whether social vulnerability modifies the long-run mortality response to tropical cyclone (TC) exposure. Using county-month mortality, TC wind exposure, temperature, population, and social vulnerability index (SVI) data for Texas from 1968 through 2018, we estimate a reduced-form panel model that combines a long-run lagged TC exposure index with SVI-group-specific treatment effects. The results show that the positive TC mortality response is concentrated among high-SVI counties. The high-SVI coefficient is positive and statistically significant, while the lower-SVI coefficient is small and statistically indistinguishable from zero. The high-minus-lower contrast is also positive, providing evidence that socially vulnerable populations experience larger mortality responses to the same TC exposure.

The validation exercises support this interpretation. Future-exposure placebo tests show no evidence that future storms predict current mortality, reducing concerns about residual pre-trends or anticipatory patterns. Exposure-randomization tests further indicate that the high-SVI estimate is unusually large relative to randomized exposure histories in most designs, while the lower-SVI estimate is not extreme under any design. Together, these checks suggest that the estimated high-SVI response is not generated mechanically by the model specification, the marginal distribution of TC exposure, county-level exposure propensity, or temporal clustering of storms.

The mortality-burden results translate these estimates into temporal, spatial, and event-level measures of TC-attributable mortality. They show that long-run TC mortality burden reflects both physical exposure and social vulnerability. Counties with repeated TC exposure and high SVI receive larger fitted mortality burdens, indicating that disaster mortality risk is not determined by hazard intensity alone. These findings have implications for equity-oriented emergency management and resilience planning: mitigation, recovery, and public-health interventions should account for the social conditions that shape communities' ability to absorb and recover from TC shocks.

This study has several limitations. First, the analysis focuses on Texas and on a 42-county estimation sample with sufficient population size and repeated TC exposure, so the estimates should not be interpreted as statewide or national mortality burdens. Second, TC exposure is measured using wind speed; rainfall, storm surge, flooding, and compound hazards are not modeled explicitly. Third, the county-level setting requires a parsimonious lag structure based on prior state-level evidence, rather than estimating a fully flexible distributed-lag response. Fourth, the SVI specification uses a binary high-versus-lower vulnerability split, which simplifies a multidimensional and continuous vulnerability construct. Finally, the reduced-form design estimates the total mortality response but does not identify the specific mechanisms linking TC exposure to delayed mortality.

Future work should extend the analysis to additional TC-exposed states, incorporate richer hazard measures such as rainfall and flooding, and examine whether results are robust to alternative vulnerability measures and more flexible heterogeneity specifications. Mechanism-focused work is also needed to distinguish the roles of healthcare access, housing damage, displacement, infrastructure disruption, financial stress, and public-resource constraints. Finally, the estimated vulnerability-adjusted mortality burdens could be integrated into resilience planning models to prioritize investments where physical hazard exposure and social vulnerability jointly produce the greatest human cost.

1068

1069 References

1070

1071 Anttila-Hughes J, Hsiang S (2013) Destruction, Disinvestment, and Death: Economic
1072 and Human Losses Following Environmental Disaster. [https://doi.org/10.2139/ssrn.](https://doi.org/10.2139/ssrn.2220501)
1073 [2220501](https://doi.org/10.2139/ssrn.2220501)

1074

1075 Berkeley Earth (n.d.) Global gridded temperature data. URL [https://berkeleyearth.](https://berkeleyearth.org/data/)
1076 [org/data/](https://berkeleyearth.org/data/), accessed 15 February 2025

1077

1078 Best K, Kerr S, Reilly A, et al (2023) Spatial regression identifies socioeconomic
1079 inequality in multi-stage power outage recovery after Hurricane Isaac. *Natural*
1080 *Hazards* 117(1):851–873. <https://doi.org/10.1007/s11069-023-05886-2>

1081

1082 Burrows K, Anderson GB, Yan M, et al (2023) Health disparities among older adults
1083 following tropical cyclone exposure in Florida. *Nature Communications* 14(1):2221.
1084 <https://doi.org/10.1038/s41467-023-37675-7>

1085

1086 Centers for Disease Control and Prevention (2024) Cdc/atsdr social vulnerabil-
1087 ity index. <https://www.atsdr.cdc.gov/place-health/php/svi/index.html>, accessed
1088 2026-02-08

1089

1090 Centers for Disease Control and Prevention (2025) Multiple cause of death data on
1091 CDC WONDER. URL <https://wonder.cdc.gov/mcd.html>, accessed 18 January 2025

1092

1093 Clark SS, Peterson SK, Shelly MA, et al (2023) Developing an equity-focused metric for
1094 quantifying the social burden of infrastructure disruptions. *Sustainable and Resilient*
1095 *Infrastructure* 8(sup1):356–369. <https://doi.org/10.1080/23789689.2022.2157116>

1096

1097 Flanagan BE, Gregory EW, Hallisey EJ, et al (2011) A Social Vulnerability Index for
1098 Disaster Management. *Journal of Homeland Security and Emergency Management*
1099 8(1):3. <https://doi.org/10.2202/1547-7355.1792>

1100

1101 Hallegatte S, Vogt-Schilb A, Rozenberg J, et al (2020) From Poverty to Disaster
1102 and Back: a Review of the Literature. *Economics of Disasters and Climate Change*
1103 4(1):223–247. <https://doi.org/10.1007/s41885-020-00060-5>

