

MITIGATING UNOBSERVED SPATIAL CONFOUNDING WHEN ESTIMATING THE EFFECT OF SUPERMARKET ACCESS ON CARDIOVASCULAR DISEASE DEATHS

BY PATRICK M. SCHNELL¹ AND GEORGIA PAPADOGEOUR²

¹*Division of Biostatistics, College of Public Health, The Ohio State University, schnell.31@osu.edu*

²*Department of Statistics, University of Florida, gpapadogeorgou@ufl.edu*

Confounding by unmeasured spatial variables has received some attention in the spatial statistics and causal inference literatures, but concepts and approaches have remained largely separated. In this paper we aim to bridge these distinct strands of statistics by considering unmeasured spatial confounding within a causal inference framework and estimating effects using outcome regression tools popular within the spatial literature. First, we show how using spatially correlated random effects in the outcome model, an approach common among spatial statisticians, does not necessarily mitigate bias due to spatial confounding, a previously published but not universally known result. Motivated by the bias term of commonly-used estimators, we propose an affine estimator which addresses this deficiency. We discuss how unbiased estimation of causal parameters in the presence of unmeasured spatial confounding can only be achieved under an untestable set of assumptions which will often be application specific. We provide a set of assumptions which describe how the exposure and outcome of interest relate to the unmeasured variables, and we show that this set of assumptions is sufficient for identification of the causal effect based on the observed data when spatial dependencies can be represented by a ring graph. We implement our method using a fully Bayesian approach applicable to any type of outcome variable. This work is motivated by and used to estimate the effect of county-level limited access to supermarkets on the rate of cardiovascular disease deaths in the elderly across the whole continental United States. Even though standard approaches return null or protective effects, our approach uncovers evidence of unobserved spatial confounding and indicates that limited supermarket access has a harmful effect on cardiovascular mortality.

1. Introduction. Over 17 million deaths were attributed to cardiovascular disease (CVD) worldwide in 2016, and the prevalence of CVD among people aged 20 or older in the United States that same year was 48% ([American Heart Association \(2019\)](#)). Poor nutrition and high body mass index are major risk factors of CVD ([American Heart Association \(2019\)](#)), and there is evidence that these risk factors are influenced by the availability of nearby supermarkets ([Laraia et al. \(2004\)](#), [Powell et al. \(2007\)](#)) which, historically, have had a higher prevalence of heart-healthy foods ([Pearce et al. \(2008\)](#), [Sallis et al. \(1986\)](#)). In a prospective cohort study of individuals who had undergone cardiac catheterization, living in an area (census tract) with low income and poor food access was associated with an increased risk of myocardial infarction or death ([Kelli et al. \(2019\)](#)).

Here, our goal is to cast this question within a causal inference framework and quantify the effect of county-level supermarket availability on the rate of CVD mortality among the elderly (65-years-old and older) in the United States. For that purpose we compile a data set including mortality, store and demographic information for all counties in the continental

Received June 2019; revised May 2020.

Key words and phrases. Causal inference, cardiovascular disease, food access, Markov random field, spatial confounding, unmeasured confounding.

United States (see also Section 2). Demographic characteristics, such as income, have been previously associated with CVD risk factors (Kelli et al. (2017)) and might be confounders of the effect of supermarket proximity and access on CVD mortality. Even though a number of demographic variables are included in our data set, the relationship of interest is possibly confounded by unobserved or difficult-to-define variables, such as regional culture relating to personal vehicles, diet and general health consciousness or state-level support for people with disabilities. Such variables could represent predictors of the exposure which influence where people live, whether or not they own a vehicle or where businesses choose to locate and could also represent predictors of the outcome such as how much people exercise, their stress levels or what food they choose to eat regardless of supermarket availability. Furthermore, these potentially unobserved demographic variables are expected to be spatially correlated, in that nearby counties are expected to have similar values of these unmeasured covariates.

Therefore, we are faced with estimating the causal effect of an exposure on an outcome using spatially-referenced, observational data and under the threat of unmeasured spatial confounding. To address this challenge, we combine spatial statistics tools and causal inference methodology within a common framework. Even though some attention has been given to causal inference topics in the spatial statistics literature (Hanks et al. (2015), Hodges and Reich (2010), Hughes and Haran (2013), Paciorek (2010)) and to spatial topics in the causal inference literature (Keele, Titiunik and Zubizarreta (2015), Papadogeorgou, Choirat and Zigler (2019), Verbitsky-Savitz and Raudenbush (2012)), there is a substantial gap in the intersection of the two fields.

In classic spatial statistics, regression models are often augmented to include spatially correlated random effects in order to “account” or “adjust” for the spatial dependence in the outcome model residuals. However, there is substantial confusion about what exactly these spatial models are capable of accounting for (Hanks et al. (2015)). In some settings, spatial mixed models are employed to estimate the relationship between an exposure and outcome without conditioning on spatial information. In this context, Hodges and Reich (2010) and Hughes and Haran (2013) proposed including a spatial random effect that is orthogonal to the exposure of interest. Other times, it is asserted that spatial models adjust for unobserved covariates which have a spatial dependence structure (Congdon (2013), Lee and Sarran (2015)). Nevertheless, the usual spatial models do not, in general, eliminate bias due to unobserved confounders, even when the residual variance components are known (Paciorek (2010)). Recently, Keller and Szpiro (2019) discuss the interpretation of estimates from regression models that progressively include spatial basis functions of higher complexity, and they conclude that increasing adjustment might even lead to bias amplification.

From a different perspective, causal inference methodology with spatial data and in the presence of unmeasured spatial confounding has been quite limited, and, to our knowledge, it has been restricted to classic causal inference tools. Within a regression discontinuity framework, Keele, Titiunik and Zubizarreta (2015) match treated to control units separated by a boundary minimizing geographical distance of matched pairs and balancing observed covariates. Relatedly, Papadogeorgou, Choirat and Zigler (2019) proposed matching treated to control units on a criterion incorporating both propensity scores and geographical distance. Although these approaches can, in some cases, address the problem of interest to spatial statisticians, they are not immediately compatible with models commonly used in spatial data analysis which are most often grounded in outcome regression. An exception is found in Thaden and Kneib (2018) where the authors propose a structural equation modeling approach treating the spatial variable as a confounder in a geoadditive model in order to eliminate bias from the unmeasured spatial variable.

In the causal inference literature, unmeasured confounding has been most often dealt with in the realm of sensitivity analysis. Sensitivity analysis is a powerful approach which aims to

quantify the robustness of estimated effects to different amounts of unmeasured confounding (Imbens and Rubin (2015), Rosenbaum (2002), Rosenbaum and Rubin (1983a), VanderWeele and Ding (2017)). However, sensitivity analysis does not directly adjust effect estimates for the presence of such confounders which is what the methodology presented here and the works referenced above aim to achieve.

In this paper we seek to bolster the bridge between spatial data analysis and causal inference. In order to do so, we consider unmeasured confounding within a formal causal inference framework and examine estimation approaches grounded on models and tools often employed by spatial statisticians. We start by focusing on continuous outcomes and linear models, studying the bias of commonly-used estimators. We propose a model-based approach to estimate the effect of a change in the exposure on an outcome of interest in the presence of unmeasured spatial confounding. Our approach is designed to easily incorporate popular tools in spatial statistics, such as hierarchical and linear mixed models. We explain that identification and estimation of the causal parameter in the presence of unmeasured confounding requires untestable assumptions regarding the unmeasured confounders and their relationship with the treatment and outcome of interest. In general, these assumptions have to be application specific, identification of the causal parameter needs to be evaluated separately for each set of assumptions and the proposed estimator would have to be adapted to alternative identifying assumptions. For continuous treatments (referred to as *exposures*), we provide one set of assumptions that is sufficient for identification of the causal exposure-response curve and one that is not. Importantly, our results illustrate that, when spatial dependencies can be represented using a ring graph, the components in our estimator involving the unmeasured confounder can be identified based solely on spatial dependencies in the observed data. Within the context of our study, we extend our approach to noncontinuous outcomes and generalized linear mixed models employing a fully-Bayesian approach, and we carefully discuss the plausibility of the causal assumptions. While our development is in the context of areal data, refinements in the context of point-referenced data are possible and are discussed where applicable.

In Section 2 we present a detailed description of our data set and present preliminary analyses using nonspatial and commonly used spatial regression models that yield suspect results. In Section 3 we define the causal estimand in terms of potential outcomes for continuous exposures and discuss commonly invoked identifiability assumptions when the observed covariates include a sufficient confounding adjustment set. The proposed methodology is introduced in Section 4 within the context of linear models. There, we rederive the result by Paciorek (2010) stating that commonly used spatial regression models do not recover the estimands of interest in the presence of unmeasured spatial confounding. Motivated by the bias of commonly used estimators, we propose the affine estimator, and we provide a set of assumptions relating the exposure and outcome to the unmeasured variables based on which the causal parameter is identifiable from observed data. In Section 5 we extend the affine estimator in the context of nonlinear models and suggest using a Bayesian approach. The estimator is compared to the currently used estimators under various generative mechanisms via simulation in Section 6. In Section 7 we discuss the plausibility of our assumptions within the context of our study, and we use the affine estimator to estimate the county-level effect of poor supermarket availability on CVD mortality. Our study illustrates the potential of the affine estimator in mitigating bias from unmeasured spatial confounders, returning effect estimates that are qualitatively different from the ones in Section 2 and more in line with subject-matter knowledge. We conclude with a discussion in Section 8.

2. County-level supermarket availability and CVD mortality. We compile a data set including mortality, store availability, demographic and behavioral data on 3093 out of 3109




FIG. 1. Percent of households with no vehicle and more than one mile from a supermarket or large grocery store (left, % HHNV1MI) and observed relative risk of CVD deaths in the 65+ age range (right, relative risk).

counties and county equivalents in the 48 contiguous states and the District of Columbia. For each areal unit (county or county equivalent), supermarket availability is defined as the proportion of housing units during 2006 that are more than one mile from the nearest supermarket or large grocery store and do not have a car, obtained from the Food Environment Atlas, June 2012 release (United States Department of Agriculture Economic Research Service (2012)). County-level population and cardiovascular disease mortality count (ICD-10 codes I00–I99) during 2007 for residents 65-years-old and over were obtained from the United States Centers for Disease Control and Prevention (CDC) WONDER query system (Centers for Disease Control (2018)). Due to privacy constraints, county death counts below 10 are censored by CDC WONDER. Figure 1 displays the exposure and crude outcome relative risk without covariate-assisted estimation or smoothing. Demographic information was acquired from the 2000 Census and includes, among others, information on urbanicity, poverty and population mobility. Covariate information also includes estimates of 2006 smoking rates derived from CDC Behavioral Risk Factor Surveillance System data (Dwyer-Lindgren et al. (2014)). Appendix A provides additional information on the data sources, data collection, and data processing pipeline, including links to the publicly-available data sets and a table including names and descriptive statistics of available covariates.

At this point we consider two common analyses investigating the relationship between limited county-level access to supermarkets on cardiovascular deaths in the elderly. For the first analysis we model CVD mortality counts as a Poisson-distributed outcome using a log link with the exposure and all covariates listed in Table A.2 as predictors. Internal standardization was implemented by using log expected death count as the offset: the population age 65+ in each county was multiplied by the overall crude rate in the same age range. All covariates were standardized, and the model was fit within the Bayesian paradigm under Gaussian priors with mean 0 and standard deviation 10 on regression coefficients. Results are presented on original scales unless otherwise noted. Samples from the posterior distribution were obtained via a Gibbs sampler with a Metropolis–Hastings block update for all regression coefficients. The Gibbs sampler was run for 10,000 iterations after a 1000-iteration burn-in. Censored outcomes were imputed subject to the known privacy constraint.

The second analysis we implemented is a common analysis method for areal spatial data. We included a spatially correlated random effect for county, U , in the linear predictor and assumed it follows a conditional autoregressive (CAR) structure (Besag (1974)), that is,

$$(1) \quad U_i | U_{-i} \sim \mathcal{N} \left[\varphi_U \sum_{j \in \partial_i} U_j / |\partial_i|, (\tau_U |\partial_i|)^{-1} \right],$$

where ∂_i and $|\partial_i|$ are the set and number of i 's neighbors, respectively. A multivariate normal representation of the distribution of \mathbf{U} is then available as

$$(2) \quad \mathbf{U} \sim \mathcal{N}[\mathbf{0}, \tau_U^{-1}(\mathbf{D} - \varphi_U \mathbf{W})^{-1}],$$

where $w_{ij} = 1$ if i and j are neighbors, and 0 otherwise and \mathbf{D} is diagonal with entries $|\partial_i|$ (Banerjee, Carlin and Gelfand (2004)). A uniform prior on $(-1, 1)$ was used for φ_U , and a gamma prior with shape and rate parameters equal to 5 was used for τ_U , jointly restricted to require the precision matrix of \mathbf{U} to be positive definite.

