Journal of the American Statistical Association

ISSN: 0162-1459 (Print) 1537-274X (Online) Journal homepage: www.tandfonline.com/journals/uasa20

A Spectral Confounder Adjustment for Spatial Regression with Multiple Exposures and Outcomes

Shih-Ni Prim, Yawen Guan, Shu Yang, Ana G Rappold, K. Lloyd Hill, Wei-Lun Tsai, Corinna Keeler & Brian J Reich

To cite this article: Shih-Ni Prim, Yawen Guan, Shu Yang, Ana G Rappold, K. Lloyd Hill, Wei-Lun Tsai, Corinna Keeler & Brian J Reich (06 Aug 2026): A Spectral Confounder Adjustment for Spatial Regression with Multiple Exposures and Outcomes, Journal of the American Statistical Association, DOI: 10.1080/01621459.2026.2714568

To link to this article: https://doi.org/10.1080/01621459.2026.2714568

A Spectral Confounder Adjustment for Spatial Regression with Multiple Exposures and Outcomes

Shih-Ni Prim[*,#], Yawen Guan[†], Shu Yang[*], Ana G Rappold[‡], K. Lloyd Hill[§], Wei-Lun Tsai[‡], Corinna Keeler[‡], and Brian J Reich[*]

[*]Department of Statistics, North Carolina State University

[†]Department of Statistics, Colorado State University

[‡]US EPA, Center for Public Health and Environmental Assessment

[§]Oak Ridge Associated Universities at US EPA, Center for Public Health and Environmental Assessment

[#]sprim@ncsu.edu

Abstract

Characterizing social vulnerability is fundamental to disaster response planning. Numerous vulnerability indicators have been developed, but they are typically not validated for their predictive power over the outcomes of interest. As with many environmental health studies where interventions are impractical or unethical, validation of social vulnerability against public health outcomes is observational and relies on spatially-dependent data. Observational studies are susceptible to bias induced by unmeasured confounders. This problem is exacerbated in spatial studies with multiple health outcomes and environmental exposure variables, as the source and magnitude of confounding bias may differ by spatial scale and across exposure/outcome pairs. We propose to mitigate confounding effects in multivariate spatial studies using a tensor-regression model that allows exposure effects to vary by exposure, outcome, and spatial scale. By extracting exposure effects that correspond to local spatial scales, we replace the strict causal assumption of no unmeasured confounders with a more realistic assumption of local unconfoundedness i.e., differences between nearby regions are unconfounded, allowing for causal interpretation. Our study of the Southern United States shows that economic resilience (Theme 1) demonstrates the strongest positive effect on diabetes and negative effect on hyperlipidemia. Our analysis reveals a concise representation of the effects of social vulnerability indices on an array of chronic health outcomes, offering interpretable epidemiological insights that adjust for multivariate spatial confounding.

Keywords: Spatial confounding; causal inference; exposure mixtures; unmeasured confounder; tensor decomposition.

1 Introduction

Large electronic health databases empower researchers to characterize how the environment and social context give rise to adverse health outcomes and to map community resilience. This task requires untangling many-to-many causal relationships that are deeply intertwined geographically. For example, poverty influences multiple diseases through pathways such as education, physical and physiological safety, and healthcare access. Consequently, poverty can drive heart failure and diabetes simultaneously, though with varying levels of severity. Because health conditions and poverty are intertwined with other spatially-varied factors, such as local environmental pollution and income, they become spatially confounded. In this paper, we offer a statistical method to study the many-to-many relationship when the assumption of no unmeasured spatial confounding is violated.

All observational studies are susceptible to missing-confounder bias, and in environmental applications, the missing confounders often have a spatial structure. Reich et al. (2006) demonstrated a profound impact of spatial confounding on the estimated regression coefficients in the models that include or exclude spatial random effects. There are many approaches to remedying this problem in the univariate setting (Paciorek 2010, Hughes and Haran 2013, Page et al. 2017, Dupont et al. 2022, Khan and Calder 2022, Marques et al. 2022, Guan et al. 2023, Adin et al. 2023, Dupont et al. 2024, Thaden and Kneib 2018, Wiecha et al. 2025, Wang and Blei 2019, Reich et al. 2021, Miao et al. 2023, Schnell and Papadogeorgou 2020, Gilbert et al. 2024, Wiecha et al. 2025). Here, we build on the spectral method in Guan et al. (2023) to relax the no unmeasured confounder assumption. Guan et al. (2023) allow the exposure effect to vary by spatial scale, and take the effect estimate for the most local scale as the causal estimate. This relaxes the stringent no unmeasured confounders assumption to the more realistic local unconfoundedness assumption.

In addition to the presence of unmeasured confounders, mixtures as exposures (i.e., many simultaneous exposures) present unique problems such as high dimensionality, collinearity, and interactions between exposures (Joubert et al. 2022). A variety of approaches has been proposed (Carrico et al. 2015, Bobb et al. 2015, Reich et al. 2020, Keil et al. 2020, Wei et al. 2020, Ferrari and Dunson 2020, Kowal et al. 2021, Gibson et al. 2021, Roy et al. 2021, Chen et al. 2022, Devick et al. 2022, Gennings et al. 2022), but they do not address spatial confounding, leaving a critical gap. As many mixture studies are concerned with epidemiological and environmental problems, the causal inference framework has also been utilized (Zigler 2021, Gennings et al. 2022, Devick et al. 2022).

Among the methods for spatial confounding, few studies include many-to-many relationships with multiple treatments and multiple responses. Note that the terms “treatment” and “exposure” are used interchangeably. For non-spatial approaches, Zheng et al. (2023) work with one treatment and Kang et al. (2025) deal with multiple treatments; both rely on factor models and use sensitivity analysis to bound the causal effects of unmeasured confounder for multiple outcomes. Miao et al. (2023) use auxiliary variables—which are not causally associated with the

response—and null treatments to find causal effects from multiple treatments when unmeasured confounders exist. For spatial data, Marques et al. (2022) investigate a single response and multiple covariates in their use of Gaussian random field to reduce spatial confounding. Urdangarin et al. (2024) extend the work of Dupont et al. (2022), who regress the covariates on the thin plate model, and the resulting residuals are used as the covariates while a penalty term for smoothness is included. Urdangarin et al. (2024) find the spectral decomposition of the spatial precision matrix and express the covariate as linear combinations of eigenvectors from frequencies assumed confounded.

We develop a novel method for modeling causal effects for spatial data with multiple exposures and outcomes in the presence of unmeasured, global spatial confounders. Motivated by Guan et al. (2023), we project spatial data into the spectral domain and use the local-unconfoundedness assumption to identify causal effects. The extension from univariate to multivariate requires a three-way tensor to quantify the effects by exposure, outcome and spatial scale. The proposed tensor structure captures complicated patterns with constraints on diminishing correlations at local scales to adhere to causal assumptions, and provides interpretable latent factors that lay bare the key relationships that drive the results. Compared to related methods, the proposed method allows for multiple exposures and outcomes unlike the univariate methods of Guan et al. (2023) and Marques et al. (2022). Compared to Urdangarin et al. (2024), the proposed method provides a comprehensive assessment of how estimates vary by spatial scale and a low-rank and interpretable estimate of the outcome/exposure structure, and avoids hard selection of the number of frequency terms that are assumed to be confounded.

The remainder of the paper proceeds as follows. The motivating data are described in Section 2. Section 3 introduces the statistical method. The method is applied to the motivating data in Section 4. Section 5 provides concluding remarks and discusses limitations. An extensive simulation study, computation, derivations, and additional results are provided in the Supplementary Materials, in which sections, tables, and figures are numbered by labels starting with "S" to distinguish from sections from the main text.

2 Research Objectives

We demonstrate the utility of the proposed method to establish a causal estimate of the strength of the effect between community resilience and adaptive capacity——characterized by the CDC's Social Vulnerability Index for Disaster Management (SVIDM)——and the incidence of prevalent chronic diseases. Although the SVIDM is a composite index designed to identify communities at risk during disasters and public health emergencies, its predictive validity regarding long-term public health burdens remains unknown. Furthermore, evaluating this relationship is heavily complicated by spatial confounding from unmeasured factors that vary geographically in a similar way to public health metrics. Because understanding this predictive power is necessary before the index can be deployed as an effective public health intervention tool, identifying the causal link between health burden and community resilience (Flanagan et al. 2011) is a critical step in operationalizing the index as a public health intervention tool. A strong predictive relationship enables public health officials to prioritize targeted interventions preemptively and allocate resources both during and in the wake of emergencies. The deployment of resources to

vulnerable areas during disasters and emergencies serves as a salient motivating example to demonstrate the value of our proposed methodology.

Composite indices are common tools for mapping community-level vulnerability and managing environmental threats. These indices synthesize multiple variables into a single metric to summarize localized contexts, such as environmental quality (Messer et al. 2014), disaster readiness (Flanagan et al. 2011), and neighborhood deprivation (Messer et al. 2006). Methodologically, researchers typically build these composite measures using factor models or principal component analysis (Reckien 2018, Greco et al. 2019). However, the most common drawback is that their predictive value is not validated. Moreover, the relationship between composite indices and public health burden is subject to confounding by unmeasured factors at multiple spatial scales, which can impact their predictive performance and lead to biased results. Therefore, we explore the relationship between health burden and a composite index of community resilience at various spatial scales to demonstrate the value of the proposed methodology and to evaluate evidence of spatial confounding.

We proceed with the following objectives:

• Test whether the relationship between SVIDM and chronic conditions varies by spatial scale, and thus caution against using results from commonly used methods (e.g. linear regression) that do not adjust for confounding.

• Estimate the causal effect of each component of SVIDM on the incidence of chronic conditions to help decision makers allocate resources on relevant SVIDM themes for a chronic outcome of interest.

• Distill the many-to-many causal inference problem down to interpretable visualizations to motivate future studies on public health policies to improve disaster preparedness.

All three aims are policy relevant: the first calls into question results from naive analyses of the data; the second frames the problem causally which is needed for decision making; and the third focuses on the dimension reduction to help policymakers extract meaningful interpretations from the analysis.

2.1 Data

The outcome data consisted of incidence rates of five prevalent chronic conditions (hypertension, chronic kidney disease, hyperlipidemia, congestive heart failure, and diabetes) by ZIP code tabulated area of residence. We obtained individual level data from the Centers for Medicare and Medicaid Services' Chronic Conditions Warehouse, which included chronic health conditions and year specific ZIP-code of residence. Each incident health condition for each unique beneficiary contributed one count to the residential ZIP-code. The counts were aggregated to their corresponding ZIP code tabulated areas (ZCTAs), and the incident rate was calculated as the total counts of incident health conditions divided by the total beneficiaries in the ZCTA each year (US Department of Health and Human Services 2023). We used the mean of the incident rates from 2017, 2018, and 2019.