1103

1104

Han Y, Ye X, Zhu C (2024) The Unequal Impact of Disasters: Assessing the Interplay Between Social Vulnerability, Public Assistance, Flood Insurance, and Migration in the U.S. <i>Urban Informatics</i> 3(1):30. https://doi.org/10.1007/s44212-024-00061-9	1105 1106 1107 1108
Kim KY, Toplu-Tutay G, Kutanoglu E, et al (2024) A stochastic optimization model for patient evacuation from health care facilities during hurricanes. <i>International Journal of Disaster Risk Reduction</i> 108:104518. https://doi.org/10.1016/j.ijdrr.2024.104518	1109 1110 1111 1112 1113
Kishore N, Marqués D, Mahmud A, et al (2018) Mortality in Puerto Rico after Hurricane Maria. <i>New England Journal of Medicine</i> 379(2):162–170. https://doi.org/10.1056/NEJMsa1803972	1114 1115 1116 1117
Markhvida M, Walsh B, Hallegatte S, et al (2020) Quantification of disaster impacts through household well-being losses. <i>Nature Sustainability</i> 3(7):538–547. https://doi.org/10.1038/s41893-020-0508-7	1118 1119 1120
National Bureau of Economic Research (2026) Mortality data: Vital statistics NCHS multiple cause of death data. https://doi.org/10.60592/s3mf-1882	1121 1122 1123
Newberry Library (n.d.) Atlas of historical county boundaries. URL https://publications.newberry.org/ahcb/ , accessed 21 February 2025	1124 1125 1126
NOAA National Centers for Environmental Information (NCEI) (2025) U.S. Billion-Dollar Weather and Climate Disasters. https://doi.org/10.25921/stkw-7w73	1127 1128 1129
Parks RM, Anderson GB, Nethery RC, et al (2021) Tropical cyclone exposure is associated with increased hospitalization rates in older adults. <i>Nature Communications</i> 12(1):1545. https://doi.org/10.1038/s41467-021-21777-1	1130 1131 1132 1133
Parks RM, Benavides J, Anderson GB, et al (2022) Association of Tropical Cyclones With County-Level Mortality in the US. <i>JAMA</i> 327(10):946–955. https://doi.org/10.1001/jama.2022.1682	1134 1135 1136 1137
Peterson SKE, Clark SS, Shelly MA, et al (2024) Assessing the household burdens of infrastructure disruptions in Texas during Winter Storm Uri. <i>Natural Hazards</i> 120(8):7065–7104. https://doi.org/10.1007/s11069-024-06480-w	1138 1139 1140
Strader SM, Walsh SJ, Linehan TM, et al (2026) Geospatial assessment of United States tropical cyclone disaster risk. <i>Natural Hazards</i> 122:301. https://doi.org/10.1007/s11069-026-08065-1	1141 1142 1143 1144
Tate E, Rahman MA, Emrich CT, et al (2021) Flood exposure and social vulnerability in the United States. <i>Natural Hazards</i> 106(1):435–457. https://doi.org/10.1007/s11069-020-04470-2	1145 1146 1147 1148 1149 1150

- 1151 Toplu-Tutay G, Hasenbein JJ, Kutanoglu E (2022) A scenario-based optimization
 1152 model for long-term healthcare infrastructure resilience against flooding. In: Pro-
 1153 ceedings of the IISE Annual Conference & Expo 2022, pp 1–6, URL [https://par.](https://par.nsf.gov/servlets/purl/10356281)
 1154 [nsf.gov/servlets/purl/10356281](https://par.nsf.gov/servlets/purl/10356281)
- 1155
 1156 Toplu-Tutay G, Hasenbein JJ, Kutanoglu E (2024) Impact of power outages depends
 1157 on who loses it: Equity-informed grid resilience planning via stochastic optimization.
 1158 Socio-Economic Planning Sciences 95:102036. [https://doi.org/10.1016/j.seps.2024.](https://doi.org/10.1016/j.seps.2024.102036)
 1159 [102036](https://doi.org/10.1016/j.seps.2024.102036)
- 1160
 1161 U.S. Census Bureau (2018) 2018 TIGER/Line shapefiles. URL [https://www.](https://www.census.gov/geographies/mapping-files/time-series/geo/tiger-line-file.html)
 1162 [census.gov/geographies/mapping-files/time-series/geo/tiger-line-file.html](https://www.census.gov/geographies/mapping-files/time-series/geo/tiger-line-file.html), accessed
 1163 21 February 2025
- 1164
 1165 Young R, Hsiang S (2024a) Mortality caused by tropical cyclones in the United States.
 1166 Nature 635(8037):121–128. <https://doi.org/10.1038/s41586-024-07945-5>
- 1167
 1168 Young R, Hsiang S (2024b) Mortality caused by tropical cyclones in the United States:
 1169 Replication data. <https://doi.org/10.5281/zenodo.10459719>

1171 **Statements and Declarations**

1173 **Funding**

1175 Gizem Toplu-Tutay received primary support through the University Graduate Con-
 1176 tinuing Fellowship and additional research support from the Bureau of Business
 1177 Research at the IC² Institute, both at The University of Texas at Austin.

1179 **Competing Interests**

1181 The authors have no relevant financial or non-financial interests to disclose.

1183 **Code and Data Availability**

1184 Data and code are archived on Zenodo at <https://doi.org/10.5281/zenodo.22097263>.

1186 **Ethics Approval**

1188 Ethics approval was not required because the study uses secondary, aggregate county-
 1189 level mortality data and other publicly available administrative data and does not
 1190 involve identifiable human participants.

1192 **Author Contributions**

1194 Gizem Toplu-Tutay contributed to the conceptualization, methodology, data curation,
 1195 formal analysis, software, visualization, writing of the original draft, and review and
 1196 editing of the manuscript. John J. Hasenbein and Erhan Kutanoglu supervised the

research and reviewed and edited the manuscript. All authors read and approved the
final manuscript.