



Contents lists available at [ScienceDirect](#)

Spatial Statistics

journal homepage: www.elsevier.com/locate/spasta



Estimating intervention effects on infectious disease control: The effect of community mobility reduction on Coronavirus spread



Andrew Giffin^{a,*}, Wenlong Gong^a, Suman Majumder^b,
Ana G. Rappold^c, Brian J. Reich^a, Shu Yang^a

^a North Carolina State University, Department of Statistics, 2311 Stinson Drive, Raleigh, NC 27607, United States of America

^b Harvard T.H. Chan School of Public Health, 677 Huntington Ave, Boston, MA 02115, United States of America

^c Environmental Protection Agency, 104 Mason Farm Road, Chapel Hill, NC 27514, United States of America

ARTICLE INFO

Article history:

Received 7 June 2021

Received in revised form 29 January 2022

Accepted 12 October 2022

Available online 21 October 2022

Keywords:

Causal inference

COVID-19

Spatiotemporal modeling

SIR model

ABSTRACT

Understanding the effects of interventions, such as restrictions on community and large group gatherings, is critical to controlling the spread of COVID-19. Susceptible–Infectious–Recovered (SIR) models are traditionally used to forecast the infection rates but do not provide insights into the causal effects of interventions. We propose a spatiotemporal model that estimates the causal effect of changes in community mobility (intervention) on infection rates. Using an approximation to the SIR model and incorporating spatiotemporal dependence, the proposed model estimates a direct and indirect (spillover) effect of intervention. Under an interference and treatment ignorability assumption, this model is able to estimate causal intervention effects, and additionally allows for spatial interference between locations. Reductions in community mobility were measured by cell phone movement data. The results suggest that the reductions in mobility decrease Coronavirus cases 4 to 7 weeks after the intervention.

Published by Elsevier B.V.

* Correspondence to: 4608 46 St NW Washington, DC 20016, United States of America.

E-mail address: giffin.andrew@gmail.com (A. Giffin).

1. Introduction

Since the Coronavirus exploded into a global pandemic, much research has been done to understand its epidemiological characteristics and quantify the effectiveness of various interventions to mitigate disease spread (Li et al., 2020; Livingston and Bucher, 2020; Dandekar and Barbastathis, 2020; Dehning et al., 2020; Cowling et al., 2020; Prem et al., 2020; Carroll and Prentice, 2021; Chen et al., 2020; Lee et al., 2021; Lyu and Wehby, 2020; Nouvellet et al., 2021; Rashed and Hirata, 2021). The traditional model for the progression of an infectious disease uses a set of differential equations to decompose the population at risk into the number of susceptible (S), infected (I), or recovered (R) individuals and provide time-dependent trajectories of these compartments. Potential effects of interventions such as school closure, work from home, social distancing, etc. can be incorporated in the SIR model for the transitions between compartments. For example, there has been substantial effort put forth to combine machine learning and/or deep learning with SIR model to make predictions of the spreading trend (Dandekar and Barbastathis, 2020; Punn et al., 2020; Magdon-Ismail, 2020; Lyu et al., 2020; Kounchev et al., 2020). Similar efforts have attempted to adopt Bayesian approaches to provide uncertainty quantification of predictions (Dehning et al., 2020; Mbuvha and Marwala, 2020; Bradley, 2020).

The SIR model and the above variants do not account for the spatial distribution of the number of susceptible, infected, or recovered people (Kermack and McKendrick, 1927). A number of attempts have been made to modify the SIR model to incorporate this spatial distribution. One strategy considers diffusion of the infected population by modeling the individual agents (Reluga, 2004; Burger et al., 2009). Another branch of the literature considers Brownian motion-like movement for subgroups of people belonging to each compartment (Chinviriyasit and Chinviriyasit, 2010; Hilker et al., 2007; Robinson et al., 2012; Wang et al., 2010); while yet another group uses cross-diffusion an assumption where the susceptible move away from an increasing gradient of the infected to bring in a spatial distribution to the SIR model (Berres and Ruiz-Baier, 2011; Milner and Zhao, 2008; Capasso and Di Liddo, 1994). Models that assume that patches of population have different infection parameters and are connected by travel or migration have also been proposed (Arino and Van den Driessche, 2003; Hyman and LaForce, 2003; Sattenspiel et al., 1995; Sattenspiel and Herring, 2003; Lee et al., 2012; Lee and Castillo-Chavez, 2015; Lee and Jung, 2015).

Our method is based on a spatial SIR model that uses an “infectious radius” to allow for spread across locations (Paeng and Lee, 2017). This method is a model discrete in time and with space partitioned into distinct areas (e.g., counties). The number of individuals in each compartment and area evolves over time using the same mechanics as the standard SIR model. A distance metric is defined between counties, allowing the number of infected in county j to account for dependence across nearby areas.

Our statistical contribution is using the SIR model as a stepping stone towards a spatiotemporal statistical model that permits causal inference without including differential-equation solvers. The result is a model with a very different flavor than the standard SIR model. While the goal of an SIR model is often to forecast the severity of a disease over time, the model presented here is designed to tease out causal effects of interventions.

Integrating causal inference into this spatial framework presents some challenges. Existing SIR models are generally used for prediction and forecasting, but not for ascertaining causal effects of interventions. Causal methods generally require adjusting for confounders that are related to both intervention and response, which allows for identification of the causal intervention effect as opposed to the observed association between intervention and response. However, if interventions at one location affect the response at other locations – a phenomenon known as interference – then the standard causal framework which adjusts only for local confounding is rendered ineffective. With a virus that spreads across space, clearly some degree of interference is involved, and a causal analysis must take this into account. This is typically done by placing assumptions on the form of interference. Commonly used assumptions on interference include “partial interference” in which the population is divided into isolated clusters, as well as “network interference” in which interference is allowed along a pre-specified network, and “spatial interference” in which interference is tempered by the distance between locations (Reich et al., 2020; Halloran and

Struchiner, 1991; Sobel, 2006; Giffin et al., 2020). This paper uses a parsimonious assumption on interference, which allows for interference from neighboring counties only. This can be thought of as a simple version of both the spatial and network interference assumptions. Intervention effects are then divided between direct (within county) interventions and indirect (between county) interventions.

The framework that we propose approximates the standard SIR model, in a way that allows for estimation of the causal effects of interventions. In particular, we use the fact that infection rates are likely smooth across space, to obtain an approximation to the rate of infection at across space and time. This approximation is a function of the intervention variables, covariates, and a spatiotemporal model, and allows us to fit the observed data and obtain intervention effect estimates.

The remainder of the paper proceeds as follows. The motivating data are described in Section 2. Section 3 introduces the statistical method and corresponding theoretical properties. The method is then evaluated using a simulation study in Section 4 where we show that the approximate model can recover causal effects of data generated from the SIR model. Section 5 applies the proposed model to estimate the effect of decreases in mobility on Coronavirus cases, finding that decreases in mobility cause a decrease in local Coronavirus cases 4–7 week later. Section 6 concludes.

2. Data description

We estimate the effect of mobility on the number of observed Coronavirus cases. There are three primary types of data for analysis: intervention (mobility reduction), response (Coronavirus cases) and covariates. All variables are defined temporally by the week (t) and spatially by the county (j). Data from March 6, 2020 through October 8, 2020 for 3,137 counties or county equivalents are used in this study.