The Social Vulnerability Index for Disaster Management was created by the Centers for Disease Control (Centers for Disease Control and Prevention; Agency for Toxic Substances and Disease Registry; Geospatial Research, Analysis, and Services Program 2018) with the intent to understand adaptive capacity for natural disasters through the lens of public health. We retrieved the 2018 SVIDM index covering the period of 2014-2018 at the ZCTA level using the `R` package `findSVI` (Xu et al. 2025). The index scores range between 0 and 1, representing the proportion of ZCTAs that are lower than or equal to the ZCTA of interest in terms of social vulnerability. Thus, higher scores signal worse abilities to respond to disasters. The index scores are structured into four themes based on 15 factors: Theme 1 describes lack of economic resilience based on measures of poverty, unemployment, income, and lack of high school diploma; Theme 2 describes household-level adaptive capacity relying on measures of number of people aged 65 or older, aged 17 or younger, number of civilian population with a disability, and number of single-parent households; Theme 3 describes fraction of minority and non-English speaking population; Theme 4 describes the housing type and transportation (multi-unit structures, mobile homes, crowding, no vehicle, group quarters).

The adjusting covariates include the population size and the indicator of urban/rural ZCTA. Population size data was downloaded from the US Census Bureau's 5-year American Community Survey (ACS-5) (US Bureau of the Census 2023). The logarithmic scale of the population is used. Urban/rural determination was obtained from 2010 data from the Integrated Public Use Microdata Series (IPUMS) National Historical Geographic Information System (NHGIS) (Manson et al. 2024). A ZCTA was considered urban if more than 50% of its population lived in urban areas. The split of rural versus urban areas was 57.2% urban and 42.8% rural. Maps of outcomes, exposures, and covariates are provided in Figures 1 and S5 show strong spatial patterns, and Tables 1 and 2 show there correlations between exposures and outcomes; these results support our multivariate spatial analysis. We focus on the South region (CDC/National Center for Health Statistics/Division of Analysis and Epidemiology 2024), which includes seventeen states (Alabama, North Carolina, Oklahoma, Arkansas, Delaware, Florida, Georgia, Kentucky, Louisiana, Maryland, Missouri, Mississippi, South Carolina, Tennessee, Texas, Virginia, West Virginia) and the District of Columbia. This region is experiencing rapid population growth and is susceptible to an array of natural disasters including hurricanes, storm surge, tornadoes, inland flooding, and extreme heat. After removing ZCTAs with less than 100 beneficiaries the dataset included 10,149 ZCTAs.

3 Statistical methods

3.1 Model in the spatial domain

Assume the spatial domain is partitioned into *S* regions, and $a_{st} = 1$ indicates that regions *s* and *t* are adjacent and $a_{st} = 0$ are otherwise. The data for region *s* consists of *R* response variables, $\{Y_{s1},...,Y_{sR}\}$, *E* exposure variables, $\{X_{s1},...,X_{sE}\}$ and *P* known confounding variables, $\{Z_{s1},...,Z_{sP}\}$. In the spatial domain, the data model is:

$$Y_{sr} = \eta_{0r} + \sum_{p=1}^{P} Z_{sp}\eta_{pr} + \sum_{e=1}^{E} X_{se}\beta_{er} + \theta_{sr} + \varepsilon_{sr}, \quad (1)$$

where η_{pr} controls the observed confounder effects, β_{er} are the exposure effects of interests, θ_{sr} are spatially-dependent residuals, and $\varepsilon_{sr} \overset{indep}{\sim} \text{Normal}(0, \tau_r^2)$ is error. While the responses are assumed to be normally distributed, the exposures can be all types—continuous, binary, categorical, etc. In the traditional spatial regression model, θ_{sr} is modeled with a spatial process independent of X_{se}. If θ_{sr} is correlated with X_{se}, it is a spatial confounder. If the spatial random effect $\theta = (\theta_{s1}, \cdots, \theta_{sR})^\top$ for $s = \{1, \cdots, S\}$ is observed, estimation of β_{er} is straightforward. Under the potential outcomes framework and the Stable Unit Treatment Value Assumption (SUTVA), if the potential outcomes follow (1), $(\beta_{1r}, \cdots, \beta_{Er})^\top$ represents the average treatment effects of the exposures on outcome *r* (Section S1). However, when the multivariate, spatial θ is unobserved and treated as a spatially-correlated error, the dependence between θ_{sr} and X_{se} leads to bias in estimating β_{er}.

Our approach for mitigating this bias makes use of the Conditional Autoregressive (CAR) model. We study the form of bias in special cases in Section S2. For an arbitrary random vector $\mathbf{V} = (V_1, ..., V_S)^\top$, the CAR model with Leroux parameterization (Leroux et al. 2000) is multivariate normal with mean zero and covariance matrix $\sigma_V^2 \{(1-\lambda_V)\mathbf{I}_S + \lambda_V \mathbf{Q}\}^{-1}$ where σ_V^2 controls the variance, $\lambda_V \in (0,1)$ determines the strength of spatial dependence, and $\mathbf{Q}$ has values $\sum_t a_{st}$ on the diagonal and $-a_{st}$ on the off-diagonals. We denote this model by $\mathbf{V} \sim \text{CAR}(\sigma_V^2, \lambda_V)$.

The multivariate spatial random effects θ_{sr} are modeled using the Linear Model of Coregionalization (LMC) (Schmidt and Gelfand 2003, Wackernagel 2003, Gelfand et al. 2005, Gelfand and Banerjee 2010, Banerjee et al. 2015). The LMC assumes the random effects for the *R* response variables can be explained by $Q \leq R$ latent spatial factors as

$$\theta_{sr} = \sum_{q=1}^{Q} A_{rq} C_{sq}, \quad (2)$$

where A_{rq} are the factor loadings and $\mathbf{C}_q = (C_{1q}, ..., C_{Sq})^\top \overset{indep}{\sim} \text{CAR}(\sigma_q^2, \lambda_C)$ for $q \in \{1, ..., Q\}$. We take *Q* to be the same as the number of responses so that the model is full rank. This multivariate spatial process includes both spatial dependence via the CAR model and cross dependence via the factor loadings. The cross-correlations have closed forms, i.e., at any location *s*, two different outcomes have correlation

$$Cor(\theta_{sr}, \theta_{sr'}) = \frac{\sum_{q=1}^{Q} \sigma_q^2 A_{rq} A_{r'q}}{\sqrt{\left(\sum_{q=1}^{Q} \sigma_q^2 A_{rq} A_{rq}\right)\left(\sum_{q=1}^{Q} \sigma_q^2 A_{r'q} A_{r'q}\right)}}.$$

To ensure identifiability, we set $A_{qq} = 1$ and $A_{rq} = 0$ for $q > r$ (Gelfand and Banerjee 2010). The remaining terms have prior distributions $A_{rq} \sim \text{Normal}(0, 0.5^2)$, so that the diagonal terms dominate.

While the multivariate CAR model in (2) accounts for multivariate spatial residual dependence, it assumes the random effects θ_{sr} are independent of the exposure variables X_{se}. The spatial model without properly accounting for the dependence between θ and X does not mitigate bias. As a result, the estimates of β_{er} can only be interpreted as causal effects when θ_{sr} and X_{se} are independent, namely, under the assumption of no unmeasured spatial confounding. In Section 3.2 we relax this global assumption.

3.2 Model in the spectral domain

We propose to estimate β_{er} in the spectral domain. By transforming a spatial variable to the spectral domain, we can decompose the spatial information to different scales. The no unmeasured confounder assumption requires θ and $\mathbf{X}$ are independent at each scale, where $\mathbf{X} = (\mathbf{X}_1, \cdots, \mathbf{X}_E)$, $\theta = (\theta_1, \cdots, \theta_R)$, and $\mathbf{X}_e$ and θ_r are vectors of length *S*. We relax this assumption by assuming that global-scale terms are susceptible to confounding whereas local-scale terms are unconfounded, i.e., local unconfoundedness (Guan et al. 2023). Our multiscale estimates can thus capture the true effect at scales where there is assumed to be no confounding.

Our spectral approach uses the graph Fourier transform to decompose the precision matrix into independent signals at different spatial scales. Let $\mathbf{Q} = \Gamma \mathbf{W} \Gamma^\top$ be the spectral decomposition of $\mathbf{Q}$ for eigenvector matrix Γ and diagonal matrix $\mathbf{W}$ with diagonal elements equal to the eigenvalues $w_1 \geq \ldots \geq w_S$. For arbitrary vector $\mathbf{V} \sim \text{CAR}(\sigma_V^2, \lambda_V)$, premultiplying by the eigenvectors projects the spatial process into the spectral domain. The spectral process is $\Gamma^\top \mathbf{V} = \mathbf{V}^* = (V_1^*, \ldots, V_S^*)^\top$ with

$$V_i^* \overset{indep}{\sim} \text{Normal}\left\{0, \sigma_V^2 / (1 - \lambda_V + \lambda_V w_i)\right\}.$$

In the spectral domain, V_i^* are independent across *i*, allowing us to model each scale separately, account for multiscale confounding, and represent weights of the eigenvectors. Larger weights on the eigenvectors that correspond to small eigenvalues result in smoother spatial process (termed global scale), while larger weights on the eigenvectors corresponding to large eigenvalues lead to rougher spatial process (termed local scale) (Reich et al. 2006).

We specify our model for adjusting for the unmeasured confounders by projecting the model in (1) to the spectral domain. For response $r \in \{1,...,R\}$, the data in the spatial domain is $\mathbf{Y}_r = (Y_{1r},...,Y_{Sr})^\top$ and the projection to the spectral domain is $\Gamma^\top \mathbf{Y}_r = \mathbf{Y}_r^* = (Y_{1r}^*,...,Y_{Sr}^*)^\top$, where Y_{ir}^* captures the variation of the outcome *r* at eigenvalue w_i. Similarly, let Z_{ip}^*, X_{ie}^* and C_{iq}^* for $p \in \{1,...,P\}$, $q \in \{1,...,Q\}$, and $e \in \{1,...,E\}$ be the projected values of the observed confounding, exposure, and spatial variables. Then (1) in the spectral domain is

$$Y_{ir}^* = \eta_{0r} + \sum_{p=1}^{P} Z_{ip}^* \eta_{pr} + \sum_{e=1}^{E} X_{ie}^* \beta_{er} + \theta_{ir}^* + \varepsilon_{ir}^*, \text{ where} \qquad (3)$$

$$\theta_{ir}^* = \sum_{q=1}^{Q} A_{rq} C_{iq}^*$$

for $i \in \{1,...,S\}$ and $\varepsilon_{ir}^* \overset{indep}{\sim} \text{Normal}(0, \tau_r^2)$. Identifiability of the regression coefficients in β is discussed in Section 3.3 and that of the other parameters are covered in Section 3.3. In the spectral domain, the latent factors are distributed as $C_{iq}^* \overset{indep}{\sim} \text{Normal}\left\{0, \sigma_q^2 / (1 - \lambda_C + \lambda_C w_i)\right\}$.