Based on these two models, we acquired what the estimated effect of a one percentage point increase in households with poor supermarket access on cardiovascular mortality would be if each model was specified correctly and was sufficiently adjusted for confounding variables. The nonspatial estimate (analysis 1) indicated that increasing poor supermarket availability is protective of CVD deaths with an estimated relative risk of 0.968 (95% CI 0.962 to 0.973). A protective effect of poor supermarket access is not consistent with either theoretical or data-driven understanding of the phenomenon, enforcing our belief that the result is at least partially due to unobserved or poorly adjusted-for confounders. The estimate from the spatial model (analysis 2) is effectively null with an estimated relative risk of 0.999 (95% CI 0.988 to 1.011). Although the spatial estimate differs from the nonspatial estimate in both location and credible interval width, we show in the next section that the spatial estimate does not necessarily mitigate bias from unobserved confounders.

3. Causal estimands and classic identifiability assumptions. Broadly speaking, the causal inference literature places substantial emphasis on defining target quantities of interest, referred to as *estimands*, and determining sufficient assumptions under which such estimands (which include unobservable quantities) are identifiable based on the observed data. We begin by defining estimands of interest following the potential outcome framework formalized by Rubin (1974) and extended to continuous exposures by Hirano and Imbens (2004). A necessary condition for an exposure Z to have an effect on an outcome Y is that Z is temporally precedent. We make the stable unit treatment value assumption (SUTVA, Rubin (1980)) which states that there is a single version of each treatment level and there is no interference between units. Based on SUTVA, we can use $Y_i(z)$ to represent the value that would have been observed at location i had it received exposure $z \in \mathcal{Z}$, where \mathcal{Z} includes all possible values of the continuous Z , and $i = 1, 2, \dots, n$. Then, $Y_i(z)$ is the *potential outcome* for location i at exposure level z , and unit i 's observed outcome Y_i is the potential outcome for the observed level of the treatment, $Y_i = Y_i(Z_i)$.

The most common estimands for continuous treatments are the *population average exposure-response curve* (PAERC) $\mu(z) = E[Y_i(z)]$, $z \in \mathcal{Z}$ and the expected rate of change in the outcome for an infinitesimal change in the exposure around z , $\mu'(z)$. Since $\mu(z)$ represents the average outcome value over the whole population had *everyone* experienced exposure z , it is clear that $\mu(z)$ includes unobserved quantities, and assumptions need to be made to ensure identifiability and to estimate it using data. The positivity and no unmeasured confounding assumptions (referred to together as the ignorability assumption) form a sufficient set of assumptions for identifiability of $\mu(z)$. Positivity states that all units can experience any $z \in \mathcal{Z}$, and the no unmeasured confounding assumption states that there exist *measured* covariates \mathbf{C} , which satisfy that, conditional on \mathbf{C} , the observed exposure Z is independent of the potential outcomes $Y(z)$, denoted as $Z \perp\!\!\!\perp Y(z)|\mathbf{C}$, $z \in \mathcal{Z}$. (See Appendix B for a discussion on identifiability of $\mu(z)$ based on these assumptions.)

Confounders \mathbf{C} are generally thought of as temporally precedent to the exposure Z and as common predictors of Z and Y , as shown in Figure 2. Since temporal order of variables matters in drawing causal conclusions, observed data are conceived as if generated in the




FIG. 2. Assumed causal diagram for the generative model. The vector C may represent a collection of multiple confounders.

following order: $[C]$, $[Z|C]$ and $[Y|Z, C]$. If the identifiability conditions of positivity and no unmeasured confounding are met in the observed data, estimation can proceed by imitating the data-generating mechanism for the exposure $Z|C$, known as the propensity score (Rosenbaum and Rubin (1983b)), the data-generating mechanism for the outcome $Y|Z, C$ (e.g., Hahn et al. (2018), Hill (2011)) or both (Hahn (1998), Robins and Rotnitzky (1995), Zhou, Elliott and Little (2019), Zigler and Dominici (2014)). In order to adhere to common approaches of spatial statistics, which emphasize analytical models imitating the outcome generative model, our primary focus is on modeling $Y|Z, C$.

Even though confounding adjustment is necessary to draw causal conclusions, C might include components that are not measured, hence violating the no unmeasured confounding assumption. Denote $C = (C^m, C^u)$, representing the measured and unmeasured components, respectively. At this point we assume that at least some of the variables in C^u are spatial and refer to Section 4.5 for a further discussion on this requirement. We refer to a variable as “spatial” if the correlation of the variable for two observations depends on their geographic locations. For areal data, like the ones in our study, this could refer to adjacency of counties. For point referenced data it could refer to the geographical distance of two points.

In this section and the next we discuss unobserved spatial confounders in the case of continuous outcomes and linear models. Focusing on this setting allows for straightforward application of theory from least squares estimation of regression coefficients and restricted maximum likelihood estimation of variance parameters. We return to non-Gaussian outcomes for the simulation study and data analysis, where we employ a fully Bayesian approach. Assume here that potential outcomes arise in the following manner:

$$(3) \quad Y_i(z) = \eta(z, C_i^m) + g(C_i^u) + \varepsilon_i(z),$$

for some function η and $\varepsilon_i(z)$ a mean zero random variable and independent of C . In (3), C^u is assumed to not interact with Z and C^m . We denote $U = g(C^u)$, representing the cumulative contribution of all unobserved covariates. Since at least some components of C^u are spatial, U also has a spatial correlation structure. Without loss of generality, we may assume $E[g(C^u)] = 0$ by absorbing any nonzero mean into the the intercept in $\eta(z, C^m)$.

4. The affine estimator in linear models: Addressing omitted variable bias of classic estimators. In this section we discuss how the classic approaches to estimation within the spatial statistics literature are biased for estimating causal quantities in the presence of unmeasured confounders and in the context of linear regression. The bias results derived below are in line with results in Paciorek (2010), and they motivate the affine estimator which is designed to explicitly remove the bias of the existing estimators. An extension to nonlinear models within the Bayesian framework is presented in Section 5.

4.1. *Omitted variable bias of ordinary and generalized least squares estimators.* Let $\mathbf{X} = (\mathbf{1}|Z|C^m)$ be the design matrix containing an intercept, the exposure Z and measured covariates, and let $\mathbf{X}_{(-z)}$ be the design matrix, including an intercept and measured confounders but not the exposure Z . For simplicity of presentation, we assume that the causal exposure response curve is linear, $\eta(z, C_i^m) = \beta_{(-z)}^\top \mathbf{x}_{i,(-z)} + \beta_z z$, which will be relaxed in Section 4.4.2. This simplification implies that $\mu'(z) = \beta_z$, corresponding to the usual linear

regression coefficient targeted in the spatial statistics literature. Using vector notation, let $\mathbf{Y} = (Y_1, Y_2, \dots, Y_n)^\top$, with \mathbf{Z} , \mathbf{U} , $\boldsymbol{\varepsilon}$ defined analogously. Then, if all of \mathbf{Y} , \mathbf{Z} , \mathbf{U} and $\mathbf{X}_{(-z)}$ were observed, estimation of the regression model

$$(4) \quad \mathbf{Y} = \mathbf{X}_{(-z)}\boldsymbol{\beta}_{(-z)} + \mathbf{Z}\beta_z + \mathbf{U} + \boldsymbol{\varepsilon}$$

would lead to consistent estimation of the causal effect through estimation of β_z .

However, the above model cannot be directly used in settings where \mathbf{U} is not measured. In (4), \mathbf{Z} and \mathbf{U} are correlated, but $\boldsymbol{\varepsilon}$ is independent of (\mathbf{Z}, \mathbf{U}) . Thus, \mathbf{U} and $\boldsymbol{\varepsilon}$ is a partition of the variability in \mathbf{Y} not due to \mathbf{Z} into one component ($\boldsymbol{\varepsilon}$) which is independent of \mathbf{Z} and one (\mathbf{U}) which is not. If \mathbf{U} is correlated with \mathbf{Z} and is omitted from the outcome regression, the *ordinary least squares* (OLS) estimator of $\boldsymbol{\beta} = (\boldsymbol{\beta}_{(-z)}^\top, \beta_z)$, $\hat{\boldsymbol{\beta}}$, will be biased. This is evident by examining the conditional expectation of $\hat{\boldsymbol{\beta}}$,

$$(5) \quad E(\hat{\boldsymbol{\beta}}|\mathbf{X}) = E[(\mathbf{X}^\top\mathbf{X})^{-1}\mathbf{X}^\top\mathbf{Y}|\mathbf{X}] = \boldsymbol{\beta} + (\mathbf{X}^\top\mathbf{X})^{-1}\mathbf{X}^\top E(\mathbf{U}|\mathbf{X}).$$

Considering $E(\hat{\boldsymbol{\beta}}) = E[E(\hat{\boldsymbol{\beta}}|\mathbf{X})]$, we see that $\hat{\boldsymbol{\beta}}$ will be biased for $\boldsymbol{\beta}$ since the second term will, in general, be nonzero for correlated \mathbf{U} , \mathbf{Z} .

When \mathbf{U} is omitted from the regression model, the component of \mathbf{U} not attributed to \mathbf{X} will be incorporated in the residuals. Since \mathbf{U} is spatially structured, residuals of the regression of \mathbf{Y} on \mathbf{X} will also be spatially correlated. In an effort to account for residual spatial correlation, spatial linear mixed models are often adopted. Typically, such models represent mechanisms similar in form to (4), but in which all right-hand-side variables are *assumed* to be independent, and some assumptions are made about the form of $\text{Var}[\mathbf{U}] = \text{Var}[\mathbf{U}|\mathbf{X}]$. These models aim to explain the spatial correlation in the residuals, and they are often effective at improving efficiency. However, they do not necessarily alleviate the omitted variable bias (Paciorek (2010)). If $\text{Var}[\mathbf{Y}|\mathbf{X}]$ (which depends on $\text{Var}[\mathbf{U}]$) is known, the *generalized least squares* (GLS) estimator of $\boldsymbol{\beta}$ is

$$(6) \quad \tilde{\boldsymbol{\beta}} = \{\mathbf{X}^\top(\text{Var}[\mathbf{Y}|\mathbf{X}])^{-1}\mathbf{X}\}^{-1}\mathbf{X}^\top(\text{Var}[\mathbf{Y}|\mathbf{X}])^{-1}\mathbf{Y},$$

with conditional expected value

$$(7) \quad E(\tilde{\boldsymbol{\beta}}|\mathbf{X}) = \boldsymbol{\beta} + \{\mathbf{X}^\top(\text{Var}[\mathbf{Y}|\mathbf{Z}])^{-1}\mathbf{X}\}^{-1}\mathbf{X}^\top(\text{Var}[\mathbf{Y}|\mathbf{X}])^{-1}E[\mathbf{U}|\mathbf{X}].$$

Therefore, even if $\text{Var}[\mathbf{Y}|\mathbf{X}]$ is known, $\tilde{\boldsymbol{\beta}}$ remains biased. This result indicates that including a spatial random effect in the regression model does not necessarily mitigate or eliminate bias arising from unmeasured spatial confounders.

4.2. *The affine estimator to account for omitted spatial variables.* The results presented above establish that spatial correlation in the outcome model residuals might arise due to spatial predictors of \mathbf{Y} , and commonly used approaches to estimate $\beta_1 = \mu'(z)$ are biased in the presence of unmeasured confounding by a spatial variable \mathbf{U} . It is now clear that mitigating bias from unmeasured spatial variables cannot be achieved based solely on an outcome regression model without making additional assumptions nor by harvesting spatial information found solely in the outcome model residuals.

An investigation of the formulas in (5) and (7) shows that bias of both least squares estimators arises from the nonzero correlation between \mathbf{U} and \mathbf{Z} , leading to a nonzero $E[\mathbf{U}|\mathbf{X}]$. Inspired by the form of the bias, we propose an estimator that includes a component that depends on \mathbf{U} . Consider the *affine estimator*,

$$(8) \quad \bar{\boldsymbol{\beta}} = \{\mathbf{X}^\top(\text{Var}[\mathbf{Y}|\mathbf{X}])^{-1}\mathbf{X}\}^{-1}\mathbf{X}^\top(\text{Var}[\mathbf{Y}|\mathbf{X}])^{-1}\{\mathbf{Y} - E[\mathbf{U}|\mathbf{X}]\},$$

which replaces \mathbf{Y} by $\mathbf{Y} - E[\mathbf{U}|\mathbf{X}]$ and is unbiased if $E[\mathbf{U}|\mathbf{X}]$ is known or, more practically, consistent if $E[\mathbf{U}|\mathbf{X}]$ is consistently estimated.

Since U is unmeasured, direct modeling of $E[U|\mathbf{X}]$ based on traditional estimation methods is *not* possible. Hence, identifiability of this component and our ability to calculate the affine estimator require additional assumptions. In Section 4.3 we provide a set of assumptions based on Gaussian Markov random field theory which pertain to the joint distribution of $(U, \mathbf{Z})|\mathbf{X}_{(-z)}$. Based on these assumptions, we discuss an approach to calculating the affine estimator in the context of restricted maximum likelihood in Section 4.4. Then, in Section 4.5 we show that these assumptions form a sufficient set for identification of the components of U on which the affine estimator is based, and, therefore, the estimation procedure is sound. The identifiability results illustrate that identification is achieved by exploiting the spatial correlation structure in the exposure and outcome model residuals which is driven by the unmeasured spatial variable.