The response data $Y_j(t)$ are new, recorded cases of Coronavirus in a given county/week. These data are taken from the publicly available Johns Hopkins University Coronavirus Resource Center, which aggregates Coronavirus case counts and provides a daily, cumulative count of cases in each county (Dong et al., 2020).

The intervention data are publicly available measures of mobility – as measured by Google devices – which have been shown to have strong associations with Coronavirus case data (Google LLC, 2020; Chang et al., 2021; Yilmazkuday, 2020). We use an aggregate mobility measure that includes the categories: “retail/recreation”, “grocery/pharmacy”, “transit station”, “workplace”, and “residential” mobility. The first four categories measure the number of visits to such locations in comparison to baseline; the “residential mobility” category is a measure of the length of stay in comparison to baseline. These data are given as percentage change from a baseline level which was taken over January 3, 2020–February 6, 2020. Intuitively, in the months since the pandemic began, “retail/recreation”, “grocery/pharmacy”, “transit station” and “workplace” mobility have decreased substantially from the baseline and “residential” mobility (i.e., amount of time spent at home) has increased. The metric that we will use as intervention $A_j(t)$ for a given county/week is the mean of percentage change from baseline of “retail/recreation”, “grocery/pharmacy”, “transit station”, “workplace”, and *negative* “residential” mobility. We note that the mean over these categories is a somewhat crude measure of mobility; however, it is used so as to reduce variability coming from any single category.

Because of privacy concerns, county/days with too few users are not provided by Google. When a subset of the categories are provided for a given county/day, the mean of the available categories is taken. When none of the categories are provided for a given count/day, we impute the missing value from any available first and second degree neighbors (with values weighted by 1 divided by the distance between centroids, where the available weights are first standardized so that they sum to 1).

In addition we include a number of static and time-varying covariates which are potentially related to virus spread or mobility. For the static covariates, we rely heavily on a curated dataset compiled by Wu et al. 2020, which examines the relationship between air pollution and COVID-19 (Wu et al., 2020a,b). Variables from this dataset include poverty rate, population density, median household income, and total population, originally taken from the 2016 American Community

Survey, as well as data on the number of hospital beds. Since this and other works have shown that PM_{2.5} (particulate matter smaller than 2.5 micrometers) may interact significantly with COVID-19, we included the county level PM_{2.5} estimates from 2016 from this dataset (Alexeeff et al., 2015; Chudnovsky et al., 2012). To this dataset we added a number of other variables. The number of ICU beds in 2020 was compiled by Environmental Systems Research Institute (ESRI) (ESRI, 2020). Additionally, from the 2018 American Community Survey (ACS) we included median age and the number of foreign born residents (US Census Bureau, 2016). From the Bureau of Labor Statistics we included the number of people employed in healthcare and social services in September 2019 (US Bureau of Labor Statistics, 2019). We also included time-varying county-level daily meteorological data, as this could inform levels of mobility and spread of disease. We use daily mean temperature (degrees Celsius), as well as relative humidity and mean dew point (degrees Celsius) from NOAA’s Global Surface Summary of the Day using the “GSODR” R package (Sparks et al., 2017). Then we obtain county-level data by predicting at the mean county latitude/longitude coordinates from the 2010 Census using thin plate spline regression from the “fields” R package (Nychka et al., 2014).

3. Main methodology

3.1. Notation

Let $Y_j(t)$ be the number of new cases of Coronavirus reported during week $t \in \{1, \dots, T\}$ and county $j \in \{1, \dots, J\}$, where $T = 31$ and $J = 3,137$. For county j , denote N_j as the population size and $\mathbf{X}_j(t)$ as a vector of covariates for week t . The direct intervention variable, $A_j(t)$, is the mobility percentage change from baseline for county j and week t . In addition to the direct intervention, we allow for an indirect (spillover) intervention $\tilde{A}_j(t)$. This variable captures the interventions received by neighboring counties. Including this term allows for interference between neighbors, since the response at county j can now be impacted by the interventions in surrounding counties. The spatial configuration of the counties is summarized by their adjacency, with $c_{jj} = 0$, $c_{jk} = 1$ if counties j and k are adjacent and $c_{jk} = 0$ otherwise. $\tilde{A}_j(t)$ is defined as the mean direct intervention over the $m_j = \sum_{k=1}^J c_{jk}$ adjacent regions: $\tilde{A}_j(t) = \sum_{k=1}^J c_{jk}A_k(t)/m_j$. Similarly, the mean of the adjacent covariate values is $\tilde{\mathbf{X}}_j(t) = \sum_{k=1}^J c_{jk}\mathbf{X}_k(t)/m_j$.

3.2. Conceptual SIR model

To motivate the proposed statistical framework, we begin by describing the discrete-time, spatial SIR model proposed in Paeng and Lee (2017). Let $S_j(t)$, $I_j(t)$ and $R_j(t)$ be the number of susceptible, infected and recovered individuals in county j at time t . These three states evolve over time but always satisfy the constraint that $N_j = S_j(t) + I_j(t) + R_j(t)$. For each county j the state evolution is similar to the classical SIR model (Kermack and McKendrick, 1927):

$$\begin{aligned} S_j(t + 1) - S_j(t) &= -\lambda_j(t), \\ I_j(t + 1) - I_j(t) &= \lambda_j(t) - \gamma I_j(t), \\ R_j(t + 1) - R_j(t) &= \gamma I_j(t), \end{aligned} \tag{1}$$

where $\gamma_j > 0$ is the recovery rate and $\lambda_j(t) = \beta \frac{S_j(t)}{N_j} \left\{ \sum_{k=1}^J W_{jk} I_k(t) \right\}$ is the rate new infections in county j and week t . Here W_{jk} is proportional to the contact rate between individuals in region j with individuals in region k .

We expand on this by allowing the infection rate $\beta_j(t) > 0$ to vary with region and time:

$$\lambda_j(t) = \beta_j(t) \frac{S_j(t)}{N_j} \left\{ \sum_{k=1}^J W_{jk} I_k(t) \right\}. \tag{2}$$

This allows us to model spatiotemporal variation in $\beta_j(t)$ as a function of the covariates $\mathbf{X}_j(t)$ and direct and indirect intervention variables $A_j(t)$ and $\tilde{A}_j(t)$. In the absence of other information on

connectivity, we simply set $W_{ij} = (1 - \phi)$ and $W_{jk} = \phi \cdot c_{jk}/m_j$ for $\phi \in [0, 1]$ so that large ϕ leads to strong spatial dependence and vice versa.

A major difficulty in using model (1) is that we do not observe the latent states $S_j(t)$, $I_j(t)$ or $R_j(t)$. We learn about these latent processes via the reported number of new cases $Y_j(t)$. Further complicating the statistical model is that there may be under-reporting and a lag time between an increase in the true and reported infection rate due to the latency of the disease. Because of these issues, we link $Y_j(t)$ and $\lambda_j(t)$ in the SIR model by the over-dispersed Poisson distribution

$$Y_j(t) | \lambda_j(t), g_j(t) \stackrel{indep}{\sim} \text{Poisson} [p \exp\{g_j(t)\} \lambda_j(t - l)], \tag{3}$$

where $p \in (0, 1)$ accounts for under-reporting, $g_j(t) \stackrel{iid}{\sim} \text{Normal}(0, \tau^2)$ captures over-dispersion, and $l \in \{0, 1, 2, \dots\}$ is the reporting lag. Setting $l > 0$ allows for cases to be observed after the actual infection, as is common in Coronavirus cases.