The distinction between confounding in the "global" versus "local" spatial scales prompts a key assumption of our proposed model. The local unconfoundedness assumption is that terms in (3) with large w_i are unaffected by spatial confounding. The matrix $\mathbf{Q}$ in the Leroux parameterization corresponds to a Laplacian matrix defined on an undirected graph defined based on neighborhood structure (Chapter 2 BiyikoÇ§u et al. 2007). The Laplacian operator on a rectangular domain with Dirichlet and Neumann boundary condition has eigenfunctions that are harmonic (sine and cosine) functions (Grebenkov and Nguyen 2013). For example, Section S2.1 provides the closed-form and visualization of the eigenfunction on a 1D path graph. In more general graph structure, the eigen functions are not in closed-form; however, Figure 2 shows that the eigenfunctions represent different spatial scales. Global confounding refers to large-spatial scale (low frequency) confounding and vice versa.

3.3 Identifiability of β in spectral domain

Unlike a standard spatial regression, we model dependence between X_{ie}^* and θ_{ie}^*. We assume a linear confounding relationship between θ_{ir}^* and X_{ie}^*: $\theta_{ir}^* \mid X_{ie}^* = \alpha_{ier} X_{ie}^* + \xi_{ir}$, where α_{ier} is the confounding effect of exposure *e* on response *r* specific to spatial scale described by eigenvalue w_i. The spatial error term, ξ_{ir}, captures the remaining variation in θ_{ir}^* after conditioning on X_{ie}^* and is independent of the non-spatial error term, ε_{ir}^*. The magnitude of α_{ier} determines the level of correlation between X_{ie}^* and θ_{ir}^*. Note that a linear confounding relationship is assumed in the projection of (1) and (2) to the spectral domain (3). In this scenario, the scale-invariant true effect β_{er} is generally inseparable from the confounding bias α_{ier} when scales are not explicitly modeled, since increasing X_{ie}^* by one increases the mean response by $\tilde{\beta}_{ier} = \beta_{er} + \alpha_{ier}$.

We use a tensor of dimension $S \times E \times R$ to denote multiscale coefficients β, which contains $\tilde{\beta}_{ier}$ at scale *i*, exposure *e*, and response *r*. To accurately estimate the coefficient tensor β, our proposed method allows for a multiscale relationship between **X** and θ. A spectral analysis of (3) estimates effects for all *S* scales. We further assume the strength of correlation diminishes at local scales, i.e. $\alpha_{ier} \to 0$ for large w_i. This **local unconfoundedness assumption** is critical in identifying $\boldsymbol{\beta}$. Thus $\hat{\beta}_{1er}$, the estimate for the multiscale coefficient $\tilde{\beta}_{1er}$, can capture the true β_{er} at eigenvalue w_1, where the correlation between X^*_{ie} and θ^*_{ir} goes to zero and

$$\tilde{\beta}_{1er} = \beta_{er}.$$

3.4 Bayesian tensor regression

The coefficient β includes effects across multiple exposures, multiple responses, and multiple scales and is thus structured as a tensor. While decomposition and the computational aspects have been the focus for tensor research (Zhang and Aeron 2017, Zhang and Xia 2018, Zhang and Han 2019), using tensor for the purpose of statistical inference has seen a recent rise in interest (Mai et al. 2021, Llosa-Vite and Maitra 2023). We employ Bayesian tensor regression methods for the regression coefficients. Since we envision $\tilde{\beta}_{ier}$ being smooth across scale *i* but not exposure *e* or response *r*, we represent the variation across *i* using *L* spline basis functions $B_{i1},...,B_{iL}$ and $\tilde{\beta}_{ier} = \sum_{l=1}^{L} B_{il}\gamma_{ler}$, where B_{il} are known B-spline basis functions. The intuition for this model is, first, assuming α_{ier} is a smooth function of w_i, the *L* spline bases are sufficient to capture its relationship with w_i. Second, the multiple exposures might influence the outcomes through common underlying biological pathways or mechanisms, or the outcomes may respond in similar ways to different exposures. A low-rank model helps identify these latent structures.

The basis coefficients γ_{ler} are modeled using the Bayesian tensor regression approach in Guhaniyogi et al. (2017). The tensor γ is written as a rank *K* tensor product of factors in each of the three dimensions, with

$$\gamma_{ler} = \sum_{k=1}^{K} T_{1lk} T_{2ek} T_{3rk}. \quad (4)$$

$\mathbf{T}_1$, $\mathbf{T}_2$, and $\mathbf{T}_3$ are matrices, whose columns are tensor margins from basis functions, exposures, and outcomes, respectively. More specifically, $\mathbf{T}_t = (T_{t1},\cdots,T_{tK}), t = 1,2,3$ and tensor margins T_{tk} are defined in Figure 3. Each tensor margin is of the dimension $\mathbb{R}^L$, $\mathbb{R}^E$, and $\mathbb{R}^R$. The latent factors T_{1lk}, T_{2ek}, and T_{3rk} are given Bayesian shrinkage priors to encourage sparsity, assuming a structure of lower than full rank exists (Section 3.5). Using canonical polyadic (CP) tensor decomposition (Kolda and Bader 2009), the sum of *K* outer products of the tensor margins approximate the tensor coefficient γ, as represented in Figure 3. Section S3 verifies that the

parameters in this model including γ are identifiable. The benefits of such decomposition are threefold. It reduces dimensionality and simplifies the model, representing essential patterns with far fewer parameters and avoiding overfitting. Further, this model captures non-linear relationships and interactions between frequency, exposures, and outcomes. Lastly, the low-rank model makes the results more interpretable and generalizable.

Kolda and Bader (2009) provide a guideline for tensor ranks: $K = \min\{LE, ER, LR\}$. In our simulations, we choose a value smaller than the above value that achieves convergence in MCMC chains. To set a range for *L*, we consider the relationship between the total number of parameters in the tensor, $K(L+E+R)$, and the number of response variables *SR*, $0 < L \leq \frac{SR}{K} - E - R.$ We fit the model with several *L* and check trace plots and estimates to choose the optimal *L*, and sensitivity tests are described in Section S8.

3.5 Shrinkage priors for tensor coefficients

The tensor regression coefficients in (4) are given horseshoe priors to shrink the regression estimates towards zero for unnecessary ranks or null exposures and outcomes. The horseshoe prior is $T_{1lk} \sim N(0, \lambda_{1k}); T_{2ek} \sim N(0, \lambda_{2e}); T_{3rk} \sim N(0, \tau^2 \lambda_{3r} \tau_r^2)$, where $\lambda_{1k}, \lambda_{2e}$, and λ_{3r} are the variances and control local shrinkage. A global shrinkage hyperparameter, τ, is embedded in T_{3rk} to shrink the coefficients in all outcomes. The variance component τ_r^2 is the variance of the random errors from ε_{ir}^* in (3) and does not affect shrinkage. Following Makalic and Schmidt (2016), we parameterize the shrinkage priors hierarchically. Each shrinkage parameter follows an inverse gamma prior, e.g. $\lambda_{jk} \mid \nu_j \sim \text{InvGamma}(\frac{1}{2}, \frac{1}{\nu_j})$ for $j \in \{1,2,3\}$. Each hyper-parameter $\nu_j \sim \text{InvGamma}(\frac{1}{2}, 1)$, resulting in the standard half-Cauchy distribution for the square roots of the shrinkage parameters.

3.6 Simulation study

We apply the methods described in Section 3 to simulated data to examine how our MSM method performs with different magnitudes of linear and nonlinear spatial confounding. Compared to the naive model that does not account for spatial confounding, MSM has over 40% of reduction in mean MSE (from 0.065 to 0.039). The mean coverage rate for MSM is 98.8%, close to the nominal rate of 95%, as opposed to 85.1% for the Naive method. More details are shown in Figure S3 and Tables S1 to S5. The results demonstrate that: 1) our proposed MSM model has low mean squared errors, 2) MSM tends to have coverage rates above the nominal rate, suggesting its conservative nature, 3) when no confounding is present, MSM's performance is similar to the model that does not allow for spatial confounding, and 4) MSM is robust to moderate nonlinear confounding. Implementation details and results are provided in Section S6.

4 Data analysis

In this section, we analyze the SVIDM and chronic conditions data described in Section 2 using the proposed MSM. We first tune and compare models and conduct checks of model fit. We then use the fitted models in the subsections below to address the study's objectives, outlined in Section 1 and reiterated below: quantify whether the relationship between SVIDM and chronic conditions varies by spatial scale, and thus establish the presence of spatial confounding; estimate the causal effect of each component of SVIDM on the incidence of chronic conditions while accounting for spatial confounding; and identify a low-rank structure that distills the many-to-many causal inference problem down to a minimal set of epidemiologically interpretable factors.

4.1 Selection and validation of the best fit model

To select the best model, we compare model performance of four combinations of the MSM method: $K = 5,10$ (tensor rank) and $L = 5,10$ (number of basis functions). We split the dataset that contains 10,149 ZCTAs into a training set (80%) and a test set (20%) and use log scores to select the best model ($K = L = 10$). The out-of-sample predictive R-squared and coverage verified the performance of the selected model (see Table 3). All models are run with 20,000 iterations, 2,000 burn-ins, and thinning rate 10. Sensitivity study with respect to different L and K is provided in Section S8. The final model was run on the full dataset with 100,000 iterations, 10,000 burn-ins, and a thinning rate of 10. Trace plots confirmed convergence for all coefficients, as shown in Section S10. Residual plots in the spatial domain (Figure 4 for diabetes and Figure S7 for all outcomes) showed no spatial patterns or outliers. In Section S11, the density plot and variogram (Figure S8) showed that residuals were normally distributed with mean around zero and have very low spatial correlations. Our analysis assumes a linear relationship with the unmeasured spatial confounders. This assumption is unverifiable, but our residual analysis in Section S11 shows no signs of non-linearity with the observed confounders, and the simulation study in Section S6 suggests robustness to moderate deviations from linearity.