4.3. *A sufficient set of assumptions.* In this section we present a set of assumptions which pertain to both the spatial and causal aspect of the affine estimator and are summarized in Table 1. We proceed with these for ease of illustration, and because they seem plausible within our study setting (see Section 7.1) but note that different or weaker assumptions for identification of $E[U|\mathbf{X}]$ are likely possible (see Sections 4.5 and 8).

4.3.1. *A Gaussian Markov random field construction of the joint distribution.* In viewing the model from a spatial perspective and to better align to the spatial modeling literature, we assume that the marginal distributions of U and ϵ are Gaussian with mean zero, independent of the measured covariates $\mathbf{X}_{(-z)}$, and that $(U, \mathbf{Z})|\mathbf{X}_{(-z)}$ is multivariate normal. We see the assumption that U is independent of $\mathbf{X}_{(-z)}$ as without loss of generality, since the same

TABLE 1

Set of assumptions based on which the causal exposure-response curve derivative is identifiable using observed data and can be estimated using the affine estimator

Causal Assumptions	
Temporal Order	The exposure is temporally precedent to the outcome.
SUTVA	No interference between units, no hidden levels of the treatment, $Y_i = Y_i(Z_i)$.
No unmeasured nonspatial confounding	$Z \perp\!\!\!\perp Y(z) \mathbf{C}^m, U, z \in \mathcal{Z}$.
Positivity*	$P(Z = z \mathbf{C}^m, U) > 0, z \in \mathcal{Z}$, which implies that:
Spatial scale restriction	Z has variation at a smaller spatial scale than that of U .
Structural Assumptions	
Outcome Additivity	The exposure and measured covariates do not interact with the unmeasured covariates.
$E[U \mathbf{Z}]$ identification	
Normality	Z, U, ϵ are jointly normal, conditional on $\mathbf{X}_{(-z)}$.
Cross-Markov	$Z_i \perp\!\!\!\perp U_{-i} U_i, \mathbf{Z}_{-i}, \mathbf{X}_{(-z)}$, where $U_{-i} = (U_1, U_2, \dots, U_{i-1}, U_{i+1}, \dots, U_n)^T, \mathbf{Z}_{-i}$ defined similarly.
Conditional correlation	$\text{Cor}(U_i, Z_i U_{-i}, \mathbf{Z}_{-i}, \mathbf{X}_{(-z)})$ is constant.
Precision matrices	The precision matrices of $U \mathbf{Z}, \mathbf{X}_{(-z)}$ and $\mathbf{Z} U, \mathbf{X}_{(-z)}$ are of CAR form.
Mean Specification	$E[Y \mathbf{X}, U]$ and $E[\mathbf{Z} \mathbf{X}_{(-z)}, U]$ are correctly specified.

*For continuous exposures, positivity can be defined in terms of the probability *density* function of Z conditional on measured and unmeasured variables.

procedure could be alternatively followed for $U - \mathbb{P}(U|\mathbf{X}_{(-z)})$, where $\mathbb{P}(U|\mathbf{X}_{(-z)})$ is the projection of U on the column space of $\mathbf{X}_{(-z)}$. This is also supported by results from simulated scenarios under which U and $\mathbf{X}_{(-z)}$ are correlated (see Section 6). Further, we recognize that joint normality might be a strong assumption, and we discuss an approach to relaxing it in Section 8. We make the following assumptions about the joint distribution of $(U, \mathbf{Z})|\mathbf{X}_{(-z)}$:

1. *Cross-Markov property*: $p(Z_i|\mathbf{Z}_{-i}, U, \mathbf{X}_{(-z)}) = p(Z_i|\mathbf{Z}_{-i}, U_i, \mathbf{X}_{(-z)})$,
2. *Constant conditional correlation*: $\text{Cor}(U_i, Z_i|\mathbf{U}_{-i}, \mathbf{Z}_{-i}, \mathbf{X}_{(-z)}) = \rho$.

The first assumption states that, conditional on measured covariates and the values of Z at all other locations, Z_i depends on U only through its value at location i , U_i . Thus, it accommodates correlation between nearby treatments, but it does not allow U_j to directly affect the value of Z_i for $i \neq j$. The second assumption states that the conditional correlation between U_i and Z_i does not vary by location. In the joint distribution of $(U, \mathbf{Z})|\mathbf{X}_{(-z)}$, these assumptions can be incorporated in the precision matrix (see Appendix C for derivations). Specifically, if

$$(9) \quad \begin{pmatrix} U \\ \mathbf{Z} \end{pmatrix} \sim \mathcal{N} \left[\begin{pmatrix} \mathbf{0} \\ \mathbf{X}_{(-z)}\boldsymbol{\gamma} \end{pmatrix}, \begin{pmatrix} \mathbf{G} & \mathbf{Q} \\ \mathbf{Q}^\top & \mathbf{H} \end{pmatrix}^{-1} \right],$$

the cross-Markov assumption is equivalent to diagonal \mathbf{Q} , and, along with the constant conditional correlation assumption, they imply that

$$(10) \quad q_{ij} = \begin{cases} -\rho\sqrt{g_{ii}h_{ii}}, & i = j, \\ 0, & i \neq j. \end{cases}$$

Given the above framework, the joint model for U and \mathbf{Z} is completed by specifying \mathbf{G} and \mathbf{H} , the precision matrices of $U|(\mathbf{Z}, \mathbf{X}_{(-z)})$ and $\mathbf{Z}|(U, \mathbf{X}_{(-z)})$, respectively, up to some *unknown* parameters that will be estimated from the data. For areal data like the ones in Section 7 we adopt conditional autoregressive structures (CAR; Besag (1974)) for \mathbf{G} and \mathbf{H} , a common assumption in standard spatial analysis models. Then, the precision matrices \mathbf{G} and \mathbf{H} are assumed to share the same neighborhood structure which is encoded in the matrices \mathbf{D} and \mathbf{W} of (2) but are allowed to differ by their precision and spatial dependence parameters (τ_U, φ_U) and (τ_Z, φ_Z) . Based on (10), the assumed CAR structure leads to $\mathbf{Q} = -\rho\sqrt{\tau_U\tau_Z}\mathbf{D}$.

In the analysis of point-referenced data, the precision matrices of $U|(\mathbf{Z}, \mathbf{X}_{(-z)})$ and $\mathbf{Z}|(U, \mathbf{X}_{(-z)})$ can be specified based on a correlation function decaying in geographical distance. In either case, since U includes all unmeasured spatial variables C^u , the correct specification of its precision matrix \mathbf{G} becomes harder for a larger number of unmeasured spatial covariates. We note again here that, since U is unmeasured, estimating the components of the joint distribution in (9) that contribute to the affine estimator *cannot* be based on traditional modeling approaches and, instead, is based on harvesting information from the spatial structure in exposure and outcome model residuals (as we see in Sections 4.4 and 4.5).

4.3.2. *Spatial scale restriction for the unmeasured spatial confounder.* In order to draw causal conclusions using our approach, the no-unmeasured and positivity assumptions still need to hold, conditional on the measured covariates C^m and the unmeasured U . The assumption of positivity implies that estimation of the causal effect of Z on Y in the presence of U is only possible if there is variability in Z within levels of U . If the spatial scale of U is smaller than that of Z , the positivity assumption will be violated, since, loosely speaking, there may be “strata” of U within which only one value of Z is possible. Therefore, from a causal perspective we assume that the spatial scales of the exposure and spatial confounder do not violate positivity of the treatment assignment conditional on the unmeasured spatial

confounder and the measured covariates. In Section 4.6 we also discuss how the spatial scale restriction is also useful in settings where spatial variables mediate the effect of interest.

The spatial scale restriction has been studied from a spatial perspective. Paciorek (2010) shows that the bias and variance of spatial model estimators depend on the relative spatial scales of the exposure and the residual, including the confounder, $\varepsilon + U$. Paciorek (2010) recommends only fitting spatial models when there is exposure variation on a spatial scale smaller than that of the unmeasured confounder, essentially ensuring positivity. From a spatial perspective, the spatial scale restriction ensures that we do not mistakenly attribute all spatial variability of the outcome residuals to the unmeasured spatial confounder when it is truly due to the exposure.

The spatial scale restriction can be enforced through the precision matrices \mathbf{G} , \mathbf{H} in (9). For geostatistical data the spatial scale of dependence is often an explicit modeling parameter, as in Paciorek (2010). For a conditional autoregressive model of areal data, the autocorrelation parameters φ_Z and φ_U do not have strict interpretations as spatial scale parameters, though the restriction $\varphi_Z < \varphi_U$ plays a similar role in identifying variance parameters.

4.4. *The affine estimator within a restricted likelihood framework.* In this section we describe estimation within a restricted likelihood framework. We do so because it allows for straightforward illustration of how model components correspond to components in the bias results of Section 4 and the affine estimator in (8). Further, it allows us to easily discuss identifiability of the model parameters in Section 4.5. In Section 5 we describe a fully-Bayesian approach to estimation, which is applicable for linear and nonlinear models, and which we follow for the remainder of this paper.

4.4.1. *Linear effect estimator.* We start by assuming the linear structure in (4). Using the conditional distribution $U|\mathbf{X}$ acquired from (9), the joint model for the observed data (integrating $U|\mathbf{X}$ out) can be factored as

$$(11) \quad \begin{aligned} Y|\mathbf{X} &\sim \mathcal{N}[\mathbf{X}\boldsymbol{\beta} - \mathbf{G}^{-1}\mathbf{Q}(\mathbf{Z} - \mathbf{X}_{(-z)}\boldsymbol{\gamma}), \mathbf{G}^{-1} + \mathbf{R}^{-1}], \\ \mathbf{Z}|\mathbf{X}_{(-z)} &\sim \mathcal{N}[\mathbf{X}_{(-z)}\boldsymbol{\gamma}, (\mathbf{H} - \mathbf{Q}^\top\mathbf{G}^{-1}\mathbf{Q})^{-1}], \end{aligned}$$

where $\mathbf{R}^{-1} = \text{Cov}(\boldsymbol{\varepsilon})$ (see Appendix D.1 for the derivation). From (11) we see that the likelihood depends on \mathbf{U} through the components of the precision matrix in (9). Note that, even though our focus is in estimating parameters of the outcome model ($\boldsymbol{\beta}$), an exposure model is also adopted to provide information on the spatial structure of \mathbf{U} . (This is related to many settings in causal inference where incorporating information from the exposure model improves estimation of causal effects (e.g., Antonelli, Parmigiani and Dominici (2019), Belloni, Chernozhukov and Hansen (2014), Wilson and Reich (2014)).)

Following a common approach to estimation for mixed models, variance parameters are estimated based on the restricted likelihood derived from (11), and the estimates are used to calculate the bias-adjusted affine estimator $\bar{\boldsymbol{\beta}}$ in (8). Defining

$$(12) \quad \begin{aligned} \mathbf{M} &= \begin{pmatrix} \mathbf{G}^{-1} + \mathbf{R}^{-1} & \mathbf{0} \\ \mathbf{0} & (\mathbf{H} - \mathbf{Q}^\top\mathbf{G}^{-1}\mathbf{Q})^{-1} \end{pmatrix}, \\ \mathbf{C} &= \begin{pmatrix} \mathbf{X} & \mathbf{G}^{-1}\mathbf{Q}\mathbf{X}_{(-z)} \\ \mathbf{0} & \mathbf{X}_{(-z)} \end{pmatrix}, \\ \mathbf{v} &= \begin{pmatrix} \mathbf{Y} + \mathbf{G}^{-1}\mathbf{Q}\mathbf{Z} \\ \mathbf{Z} \end{pmatrix} \quad \text{and} \\ \boldsymbol{\theta} &= \begin{pmatrix} \boldsymbol{\beta} \\ \boldsymbol{\gamma} \end{pmatrix}, \end{aligned}$$

we can write the joint distribution of (\mathbf{Y}, \mathbf{Z}) as

$$(13) \quad f(\mathbf{Y}, \mathbf{Z}|\boldsymbol{\beta}, \boldsymbol{\gamma}) \propto |\mathbf{M}|^{-\frac{1}{2}} \exp\left[-\frac{1}{2}(\mathbf{v} - \mathbf{C}\boldsymbol{\theta})^\top \mathbf{M}^{-1}(\mathbf{v} - \mathbf{C}\boldsymbol{\theta})\right]$$

and the restricted likelihood as

$$(14) \quad RL \propto [|\mathbf{M}| \cdot |\mathbf{C}^\top \mathbf{M}^{-1} \mathbf{C}|]^{-1/2} \exp\left[-\frac{1}{2} \mathbf{v}^\top (\mathbf{M}^{-1} - \mathbf{M}^{-1} \mathbf{C} (\mathbf{C}^\top \mathbf{M}^{-1} \mathbf{C})^{-1} \mathbf{C}^\top \mathbf{M}^{-1}) \mathbf{v}\right].$$

If $\widehat{\mathbf{M}}$, $\widehat{\mathbf{C}}$, and $\widehat{\mathbf{v}}$ are maximizers of the restricted likelihood in (14), we calculate $(\bar{\boldsymbol{\beta}}, \bar{\boldsymbol{\gamma}}) = (\widehat{\mathbf{C}}^\top \widehat{\mathbf{M}}^{-1} \widehat{\mathbf{C}})^{-1} \widehat{\mathbf{C}}^\top \widehat{\mathbf{M}}^{-1} \widehat{\mathbf{v}}$. (Readers interested in the REML approach can find additional information in Appendix A.)