3.3. Gaussian approximation

The SIR model in Section 3.2 is an elegant way to model and forecast the spread of a disease through a population. However, the solution to the difference equations is complex, which hinders the ability to estimate model parameters in a statistical manner. Previously methods have been proposed to mimic the mechanistic model to provide realistic forecasts, e.g., [Buckingham-Jeffery et al. \(2018\)](#). Here, we propose an approximation to the SIR model to allow for estimation of intervention effects on local infection rates.

The key insight is that the rate of new infections in (2) can be re-expressed as

$$\lambda_j(t) = \beta_j(t) \exp\{\theta_j(t)\} + v_j(t), \tag{4}$$

$$\exp\{\theta_j(t)\} = S_j(t) I_j(t) / N_j, \tag{5}$$

$$v_j(t) = \beta_j(t) \phi \frac{S_j(t)}{N_j} \sum_{k=1}^J \frac{c_{jk}}{m_j} \{I_k(t) - I_j(t)\}. \tag{6}$$

Both $\theta_j(t)$ and $v_j(t)$ in (4) are latent spatial and temporal processes due to the unobserved $S_j(t)$ and $I_j(t)$, which we approximate by spatiotemporal models.

We model $\theta_j(t)$ as a separable spatiotemporal continuous autoregressive model ([Stein, 2005](#)). Denote $\theta(t) = \{\theta_1(t), \dots, \theta_j(t)\}$ and $\boldsymbol{\theta} = \{\theta(1)^T, \dots, \theta(T)^T\}^T$. Then the spatiotemporal conditional autoregressive model (STCAR) for $\boldsymbol{\theta}$ is multivariate normal with mean zero and covariance $\sigma^2 \Sigma(\rho_t) \otimes \Omega(\rho_s)$. Spatial dependence is governed by $\Sigma(\rho_s) = (\mathbf{M} - \rho_s \mathbf{C})^{-1}$, where \mathbf{M} is diagonal with diagonal elements m_j and the (j, k) element of \mathbf{C} is 1 if sites j and k are adjacent and zero otherwise. If $\boldsymbol{\theta}(t) \sim \text{Normal}\{0, \sigma^2 \Sigma(\rho_s)\}$, then $\theta_j(t) | \theta_k(t)$ for all $k \neq j \sim \text{Normal}(\rho_s \sum_{k \sim j} \theta_k(t) / m_j, \sigma^2 / m_j)$, so that ρ_s determines the strength of spatial dependence and σ is the scale parameter. We denote this model as $\boldsymbol{\theta}(t) \sim \text{CAR}(\sigma, \rho_s)$. Similarly, temporal dependence is governed by $\Omega(\rho_t)$, which has the same form as $\Sigma(\rho_s)$ except for temporal adjacency structure with time t having $t - 1$ and $t + 1$ considered neighbors. We refer to this as the STCAR(σ, ρ_s, ρ_t) model.

The second term in (4), $v_j(t)$, sums over the m_j regions adjacent to region j , and is a function of the local differences, $I_k(t) - I_j(t)$. Assuming the number of infected individuals varies smoothly across space, these local differences should be small. Therefore, one approximation we consider is simply setting $v_j(t) = 0$ for all j and t . For cases where these terms cannot be removed, we note that because they are functions of local differences they should be less spatially correlated than the spatial process itself, and thus a second approximation is $\log\{v_j(t)\} \stackrel{iid}{\sim} \text{Normal}(\mu_v, \sigma_v^2)$. This form implicitly forces the $v_j(t)$ term in (7) to be positive. This approximation is justified because this term will be close to zero for spatially-correlated processes, as can be seen in (6). Approximating $v_j(t)$ as Gaussian is another option, but this would require additional constraints to ensure $\lambda_j(t) > 0$. Modeling $v_j(t)$ as log-normal solves this issue and provides accurate results as shown in Section 4.

A crucial point is that neither our model for $\theta_t(\mathbf{s})$ nor $v_t(\mathbf{s})$ attempt to retain the mechanistic properties of $I_j(t)$ and $S_j(t)$. Therefore, this model would not provide reliable forecasts about future

disease spread. However, forecasting is not our objective. Our aim is to provide a computationally feasible method to use observed data to estimate effects of intervention variables on local infection rates. The following section verifies that the approximation indeed has a causal interpretation under standard assumptions from causal inference, and proposes an estimation procedure for these effects.

For the final component, modeling the spatiotemporal infection rate $\beta_j(t)$, we regress $\log\{\beta_j(t)\}$ on covariates $\mathbf{X}_j(t)$ and $\tilde{\mathbf{X}}_j(t)$ as well as mobility variables $A_j(t)$ and $\tilde{A}_j(t)$ as

$$\eta_j(t) = \log\{\beta_j(t)\} = \alpha_0 + \mathbf{X}_j(t)^\top \boldsymbol{\alpha}_1 + \tilde{\mathbf{X}}_j(t)^\top \boldsymbol{\alpha}_2 + A_j(t)\delta_1 + \tilde{A}_j(t)\delta_2$$

where δ_1 and δ_2 quantify the direct and indirect (spillover) causal effects of mobility on infection rate. The covariate vectors $\mathbf{X}_j(t)$ and $\tilde{\mathbf{X}}_j(t)$ can include both the original covariates or covariates derived as propensity scores (Section 3.4).

Since it is impossible to identify the reporting rate p in (3) and intercept terms α_0 and μ_v we fit the final model

$$Y_j(t) | g_j(t), \boldsymbol{\theta}, v_j(t) \stackrel{\text{indep}}{\sim} \text{Poisson} \left[\exp\{g_j(t) + \eta_j(t-l) + \theta_j(t-l)\} + \exp\{\tilde{v}_j(t)\} \right] \tag{7}$$

$$\eta_j(t) = \alpha_0 + \mathbf{X}_j(t)^\top \boldsymbol{\alpha}_1 + \tilde{\mathbf{X}}_j(t)^\top \boldsymbol{\alpha}_2 + A_j(t)\delta_1 + \tilde{A}_j(t)\delta_2 \tag{8}$$

$$g_j(t) \stackrel{\text{iid}}{\sim} \text{Normal}(0, \tau^2)$$

$$\boldsymbol{\theta} \sim \text{STCAR}(\sigma^2, \rho_s, \rho_t)$$

$$\tilde{v}_j(t) \stackrel{\text{iid}}{\sim} \text{Normal}(\mu_v, \sigma_v^2).$$

The final term in (7) combines the overdispersion term $g_j(t)$ and the nugget term $v_j(t)$. When the dimensions of the covariates are large, we propose dimension reduction for $\log\{\beta_j(t)\}$ using the generalized propensity score (Imbens, 2000) in Section 3.4. To complete the Bayesian model, we specify noninformative independent $\text{Normal}(0, 10^2)$ priors for $\alpha_0, \delta_1, \delta_2, \mu_v$ and the elements of $\boldsymbol{\alpha}_1$ and $\boldsymbol{\alpha}_2$, and noninformative priors for covariance parameters $\sigma^2, \tau^2, \sigma_v^2 \sim \text{InvGamma}(0.1, 0.1)$ and $\rho_s, \rho_t \sim \text{Uniform}(0, 1)$.