To highlight the difference between our proposed multiscale method and methods that either lack spatial information or do not allow effects to vary by scale, we compare the results with those from the univariate spectral models (USM), a naive spatial model ("naive"; see Section S6.2.1) and a non-spatial ordinary least squares (OLS). The USM model used the same settings (*K*, *L*, iterations, burn-ins, and thinning rate) as MSM.

4.2 The summary effect estimates across spatial scales

The relationship between SVIDM and chronic conditions varied by spatial scale, strongly indicating that the results are sensitive to spatial confounding. The size and uncertainty associated with the coefficient estimates at all eigenvalues (spectral frequencies) are given in Figure 5. Focusing on the MSM model across outcomes and SVIDM themes, we note that the global spatial scale estimates—which correspond to smaller eigenvalues—differ markedly from local scale estimates with larger eigenvalues. In most treatment-outcome pairs, the global scale estimate of the effect coefficient (eigenvalues < 2) was biased away from the null. This implies

that the models that do not account for unmeasured confounding at larger spatial scales are biased away from the null. The effect of socioeconomic measures on kidney disease, hyperlipidemia, heart failure, and diabetes did not reduce to null (at larger eigenvalues) but were markedly reduced in comparison to the global effects at lower eigenvalues. In contrast, the effect of Theme 4 scores describing housing type and transportation for kidney disease was reduced to null at a global scale (smaller eigenvalues).

The discrepancy between local and global estimates can be quantified by the posterior probabilities of the mean effect at the lowest frequency being greater than those at the highest frequency (Table 4). Treatment-outcome pairs corresponding to Theme 1, 2 and 4 have very different estimates at the most local and most global spatial scales, strongly suggesting the presence of spatial confounding. In contrast, Theme 3 coefficient estimates for racial and ethnic minority measures from the most local and most global scales have the largest percentages of overlap and stand out as not having overall trends for the outcomes.

The multiscale coefficients also highlight how spatial scale of analysis could affect the result. As explained in Section S2.1, derived from the eigen decomposition of the adjacency matrix's precision matrix, eigenvalues correspond to spatial scales: the larger the eigenvalues, the more local the scales. Though eigenvalues cannot be uniquely translated into specific physical distances, a rough correspondence can be provided for illustrative purposes. Among the pairwise distances of ZCTAs, the 10^{th} percentile is around 283 kilometers, and the 90^{th} percentile is around 1638 kilometers. For context, the mean width of a state is around 400 kilometers. These numbers and the curves in Figure 5 reveal that very different conclusions can be drawn, depending on whether data come from the census tract, ZCTA, county, state, or region level. If the information only comes from the regional level, for example, no local information is available, and the results can be biased away from null.

4.3 Causal effect estimates

Despite the evidence of confounding at the global spatial scale, a number of outcome-treatment pairs have a significant causal effect at the local scale, indicated by the corresponding posterior probabilities of coefficients being different from zero are close to 0 or 1 (the pairs in red or blue in Table 5). Economic resilience (Theme 1) demonstrates a positive effect on diabetes, chronic kidney disease, and congestive heart failure but a negative effect on hyperlipidemia. Furthermore, vulnerable household characteristics (Theme 2), including age extremes, disability, and single-parent structures, has a 90% probability of having a positive effect on diabetes rates. Housing type and transportation (Theme 4) shows positive effects on chronic kidney disease and congestive heart failure. The effects are not universal, however, as other pairs have less extreme posterior probabilities.

To additionally demonstrate the utility of the multiscale model, we contrast MSM coefficients with those from the Univariate (USM), Naive, and OLS models. The multivariate (MSM) and univariate (USM) coefficient estimates were similar in magnitude across all outcomes, but the MSM model had advantageously smaller credible intervals, especially for local spatial scales/large eigenvalues (Figure 6).

Figure 5 shows how the global spatial scales can nudge the estimates towards biased for the naive model and OLS. Although MSM and naive models' credible intervals generally overlap, the naive model takes in all information and its credible intervals might be overly narrow, resulting in false positive results. The nudge toward bias is particularly evident for measures in Theme 2 (household composition and disability) where OLS estimates best compare to the global MSM estimates. As shown in Figure S3, the naive model has the lowest mean standard deviation, but the coverage rate can be lower than nominal levels, especially when spatial confounding is strong. Similarly, OLS takes in information from all locations, but unlike the naive model, it does not consider spatial locations. Its results also have smaller uncertainty and can be influenced by spatial confounding to a larger degree, which is particularly evident in plots for household composition and disability.

A forest plot in Figure 6 shows the coefficient estimates with 95% credible intervals for MSM, USM, naive model, and OLS. This is essentially the same as Figure 5, except the curves are sliced for the largest eigenvalues (local effects). For MSM, the intervals that do not cross the zero line are the same as those colored blue or red in Table 5. The estimates from the naive and OLS models show how adjusting for spatial relationships moves the estimates closer to null, but the naive model, nonetheless, can still be over-confident, as shown in Figure S3. As discussed above, MSM and USM have the largest uncertainty because the spectral model breaks down the estimates by spatial scales, and the estimates come from the most local spatial scale. The separation of signals filters out confounded information and calibrates uncertainty, as reflected in the wider intervals. Further, the comparison between MSM and USM highlights the efficiency gained by jointly modeling all exposures and outcomes. All credible intervals from MSM are narrower than those from USM, and significant associations—such as for theme 1 and CKD—can be overlooked due to the too-wide credible interval of USM.

The difference among these four models points to important implications for real world data. For epidemiological studies, data analysis—even when spatial information is available—univariate linear models are widely used. Figure 6 shows how much results from univariate OLS can deviate from those from spatial and spectral models. Further, the spatial "naive" method can be overly confident, when unmeasured confounding is present. The results thus reflect the danger of basing policies on potentially false positive results, if OLS and naive spatial models are the only ones used. Spectral models can account for unmeasured confounding present in global spatial scales; moreover, if no confounding exists, spectral methods will give results that agree with the naive model. In other words, the adjustment provided by MSM will not result in false conclusions if no confounding exists. The multivariate model also provides efficiency gain, especially for correlated outcomes, commonly seen in epidemiological studies.

4.4 Sensitivity analysis for bias

Our analysis relies on the unverifiable assumption of local unconfoundedness. Using methods motivated by Cinelli and Hazlett (2020), we conduct a sensitivity analysis to gauge the strength of a missing local confounder that would alter our conclusions. We assume there are $N_U = \max\{R, E\}$ latent confounders, and let ρ determine the strength of confounding. Specifically, we assume confounding in the spectral domain has the form

$$\mathbf{X}^* = \left(\sqrt{1-\rho^2}\mathbf{U}_X + \rho\mathbf{U}_C\right)\mathbf{B}_X^T$$
$$\mathbf{\Theta}^* = \left(\sqrt{1-\rho^2}\mathbf{U}_\Theta + \rho\mathbf{U}_C\right)\mathbf{B}_\Theta^T,$$

where $\mathbf{U}_X$, $\mathbf{U}_\Theta$, and $\mathbf{U}_C$ are mutually independent $S \times N_U$ matrices with independent standard normal elements. To preserve the cross-correlation between exposures and spatial random effects, $\mathbf{B}_X$ and $\mathbf{B}_\Theta$ are selected so that $\mathbf{B}_\Theta\mathbf{B}_\Theta^T$ and $\mathbf{B}_X\mathbf{B}_X^T$ match the sample covariance of the posterior mean of the spatial random effects, $\hat{\Theta}^*$, and observed exposure variables, $\mathbf{X}_0^*$, respectively, with columns of zeros added so that both $\mathbf{B}_X$ and $\mathbf{B}_\Theta$ have N_U columns. In this formulation, ρ is both the correlation between the latent confounder $\mathbf{U}_C\mathbf{B}_X^T$ and exposure $\mathbf{X}^*$ and the latent confounder $\mathbf{U}_C\mathbf{B}_\Theta^T$ and spatial random effect Θ^*. The resulting bias for exposure *e* and response *r* is

$$\alpha_{ier} = \frac{\text{Cov}(\Theta_{ir}^*, X_{ie}^*)}{\text{Var}(X_{ie}^*)} = \rho^2 \frac{(\mathbf{B}_\Theta^\top \mathbf{B}_X)_{re}}{(\mathbf{B}_X^\top \mathbf{B}_X)_{ee}}.$$

Derivations are provided in Section S12. To focus on the local spatial scale, we compute $\mathbf{B}_\Theta$ and $\mathbf{B}_X$ using only the $\hat{\Theta}$ and $\mathbf{X}_0^*$ for the one hundred most local terms.

Figure 7 shows how bias increases with ρ for four exposure-outcome combinations that have significant effect in MSM. The critical value of ρ is the value for which the bias equals the lower boundary of the credible intervals (dashed line) of the effect estimate, as shifting the interval by this value of the bias would place zero in the credible interval. For the combinations Theme 1 and CKD and Theme 4 and CHF, the effects remain significant only for $\rho < 0.069$ and $\rho < 0.117$, respectively. For Theme 1 and Hyperlipidemia and Theme 1 and Diabetes, the effects remain significant for greater values of ρ (0.507 and 0.368, respectively). As a benchmark to interpret these results, the absolute correlation over the largest one hundred eigenvalues between the observed confounders and exposures ranges between 0.02 and 0.20. Thus, the causal conclusions are robust for the associations between Theme 1 and Hyperlipidemia and Theme 1 and Diabetes but fragile for those between Theme 1 and CKD and Theme 4 and CHF.

4.5 Summarizing the tensor structure

Lastly, to answer our motivating question regarding the latent structure between variables, we use the tensor estimates to decompose the health effects. To provide a succinct representation of this structure, we use the `R` package `rTensor` (Li et al. 2018) to find a rank-2 canonical polyadic tensor decomposition of the posterior mean of the tensor $\hat{\boldsymbol{\gamma}}$ in (4). Each curve in Figure 8 represents a tensor margin (see Figure 3). For basis functions, term with index near one capture global features and indices near $L = 10$ capture local features. For tensor margins representing basis functions, the first factor (green) captures global trends. For tensor margins representing

exposures, the first factor is positive for Theme 1 (lack of economic resilience) and Theme 2 (household-level adaptive capacity) and negative for Theme 3 (fraction of minority) and Theme 4 (housing type and transportation). The green tensor margins for basis function and exposure together explain why the coefficient estimates for global scales are positive for Themes 1 and 2 and negative for Themes 3 and 4 (Figure 5). Factor 2 (purple) dominates for the local-scale trends of interest and Theme 1 has the highest loading and thus explains the strongest effects in Theme 1, shown in Figure 6. For this factor and thus the final estimates in Figure 6, the responses have similar loadings except for hyperlipidemia, which is also reflected in Figure 6. The tensor margins for outcomes indicate that the exposure effects for chronic kidney disease and congestive heart failure are largely similar, while the effects for other outcomes differ. The complexity in data with multiple exposures and multiple outcomes makes it difficult to parse complex structures. By using a tensor structure, MSM offers a way to construct interpretable latent structures, enabling understanding of the relationship among different dimensions.