The restricted likelihood formulation allows us to make illuminating connections between our approach, the mixed effects models often used in spatial statistics and the bias results of existing approaches in Section 4. If $\rho = 0$, the matrix \mathbf{Q} is zero, the model in (9) reduces to the case where \mathbf{U}, \mathbf{Z} are independent and the restricted likelihood estimation method would lead to the estimator $\widehat{\boldsymbol{\beta}}$. A nonzero correlation ρ between \mathbf{U} and \mathbf{Z} leads to the inclusion of the $-\mathbf{G}^{-1} \mathbf{Q}$ component in the coefficient of \mathbf{Z} , corresponding to the bias correction term $E[\mathbf{U}|\mathbf{X}] = -\mathbf{G}^{-1} \mathbf{Q}(\mathbf{Z} - \mathbf{X}_{(-z)} \boldsymbol{\gamma})$.

4.4.2. *Semiparametric effect estimator.* To better accommodate continuous exposures, we can flexibly model the exposure-response relationship using penalized regression splines. Penalized regression splines may be represented as linear mixed models (Ruppert, Wand and Carroll (2003)), allowing for the linear effect-assumption in (4) to be relaxed to

$$(15) \quad \mathbf{Y} = \mathbf{1}\beta_0 + \mathbf{f}(\mathbf{Z}) + \mathbf{X}_{(-z)} \boldsymbol{\beta}_{(-z)} + \mathbf{U} + \boldsymbol{\varepsilon},$$

where $\mathbf{f}(\mathbf{Z}) = (f(Z_1), \dots, f(Z_n))^\top$ and f being a smooth function of Z . Our chosen radial basis penalized spline model for f may then be written as

$$(16) \quad \bar{f}(z) = \sum_{a=1}^A \beta_a z^a + \sum_{k=1}^K l_k |z - \xi_k|^A,$$

where A is the degree of the spline and the ξ_k are prespecified knots. Letting

$$(17) \quad \mathbf{X} = \begin{pmatrix} 1 & z_1^1 & \dots & z_1^A \\ \vdots & \vdots & \ddots & \vdots \\ 1 & z_n^1 & \dots & z_n^A \end{pmatrix}, \quad \mathbf{L} = \begin{pmatrix} |z_1 - \xi_1|^A & \dots & |z_1 - \xi_K|^A \\ \vdots & \ddots & \vdots \\ |z_n - \xi_1|^A & \dots & |z_n - \xi_K|^A \end{pmatrix}$$

and $\mathbf{V} = \psi^{-1} \mathbf{L} \mathbf{L}^\top + \mathbf{G}^{-1} + \mathbf{R}^{-1}$, where $\psi > 0$ is a roughness penalty, the restricted likelihood is as in (14) with updated components \mathbf{C} and \mathbf{M} . Letting $\mathbf{T} = (\mathbf{X} \mathbf{L})$ and \mathbf{A} be a diagonal matrix with the first $A + 1$ elements equal to 0 and the rest equal to 1 (corresponding to penalization of the $\boldsymbol{\beta}$ and \mathbf{l} , respectively), the estimate of $\boldsymbol{\theta} = (\beta_0, \beta_1, \dots, \beta_A, l_1, l_2, \dots, l_K)$ is $\bar{\boldsymbol{\theta}} = (\mathbf{T}^\top \bar{\mathbf{V}}^{-1} \mathbf{T} + \psi \mathbf{A})^{-1} \mathbf{T}^\top \bar{\mathbf{V}}^{-1} (\mathbf{Y} - \bar{\mathbf{B}} \mathbf{Z})$.

4.5. *Identifiability of parameters.* It is evident from the form of the affine estimator in (8) and the restricted likelihood in (14) that calculating the estimator depends on being able to estimate components of the relationship of the unmeasured confounder with the exposure and outcome of interest. Therefore, it is natural to wonder whether these components can be identified and, if so, which assumptions are key in driving identifiability and which can be relaxed. We provide two examples: one in which identifiability can be proved analytically and one in which model components are not identifiable. The details of how identifiability is achieved or lost in these examples illuminate the key assumptions and their roles.




FIG. 3. Graph representation of ring with the cross-Markov property and dependence parameters.

4.5.1. *Identifiability of model components for a ring graph.* The first example is one in which identifiability can be analytically proved. We consider the setting without measured variables since including them complicates the notation without providing any additional insight, and results extend trivially. We also assume that the spatial dependence can be represented in a ring graph of n locations depicted in Figure 3 with CAR specifications for $U|Z$ and $Z|U$. This spatial structure and model yields the precision matrix

$$(18) \quad \mathbf{G}_n = \tau_U \begin{pmatrix} 2 & -\varphi_U & & -\varphi_U \\ -\varphi_U & 2 & -\varphi_U & \\ & \ddots & \ddots & \ddots \\ & & -\varphi_U & 2 & -\varphi_U \\ -\varphi_U & & & -\varphi_U & 2 \end{pmatrix}$$

for $U|Z$ and similar for \mathbf{H}_n for $Z|U$. For simplicity, we assume that $E[Z] = \mathbf{0}$ (marginally over U) and that $E[Y|Z, U] = \beta_Z Z + U$. We present two results: the first discussing parameter identifiability based on Z alone (Theorem 1), and the second based on (Y, Z) (Theorem 2). Importantly, these identifiability results allow for U to be completely unmeasured.

THEOREM 1. *In the scenario defined in this section, it can be determined whether or not $\rho\varphi_U = 0$ by observing Z . Further, if $\rho\varphi_U \neq 0$, then $(\tau_Z, \varphi_Z, \varphi_U, |\rho|)$ is also identifiable by observing Z .*

PROOF. The proof relies on a few matrix lemmas which are stated and given proof outlines in Appendix E.1. We have that $\text{Prec}[Z] = \tau_Z[\mathbf{H}_n - 4\rho\mathbf{G}_n^{-1}]$ (marginally over U), and

$$(19) \quad \lim_{n \rightarrow \infty} \text{Prec}[Z]_{ij} = \begin{cases} \tau_Z \left[2 - 2\rho^2 \frac{1}{\sqrt{1 - \varphi_U^2}} \right], & i = j, \\ \tau_Z \left[-\varphi_Z - 2\rho^2 \frac{\varphi_U}{\sqrt{1 - \varphi_U^2} (1 + \sqrt{1 - \varphi_U^2})} \right], & |i - j| = 1, \\ \tau_Z \left[0 - 2\rho^2 \frac{\varphi_U^{|i-j|}}{\sqrt{1 - \varphi_U^2} (1 + \sqrt{1 - \varphi_U^2})^{|i-j|}} \right], & |i - j| > 1. \end{cases}$$

First, note that, for any lag l , the number of pairs of locations with $|i - j| = l$ grows linearly with n . It can be determined whether or not $\rho\varphi_U = 0$ because $\lim_{n \rightarrow \infty} \text{Prec}[Z]_{ij} = 0$ for all (i, j) such that $|i - j| > 1$ if and only if $\rho\varphi_U = 0$. If $\rho\varphi_U \neq 0$, then for $|i - j| > 1$ and $|i' - j'| = |i - j| + 1$ the ratio $\lim_{n \rightarrow \infty} \text{Prec}[Z]_{ij} / \text{Prec}[Z]_{i'j'}$ depends only on φ_U and is bijective, thus φ_U is identified. With φ_U identified, the three cases in (19) form a system of equations solvable for $(\tau_Z, \varphi_Z, |\rho|)$. \square

Based on Theorem 1, if $\rho\varphi_U = 0$, some parameters are not identifiable by only observing \mathbf{Z} . That is, because when $\rho = 0$, Z and U are uncorrelated, hence Z cannot provide any information on U ; when $\varphi_U = 0$, the unmeasured variable is not spatial and the variability in Z cannot be decomposed accordingly (a situation we examine closely in Section 4.5.2). In contrast, when $\rho\varphi_U \neq 0$, a number of spatial parameters are identified based solely on the vector of treatments \mathbf{Z} . For intuition about why the unobserved variable’s spatial dependence parameter φ_U (but not its precision τ_U) can be identified by examining the behavior of $\text{Prec}[\mathbf{Z}]$ away from the tridiagonal, recall that, off of the tridiagonal, i and j are not neighbors and that the precision in (19) tells us about the strength of dependence between Z_i and Z_j , given the value of \mathbf{Z} at all other locations. Our cross-Markov assumption states that, *conditionally on U* , the values of \mathbf{Z} at nonneighboring locations are independent, given the values of \mathbf{Z} at other locations. If Z_i and Z_j are not independent conditional only on \mathbf{Z} (and not on U), this dependence has to arise through paths in Figure 3 that pass through U . The strength of dependence is a function of ρ , the strength of connection between U and Z at a given location (which is independent of the distance between i and j) and φ_U which determines how quickly the dependence of U_i, U_j attenuates with distance $|i - j|$. These dependencies are graphically represented in Figure 3 where the dependence between Z_{i-1} and Z_{i+1} , conditional on \mathbf{Z} at other locations but marginally over U , arises from the paths $Z_{i-1} - U_{i-1} - U_i - U_{i+1} - Z_{i+1}$ and $Z_{i-1} - U_{i-1} - U_{i-2} - \dots - U_1 - U_n - \dots - U_{i+1} - Z_{i+1}$ and the dependency due to the latter path diminishes as n becomes large. We can identify φ_U off of the tridiagonal by examining this attenuation, and ρ is separable from τ_Z only when examining the tridiagonal as well.

THEOREM 2. *The parameter φ_U is identifiable by observing (\mathbf{Y}, \mathbf{Z}) . Further, if $\varphi_U \neq 0$, then the full parameter $(\beta_Z, \tau_Z, \varphi_Z, \tau_U, \varphi_U, \rho, \tau_\varepsilon)$ is identifiable by observing (\mathbf{Y}, \mathbf{Z}) .*

PROOF. Note that $E[\mathbf{Y}|\mathbf{Z}]$ is identified irrespective of $\text{Var}[\mathbf{Y}|\mathbf{Z}]$, and that, by Theorem 1, we can identify whether $\rho\varphi_U = 0$ by observing \mathbf{Z} .

We start by showing that φ_U is identifiable by observing (\mathbf{Y}, \mathbf{Z}) . Since $\text{Var}[\mathbf{Y}|\mathbf{Z}] = \mathbf{G}_n^{-1} + \tau_\varepsilon^{-1}\mathbf{I}_n$, by noting the similarity between the expressions for $\text{Var}[\mathbf{Y}|\mathbf{Z}]$ and $\text{Prec}[\mathbf{Z}]$ in (19), we have

$$(20) \quad \lim_{n \rightarrow \infty} \text{Var}[\mathbf{Y}|\mathbf{Z}]_{ij} = \tau_U^{-1} \frac{1}{2\sqrt{1 - \varphi_U^2}} \left(\frac{\varphi_U}{1 + \sqrt{1 - \varphi_U^2}} \right)^{|i-j|}, \quad i \neq j.$$

Therefore, $\lim_{n \rightarrow \infty} \text{Var}[\mathbf{Y}|\mathbf{Z}]_{ij} = 0$ for $i \neq j$ if and only if $\varphi_U = 0$. Since we can identify $\text{Var}[\mathbf{Y}|\mathbf{Z}]$, φ_U is identifiable by $\lim_{n \rightarrow \infty} \text{Var}[\mathbf{Y}|\mathbf{Z}]_{ij} / \text{Var}[\mathbf{Y}|\mathbf{Z}]_{i'j'}$ for $|i - j| > 0$ and $|i' - j'| = |i - j| + 1$.

Next, assume that $\varphi_U \neq 0$, and we show that the remaining parameters are identifiable by observing (\mathbf{Y}, \mathbf{Z}) . We first note that, since $E[\mathbf{Y}|\mathbf{Z}] = (\beta_Z - \rho\sqrt{\tau_Z/\tau_U})\mathbf{Z}$, the combined coefficient of \mathbf{Z} , $\beta_Z - \rho\sqrt{\tau_Z/\tau_U}$ is identifiable. Since we can identify whether $\rho\varphi_U \neq 0$ from observing \mathbf{Z} (Theorem 1) and since we have here that $\varphi_U \neq 0$, we can identify whether $\rho = 0$ which allows us to consider the cases where $\rho = 0$ and $\rho \neq 0$, separately.