3.4. The potential outcomes framework

In this section, we provide a causal interpretation of δ_1 and δ_2 under the potential outcomes framework (Rubin, 1974). To ease discussion, let $Y_j(t)$ be the total number of cases, rather than reported number (i.e., $p = 1$). We use the potential outcomes framework to define the causal effect of mobility on infection rate. We use overbar to denote all history. Let $\bar{\mathbf{a}}_j(t) = (a_j(1), \dots, a_j(t))^\top$ be the trajectory of mobility level at region j through time t . Let $\bar{\mathbf{a}}(t) = (\bar{\mathbf{a}}_1(t), \dots, \bar{\mathbf{a}}_j(t))$ be the trajectory of mobility level for all regions through time t . Define $Y_j^{\bar{\mathbf{a}}(t)}(t)$ to be the (possibly counterfactual) number of new cases in region j at time t had all the regions controlled the mobility level at $\bar{\mathbf{a}}(t)$ through time t .

Assume the potential outcome model for the SIR model with $\theta_j(t) = \log\{S_j(t)I_j(t)/N_j\}$ and $v_j(t) = \beta_j(t) \frac{S_j(t)}{N_j} \sum_{k=1}^J \frac{c_{jk}}{m_j} \{I_k(t) - I_j(t)\}$ as in (5) and (6), and

$$Y_j^{\bar{\mathbf{a}}(t)} \sim \text{Poisson}\{\lambda_j^{\bar{\mathbf{a}}(t)}\},$$

$$\lambda_j^{\bar{\mathbf{a}}(t)}(t) = \beta_j^{\bar{\mathbf{a}}(t)}(t) \exp\{\theta_j(t)\} + v_j(t), \tag{9}$$

$$\log\{\beta_j^{\bar{\mathbf{a}}(t)}(t)\} = a_j(t)\delta_1 + \tilde{a}_j(t)\delta_2 + h\{\bar{\mathbf{X}}(t), \bar{\mathbf{S}}(t-1), \bar{\mathbf{I}}(t-1)\}.$$

Consider two regimes $\bar{\mathbf{a}}(t)$ and $\bar{\mathbf{a}}'(t)$, where all components are the same except for $a_j(t) = a'_j(t) + 1$, and $v_j(t) = 0$. Then under model (9),

$$\frac{E\{Y_j^{\bar{\mathbf{a}}(t)}(t) \mid \bar{\mathbf{X}}(t), \bar{\mathbf{S}}(t-1), \bar{\mathbf{I}}(t-1)\}}{E\{Y_j^{\bar{\mathbf{a}}'(t)}(t) \mid \bar{\mathbf{X}}(t), \bar{\mathbf{S}}(t-1), \bar{\mathbf{I}}(t-1)\}} = \exp(\delta_1),$$

where $\bar{\mathbf{X}}(t)$ is the full history of covariates and intervention through time t , and $\bar{\mathbf{S}}(t-1)$ and $\bar{\mathbf{I}}(t-1)$ are the full history of susceptible and infected, excluding time t . (Because $S_j(t-1) + I_j(t-1) +$

$R_j(t - 1) = N_j$, observing $\bar{\mathbf{S}}(t - 1)$ and $\bar{\mathbf{I}}(t - 1)$ implies $\bar{\mathbf{R}}(t - 1)$.) Then $\exp(\delta_1)$ entails the risk ratio of new cases by increasing mobility level by 1 unit locally (direct effect). If $v_j(t)$ is nonzero but small in comparison to $\beta_j^{\bar{\mathbf{a}}(t)}(t) \exp\{\theta_j(t)\}$, this should hold as a good approximation.

Similarly, if we consider two regimes $\bar{\mathbf{a}}(t)$ and $\bar{\mathbf{a}}'(t)$ where all components are the same except for $\bar{a}_j(t) = \bar{a}'_j(t) + 1$, and $v_j(t) = 0$, then

$$\frac{E\{Y_j^{\bar{\mathbf{a}}(t)}(t) \mid \bar{\mathbf{X}}(t), \bar{\mathbf{S}}(t - 1), \bar{\mathbf{I}}(t - 1)\}}{E\{Y_j^{\bar{\mathbf{a}}'(t)}(t) \mid \bar{\mathbf{X}}(t), \bar{\mathbf{S}}(t - 1), \bar{\mathbf{I}}(t - 1)\}} = \exp(\delta_2)$$

encodes the risk ratio of new cases by increasing mobility level by 1 unit non-locally (indirect effect).

For the effect from neighbor k , consider two regimes $\bar{\mathbf{a}}(t)$ and $\bar{\mathbf{a}}'(t)$, where all components are the same, except for $a_k(t)$ and $a'_k(t) + 1$, and $v_j(t) = 0$. Then

$$\frac{E\{Y_j^{\bar{\mathbf{a}}(t)}(t) \mid \bar{\mathbf{X}}(t), \bar{\mathbf{S}}(t - 1), \bar{\mathbf{I}}(t - 1)\}}{E\{Y_j^{\bar{\mathbf{a}}'(t)}(t) \mid \bar{\mathbf{X}}(t), \bar{\mathbf{S}}(t - 1), \bar{\mathbf{I}}(t - 1)\}} = \exp(\delta_2 c_{jk} / m_j),$$

which encodes the indirect effect of area k on area j .

To link the potential variables and the observed variables and identify the direct and indirect effects, we invoke the following assumptions.

Assumption 1 (Interference). Given the past information through t , the intervention $\mathbf{A}(t)$ affects the potential infection rate at time t and region j only through $A_j(t)$ and $\bar{A}_j(t)$.

Assumption 2 (Causal Consistency). $Y_j(t) = Y_j^{\bar{\mathbf{A}}(t)}(t), \forall t, j$.

Assumption 3 (Ignorability of Intervention Process Variables). $A_j(t), \bar{A}_j(t) \perp\!\!\!\perp Y_j^{\bar{\mathbf{a}}(t)}(t) \mid \{\bar{\mathbf{X}}(t), \bar{\mathbf{S}}(t - 1), \bar{\mathbf{I}}(t - 1)\}, \forall t, j$.

We further define a generalized propensity score for the direct and indirect effects

$$\bar{e}_{jt} \{\kappa_1, \kappa_2; \bar{\mathbf{X}}(t), \bar{\mathbf{S}}(t - 1), \bar{\mathbf{I}}(t - 1)\} = f\{A_j(t) = \kappa_1, \bar{A}_j(t) = \kappa_2, \mid \bar{\mathbf{X}}(t), \bar{\mathbf{S}}(t - 1), \bar{\mathbf{I}}(t - 1)\}.$$

In practice, we often assume that the direct and indirect interventions are independent, and can write this score as the two separate components

$$e_{jt} \{\kappa; \bar{\mathbf{X}}(t), \bar{\mathbf{S}}(t - 1), \bar{\mathbf{I}}(t - 1)\} = f\{A_j(t) = \kappa \mid \bar{\mathbf{X}}(t), \bar{\mathbf{S}}(t - 1), \bar{\mathbf{I}}(t - 1)\},$$

$$\bar{e}_{jt} \{\kappa; \bar{\mathbf{X}}(t), \bar{\mathbf{S}}(t - 1), \bar{\mathbf{I}}(t - 1)\} = \bar{f}\{\bar{A}_j(t) = \kappa \mid \bar{\mathbf{X}}(t), \bar{\mathbf{S}}(t - 1), \bar{\mathbf{I}}(t - 1)\},$$

denoted by $e_{jt} \{\kappa; \bar{\mathbf{X}}(t), \bar{\mathbf{S}}(t - 1), \bar{\mathbf{I}}(t - 1)\}$ and $\bar{e}_{jt} \{\kappa; \bar{\mathbf{X}}(t), \bar{\mathbf{S}}(t - 1), \bar{\mathbf{I}}(t - 1)\}$, respectively.