5 Discussion

We introduce a multivariate spectral model (MSM) for areal data that estimates the effects of multiple exposures on multiple outcomes in the presence of spatial confounding. The model relaxes the no unmeasured confounder assumption and permits global confounding. As the simulation study demonstrates, when the conditions hold, the MSM method has low MSE and achieves conservative coverage rates that corrects over-confidence—reflected in the too-narrow credible intervals—from single-scale methods. Nonetheless, being a parametric model, the proposed MSM has its limitations in the pre-specified model. A semiparametric or nonparametric framework could provide more flexibility and robustness in the case that the model is misspecified. Further, unmeasured local confounding is unverifiable, and our proposed method does not account for spillover and interference. Spillover effects are possible if neighboring, economically vulnerable communities influence local health outcomes. For instance, residents may travel to adjacent areas to access public resources, thereby diminishing the supply available to local residents. We leave these limitations for future research. Finally, we note that our results and interpretations are limited to the Southeastern US and require external validation before generalizing.

Data analysis shows the associations between disaster resilience index in four themes and the incidence of five chronic diseases. The 95% credible intervals show that ZCTA areas lacking economic resilience (high scores for Theme 1) are more likely to have higher incident rates of chronic kidney disease and diabetes but lower rates of hyperlipidemia, while lacking disaster resilience in terms of housing type and transportation (high scores for Theme 4) are associated with higher incident rates of congestive heart failure. Besides being able to provide unbiased causal estimate effects with adjustment for unmeasured spatial confounders, MSM offers efficiency gain, provides interpretable findings, and encourages the consideration of the role of spatial scales in the effect estimates.

Funding

Research reported in this publication was partially supported by the National Heart, Lung, And Blood Institute of the National Institutes of Health under Award Number T32HL079896, National Institutes of Health grants R01ES031651-01 and R01ES036270-01A1, and National Science Foundation grant DMS2152887. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. This project was supported by an appointment (of S.E.C.) to the Research Participation Program at the Center for Public Health and Environmental Assessment, U.S. Environmental Protection Agency (EPA), administered by the Oak Ridge Institute for Science and Education (ORISE) through an interagency agreement between the U.S. Department of Energy and the U.S. EPA.

Disclosure statement

The authors have no conflicts of interest. Although this work has been reviewed for publication by the U.S. EPA, it does not necessarily reflect the views and policies of the agency. Mention of trade names or commercial products does not constitute endorsement or recommendation for use.

Data Availability Statement

All available data and code are provided in our public git repository: https://github.com/snprim/multivar_spectral_confounder.

References

Adin, A., Goicoa, T., Hodges, J. S., Schnell, P. M. and Ugarte, M. D. (2023), 'Alleviating Confounding in Spatio-Temporal Areal Models With an Application on Crimes Against Women in India', *Statistical Modelling* **23**(1), 9–30.

Banerjee, S., Carlin, B. P. and Gelfand, A. E. (2015), *Hierarchical Modeling and Analysis for Spatial Data*, CRC Press.

BiyikoÇ§u, T., Leydold, J. and Stadler, P. F. (2007), *Laplacian Eigenvectors of Graphs*, Springer Berlin, Heidelberg.

Bobb, J. F., Valeri, L., Claus Henn, B., Christiani, D. C., Wright, R. O., Mazumdar, M., Godleski, J. J. and Coull, B. A. (2015), 'Bayesian Kernel Machine Regression for Estimating the Health Effects of Multi-Pollutant Mixtures', *Biostatistics* **16**(3), 493–508.

Carrico, C., Gennings, C., Wheeler, D. C. and Factor-Litvak, P. (2015), 'Characterization of Weighted Quantile Sum Regression for Highly Correlated Data in a Risk Analysis Setting', *Journal of Agricultural, Biological, and Environmental Statistics* **20**(1), 100–120.

CDC/National Center for Health Statistics/Division of Analysis and Epidemiology (2024), 'Geographic division or region', https://www.cdc.gov/nchs/hus/sources-definitions/geographic-region.htm.

Centers for Disease Control and Prevention; Agency for Toxic Substances and Disease Registry; Geospatial Research, Analysis, and Services Program (2018), 'CDC/ATSDR Social Vulnerability Index 2018 Database U.S.', https://www.atsdr.cdc.gov/placeandhealth/svi/data_documentation_download.html.

Chen, H. Y., Li, H., Argos, M., Persky, V. W. and Turyk, M. E. (2022), 'Statistical Methods for Assessing the Explained Variation of a Health Outcome by a Mixture of Exposures', *International Journal of Environmental Research and Public Health* **19**(5), 2693.

Cinelli, C. and Hazlett, C. (2020), 'Making sense of sensitivity: Extending omitted variable bias', *Journal of the Royal Statistical Society Series B: Statistical Methodology* **82**(1), 39–67.

Devick, K. L., Bobb, J. F., Mazumdar, M., Claus Henn, B., Bellinger, D. C., Christiani, D. C., Wright, R. O., Williams, P. L., Coull, B. A. and Valeri, L. (2022), 'Bayesian Kernel Machine Regression-Causal Mediation Analysis', *Statistics in Medicine* **41**(5), 860–876.

Dupont, E., Marques, I. and Kneib, T. (2024), 'Demystifying Spatial Confounding', *arXiv:2309.16861* .

Dupont, E., Wood, S. N. and Augustin, N. H. (2022), 'Spatial+: A Novel Approach to Spatial Confounding', *Biometrics* **78**(4), 1279–1290.

Ferrari, F. and Dunson, D. B. (2020), 'Identifying Main Effects and Interactions Among Exposures Using Gaussian Processes', *The Annals of Applied Statistics* **14**(4), 1743–1758.

Flanagan, B. E., Gregory, E. W., Hallisey, E. J., Heitgerd, J. L. and Lewis, B. (2011), 'A Social Vulnerability Index for Disaster Management', *Journal of Homeland Security and Emergency Management* **8**(1).

Gelfand, A. and Banerjee, S. (2010), Multivariate Spatial Process Models, *in* A. Gelfand, P. Diggle, M. Fuentes and P. Guttorp, eds, 'Handbook of Spatial Statistics', CRC Press, pp. 495–515.

Gelfand, A. E., Banerjee, S. and Gamerman, D. (2005), 'Spatial Process Modelling for Univariate and Multivariate Dynamic Spatial Data', *Environmetrics* **16**(5), 465–479.

Gennings, C., Svensson, K., Wolk, A., Lindh, C., Kiviranta, H. and Bornehag, C.-G. (2022), 'Using Metrics of a Mixture Effect and Nutrition From an Observational Study for Consideration Towards Causal Inference', *International Journal of Environmental Research and Public Health* **19**(4), 2273.

Gibson, E. A., Rowland, S. T., Goldsmith, J., Paisley, J., Herbstman, J. B. and Kiourmourtzoglou, M.-A. (2021), 'Bayesian Non-Parametric Non-Negative Matrix Factorization for Pattern Identification in Environmental Mixtures', *arXiv:2109.12164* .

Gilbert, B., Datta, A., Casey, J. A. and Ogburn, E. L. (2024), 'A Causal Inference Framework for Spatial Confounding', *arXiv:2112.14946* .

Grebenkov, D. S. and Nguyen, B.-T. (2013), 'Geometrical Structure of Laplacian Eigenfunctions', *SIAM Review* **55**(4), 601–667.

Greco, S., Ishizaka, A., Tasiou, M. and Torrisi, G. (2019), 'On the Methodological Framework of Composite Indices: A Review of the Issues of Weighting, Aggregation, and Robustness', *Social Indicators Research* **141**, 61–94.

Guan, Y., Page, G. L., Reich, B. J., Ventrucci, M. and Yang, S. (2023), 'Spectral Adjustment for Spatial Confounding', *Biometrika* **110**(3), 699–719.

Guhaniyogi, R., Qamar, S. and Dunson, D. B. (2017), 'Bayesian Tensor Regression', *The Journal of Machine Learning Research* **18**(79), 1–31.

Hughes, J. and Haran, M. (2013), 'Dimension Reduction and Alleviation of Confounding for Spatial Generalized Linear Mixed Models', *Journal of the Royal Statistical Society Series B: Statistical Methodology* **75**(1), 139–159.

Joubert, B. R., Kioumourtzoglou, M.-A., Chamberlain, T., Chen, H. Y., Gennings, C., Turyk, M. E., Miranda, M. L., Webster, T. F., Ensor, K. B., Dunson, D. B. and Coull, B. A. (2022), 'Powering Research Through Innovative Methods for Mixtures in Epidemiology (PRIME)

Program: Novel and Expanded Statistical Methods', *International Journal of Environmental Research and Public Health* **19**(3), 1378.

Kang, S., Franks, A., Audirac, M., Braun, D. and Antonelli, J. (2025), 'Partial Identification and Unmeasured Confounding with Multiple Treatments and Multiple Outcomes', *arXiv:2311.12252* .

Keil, A. P., Buckley, J. P., O'Brien, K. M., Ferguson, K. K., Zhao, S. and White, A. J. (2020), 'A Quantile-Based g-Computation Approach to Addressing the Effects of Exposure Mixtures', *Environmental Health Perspectives* **128**(4), 47004.

Khan, K. and Calder, C. A. (2022), 'Restricted Spatial Regression Methods: Implications for Inference', *Journal of the American Statistical Association* **117**(537), 482–494.

Kolda, T. G. and Bader, B. W. (2009), 'Tensor Decompositions and Applications', *SIAM Review* **51**(3), 455–500.

Kowal, D. R., Bravo, M., Leong, H., Bui, A., Griffin, R. J., Ensor, K. B. and Miranda, M. L. (2021), 'Bayesian Variable Selection for Understanding Mixtures in Environmental Exposures', *Statistics in Medicine* **40**(22), 4850–4871.

Leroux, B. G., Lei, X. and Breslow, N. (2000), Estimation of Disease Rates in Small Areas: A New Mixed Model for Spatial Dependence, *in* M. E. Halloran and D. Berry, eds, 'Statistical Models in Epidemiology, the Environment, and Clinical Trials', Springer New York, New York, NY, pp. 179–191.