If $\rho = 0$, the combined coefficient of \mathbf{Z} is equal to β_Z , and β_Z is identified (which is expected since for $\rho = 0$ there is no confounding by U). Additionally, we can return to (19) to identify (τ_Z, φ_Z) from the first two cases. Finally, with $\varphi_U \neq 0$ identified, τ_U can be identified from the off-diagonal elements of $\text{Var}[\mathbf{Y}|\mathbf{Z}]$ in (20).

If $\rho \neq 0$, we can identify $(\tau_Z, \varphi_Z, \varphi_U, |\rho|)$ by Theorem 1. Recall that $E[\mathbf{Y}|\mathbf{Z}] = (\beta_Z - \mathbf{G}_n^{-1}\mathbf{Q}_n)\mathbf{Z}$ and $-\mathbf{G}_n^{-1}\mathbf{Q}_n = 2\rho\sqrt{\frac{\tau_Z}{\tau_U}}(\tau_U\mathbf{G}_n^{-1})$, where $\tau_U\mathbf{G}_n^{-1}$ does not depend on τ_U . Since $\varphi_U \neq 0$, \mathbf{G}_n is not a scalar matrix, and the terms \mathbf{Z} and $\tau_U\mathbf{G}_n^{-1}\mathbf{Z}$ are known and not collinear.

Thus, we can separately identify their coefficients β_Z and $\rho\sqrt{\tau_Z/\tau_U}$. Since τ_U and τ_Z are both positive and $(\tau_Z, |\rho|)$ has been previously identified, we can identify $\text{sign}(\rho)$, and τ_U .

In both cases ($\rho = 0$ or $\rho \neq 0$), τ_ε can be identified from the diagonal elements of $\text{Var}[\mathbf{Y}|\mathbf{Z}]$. \square

Theorems 1 and 2 combined establish that, when spatial dependencies can be represented using a ring, and the unmeasured variable is spatial ($\varphi_U \neq 0$) and is truly a confounder ($\rho \neq 0$), the effect of the exposure \mathbf{Z} , β_Z , is identifiable.

The above results and the details of the proofs indicate that the cross-Markov property is the critical assumption in identifying β_Z . This is not to say that we can remove the more parametric or the constant conditional correlation assumptions but rather that these assumptions are most likely not the only or weakest ones allowing for identification. In turn, the cross-Markov assumption can be viewed as a relaxation of the usual assumption of no unobserved confounders, expressing that confounding by unobserved spatial variables is local. To see this, the assumption may be rewritten as $Z_i \perp\!\!\!\perp U_{-i} | (\mathbf{Z}_{-i}, U_i, \mathbf{X}_{(-z)})$, expressing that conditional on measured covariates, the value of U at location i and the value of Z everywhere else, U_k with $k \neq i$ is not a predictor of Z_i and, therefore, does not confound the relationship between Z_i and Y_i .

Although we have proved these asymptotic results for a ring of locations, the key requirement on the structure of a graph with large connected components is that there are enough pairs of locations at varying lags. This requirement does not seem problematic for, for example, counties in the United States, as in our study in Section 7.

4.5.2. *A nonsufficient set of assumptions: Unmeasured nonspatial confounders.* Within framework (3) any set of assumptions that suffice for identification of $E[U|\mathbf{X}]$ would also allow for identification of the causal parameter $\mu'(z)$. In Section 4.5.1 we showed that, for adjacency structures described as a ring of growing size, identifiability is achieved when the unmeasured confounder is spatial ($\varphi_U \neq 0$). Here, we establish that identifiability is lost when the unmeasured confounder is not spatial ($\varphi_U = 0$) under any adjacency structure.

Assume that U_i are independent and identically distributed random variables (exhibiting no spatial structure) with $\mathbf{G} = \tau_U \mathbf{I}$. Further, assume that \mathbf{Z} and ε are not spatially structured with precision matrices $\mathbf{H} = \tau_Z \mathbf{I}$ and $\mathbf{R} = \tau_\varepsilon \mathbf{I}$. Then, the parameter vector $(\tau_U, \tau_Z, \tau_\varepsilon, \rho)$ in the restricted likelihood is reduced to $(\sigma^2, \varphi) = (\tau_U^{-1} + \tau_\varepsilon^{-1}, \tau_Z(1 - \rho^2))$, and the parameters in $(\tau_U, \tau_\varepsilon)$ and in (τ_Z, ρ) are not separately identifiable (the mathematical derivations are included in Appendix E.2). Thus, when there is no spatial structure, $\beta = (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top (\mathbf{Y} - \rho\sqrt{\frac{\tau_Z}{\tau_U}} \mathbf{Z})$ is not identifiable either.

This result is intuitively obvious. If U is not spatially structured, there is no information in the observed data to differentiate outcome model residuals' variability due to U from that due to ε and, similarly, nothing to differentiate intrinsic variability in Z from variability due to U . In such case, $E[U|\mathbf{X}]$ is not identifiable based on observed data indicating that adjustment for U is not possible if the unmeasured confounders do not exhibit spatial structure. This is in line with recent work showing that latent variable approaches cannot be used to acquire identifiability of causal parameters without additional assumptions (D'Amour (2019), Ogburn, Shpitser and Tchetgen Tchetgen (2019)).

4.6. *Spatially correlated mediating variables.* From model (11) and the form of the restricted likelihood in (14), it is evident that information about the elements (\mathbf{G}, \mathbf{Q}) in the bias correction term $-\mathbf{G}^{-1} \mathbf{Q}(\mathbf{Z} - \mathbf{X}_{(-z)} \boldsymbol{\gamma})$ is found in the spatial variability of exposure and outcome models' residuals. However, the estimation procedure cannot differentiate between spatial structure arising from spatial confounders (temporally precedent of Z) or variables

found on the causal pathway between Z and Y (mediators). If Z and Y are measured within a small time window, it may be reasonable to assume that there are no spatial covariates mediating the effect of Z on Y , and, for that reason, our estimates correspond to estimates of $\beta_z = \mu'(z)$.

On the other hand, in the presence of spatial intermediate variables, estimates using the affine estimator might more closely resemble the direct effect of Z on Y , not due to changes to the spatial mediators (Baron and Kenny (1986)). In this setting the spatial scale restriction provides some protection against adjustment for spatial mediators since variation in spatial scales smaller than that of the exposure is not adjusted for. Therefore, the spatial scale restriction allows for unmeasured spatial confounder bias mitigation while protecting us from adjusting for variables on the causal pathway between Z and Y .

5. The affine estimator in nonlinear settings: A Bayesian implementation. The REML framework has allowed us to analytically investigate bias and identifiability by allowing us to integrate out the distribution of the unmeasured confounder from the observed data likelihood. However, such analytical approach is less applicable to noncontinuous outcomes and nonlinear models. Here, we extend the affine estimator to nonlinear models within the Bayesian framework.

5.1. *Estimation of causal parameters with nonlinear models.* When $\mu(z) = E[Y(z)]$ is the estimand of interest, and the outcome model specifies $E(Y|Z, \mathbf{C}^m, U)$, $\mu(z)$ can be written as $E_{\mathbf{C}^m, U}[E(Y|Z = z, \mathbf{C}^m, U)]$ under assumptions. This is often referred to as the g-formula or g-computation (Robins (1986)). If the outcome model is linear without exposure-covariate interactions, the regression coefficient for the exposure can be directly interpreted as an estimate of the causal quantity $\mu'(z)$. However, the exposure’s coefficient cannot be directly causally interpreted in the case of nonlinear models, even when the model is correctly specified. For example, in the context of logistic regression with binary outcomes, the coefficient of the exposure β_z is not equal to $\mu'(z)$, and a linear specification of the exposure-response relationship in the linear predictor of the logistic regression does *not* imply a linear $\mu(z)$. Therefore, in nonlinear outcome models we need to proceed with care when translating estimated coefficients to estimates of causal quantities, and an integration step (over the distribution of confounders in the target population) needs to be employed in order to acquire estimates of $\mu(z)$ from a nonlinear model fit.

Poisson models for count outcomes, like the one in our study, are an exception. In such models the parameter β_z (or $\exp(\beta_z)$, which is often used in Poisson models), can be interpretable as causal but for an estimand that is slightly different from $\mu'(z)$. Let P_i be the population at risk at location i , and $Y_i(z)$ be the potential outcome at location i if the exposure was set to z . Consider the PAERC defined in terms of the *standardized* outcome rate as $E[P^{-1}Y(z)]$. Under a structural model similar to (3) for linear η , assume that $[Y(z)|\mathbf{X}_{(-z)}, U] \sim \text{Poisson}(P \exp\{\beta_z z + \beta_{(-z)}^\top \mathbf{X}_{(-z)} + U\})$. Then, the PAERC can be written as

$$(21) \quad E_{\mathbf{X}_{(-z)}, U}[\exp\{\beta_z z + \beta_{(-z)}^\top \mathbf{X}_{(-z)} + U\}] = \exp\{\beta_z z\} E_{\mathbf{X}_{(-z)}, U}[\exp\{\beta_{(-z)}^\top \mathbf{X}_{(-z)} + U\}].$$

From (21) the coefficient β_z can be interpreted as the log-relative standardized rate for a one-unit exposure change, $\log\{E[P^{-1}Y(z + 1)]/E[P^{-1}Y(z)]\}$. Equivalently, β_z can be interpreted as the instantaneous effect of the exposure in the relative scale as $\partial\{\log E[P^{-1}Y(z)]\}/\partial z$. Note that this is substantially different from the standard interpretation of estimated coefficients in Poisson models, and β_z *cannot* be used as an estimate of $\log\{E[Y(z + 1)/Y(z)]\}$.

5.2. *Bayesian implementation of the affine estimator.* The presentation above indicates that, for nonlinear models where noncollapsibility is an issue and estimated coefficients do not always estimate causal quantities, estimation might require an explicit model for Y conditional on measured covariates *and* the unmeasured component. This is straightforwardly achieved for the affine estimator within a Bayesian implementation, for which U is viewed as a missing variable that is iteratively imputed through a Gibbs sampler. Therefore, placing the affine estimator within the Bayesian paradigm does not require marginalization over U which allows for estimation outside the realm of linear regression. This is exploited in Sections 6 and 7 where we consider a count outcome and a Poisson model with the log link and linear predictor $\eta_i = o_i + \mathbf{x}_i^\top \boldsymbol{\beta} + u_i$, where o_i is the offset in the usual sense and sampling from the conditional posterior distribution of $\boldsymbol{\beta}$ can be performed without modifying standard algorithms for Poisson regression.

Apart from its generalizability to nonlinear models, the Bayesian approach has a number of additional benefits over the REML approach, including computational gains. The Bayesian implementation is computationally more efficient thanks to the conditional nature of Gibbs sampling which allows us to take advantage of sparsity in $\mathbf{P} = \text{Prec}[\mathbf{U}, \mathbf{Z}]$. For example, the log full conditional density of the dependence parameters is (up to an additive constant)

$$(22) \quad \frac{1}{2} [\mathbf{P} - \mathbf{u}^\top \mathbf{G} \mathbf{u} - 2\mathbf{u}^\top \mathbf{Q}(\mathbf{z} - \mathbf{X}_{(-z)} \boldsymbol{\gamma}) - (\mathbf{z} - \mathbf{X}_{(-z)} \boldsymbol{\gamma})^\top \mathbf{H}(\mathbf{z} - \mathbf{X}_{(-z)} \boldsymbol{\gamma})] \\ + \log p(\tau_U, \tau_Z, \varphi_U, \varphi_Z, \rho),$$

where \mathbf{P} , \mathbf{G} , \mathbf{H} and \mathbf{Q} are sparse matrices depending on $(\tau_U, \tau_Z, \varphi_U, \varphi_Z, \rho)$ and $p(\tau_U, \tau_Z, \varphi_U, \varphi_Z, \rho)$ is the prior for these parameters. In contrast, the REML approach requires inverting \mathbf{G} , \mathbf{R} , $(\mathbf{H} - \mathbf{Q}^\top \mathbf{G}^{-1} \mathbf{Q})$, \mathbf{M} and $\mathbf{C}^\top \mathbf{M} \mathbf{C}$ in (12) and (14) at each evaluation of the restricted likelihood.

Another advantage of the Bayesian approach is that it is easier to incorporate non-Gaussian exposures by distinguishing between \mathbf{Z} in (9) and the exposure parameterization in the outcome model. For example, in our analysis of the food access data in Section 7 we replace Z with $\log Z$ in (9) to make the assumption of joint normality more plausible, while using Z in the outcome model to retain the desired interpretation of regression coefficients on the original percentage point scale. Of course, that comes with the caveat that our assumptions are now based on the transformation of the exposure variable. The exposure model and outcome model may be further decoupled by, for example, assuming joint normality of U and a latent variable in a probit model of a binary exposure.