Lastly we require a positivity assumption that ensures all regions can take any plausible mobility level. This is formulated based on the propensity score:

Assumption 4 (Positivity). For all $\bar{\mathbf{X}}(t), \bar{\mathbf{S}}(t - 1), \bar{\mathbf{I}}(t - 1)$ with $f\{\bar{\mathbf{X}}(t), \bar{\mathbf{S}}(t - 1), \bar{\mathbf{I}}(t - 1)\} > 0, \bar{e}_{jt} \{\kappa_1, \kappa_2; \bar{\mathbf{X}}(t), \bar{\mathbf{S}}(t - 1), \bar{\mathbf{I}}(t - 1)\} > 0, \forall t, \kappa_1, \kappa_2$ where $f(\cdot)$ denotes the generic probability function; i.e., it is a probability density function for a continuous variable and a probability mass function for a discrete variable.

Under **Assumptions 2** and **3**, the induced model from (9) for $Y_j(t)$ is $Y_j(t) \sim \text{Poisson}\{\lambda_j(t)\}$. One can fit the induced model with the observed data to estimate the causal parameters. Because $\bar{\mathbf{X}}, \bar{\mathbf{I}}$, and $\bar{\mathbf{S}}$ are high-dimensional, fitting the above model directly may become a daunting task. We consider the generalized propensity score approach for dimension reduction. Under **Assumption 3**, Imbens' generalized propensity score approach, we can show that

$$A_j(t), \bar{A}_j(t) \perp\!\!\!\perp \{\bar{\mathbf{X}}(t), \bar{\mathbf{S}}(t - 1), \bar{\mathbf{I}}(t - 1)\} \mid \bar{e}_{jt} \{\kappa_1, \kappa_2; \bar{\mathbf{X}}(t), \bar{\mathbf{S}}(t - 1), \bar{\mathbf{I}}(t - 1)\}, \forall \kappa_1, \kappa_2. \text{ (Balancing)}$$

In practice, we model the two generalized propensity scores using linear regressions of $A_j(t)$ and $\bar{A}_j(t)$ onto covariates and past interventions and responses (Giffin et al., 2020). Then, their

distributions for a given j and t (i.e., their generalized propensity scores values) can be completely summarized with a single sufficient statistic: the estimated conditional mean (i.e., the fitted value from the regression). For this reason, henceforth $e_j(t)$ and $\tilde{e}_j(t)$ (without the index for κ) will refer to these sufficient conditional means, for which the result (Balancing) applies. This shows that the generalized propensity score serves as a balancing score; i.e., at the same level of the propensity score, the distribution of the history of confounding variables are the same across different intervention levels.

Allowing \mathbf{X} and $\tilde{\mathbf{X}}$ to include propensity score components, the induced model for $Y_j(t)$ given the intervention and generalized propensity score is

$$E\{Y_j(t) \mid \bar{\mathbf{A}}(t), \bar{\mathbf{e}}(t), \tilde{\mathbf{V}}(t-1)\} = \exp\{A_j(t)\delta_1 + \tilde{A}_j(t)\delta_2 + \alpha_1 + \mathbf{X}_j(t)^\top \alpha_2 + \tilde{\mathbf{X}}_j(t)^\top \alpha_3\} \cdot \exp\{\theta_j(t)\} + v_j(t).$$

The derivation is provided in the Supporting Information. Therefore, we can fit the model $Y_j(t) \sim \text{Poisson}\{\lambda_j^e(t)\}$, where

$$\begin{aligned} \lambda_j^e(t) &= \beta_j^e(t) \exp\{\theta_j(t)\} + v_j(t), \\ \beta_j^e(t) &= \exp\{A_j(t)\delta_1 + \tilde{A}_j(t)\delta_2 + \alpha_1 + \mathbf{X}_j(t)^\top \alpha_2 + \tilde{\mathbf{X}}_j(t)^\top \alpha_3\}, \end{aligned} \tag{10}$$

To model the propensity of A_t given the past history, we require some structural assumption to reduce the dimension of the historical variables. We posit a model for $A_j(t)$ adjusting for its direct causes $A_j(t-1)$, $A_j(t-2)$, $X_j(t)$, $X_j(t-1)$ and $Y_j(t-1)$ to get $e_j(t)$. $\tilde{A}_j(t)$ can be modeled similarly with the neighbor-averaged variables $\tilde{A}_j(t-2)$, $\tilde{X}_j(t)$, $\tilde{X}_j(t-1)$ and $\tilde{Y}_j(t-1)$. In practice, simply taking the average over neighbors of $e_j(t)$ gives similar estimates for $\tilde{e}_j(t)$.

4. Simulation study

In this section we conduct a simulation study to evaluate the statistical properties of the causal effects that stem from the approximate SIR models. Data are generated from the full SIR model in Section 3.2 and analyzed using the approximated spatial model in Section 3.3. The objectives are to determine when the approximations give unbiased point estimation and valid interval estimation, and to compare different approximations using these criteria. The code required to run the simulation study is provided on the Harvard Dataverse (Giffin, 2022).

4.1. Methods

We generate data on a 15×15 square grid of $J = 225$ regions for $T = 30$ time steps. Rook neighbors are considered adjacent. Each region has population $N_j = 100,000$ and initial states are generated as $I_j(1) = 100 \exp\{U_j\}$, $R_j(1) = 0$ and $S_j(1) = N_j - I_j(1)$, where $(U_1, \dots, U_J) \sim \text{CAR}(1, \rho_s)$. A confounding variable $X_j(t)$ is generated from the STCAR(1, ρ_s, ρ_t) distribution and the intervention variable is generated as $A_j(t) = \rho_x X_j(t) + \sqrt{1 - \rho_x^2} E_j(t)$ where $E_j(t)$ is also generated from the STCAR(1, ρ_s, ρ_t) distribution. The latent states for times $t \in \{2, \dots, T\}$ are generated from the full mechanistic model given by (1) and (2) with recovery rate $\gamma = 0.1$ and infection rate

$$\log\{\beta_j(t)\} = \alpha_0 + X_j(t)\alpha_1 + \tilde{X}_j(t)\alpha_2 + A_j(t)\delta_1 + \tilde{A}_j(t)\delta_2$$

for $\alpha_0 = -3$, $\alpha_1 = 0.5$, $\alpha_2 = 0.3$, $\delta_1 = 0.5$ and $\delta_2 = 0.2$. The data are then generated as $Y_j(t) \sim \text{Poisson}\{p\lambda_j(t-l)\}$ where the reporting rate is $p = 0.5$ and the lag is $l = 2$.