Li, J., Bien, J. and Wells, M. T. (2018), 'rTensor: An R package for multidimensional array (tensor) unfolding, multiplication, and decomposition', *Journal of Statistical Software* **87**(10), 1–31.

Llosa-Vite, C. and Maitra, R. (2023), 'Reduced-Rank Tensor-on-Tensor Regression and Tensor-Variate Analysis of Variance', *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(2), 2282–2296.

Mai, Q., Zhang, X., Pan, Y. and Deng, K. (2021), 'A Doubly Enhanced EM Algorithm for Model-Based Tensor Clustering', *Journal of the American Statistical Association* **117**(540), 2120–2134.

Makalic, E. and Schmidt, D. F. (2016), 'A Simple Sampler for the Horseshoe Estimator', *IEEE Signal Processing Letters* **23**(1), 179–182.

Manson, S., Schroeder, J., Riper, D. V., Knowles, K., Kugler, T., Roberts, F. and Ruggles, S. (2024), 'IPUMS National Historical Geographic Information System: Version 19.0 [dataset]'. Minneapolis, MN: IPUMS.

Marques, I., Kneib, T. and Klein, N. (2022), 'Mitigating Spatial Confounding by Explicitly Correlating Gaussian Random Fields', *Environmetrics* **33**(5), e2727.

Messer, L. C., Jagai, J. S., Rappazzo, K. M. and Lobdell, D. T. (2014), 'Construction of an Environmental Quality Index for Public Health Research', *Environmental Health* **13**.

Messer, L. C., Laraia, B. A., Kaufman, J. S., Eyster, J., Holzman, C., Culhane, J., Elo, I., Burke, J. G. and O'Campo, P. (2006), 'The Development of a Standardized Neighborhood Deprivation Index', *Journal of Urban Health* **83**, 1041–1062.

Miao, W., Hu, W., Ogburn, E. L. and Zhou, X.-H. (2023), 'Identifying Effects of Multiple Treatments in the Presence of Unmeasured Confounding', *Journal of the American Statistical Association* **118**(543), 1953–1967.

Paciorek, C. J. (2010), 'The Importance of Scale for Spatial-Confounding Bias and Precision of Spatial Regression Estimators', *Statistical Science* **25**(1), 107–125.

Page, G. L., Liu, Y., He, Z. and Sun, D. (2017), 'Estimation and Prediction in the Presence of Spatial Confounding for Spatial Linear Models', *Scandinavian Journal of Statistics* **44**(3), 780–797.

Reckien, D. (2018), 'What is in an Index? Construction Method, Data Metric, and Weighting Scheme Determine the Outcome of Composite Social Vulnerability Indices in New York City', *Regional Environmental Change* **18**, 1439–1451.

Reich, B. J., Guan, Y., Fourches, D., Warren, J. L., Sarnat, S. E. and Chang, H. H. (2020), 'Integrative Statistical Methods for Exposure Mixtures and Health', *The Annals of Applied Statistics* **14**(4), 1945–1963.

Reich, B. J., Hodges, J. S. and Zadnik, V. (2006), 'Effects of Residual Smoothing on the Posterior of the Fixed Effects in Disease-Mapping Models', *Biometrics* **62**(4), 1197–1206.

Reich, B. J., Yang, S., Guan, Y., Giffin, A. B., Miller, M. J. and Rappold, A. (2021), 'A Review of Spatial Causal Inference Methods for Environmental and Epidemiological Applications', *International Statistical Review* **89**(3), 605–634.

Roy, A., Lavine, I., Herring, A. H. and Dunson, D. B. (2021), 'Perturbed Factor Analysis: Accounting for Group Differences in Exposure Profiles', *The Annals of Applied Statistics* **15**(3), 1386–1404.

Schmidt, A. M. and Gelfand, A. E. (2003), 'A Bayesian Coregionalization Approach for Multivariate Pollutant Data', *Journal of Geophysical Research: Atmospheres* **108**(D24).

Schnell, P. and Papadogeorgou, G. (2020), 'Mitigating Unobserved Spatial Confounding When Estimating the Effect of Supermarket Access on Cardiovascular Disease Deaths', *The Annals of Applied Statistics* **14**(4), 2069–2095.

Thaden, H. and Kneib, T. (2018), 'Structural Equation Models for Dealing With Spatial Confounding', *The American Statistician* **72**(3), 239–252.

Urdangarin, A., Goicoa, T., Kneib, T. and Ugarte, M. D. (2024), 'A Simplified Spatial+ Approach to Mitigate Spatial Confounding in Multivariate Spatial Areal Models', *Spatial Statistics* **59**, 100804.

US Bureau of the Census (2023), 'American Community Survey 5-Year Data (2009-2021)', https://www.census.gov/data/developers/data-sets/acs-5year.html.

US Department of Health and Human Services (2023), 'Chronic Conditions Data Warehouse', https://www2.ccwdata.org/web/guest/condition-categories-chronic/.

Wackernagel, H. (2003), *Multivariate Geostatistics*, Springer Berlin, Heidelberg.

Wang, Y. and Blei, D. M. (2019), 'The Blessings of Multiple Causes', *Journal of the American Statistical Association* **114**(528), 1574–1596.

Wei, R., Reich, B. J., Hoppin, J. A. and Ghosal, S. (2020), 'Sparse Bayesian Additive Nonparametric Regression with Application to Health Effects of Pesticides Mixtures', *Statistica Sinica* **30**(1), 55–79.

Wiecha, N., Hoppin, J. A. and Reich, B. J. (2025), 'Two-stage estimators for spatial confounding with point-referenced data.', *Biometrics* **81**(3).

Xu, H., Li, R. and Bilal, U. (2025), *findSVI: Calculate Social Vulnerability Index for Communities*. R package version 0.2.0.

Zhang, A. and Han, R. (2019), 'Optimal Sparse Singular Value Decomposition for High-Dimensional High-Order Data', *Journal of the American Statistical Association* **114**(528), 1708–1725.

Zhang, A. and Xia, D. (2018), 'Tensor SVD: Statistical and Computational Limits', *IEEE Transactions on Information Theory* **64**(11), 7311–7338.

Zhang, Z. and Aeron, S. (2017), 'Exact Tensor Completion Using t-SVD', *IEEE Transactions on Signal Processing* **65**(6), 1511–1526.

Zheng, J., Wu, J., D'Amour, A. and Franks, A. (2023), 'Sensitivity to Unobserved Confounding in Studies with Factor-Structured Outcomes', *Journal of the American Statistical Association* **119**(547), 2026–2037.

Zigler, C. M. (2021), 'Invited Commentary: the Promise and Pitfalls of Causal Inference With Multivariate Environmental Exposures', *American Journal of Epidemiology* **190**(12), 2658–2661.

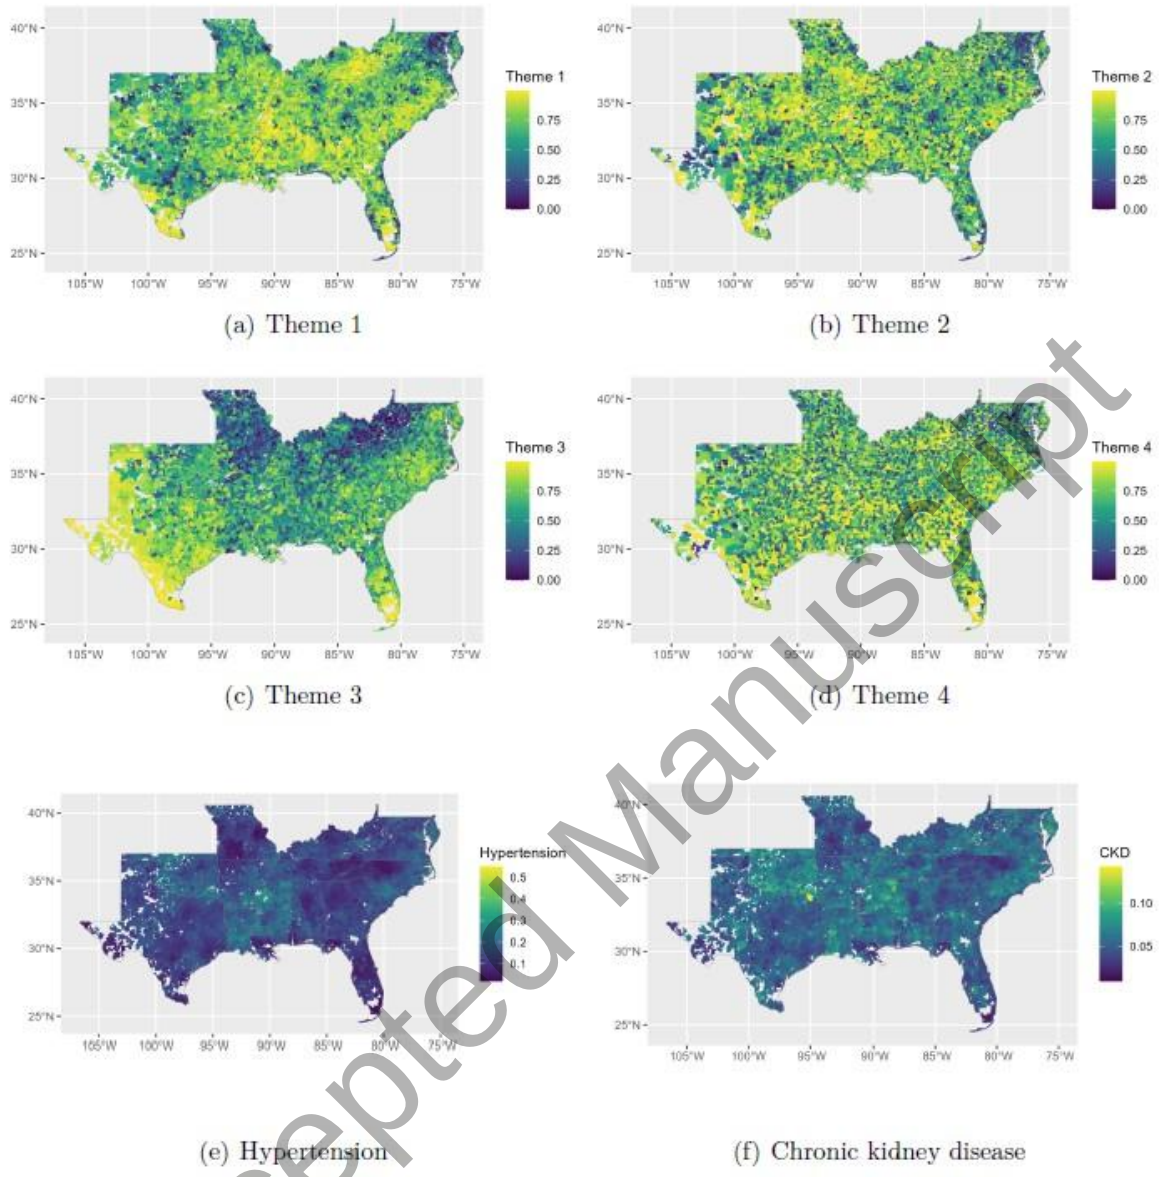

Figure 1: **Maps for exposures and outcomes**: Social Vulnerability Index for Disaster Management themes 1-4 in southern states are in (a)-(d). SVIDM ranges between 0 and 1, with the value indicating the proportion of ZCTAs with lower or equal social vulnerability. Two of the five chronic health outcomes in southern states are in (e) and (f). The maps show incident rates—the proportion of beneficiaries with the condition in a ZCTA.