5.3. *A regularization prior on the precision matrix of (U, Z) .* The estimation of the joint precision matrix \mathbf{P} of (U, Z) is critical in mitigating bias due to the unobserved spatial confounder U . However, the present setting is a “low-information” one, as we neither observe U directly nor obtain independent replicates. In such settings the restricted likelihood may have maxima at the boundary of allowed values. For example, Chung, Rabe-Hesketh and Choi (2013) noted that it is not unusual in random effects metaanalysis for the REML estimate of the between-study standard deviation to be zero and suggested regularizing the REML estimate by multiplying the restricted likelihood by a weakly-informative gamma prior for the between-study variance. Along another thread, Won et al. (2013) considered estimating the covariance matrix in high-dimensional settings where maximum likelihood estimates of such covariance matrices are often ill conditioned and cannot be inverted accurately. They propose a constrained maximum likelihood approach using the constraint $\kappa(\boldsymbol{\Sigma}) \leq \kappa_{\max}$, where $\kappa(\boldsymbol{\Sigma})$ is the condition number (the ratio of the largest to smallest eigenvalue) and κ_{\max} is prespecified. They note that this optimization is equivalent to maximizing the likelihood times an exponential prior on $\kappa(\boldsymbol{\Sigma})$ left-truncated at 1.

We have observed that this problem manifests in the fully Bayesian implementation as occasional failure of the MCMC sampler to converge. For our purposes we adopt a truncated exponential prior for $\kappa(\mathbf{P})$ with rate 1/10 and range $(1, \infty)$ which directly addresses the ill-conditioning problem. With this specification the difference in log-prior density between $\kappa(\mathbf{P}) = 1$ and $\kappa(\mathbf{P}) = 100$ is 9.9.

6. Simulation study.

6.1. *Linear effect.* We perform simulations to compare the affine estimator to the non-spatial and spatial random effect estimators under several generative models (GMs). Under all GMs, we consider a single measured covariate, X , generated uniformly on $(-1/2, 1/2)$. We assume that the mean of $\mathbf{Z}|\mathbf{X}_{(-z)}$ is X . Four GMs reflect $(U, \mathbf{Z})|\mathbf{X}_{(-z)}$ generation according to (9) and (10), with $U|\mathbf{Z}, \mathbf{X}_{(-z)}$ and $\mathbf{Z}|U, \mathbf{X}_{(-z)}$ being one-dimensional CAR models. The within-variable dependence parameters are denoted by φ_U and φ_Z , and precision parameters by $\tau_U = \tau_Z = 1$. The first model we consider is the unconfounded GM (GM 1), where U and \mathbf{Z} are independent of each other ($\rho = 0$) but still spatially structured with $\varphi_U = 0.5$ and $\varphi_Z = 0.2$. The unmeasured variable U is still predictive of the outcome, hence inducing spatial correlation in the observed outcomes. For the remaining three CAR models, we specify cross-variable dependence ($\rho = 0.3$), and vary the within-variable dependence parameters (φ_U, φ_Z) at $(0.5, 0.2)$ for GM 2, representing a confounder at a larger spatial scale than the exposure, $(0.2, 0.5)$ for GM 3, representing a confounder at a smaller spatial scale than the exposure violating our causal assumptions, and $(0.35, 0.35)$ for GM 4, where confounder and exposure vary at the same spatial scale. The fifth and sixth GMs represent situations in which the analysis model is misspecified. For GM 5, U was generated such that its *marginal* distribution is a one-dimensional CAR model with $\varphi_U = 0.5$ and $\tau_U = 1$ and $\mathbf{Z}|U, \mathbf{X}_{(-z)} \sim \mathcal{N}[U + X, \mathbf{I}]$. Therefore, in this GM the model (9) is misspecified, in that \mathbf{G} does not describe the true precision matrix of $U|\mathbf{Z}, \mathbf{X}_{(-z)}$, and the assumption of constant conditional correlation is violated. However, the precision matrix of $\mathbf{Z}|U, \mathbf{X}_{(-z)}$ is still correctly specified, and the cross-Markov property holds. For GM 6, $\tan(U)$ takes the place of U in (9), so that the joint normality assumption on $(U, \mathbf{Z})|\mathbf{X}_{(-z)}$ is violated. In all six GMs the potential outcomes are generated as $Y_i(z) \stackrel{iid}{\sim} \text{Poisson}[\exp\{z + X_i + U_i\}]$.

Under each GM we generate 500 data sets of size $n = 300$ and fit the nonspatial, spatial and affine estimators. When assumptions on the forms of variances are required, we assume CAR structures, and for the affine estimator we assume that \mathbf{Q} is of the Markov form (10). Linear predictor models for the exposure and outcome are correctly specified. For the spatial and affine estimators we evaluate variations with and without the restriction that $\varphi_Z < \varphi_U$ discussed in Section 4.3.2, with the restricted estimators denoted by (-RS). For regression coefficients we used Gaussian priors with mean 0 and standard deviation 10. For the Spatial-RS, Affine, and Affine-RS estimators we used the regularization prior discussed in Section 5.3. This prior couples the distributions of U and \mathbf{Z} , which is not usually a feature in spatial analyses, and so is not used for the nonspatial and spatial estimators. Instead, for the non-spatial and unrestricted spatial estimators flat priors were used for all variance parameters. In all cases the precision matrix \mathbf{P} was restricted to be positive definite. Due to the computational cost of computing the condition number of the precision matrix \mathbf{P} when evaluating the prior for variance parameters, we used an approximation to the condition number acquired by the same model and parameter values on a four-location, one-dimensional ring instead of the 300-location line (the first and last locations are also neighbors). Posterior samples were drawn using 10,000 Gibbs sampler iterations after 1000 burn-in iterations.

Table 2 displays the simulation results in terms of bias, standard deviation and root mean squared error (RMSE) of posterior means across data sets and empirical coverage of 95%

TABLE 2

Simulation results from 500 data sets of size $n = 300$. The -RS suffix indicates estimators with the restriction $\varphi_Z \leq \varphi_U$

	Mechanism	Estimator	Bias	Std. Err.	RMSE	95% CI Coverage
GM 1	Unconfounded	Nonspatial	0.00	0.14	0.14	0.62
		Spatial	0.01	0.12	0.12	0.93
		Spatial-RS	0.02	0.12	0.12	0.93
		Affine	0.03	0.27	0.27	0.98
		Affine-RS	0.03	0.29	0.29	0.96
GM 2	Large-scale confounder	Nonspatial	0.37	0.15	0.40	0.02
		Spatial	0.37	0.12	0.39	0.05
		Spatial-RS	0.37	0.12	0.39	0.04
		Affine	0.24	0.34	0.41	0.94
		Affine-RS	0.14	0.30	0.33	0.95
GM 3	Large-scale exposure	Nonspatial	0.33	0.14	0.36	0.03
		Spatial	0.35	0.11	0.36	0.03
		Spatial-RS	0.34	0.10	0.35	0.04
		Affine	0.33	0.25	0.42	0.91
		Affine-RS	0.23	0.36	0.42	0.83
GM 4	Same scales	Nonspatial	0.34	0.15	0.38	0.03
		Spatial	0.35	0.11	0.37	0.05
		Spatial-RS	0.35	0.11	0.37	0.06
		Affine	0.29	0.27	0.39	0.93
		Affine-RS	0.18	0.32	0.36	0.89
GM 5	Non-constant conditional correlation	Nonspatial	0.35	0.11	0.37	0.00
		Spatial	0.37	0.06	0.37	0.00
		Spatial-RS	0.37	0.06	0.37	0.00
		Affine	0.38	0.27	0.47	0.57
		Affine-RS	0.22	0.20	0.30	0.75
GM 6	Non-normal joint distribution	Nonspatial	0.24	0.08	0.26	0.06
		Spatial	0.24	0.08	0.26	0.14
		Spatial-RS	0.24	0.08	0.26	0.15
		Affine	0.18	0.18	0.26	0.94
		Affine-RS	0.08	0.17	0.19	0.96

equal-tail credible intervals. We first note that there is minimal difference between the spatial and spatial-RS estimators, even when the spatial scale assumption is violated in the large-scale exposure scenario. Additionally, the spatial and spatial-RS estimators have similar biases to the nonspatial estimator. For the unconfounded GM 1, all estimators are unbiased, all spatial estimators have approximately correct confidence interval coverage but both affine estimators have much larger standard errors and, therefore, RMSE. As expected due to misspecification of the dependence structure, the posterior distributions from the nonspatial models are too concentrated and therefore the credible intervals are anticonservative. In all GMs with confounding (2–6), the affine-RS estimator mitigates bias relative to the nonspatial estimator, whereas the spatial and spatial-RS estimators do not. The unrestricted affine estimator generally mitigates bias to a lesser extent, especially in the large-scale exposure, same-scales and nonconstant conditional correlation scenarios. When the restricted scale assumption is correct, the affine-RS estimator has a smaller standard error than the unrestricted affine estimator. Additionally, the affine-RS estimator has a smaller RMSE than all other estimators, except in the case of a large-scale exposure where its scale restriction is false. In the pres-

ence of unobserved confounding, both affine estimators have credible interval coverage rates far superior to the other estimators. Both have approximately nominal coverage rates in the independent, large-scale confounder, and nonnormal joint distribution scenarios. Analogous simulation results for maximum a posteriori (restricted likelihood multiplied by a prior) estimation in the Gaussian outcome are available in Appendix G.

Although our model specifies that U is independent of X , we conducted a smaller simulation for a scenario in which the large-scale unmeasured confounder in GM 2 is correlated with the measured confounder, by specifying that $E[U|X] = X$. We simulated 100 data sets and fit the same models (i.e., without a mean model for U) as in the previous simulations. The results in estimating β_z were similar to those under GM 2 (Appendix F), but the estimates of β_x were biased upward (not shown). This is expected since the part of U that is correlated with X is captured and adjusted for with the inclusion of X in the outcome model, and the affine estimator targets the component of U that is orthogonal to X .

6.2. Nonlinear effect. The bias-variance trade-off observed between the spatial and affine estimators in the linear case was also observed for a nonlinear effect. We generated 500 data sets of size 300 where U, Z are generated from (9) with $(\tau_U, \varphi_U, \tau_Z, \varphi_Z, \rho) = (1, 0.5, 1, 0.2, 0.3)$, and Y is a Poisson variable with log link and linear predictor in the form of the right-hand side of (15), with $(\beta_0, \beta_x) = (0, 0)$ and $f(z) = 2/(1 + e^{-6z}) - 1$. Therefore, the true effect curve is an antisymmetric sigmoid with asymptotes -1 and 1 . We fit the restricted-scale semiparametric spatial and constrained affine estimators using a penalized cubic spline model with a radial basis and used the same priors as in the linear simulations. Inference was based on 10,000 posterior draws after 5000 burn-in iterations.

Figure 4 displays a graphical summary of the simulation results. For the most part, both estimators capture the general shape of the mean response curve. However, the spatial estimator is biased toward more extreme estimates as the exposure deviates from 0, and this bias is mitigated by the constrained affine estimator. On the other hand, the constrained affine estimator exhibits substantially greater variability, especially for exposure ranges with limited




FIG. 4. *Nonlinear effect simulation results. Mean and pointwise 95% sampling intervals of posterior mean log population average exposure-response curve from 500 data sets of size $n = 300$. True log PAERC indicated by solid black curve.*

available data (away from an exposure value of 0). There also appears to be an asymmetry in that, when the true log PAERC is negative, the bias of the spatial estimator is more pronounced, and the constrained affine estimator is more effective at mitigating bias than when the log PAERC is positive. This asymmetry is likely due to the nonlinear relationship between the model's linear predictor and the expected outcome according to the Poisson likelihood.

7. Estimating the county-level effects of poor supermarket availability on CVD mortality. Here, we use the affine estimator in order to estimate the county-level effect of poor supermarket availability on CVD mortality. We consider the affine estimator with and without the spatial scale restriction, and we also consider the spatial random effect estimator with the spatial scale restriction, an extension to the results shown in Section 2. In the outcome model we include the exposure on the percentage point scale to aid in interpretation of its coefficient. However, we replace \mathbf{Z} with $\log \mathbf{Z}$ in the joint model (9) to better satisfy the condition of joint normality. For the spatial-RS model we restrict ρ to be zero, and for the spatial-RS and affine-RS models we apply the constraint $\varphi_U \geq \varphi_Z$. Note that the prior distribution for the spatial-RS model differs substantially from that of the unconstrained spatial model. In all three models we use a similar approximation to the condition number prior of Section 5.3 that was used in the simulation study: rather than computing the condition number on the full joint precision matrix \mathbf{P} , we use the analog of \mathbf{P} derived from a 4×4 regular grid. Posterior distributions from all models were simulated by retaining 10,000 Gibbs sampler iterations after 1000 burn-in iterations. The affine-RS model took approximately 1.5 hours to fit on a laptop for a sample size of $n = 3,093$.

In Section 7.1 we examine the assumptions underlying the affine estimator, and in Section 7.2 we report summaries of the posterior distribution of the causal effect estimates.