We compare six simulation scenarios defined by the spatial (ρ_s) and temporal (ρ_t) dependence of the intervention variable $A_j(t)$, the strength of confounding (ρ_x) and the strength of spatial dependence (ϕ) in the SIR model:

1. Base model: $\rho_s = 0.9$, $\rho_t = 0.5$, $\rho_x = 0.5$ and $\phi = 0.4$
2. Strong spatial dependence in $A_j(t)$: $\rho_s = 0.99$, $\rho_t = 0.5$, $\rho_x = 0.5$ and $\phi = 0.4$
3. Strong temporal dependence in $A_j(t)$: $\rho_s = 0.3$, $\rho_t = 0.9$, $\rho_x = 0.5$ and $\phi = 0.4$

Table 1

Spatiotemporal model simulation study results. The six data-generating scenarios are given in Section 4.1. The two recommended models, the full model (“Full”) and the model without a nugget term (“No nugget”), are given in bold; and are compared to the model without the generalized propensity score (“No PS”) and the model without spatial dependence (“Non-spatial”) in terms of bias and empirical coverage of 90% prediction intervals, separately for the direct (δ_1) and indirect (δ_2) effects. Bias is multiplied by 100 and standard errors are given in subscripts.

| Scenario | Method | Direct effect | | Indirect effect | |
|---|------------------|-----------------------------|------------------------|------------------------------|------------------------|
| | | Bias | 90% Cov. | Bias | 90% Cov. |
| 1. Base Model | Full | 0.25 _{0.25} | 90 ₃ | 0.36 _{0.45} | 93 ₃ |
| | No nugget | 0.12 _{0.24} | 94 ₂ | 0.36 _{0.46} | 94 ₂ |
| | No PS | 2.18 _{0.25} | 74 ₄ | 2.19 _{0.53} | 79 ₄ |
| | Non-spatial | 0.52 _{0.25} | 96 ₂ | -0.41 _{0.64} | 80 ₄ |
| 2. Strong spatial dependence in $A_j(t)$ | Full | 0.79 _{0.24} | 92 ₃ | 0.25 _{0.44} | 91 ₃ |
| | No nugget | 0.60 _{0.23} | 91 ₃ | 0.09 _{0.43} | 92 ₃ |
| | No PS | 2.87 _{0.23} | 65 ₅ | 1.64 _{0.54} | 78 ₄ |
| | Non-spatial | 1.59 _{0.26} | 99 ₁ | -1.02 _{0.71} | 73 ₄ |
| 3. Strong temporal dependence in $A_j(t)$ | Full | 1.16 _{0.47} | 95 ₂ | 0.62 _{0.96} | 93 ₃ |
| | No nugget | 1.06 _{0.47} | 94 ₂ | 0.73 _{0.94} | 93 ₃ |
| | No PS | 12.71 _{0.45} | 2 ₁ | 14.28 _{1.01} | 20 ₄ |
| | Non-spatial | 1.78 _{0.53} | 97 ₂ | -0.06 _{1.29} | 83 ₄ |
| 4. Strong spatiotemporal dependence in $A_j(t)$ | Full | 2.15 _{0.53} | 86 ₄ | 3.63 _{1.04} | 86 ₃ |
| | No nugget | 2.18 _{0.53} | 88 ₃ | 3.43 _{1.02} | 86 ₃ |
| | No PS | 13.59 _{0.51} | 4 ₂ | 17.37 _{1.26} | 22 ₄ |
| | Non-spatial | 4.44 _{0.62} | 94 ₂ | 2.15 _{1.63} | 72 ₅ |
| 5. Strong confounding | Full | 0.97 _{0.46} | 92 ₃ | -0.86 _{1.01} | 91 ₃ |
| | No nugget | 0.77 _{0.46} | 94 ₂ | -0.90 _{0.98} | 92 ₃ |
| | No PS | 2.99 _{0.50} | 78 ₄ | 1.44 _{1.09} | 82 ₄ |
| | Non-spatial | 1.33 _{0.49} | 95 ₂ | -0.31 _{1.37} | 78 ₄ |
| 6. Weak SIR spatial dependence | Full | 0.28 _{0.25} | 92 ₃ | 0.67 _{0.48} | 96 ₂ |
| | No nugget | 0.22 _{0.25} | 93 ₃ | 0.55 _{0.48} | 96 ₂ |
| | No PS | 2.53 _{0.26} | 74 ₄ | 1.91 _{0.56} | 82 ₄ |
| | Non-spatial | 0.64 _{0.27} | 97 ₂ | -0.30 _{0.67} | 80 ₄ |

- 4. Strong spatiotemporal dependence $A_j(t)$: $\rho_s = 0.9, \rho_t = 0.9, \rho_x = 0.5$ and $\phi = 0.4$
- 5. Strong confounding: $\rho_s = 0.9, \rho_t = 0.5, \rho_x = 0.9$ and $\phi = 0.4$
- 6. Weak SIR spatial dependence: $\rho_s = 0.9, \rho_t = 0.5, \rho_x = 0.5$ and $\phi = 0.2$

For each scenario we generate 100 datasets.

The propensity score $e_j(t)$ is the fitted value from a least squares regression of $A_j(t)$ onto $A_j(t - 1), A_j(t - 2), X_j(t), X_j(t - 1)$ and $Y_j(t - 1)$. We then model $\log\{\beta_j(t)\}$ as a linear combination of $A_j(t), \tilde{A}_j(t), X_j(t), \tilde{X}_j(t), e_j(t), \tilde{e}_j(t), e_j(t)^2, \tilde{e}_j(t)^2$ and $e_j(t)\tilde{e}_j(t)$. For each simulated dataset we fit the full model in (7) with the propensity scores included as covariates (“Full”), and compare this model to the full model but with $v_j(t) = 0$ (“No nugget”), the full model but without the propensity scores $e_j(t) = \tilde{e}_j(t) = 0$ (“No PS”) and the full model but without spatial dependence $\rho_s = 0$ (“Non-spatial”). For the non-spatial model we also set $v_j(t) = 0$ because the MCMC algorithm did not converge with these terms included. The models are fit using responses from time points $t \in \{5, \dots, n_t\}$ to accommodate lagged predictors. We fit the model using MCMC with 10,000 iterations and the first 2,000 samples are discarded as burn-in to approximate the posterior median and 95% interval of the direct (δ_1) and indirect (δ_2) effects.

4.2. Results

The results are given in Table 1. Broadly speaking the Full model and the No Nugget model perform well. Giving similar precision and coverage results under all scenarios, they generally outperform the other models in terms of coverage and bias. Moreover, their coverage rates remain




Fig. 1. Simulation results for the full model. The boxplots represent the sampling distribution of the posterior estimator of the direct (δ_1) and indirect (δ_2) effects for the full model. The horizontal lines are the true values and the values above the boxplots are the percentage of datasets for which the posterior 90% interval excludes zero.