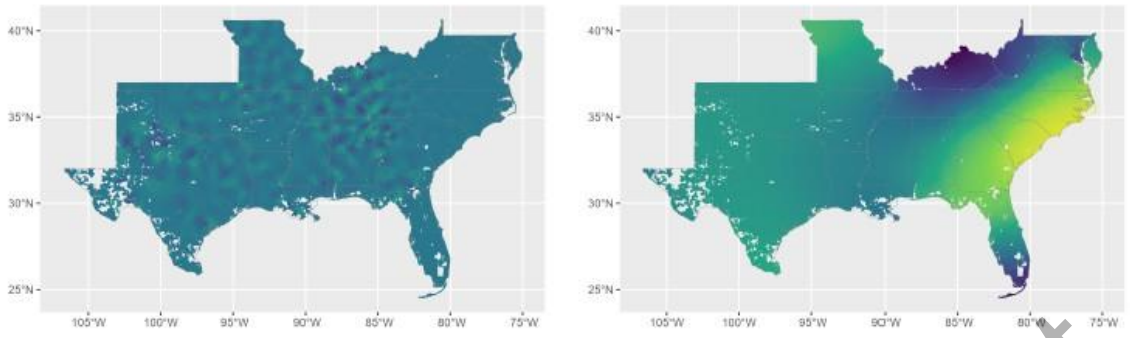

(a) Local scale (b) Global scale

Figure 2: **Spatial scales in the South region of the US:** The eigenvalues range from 0 to 23.12. Panel (a) shows a local spatial scale that corresponds to eigenvalue 1.78, and Panel (b) shows a global spatial scale that corresponds to eigenvalue 0.006.

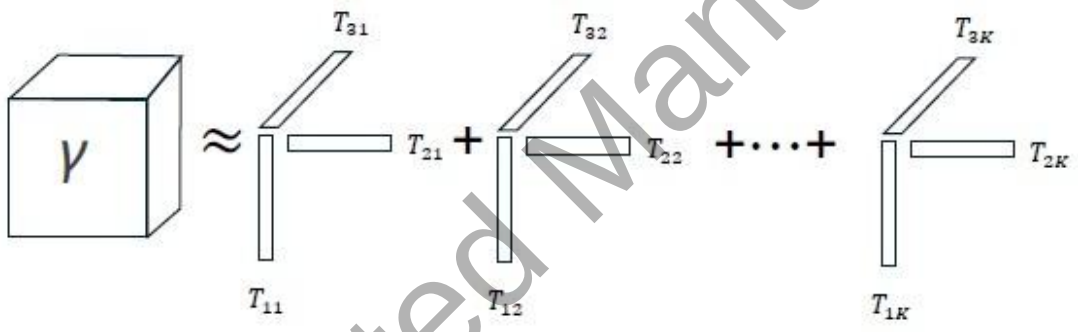

Figure 3: CP decomposition: $\mathbf{T}_1$, $\mathbf{T}_2$, and $\mathbf{T}_3$ are matrices, which are formed by columns that are tensor margins: T_{1k}, T_{2k}, T_{3k} for $k \in \{1, \cdots, K\}$. Here $T_{1k} = (T_{11k}, \cdots, T_{1Lk})^\top, T_{2k} = (T_{21k}, \cdots, T_{2Ek})^\top, T_{3k} = (T_{1rk}, \cdots, T_{1Rk})^\top$. The outer product of three tensor margins, each of which a vector of lengths *L*, *E*, and *R*, respectively, is a $L \times E \times R$ tensor. The sum of *K* such tensors approximates the tensor γ.

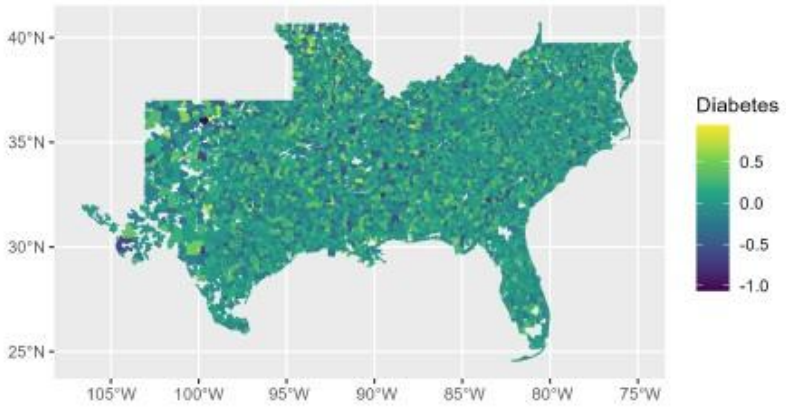

Figure 4: **Residual plot in the spatial domain for diabetes**: No apparent patterns and outliers are present.

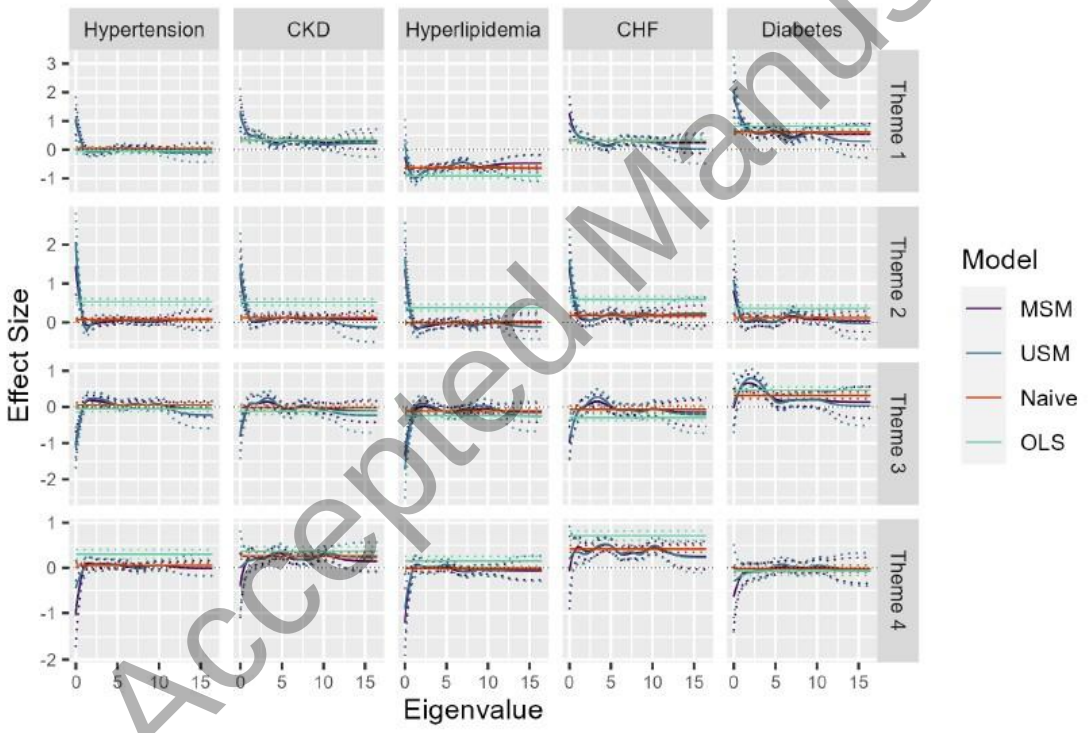

Figure 5: **Posterior estimates by spatial scale**: The posterior means for all frequencies and the 95% credible intervals are shown below. Larger eigenvalues correspond to local scales, and smaller eigenvalues global scales. The four themes for the disaster resilience index are economic resilience (Theme 1), household composition (Theme 2), fraction of minority (Theme 3), and housing type and transportation (Theme 4). For each Theme, higher scores indicate less resilience.

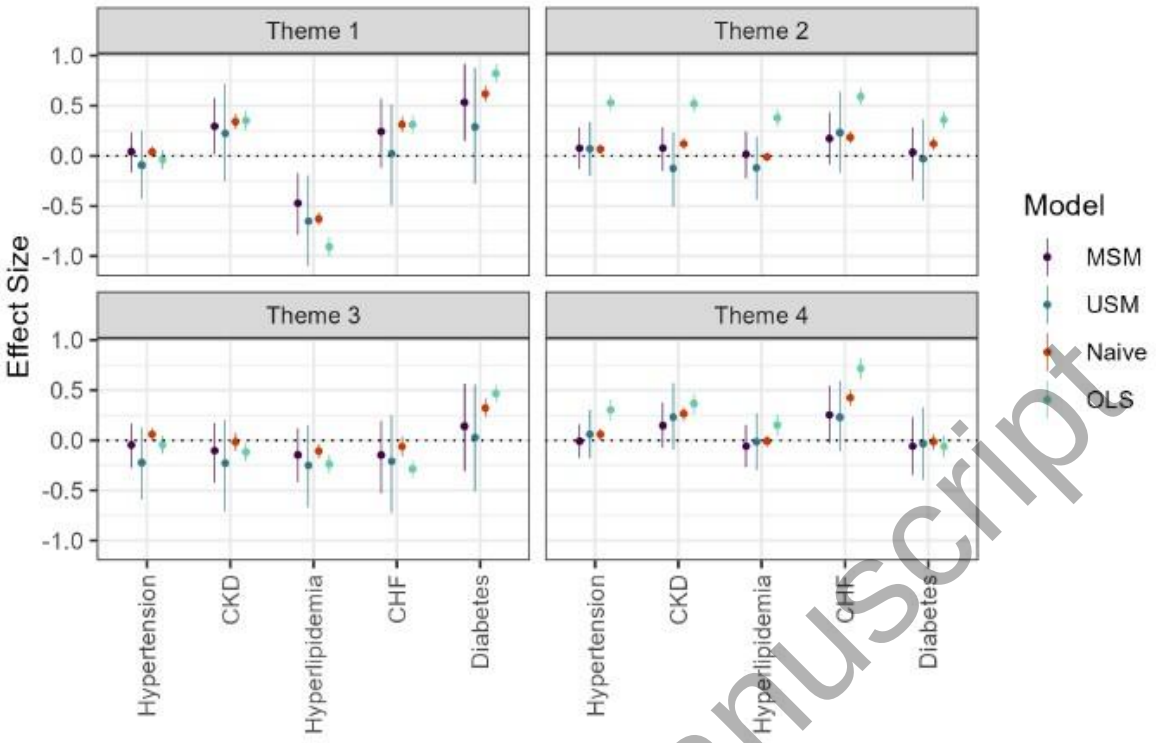

Figure 6: **Forest plot for estimated coefficients**: MSM is our multivariate spectral model, and USM is the univariate version. The naive model is spatial but does not allow coefficients to vary by scales. The OLS model is non-spatial and treats each location independently. MSM, more efficient than USM, has larger uncertainty than naive and OLS models but provides unbiased estimates, assuming local unconfoundedness.