7.1. Examining the plausibility of the assumptions in the context of our study. A number of assumptions, previously presented in Table 1, are necessary to identify the causal effect of interest in the presence of unmeasured spatial confounders.

Temporal ordering and SUTVA are standard assumptions in causal inference and are necessary to define our causal effect. The temporal ordering of the exposure and outcome is immediate satisfied since the exposure data were compiled from 2000 and 2006 data sets, while the mortality outcome data were compiled from 2007 reports. SUTVA is expected to hold, at least approximately, since we can assume that the county-level effect of poor supermarket access on CVD mortality is due to individual-level causal effects, and the home address county listed on death certificates in 2007 corresponds well to the deceased person's county of residence from 2000 to 2006.

Outcome additivity and the appropriateness of the joint normality and CAR assumptions for $(U, \mathbf{Z}) | \mathbf{X}_{(-z)}$ are modeling assumptions that may be, at least partially, addressed via standard diagnostics. Maps of Pearson residuals, based on posterior mean parameters and plots of those residuals versus linear predictors, indicated no visually apparent residual spatial correlation or nonlinearity in either the log exposure or outcome models in the affine-RS approach. A scatterplot of the joint distribution of the residual log exposure, after adjusting for covariates vs. mean, imputed confounder U appeared Gaussian.

The cross-Markov and constant conditional correlation assumptions are assumptions about the relationship between unobserved confounders and the exposure to which standard diagnostics are not applicable. The plausibility of these assumptions depends heavily on the application and hypothesized confounders. For example, consider an unmeasured variable representing cultural preference toward purchasing prepared food from restaurants vs. cooking at home. This variable might act as a confounder in our study, since an increase in such a preference could both depress the demand for and availability of supermarkets and might drive




FIG. 5. Graph representation of the cross-Markov property $p(Z_i | \mathbf{Z}_{-i}, \mathbf{U}) = p(Z_i | \mathbf{Z}_{-i}, U_i)$.

food choices independent of supermarket availability. The cross-Markov property for this variable (an illustration of which is shown in Figure 5) allows for such cultural preferences to have complex dependence structures across locations. However, grocery store accessibility within a county i , Z_i , is only allowed to depend on such cultural preferences only through its value within the county (conditional on the grocery accessibility in all other locations). This assumption is reasonable for large counties where the food culture in neighboring counties does not directly influence the demand for (and eventual availability of) supermarkets, except through its correlation with the food culture within the county itself. Given the cross-Markov property, the constant conditional correlation assumption implies that the strength of the relationship between this aspect of food culture and supermarket availability is constant (conditional on other observed variables in $\mathbf{X}_{(-z)}$). This can be seen by noting that

$$(23) \quad E[Z_i | \mathbf{Z}_{-i}, U_i] = \mathbf{x}_{(-z),i}^\top \boldsymbol{\gamma} + \frac{\varphi_Z}{|\partial_i|} \sum_{j \in \partial_i} (Z_j - \mathbf{x}_{(-z),j}^\top \boldsymbol{\gamma}) + \rho \sqrt{\frac{\tau_U}{\tau_Z}} U_i.$$

Thus if both the cross-Markov and constant conditional correlation hold, $\sqrt{\tau_U/\tau_Z} U_i$ above behaves like an additive predictor of Z_i with regression coefficient ρ . A plot of residuals from the regression implied by (23) did not indicate any departures from linearity.

An assumption on the spatial scales of unobserved confounders is also critical for reliable identification of causal effects. We can evaluate this assumption within the model by examining the joint posterior distribution of (φ_U, φ_Z) . In our case the posterior distribution of $\varphi_U - \varphi_Z$ from the unconstrained affine model was approximately Gaussian with mean 0.091, standard deviation 0.024 and first percentile 0.030. The posterior from the affine-RS model was similar, indicating that the assumption is satisfied within the scope of the model.

7.2. *Estimating the effect of poor supermarket availability on CVD mortality.* Figure 6 displays the posterior distribution of the exponentiated exposure coefficient from the nonspatial, spatial, spatial-RS, affine and affine-RS models. As noted in Section 5.1, this quantity can be interpreted as the relative expected risk of CVD mortality among the population in the 65+ age range due to a one percentage point increase in poor supermarket access in a randomly-chosen county. The agreement of the posterior distribution of $\varphi_U - \varphi_Z$ in the spatially restricted and unrestricted models implies that the posterior densities from the spatial and spatial-RS models closely coincide (both posterior geometric means 0.999, 95% CIs (0.988, 1.011)), as do those from the affine and affine-RS models (posterior geometric mean 1.005, 95% CI (0.993, 1.018) and 1.005 (0.992, 1.018), respectively). As we saw in Section 2, the nonspatial model reports a definitive, protective effect of poor supermarket availability (posterior mean 0.968, 95% CI (0.962, 0.973)). In contrast, all four spatial models return a smaller and potentially null effect. The spatial and spatial-RS models return approximately null effects with posterior probability of a relative risk smaller than one equal to 0.54 and 0.53, respectively. In contrast, the affine and affine-RS models estimate that poor supermarket access might have a harmful effect on CVD mortality, with posterior probability of a relative risk greater than one equal to 0.8 and 0.79, respectively.




FIG. 6. Posterior densities of county-level relative risk of CVD mortality in the 65+ age range due to a one percentage point increase in the proportion of households with no vehicle and more than one mile from a supermarket or large grocery store.

The change in the point estimate between the spatial random-effect models and the affine models is largely attributable to the posterior distribution of ρ . This distribution is skewed slightly left, with a posterior mean of -0.020 and 95% CI $(-0.040, -0.001)$, indicating confounding by the latent variable U , even though the posteriors of $\exp(\beta_Z)$ from the spatial(-RS) and affine(-RS) models overlap. Among posterior draws from the affine-RS model, the correlation between ρ and $\exp(\beta_Z)$ is -0.448 , indicating that unobserved spatial confounding is likely to bias the spatial estimator downward.

We also fit semiparametric versions of each model following the approach in Section 4.4.2. Due to the nonuniform distribution of exposures, we used a truncated cubic basis for the penalized spline. We retained 10,000 MCMC iterations after 5000 burn-in iterations. The affine-RS model took approximately two hours to fit. Results are shown in Figure 7 and are in broad agreement with those from the generalized linear models. The nonspatial model indicates a protective effect of increasing the proportion of households with poor supermarket access on CVD mortality across the observed exposure range. In contrast, the spatial and affine models indicate null and weakly harmful effects, respectively. In all models the posterior geometric mean of PAERC indicates a protective effect at extremely high levels of poor supermarket access (above 15%), though this is likely an artifact of extrapolation of the spline basis since data are sparse in that range and the credible bands widen dramatically.

8. Discussion. By positing a joint model for the exposure and unmeasured spatially-correlated confounders, we were able to extend commonly-used spatial data analysis tools to mitigate bias due to such confounders. In contrast to existing approaches that indicate a potentially protective effect of poor supermarket access on cardiovascular disease deaths among the elderly at the county level, the proposed approach leads to estimated effects and inferences that are more plausible and in line with subject-matter knowledge which indicate that poor supermarket access is likely to be harmful on cardiovascular health. Our study also contributes to the literature on racial and socioeconomic disparities. Much of the recent research on food access and health makes reference to *food deserts*, census tracts with low




FIG. 7. Posterior geometric mean and pointwise 95% credible bounds for the population average exposure-response curve. Rugs indicate observed exposures. Results from the constrained spatial and unconstrained affine models are similar to those from the unconstrained spatial and constrained affine models, respectively.

income and poor food access (United States Department of Agriculture Economic Research Service (2012)). At the neighborhood level, supermarkets are more prevalent in high-income neighborhoods and in neighborhoods with predominantly white residents compared to those whose residents have lower incomes or are black (Morland et al. (2002)).

At the same time, there exist limitations to our study that extend outside the scope of unobserved confounders. County-level effects of supermarket access can be extremely relevant for policy making and local planning, but they cannot be directly translated to individual-level effects that may also be of substantial scientific interest. Additionally, although outcomes were measured in the 65+ age range, the exposure and covariates were generally measured across all age ranges. Finally, it is unlikely that effects of poor supermarket access would manifest in differences in mortality in the short term (e.g., the following year). Instead, effects are more likely to be cumulative over timespans on the order of many years.

From a statistical perspective, our approach to mitigating bias from unobserved spatial confounders is rooted in the causal inference framework and exploits spatial statistics tools that can be used to *directly adjust* for structured unmeasured confounding. We hope that our work contributes to the growing bridge between spatial data analysis and causal inference. The methodology is intended to be amenable to researchers accustomed to the usual spatial statistics literature but could, potentially, be useful in situations calling for mixed models, more generally, with appropriate modifications. For that reason the proposed approach may be widely applicable to scenarios in spatial statistics, time-series analyses and spatiotemporal settings.

One of the key assumptions in drawing causal conclusions is that of positivity. In the presence of spatial confounders and for positivity to hold, the spatial scale of the confounder must be larger than that of the treatment. Checking the robustness of estimated effects using the affine estimator with and without the spatial scale restriction can provide intuition on the plausibility of positivity due to unmeasured spatial confounders. This should, however, be employed with care, if spatial mediating variables are expected to exist.

A natural question that arises is, whether and at what occasions methodology, which directly adjusts for unmeasured confounding, should be preferred over classical sensitivity analysis. We find that in settings with structured data, such as spatial and temporal data, unmeasured confounders will also be expected to be structured. In those situations we find that methodology that directly adjusts for these variables can provide more accurate effect es-

timates, and strengthen the causal claim of scientific conclusions. An interesting line of future research could extend sensitivity analysis to settings where methods, like the one presented here, alleviate bias from structured confounders, and sensitivity of results in the presence of unstructured covariates is evaluated.

Based on the structural model (3), we discussed a set of assumptions that allowed for identification of the bias correction term using only the observed data while providing a simple expression for the expected value of the confounder conditional on the exposure: joint normality of the spatial confounder and exposure of interest, the cross-Markov and constant conditional correlation assumptions. There are several ways in which these assumptions could be relaxed. First, joint normality may be relaxed by assuming joint normality of an underlying random effect process, with the realizations of both the exposure and covariates arising from other distributions. For example, a latent probit model could be used to assess the effects of binary exposures. Furthermore, it may be possible to use more flexible, non-parametric alternatives to the multivariate normal form of the random effect structure, such as spatial Dirichlet processes, both in point-referenced (Gelfand, Kottas and MacEachern (2005)) and areal (Kottas, Duan and Gelfand (2008)) data. However, questions of identifiability in less parametric settings will likely be difficult to answer. The cross-Markov assumption may be relaxed by, for example, treating the joint distribution of the exposure and confounder as a multivariate conditional autoregressive process (Gelfand and Vounatsou (2003)) and expanding the allowable neighbor relations. The constant conditional correlation assumption may be relaxed by allowing the conditional correlation to vary smoothly in space or based on the number of neighboring locations. In any case, formal treatment of general requirements on the: (a) spatial dependence structure (such as the ring graph in Section 4.5.1), including (b) the cross-Markov structure specifying the conditional dependence between the exposure and the unmeasured confounder and (c) the conditional correlation between the unmeasured confounder and exposure, allowing for identification of the causal effect in the presence of unmeasured confounding is an interesting topic of future research. On a more technical note, priors that place a positive probability on $\tau_U^{-1} = 0$ could be allowed.

Furthermore, even though the structural model in (3) allows for arbitrary interactions among the exposure and measured covariates (allowing, e.g., for treatment effect heterogeneity), it assumes that there are no interactions between measured and unmeasured covariates. The extent to which this assumption can be relaxed is an interesting line of future work, especially in the light of recent results in causal inference for unmeasured confounding (D'Amour (2019), Ogburn, Shpitser and Tchetgen Tchetgen (2019)). Therefore, an interesting question that arises is, "Can we harvest the spatial information of the data to mitigate bias from unmeasured confounders without imposing structural assumptions?"

We consider this to be the most pressing topic of future study: what are the general conditions under which causal effects are identifiable in the presence of unmeasured spatial confounding and to what extent is bias mitigation robust to model misspecification? Within the context of each study, researchers would need to verify whether the set of reasonable (within their context) assumptions suffices for identification of the bias correction term, while we have illustrated situations in which the causal effect is and is not identifiable.

Acknowledgments. We would like to thank Dr. Lucas R.F. Henneman for his contributions to the original version of the manuscript. We would also like to thank Dr. Jim Hodges for his valuable thoughts and input during various stages of this paper. The work presented here was partially supported by grants UL1TR002733 and R01MH118927 from the National Institutes of Health.

SUPPLEMENTARY MATERIAL

Supplement to “Mitigating unobserved spatial confounding when estimating the effect of supermarket access on cardiovascular disease deaths” (DOI: [10.1214/20-AOAS1377SUPPA](https://doi.org/10.1214/20-AOAS1377SUPPA); .pdf). (Schnell and Papadogeorgou (2020a)) PDF file containing detailed data descriptions, derivations, and additional simulation results.