Table 2

Simulation study results for model fit. The data-generating scenarios are given in Section 4.1. We compare the full model (“Full”), the model without a nugget term (“No nugget”), the model without the generalized propensity score (“No PS”) and the model without spatial dependence (“Non-spatial”). Estimation of the log mean $l_j(t) = \log[E\{Y_j(t)\}]$ is measured using the correlation between the true and estimated $l_j(t)$ and empirical coverage of 90% intervals. All values are multiplied by 100 and standard errors are given in subscripts.

| Scenario | Correlation | | | | Coverage | | | |
|----------|---------------------|---------------------|---------------------|---------------------|---------------------|---------------------|---------------------|---------------------|
| | Full | No nugget | No PS | Non-spatial | Full | No nugget | No PS | Non-spatial |
| 1 | 93.5 _{0.1} | 93.5 _{0.1} | 93.5 _{0.1} | 88.4 _{0.1} | 94.4 _{0.1} | 94.4 _{0.1} | 94.4 _{0.1} | 93.0 _{0.1} |
| 2 | 96.4 _{0.1} | 96.4 _{0.1} | 96.4 _{0.1} | 92.9 _{0.2} | 94.5 _{0.1} | 94.5 _{0.1} | 94.5 _{0.1} | 93.0 _{0.2} |
| 3 | 94.9 _{0.1} | 95.0 _{0.1} | 94.9 _{0.1} | 91.0 _{0.2} | 94.5 _{0.1} | 94.5 _{0.1} | 94.5 _{0.1} | 92.7 _{0.2} |
| 4 | 96.9 _{0.1} | 97.0 _{0.1} | 96.9 _{0.1} | 93.8 _{0.3} | 93.9 _{0.1} | 94.2 _{0.1} | 94.1 _{0.1} | 92.5 _{0.3} |
| 5 | 94.5 _{0.1} | 94.5 _{0.1} | 94.5 _{0.1} | 90.4 _{0.1} | 94.3 _{0.1} | 94.3 _{0.1} | 94.3 _{0.1} | 93.0 _{0.1} |
| 6 | 92.0 _{0.1} | 91.9 _{0.1} | 92.0 _{0.1} | 87.9 _{0.1} | 93.5 _{0.1} | 93.5 _{0.1} | 93.5 _{0.1} | 93.1 _{0.1} |

reasonably close to their 90% target, although they appear to have a small, positive bias. The Non-spatial model performs worse than the Full and No Nugget models in both precision and coverage, but it remains competitive. The No PS model performs the worst, with extremely poor performance for the strong temporal/spatiotemporal dependence scenarios. The four models seem to have the most difficulty estimating effects when there is strong temporal/spatiotemporal dependence in $A_j(t)$ or strong confounding. All models appear to estimate the direct effects better than the indirect effects. Fig. 1 provides the sampling distribution and power of the direct and indirect effect estimators. These plots show that for the simulation study the estimates have the correct sign with probability near one and have reasonable power in most cases. Therefore, in these synthetic cases the algorithm performs well.

Table 2 compares the model fit in terms of the log mean response $l_j(t) = \log[E\{Y_j(t)\}]$, by giving the correlation (across space and time) between the true and estimated $l_j(t)$ and the coverage of credible intervals. The Full model and the models without nugget or propensity score fit equally well while the non-spatial model shows some degradation in fit. However, the differences in fit are much less dramatic than the difference in bias and covariance for the direct and indirect effects δ_1 and δ_2 , as expected.




Fig. 2. Posterior distributions for direct indirect intervention effects for Full model across lags 0–7 weeks. Estimates correspond to the expected percentage increase in cases from an increase in mobility 50% above baseline: $100 \cdot \{\exp(50 \cdot \delta_i) - 1\}$, $i = 1, 2$. The direct effect estimates for lags 3–5 are 3.5, 15.7, and 24.8, with the latter two significant at the $\alpha = 0.05$ level. The indirect estimates for lags 3–5 are -9.6 , -7.6 , and 0.3 , though none are significant.

5. Estimating the causal effect of community mobility reduction on Coronavirus spread

We implement four models on the real data: A Full model, a No Nugget model, a No PS model, and a Non-spatial model. The models are fit using responses from time points $t \in \{8, \dots, 31\}$ (April 24, 2020 through October 8, 2020) to accommodate lagged predictors. The direct propensity score is estimated with least squares regressions of $A_j(t)$ onto the previous intervention $A_j(t - 1)$, the local covariates $X_j(t)$ and $X_j(t - 1)$, the number of weeks since the first case in county j (enters both linearly and quadratically), time t (enters both linearly and quadratically), a time and intervention interaction $t \cdot A_j(t - 1)$, the baseline level of mobility $A_j(1)$, and $\log\{Y_j(t - 1)\} - \log(N_j)$. The propensity score $e_j(t)$ is set to the fitted values of the resulting model. The propensity of the spillover intervention is estimated as the local average of its neighboring direct scores: $\tilde{e}_j(t) = \sum_{k \sim j} e_k(t) / m_j$. The direct and indirect propensity scores are added to \mathbf{X} and $\tilde{\mathbf{X}}$, respectively, as in (8), along with l -lagged interventions and covariates. Each model is implemented with lags of $l = 0, \dots, 7$. The No Nugget, No PS, and Non-spatial models were run for 100,000 MCMC iterations, with a burn-in tuning period of 20,000 iterations. The Full model required more iterations to converge, so these values were increased to 200,000 and 40,000, respectively. The 200,000 iteration Full model took approximately 30 h to run using 4 cores of a Dell R7425 Dual Processor AMD Epyc 2.2 GHz cluster, with 512 GB RAM and running 64Bit Ubuntu Linux Version 18.04. Finally, the code and data required to run the primary analysis are provided on the Harvard Dataverse. The code required to run the simulation study is provided on the Harvard Dataverse (Giffin, 2022).

5.1. Effect estimation

Table 3 gives the estimated posteriors for the direct and spillover effects for the four different models across lags $l = 0, \dots, 7$, and Fig. 2 shows these results graphically for the Full model. Rather

Table 3
Posterior estimates for direct and indirect treatment effects for lags 0–7 weeks for all models.
 Estimates correspond to the expected percentage increase in cases from an increase in mobility 50% above baseline ($100 \cdot (\exp(50 \cdot \delta) - 1)$). 95% credible intervals are shown in parentheses, and estimates significant at the 95% level are denoted with *.