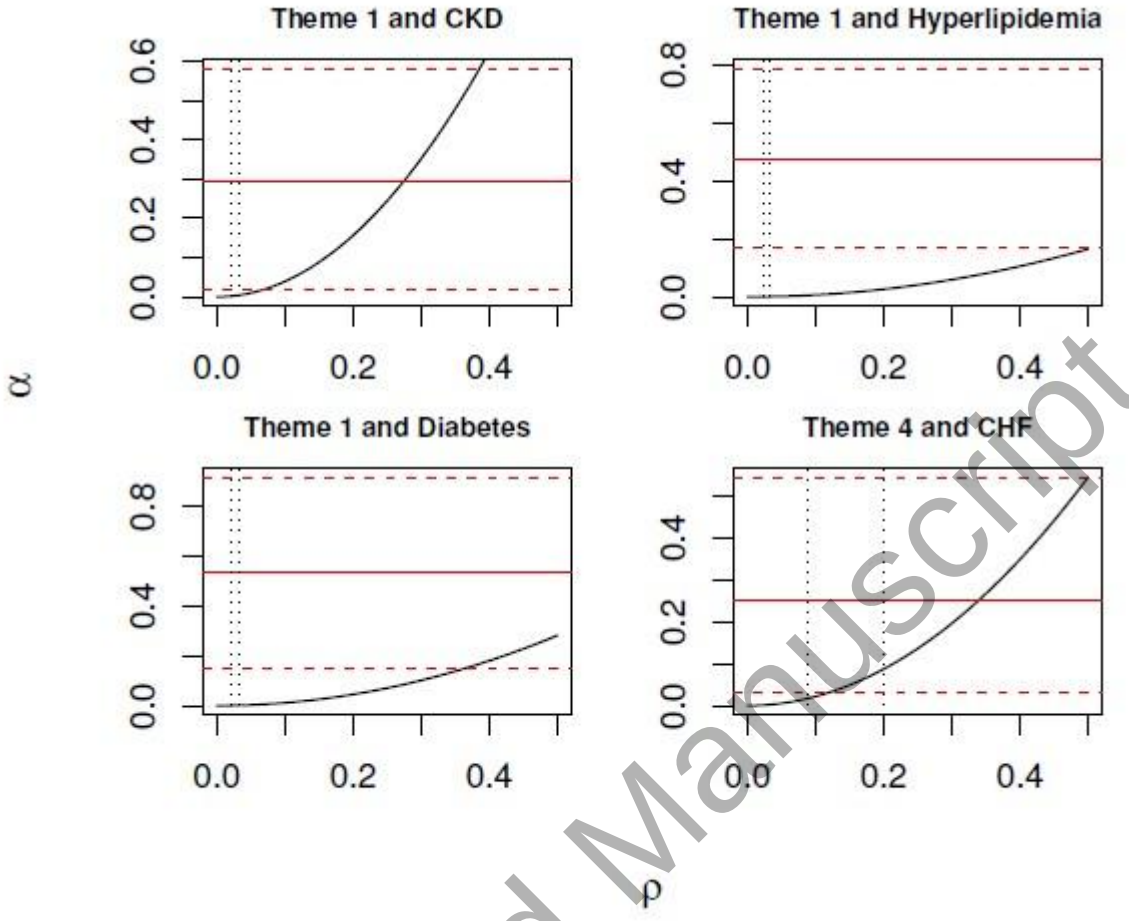

Figure 7: **Sensitivity analysis**: The red solid and dotted lines indicate the posterior means and 95% credible intervals, respectively, from MSM. The black curve is the absolute bias that would be induced by a missing local confounding variable plotted against its correlation ρ with the exposure and response, and the vertical lines indicate correlation between the exposure and two observed confounders from the one hundred most local spatial scales

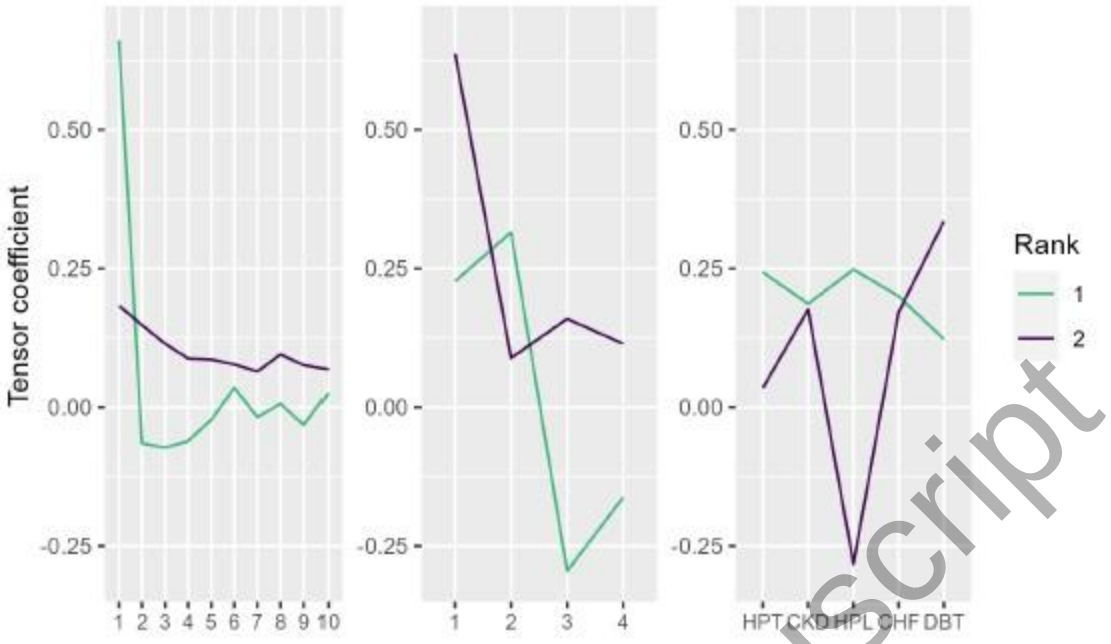

Figure 8: **Tensor estimates**: The curves show one rank-2 CP decomposition of the posterior means of the low-rank tensor that is of the dimension $L \times E \times R$, as presented in (4).For basis functions (left panel), the x-axis shows the ten basis functions that represent weights from the most global (1) to the most local (10). The x-axes for exposures (middle panel) and outcomes (right panel) include the four themes of the disaster resilience index and the five health conditions. Each curve represents a tensor margin and shows the relationship among basis functions, exposures, and outcomes.

Table 1: **Summary statistics of response variables**: The mean and range are from incident rates—the proportion of beneficiaries with the condition in a ZCTA.

			Correlation				
	Mean	Range	Hypertension	CKD	Hyperlipidemia	CHF	Diabetes
Hypertension	0.130	0.530	1.000	0.718	0.765	0.621	0.685
CKD	0.051	0.131	0.718	1.000	0.584	0.625	0.648
Hyperlipidemia	0.110	0.394	0.765	0.584	1.000	0.425	0.527
CHF	0.028	0.087	0.621	0.625	0.425	1.000	0.498
Diabetes	0.032	0.137	0.685	0.648	0.527	0.498	1.000

Table 2: Cross correlations of SVIDM

	Theme 1	Theme 2	Theme 3	Theme 4
Theme 1	1.000	0.538	0.233	0.404
Theme 2	0.538	1.000	0.019	0.249
Theme 3	0.233	0.019	1.000	0.588
Theme 4	0.404	0.249	0.588	1.000

Table 3: Predictive validity assessment with 80% training data and 20% test data: The predictive R^2 shows the squared correlation between predictive and true responses, and coverage shows the percentage of observed responses in the 95% credible intervals.

Model	Predictive R^2	Coverage
MSM ($K=5, L=5$)	0.216	0.950
MSM ($K=5, L=10$)	0.248	0.950
MSM ($K=10, L=5$)	0.154	0.951
MSM ($K=10, L=10$)	0.278	0.951
Naive ($L=1$)	0.065	0.950

Table 4: **Probabilities of coefficient estimates at the lowest frequency greater than those at the highest frequency**: Values close to one (in red) or zero (in blue) mean that coefficient estimates vary by spatial scales, which suggests likely spatial confounding. As shown below, strong spatial confounding is suspected for most exposure/outcome pairs.

	Hypertension	CKD	Hyperlipidemia	CHF	Diabetes
Theme 1	1.00	1.00	0.68	1.00	1.00
Theme 2	1.00	1.00	1.00	1.00	0.99
Theme 3	0.00	0.00	0.00	0.00	0.39
Theme 4	0.08	0.12	0.01	0.22	0.13

Table 5: Posterior probability of effect size being different from zero: Positive coefficients suggest that higher scores are associated with higher incidence of chronic health conditions (worse health outcomes), and negative coefficients suggest otherwise. Posterior probability over 0.90 (in red) and under 0.10 (in blue) suggest strong evidence for causal effect to be different from zero; positive or negative effect respectively.

	Hypertension	CKD	Hyperlipidemia	CHF	Diabetes
Theme 1	0.63	0.97	0.00	0.90	1.00
Theme 2	0.73	0.76	0.54	0.90	0.61
Theme 3	0.34	0.26	0.11	0.23	0.77
Theme 4	0.56	0.92	0.35	0.96	0.36