Supplement data and code to “mitigating unobserved spatial confounding when estimating the effect of supermarket access on cardiovascular disease deaths” (DOI: [10.1214/20-AOAS1377SUPPB](https://doi.org/10.1214/20-AOAS1377SUPPB); .zip). (Schnell and Papadogeorgou (2020b)) Data sets and code for reproducing simulations and analyses.

REFERENCES

- AMERICAN HEART ASSOCIATION (2019). Heart disease and stroke statistics—2019 update: A report from the American Heart Association. *Circulation* **139** e56–e528.
- ANTONELLI, J., PARMIGIANI, G. and DOMINICI, F. (2019). High-dimensional confounding adjustment using continuous spike and slab priors. *Bayesian Anal.* **14** 825–848. MR3960772 <https://doi.org/10.1214/18-BA1131>
- BANERJEE, S., CARLIN, B. P. and GELFAND, A. E. (2004). *Hierarchical Modeling and Analysis for Spatial Data*. CRC Press, Boca Raton, FL. MR3362184
- BARON, R. M. and KENNY, D. A. (1986). The moderator-mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *J. Pers. Soc. Psychol.* **51** 1173–1182.
- BELLONI, A., CHERNOZHUKOV, V. and HANSEN, C. (2014). Inference on treatment effects after selection among high-dimensional controls. *Rev. Econ. Stud.* **81** 608–650. MR3207983 <https://doi.org/10.1093/restud/rdt044>
- BESAG, J. (1974). Spatial interaction and the statistical analysis of lattice systems. *J. Roy. Statist. Soc. Ser. B* **36** 192–236. MR0373208
- CENTERS FOR DISEASE CONTROL (2018). Underlying Cause of Death 1999–2017 on CDC WONDER Online Database. <https://wonder.cdc.gov/ucd-icd10.html>.
- CHUNG, Y., RABE-HESKETH, S. and CHOI, I.-H. (2013). Avoiding zero between-study variance estimates in random-effects meta-analysis. *Stat. Med.* **32** 4071–4089. MR3102435 <https://doi.org/10.1002/sim.5821>
- CONGDON, P. (2013). Assessing the impact of socioeconomic variables on small area variations in suicide outcomes in England. *Int. J. Environ. Res. Public Health* **10** 158–177.
- D’AMOUR, A. (2019). On multi-cause causal inference with unobserved confounding: Counterexamples, impossibility, and alternatives. In *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics (AISTATS)* **89**.
- DWYER-LINDGREN, L., MOKDAD, A. H., SREBOTNJAK, T., FLAXMAN, A. D., HANSEN, G. M. and MURRAY, C. J. (2014). Cigarette smoking prevalence in US counties: 1996–2012. *Popul. Health Metr.* **12** 5. <https://doi.org/10.1186/1478-7954-12-5>
- GELFAND, A. E., KOTTAS, A. and MACÉACHERN, S. N. (2005). Bayesian nonparametric spatial modeling with Dirichlet process mixing. *J. Amer. Statist. Assoc.* **100** 1021–1035. MR2201028 <https://doi.org/10.1198/016214504000002078>
- GELFAND, A. E. and VOUNATSOU, P. (2003). Proper multivariate conditional autoregressive models for spatial data analysis. *Biostatistics* **4** 11–15.
- HAHN, J. (1998). On the role of the propensity score in efficient semiparametric estimation of average treatment effects. *Econometrica* **66** 315–331. MR1612242 <https://doi.org/10.2307/2998560>
- HAHN, P. R., CARVALHO, C. M., PUELZ, D. and HE, J. (2018). Regularization and confounding in linear regression for treatment effect estimation. *Bayesian Anal.* **13** 163–182. MR3737947 <https://doi.org/10.1214/16-BA1044>
- HANKS, E. M., SCHLIEP, E. M., HOOTEN, M. B. and HOETING, J. A. (2015). Restricted spatial regression in practice: Geostatistical models, confounding, and robustness under model misspecification. *Environmetrics* **26** 243–254. MR3340961 <https://doi.org/10.1002/env.2331>
- HILL, J. L. (2011). Bayesian nonparametric modeling for causal inference. *J. Comput. Graph. Statist.* **20** 217–240. Supplementary material available online. MR2816546 <https://doi.org/10.1198/jcgs.2010.08162>
- HIRANO, K. and IMBENS, G. W. (2004). The propensity score with continuous treatments. In *Applied Bayesian Modeling and Causal Inference from Incomplete-Data Perspectives*. Wiley Ser. Probab. Stat. 73–84. Wiley, Chichester. MR2134803 <https://doi.org/10.1002/0470090456.ch7>
- HODGES, J. S. and REICH, B. J. (2010). Adding spatially-correlated errors can mess up the fixed effect you love. *Amer. Statist.* **64** 325–334. MR2758564 <https://doi.org/10.1198/tast.2010.10052>

- HUGHES, J. and HARAN, M. (2013). Dimension reduction and alleviation of confounding for spatial generalized linear mixed models. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **75** 139–159. MR3008275 <https://doi.org/10.1111/j.1467-9868.2012.01041.x>
- IMBENS, G. W. and RUBIN, D. B. (2015). *Causal Inference—for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge Univ. Press, New York. MR3309951 <https://doi.org/10.1017/CBO9781139025751>
- KEELE, L., TITIUNIK, R. and ZUBIZARRETA, J. R. (2015). Enhancing a geographic regression discontinuity design through matching to estimate the effect of ballot initiatives on voter turnout. *J. Roy. Statist. Soc. Ser. A* **178** 223–239. MR3291769 <https://doi.org/10.1111/rssa.12056>
- KELLER, J. P. and SZPIRO, A. A. (2019). Selecting a scale for spatial confounding adjustment. arXiv preprint [arXiv:1909.11161](https://arxiv.org/abs/1909.11161).
- KELLI, H. M., HAMMADAH, M., AHMED, H., KO, Y.-A., TOPEL, M., SAMMAN-TAHHAN, A., AWAD, M., PATEL, K., MOHAMMED, K. et al. (2017). Association between living in food deserts and cardiovascular risk. *Circulation: Cardiovascular Quality and Outcomes* **10** e003532.
- KELLI, H. M., KIM, J. H., SAMMAN TAHHAN, A., LIU, C., KO, Y.-A., HAMMADAH, M., SULLIVAN, S., SANDESARA, P., ALKHODER, A. A. et al. (2019). Living in food deserts and adverse cardiovascular outcomes in patients with cardiovascular disease. *Journal of the American Heart Association* **8** e010694.
- KOTTAS, A., DUAN, J. A. and GELFAND, A. E. (2008). Modeling disease incidence data with spatial and spatio-temporal Dirichlet process mixtures. *Biom. J.* **50** 29–42. MR2414636 <https://doi.org/10.1002/bimj.200610375>
- LARAIA, B. A., SIEGA-RIZ, A. M., KAUFMAN, J. S. and JONES, S. J. (2004). Proximity of supermarkets is positively associated with diet quality index for pregnancy. *Prev. Med.* **39** 869–875.
- LEE, D. and SARRAN, C. (2015). Controlling for unmeasured confounding and spatial misalignment in long-term air pollution and health studies. *Environmetrics* **26** 477–487. MR3415567 <https://doi.org/10.1002/env.2348>
- MORLAND, K., WING, S., ROUX, A. D. and POOLE, C. (2002). Neighborhood characteristics associated with the location of food stores and food service places. *Am. J. Prev. Med.* **22** 23–29.
- OGBURN, E. L., SHPITSER, I. and TCHETGEN TCHETGEN, E. J. (2019). Comment on “Blessings of multiple causes” [MR4047282]. *J. Amer. Statist. Assoc.* **114** 1611–1615. MR4047286 <https://doi.org/10.1080/01621459.2019.1689139>
- PACIOREK, C. J. (2010). The importance of scale for spatial-confounding bias and precision of spatial regression estimators. *Statist. Sci.* **25** 107–125. MR2741817 <https://doi.org/10.1214/10-STS326>
- PAPADOGEORGOU, G., CHOIRAT, C. and ZIGLER, C. M. (2019). Adjusting for unmeasured spatial confounding with distance adjusted propensity score matching. *Biostatistics* **20** 256–272. MR3922132 <https://doi.org/10.1093/biostatistics/kxx074>
- PEARCE, J., HISCOCK, R., BLAKELY, T. and WITTEN, K. (2008). The contextual effects of neighbourhood access to supermarkets and convenience stores on individual fruit and vegetable consumption. *J. Epidemiol. Community Health* **62** 198–201.
- POWELL, L. M., AULD, M. C., CHALOUPKA, F. J., O’MALLEY, P. M. and JOHNSTON, L. D. (2007). Associations between access to food stores and adolescent body mass index. *Am. J. Prev. Med.* **33** S301–S307.
- ROBINS, J. (1986). A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. *Math. Model.* **7** 1393–1512. MR0877758 [https://doi.org/10.1016/0270-0255\(86\)90088-6](https://doi.org/10.1016/0270-0255(86)90088-6)
- ROBINS, J. M. and ROTNITZKY, A. (1995). Semiparametric efficiency in multivariate regression models with missing data. *J. Amer. Statist. Assoc.* **90** 122–129. MR1325119
- ROSENBAUM, P. R. (2002). *Observational Studies*, 2nd ed. *Springer Series in Statistics*. Springer, New York. MR1899138 <https://doi.org/10.1007/978-1-4757-3692-2>
- ROSENBAUM, P. R. and RUBIN, D. B. (1983a). Assessing sensitivity to an unobserved binary covariate in an observational study with binary outcome. *J. Roy. Statist. Soc. Ser. B* **45** 212–218.
- ROSENBAUM, P. R. and RUBIN, D. B. (1983b). The central role of the propensity score in observational studies for causal effects. *Biometrika* **70** 41–55. MR0742974 <https://doi.org/10.1093/biomet/70.1.41>
- RUBIN, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *J. Educ. Psychol.* **66** 688–701.
- RUBIN, D. B. (1980). Randomization analysis of experimental data: The Fisher randomization test comment. *J. Amer. Statist. Assoc.* **75** 591–593.
- RUPPERT, D., WAND, M. P. and CARROLL, R. J. (2003). *Semiparametric Regression*. *Cambridge Series in Statistical and Probabilistic Mathematics* **12**. Cambridge Univ. Press, Cambridge. MR1998720 <https://doi.org/10.1017/CBO9780511755453>
- SALLIS, J. F., NADER, P. R., RUPP, J. W., ATKINS, C. J. and WILSON, W. C. (1986). San Diego surveyed for heart-healthy foods and exercise facilities. In *Public Health Reports* (1974-) 216–219.
- SCHNELL, P. M. and PAPADOGEORGOU, G. (2020a). Supplement to “Mitigating unobserved confounding when estimating the effect of supermarket access on cardiovascular disease deaths.” <https://doi.org/10.1214/20-AOAS1377SUPPA>

- SCHNELL, P. M. and PAPADOGEORGOU, G. (2020b). Supplement data and code to “Mitigating unobserved confounding when estimating the effect of supermarket access on cardiovascular disease deaths.” <https://doi.org/10.1214/20-AOAS1377SUPPB>
- THADEN, H. and KNEIB, T. (2018). Structural equation models for dealing with spatial confounding. *Amer. Statist.* **72** 239–252. MR3836447 <https://doi.org/10.1080/00031305.2017.1305290>
- UNITED STATES DEPARTMENT OF AGRICULTURE ECONOMIC RESEARCH SERVICE (2012). Food Environment Atlas. <https://www.ers.usda.gov/data-products/food-environment-atlas>.
- VANDERWEELE, T. J. and DING, P. (2017). Sensitivity analysis in observational research: Introducing the E-value. *Ann. Intern. Med.* **167** 268–274. <https://doi.org/10.7326/M16-2607>
- VERBITSKY-SAVITZ, N. and RAUDENBUSH, S. W. (2012). Causal inference under interference in spatial settings: A case study evaluating community policing program in Chicago. *Epidemiol. Methods* **1** 105–130.
- WILSON, A. and REICH, B. J. (2014). Confounder selection via penalized credible regions. *Biometrics* **70** 852–861. MR3295746 <https://doi.org/10.1111/biom.12203>
- WON, J.-H., LIM, J., KIM, S.-J. and RAJARATNAM, B. (2013). Condition-number-regularized covariance estimation. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **75** 427–450. MR3065474 <https://doi.org/10.1111/j.1467-9868.2012.01049.x>
- ZHOU, T., ELLIOTT, M. R. and LITTLE, R. J. A. (2019). Penalized spline of propensity methods for treatment comparison. *J. Amer. Statist. Assoc.* **114** 1–19. MR3941229 <https://doi.org/10.1080/01621459.2018.1518234>
- ZIGLER, C. M. and DOMINICI, F. (2014). Uncertainty in propensity score estimation: Bayesian methods for variable selection and model-averaged causal effects. *J. Amer. Statist. Assoc.* **109** 95–107. MR3180549 <https://doi.org/10.1080/01621459.2013.869498>