| Model | Lag (weeks) | Direct effect | Indirect effect |
|-------------|-------------|---------------------|-----------------------|
| Full | 0 | 2.0 (-9.0, 14.1)* | -12.9 (-35.1, 17.0)* |
| | 1 | 1.0 (-9.5, 12.6)* | -28.0 (-45.7, -4.6)* |
| | 2 | 4.9 (-5.8, 16.8)* | -9.4 (-31.8, 22.1)* |
| | 3 | 3.5 (-7.8, 16.4)* | -9.6 (-32.5, 20.0)* |
| | 4 | 15.7 (2.8, 29.9)* | -7.6 (-33.1, 28.3)* |
| | 5 | 24.8 (11.0, 40.2)* | 0.3 (-25.3, 34.6)* |
| | 6 | 25.7 (12.3, 40.7)* | -9.9 (-33.2, 21.0)* |
| | 7 | 24.3 (11.0, 39.1)* | 24.2 (-7.5, 66.5)* |
| No Nugget | 0 | -0.8 (-11.6, 11.7)* | -21.7 (-41.1, 4.1)* |
| | 1 | -1.5 (-11.8, 9.8)* | -28.5 (-46.4, -4.6)* |
| | 2 | 3.6 (-7.2, 15.6)* | -15.4 (-35.7, 12.1)* |
| | 3 | 0.3 (-10.4, 12.3)* | -19.3 (-38.5, 6.2)* |
| | 4 | 12.3 (-0.3, 28.0)* | -13.7 (-35.5, 15.4)* |
| | 5 | 22.7 (9.1, 37.8)* | 3.1 (-22.8, 37.7)* |
| | 6 | 22.2 (8.7, 37.2)* | -18.0 (-38.8, 11.4)* |
| | 7 | 22.8 (9.2, 37.7)* | 2.3 (-22.9, 35.7)* |
| No PS | 0 | -2.7 (-8.0, 2.8)* | 2.9 (-9.4, 17.1)* |
| | 1 | 2.1 (-3.1, 7.7)* | 8.7 (-5.0, 24.1)* |
| | 2 | 6.6 (0.9, 12.9)* | 20.1 (5.1, 39.2)* |
| | 3 | 11.9 (5.6, 18.7)* | 32.5 (12.3, 57.3)* |
| | 4 | 15.6 (9.2, 22.3)* | 32.7 (15.1, 59.5)* |
| | 5 | 15.3 (9.6, 21.6)* | 27.6 (10.4, 48.4)* |
| | 6 | 12.4 (6.7, 18.4)* | 20.3 (5.6, 37.4)* |
| | 7 | 10.8 (5.2, 16.7)* | 21.0 (6.3, 37.2)* |
| Non-spatial | 0 | -2.6 (-17.0, 13.8)* | 7.7 (-12.5, 32.0)* |
| | 1 | -5.9 (-19.8, 10.1)* | -46.4 (-57.1, -34.5)* |
| | 2 | -0.9 (-14.9, 15.5)* | -25.7 (-39.3, -9.4)* |
| | 3 | -4.3 (-17.7, 11.5)* | -37.2 (-48.3, -23.5)* |
| | 4 | 9.8 (-6.1, 28.8)* | -21.6 (-35.2, -4.2)* |
| | 5 | 25.1 (7.3, 45.5)* | 35.8 (13.3, 62.4)* |
| | 6 | 19.3 (3.0, 38.0)* | -10.5 (-24.3, 6.2)* |
| | 7 | 24.2 (6.8, 44.5)* | 6.1 (-11.1, 26.5)* |

than present the estimated values for δ_1 and δ_2 , the estimates shown are $100 \cdot \{\exp(50 \cdot \delta_i) - 1\}$, $i = 1, 2$, and correspond to the expected percentage increase in cases from an increase in mobility 50% above baseline. Because it is not clear that the observed data can identify the correct lag, we provide model results for all lags, and speculate that lags of 3–5 weeks are most appropriate to see effects of changes in mobility. The Centers for Disease Control and Prevention guidelines suggest that COVID symptoms appear 2–14 days after virus exposure (Centers for Disease Control and Prevention, 2020). Given this, 3–5 weeks would allow the virus to spread through roughly several generations of people – which would be enough to see changes in county-level case counts.

The key trend evident in these results is that the direct effect appears to be positive and significantly different from zero for higher lags. The Full model gives positive significant estimates on δ_1 for lags 4–7 weeks; the No Nugget for lags 5–7 weeks; the No PS model for lags 2–7 weeks; and the Non-spatial model for lags 5–7. The size of this effect differs between models/lags, but the Full model for lags 5–7 predicts that a mobility level 50 percent above baseline will increase Coronavirus cases by between 11 and 41 percent. This trend is demonstrated in the grouping of posteriors above zero in the top sub-figure of Fig. 2. This provides evidence that a decrease in local mobility should decrease local Coronavirus cases roughly 4–7 weeks after the change in mobility.

The indirect effect estimates are more variable with wider credible intervals, as illustrated in Fig. 2. For the Full model, a lag of 7 gives a positive but non-significant estimate for the indirect effect, although the indirect effect posteriors for other lags seem centered at zero, with the exception

Table 4
Posterior estimates for coefficients in No Nugget, lag 5 model. 95% credible intervals are shown in parentheses, and estimates significant at the 95% level are denoted with *. Coefficients are not transformed as in Table 3.

| Variable | Estimate (95% interval) |
|--------------------------------------|---------------------------|
| Intercept | -7.43 (-7.68, -7.18)* |
| $A_j(t)$ | 0.004 (0.002, 0.006)* |
| $\tilde{A}_j(t)$ | 0.0003 (-0.0055, 0.006) |
| $e_j(t)$ | -0.0011 (0.0007, 0.0054)* |
| $\tilde{e}_j(t)$ | 0.006 (0.001, 0.013)* |
| Mean temperature | -0.008 (-0.016, 0.000)* |
| Relative Humidity | -0.009 (-0.011, -0.006)* |
| Dew point | 0.045 (0.033, 0.057)* |
| PM _{2.5} | 0.082 (0.066, 0.097)* |
| Poverty rate | 1.07 (0.88, 1.26)* |
| Population density | -0.035 (-0.043, -0.026)* |
| Median household income (thousands) | -0.002 (-0.003, -0.001)* |
| Population (millions) | -0.048 (-0.066, -0.030)* |
| # hospital beds | 5.95 (3.64, 8.29)* |
| # ICU beds | 24.01 (6.99, 40.94)* |
| Median age | -0.021 (-0.023, -0.018)* |
| Proportion of foreign-born residents | 4.16 (3.97, 4.35)* |
| Proportion of health-employed | 0.82 (0.59, 1.04)* |
| ρ_s | 0.9949 (0.9945, 0.9950)* |
| ρ_t | 0.452 (0.441, 0.463)* |

of lag 1 which is significant and negative. For lag 1 the No Nugget model also shows this counter-intuitive result. The No PS model shows a number of positive and significant indirect effects, and the Non-spatial model paradoxically shows several significant indirect effects – some positive and some negative. It is not entirely clear why there is more variability in the indirect effects or why we see these counter-intuitive results. Some of these could simply be spurious significance resulting from estimating many lags. Another possibility is that our measure of connectivity does not precisely measure the levels of contact between different counties. This could explain the wide intervals, as the connectivity is over-estimated between some counties and under-estimated between others. In any event, the indirect effect story appears less clear cut than that of the direct effect.

Table 4 gives parameter estimates from the No Nugget/Lag 5 model. In addition to the intervention effects and accounting for the propensity scores, the covariate coefficients tell an interesting story about their association with Coronavirus case counts. For example, expected Coronavirus case counts are decreasing in median age of population, population density, county population, and relative humidity. Alternatively, Coronavirus case counts tend to increase with the poverty rate, PM_{2.5}, number of hospital beds, number of ICU beds, and percentage healthcare-employed residents.

6. Discussion

This analysis uses a spatiotemporal model to provide insights into how mobility affects the spread of Coronavirus. The analysis uses a propensity score framework to obtain causal estimates of this effect, while allowing for potential interference between nearby counties. We show via simulation study that the causal estimates from our computationally-efficient approximate model have good statistical properties for data generated from a mechanistic process. In particular, the data suggest that decreases in mobility appear to cause a statistically significant decrease on the number of local Coronavirus cases after roughly 5 weeks. This is significant, as such a long lag might be overlooked in a naive analysis of the relationship between mobility and Coronavirus.

The core of our contribution is developing a novel method which allows for the recovery causal treatment effect estimates from an SIR model – a class of models which is primarily used for prediction and forecasting rather than causal inference. This theoretical gap is bridged using the potential outcomes framework and propensity scores. Because many infectious disease models are predictive in nature, this provides an importance new inferential tool for researchers.

There are several noteworthy limitations of this analysis. The intervention data are somewhat crude measures of mobility. Taken from a small minority of smartphone users, they provide a limited view into the distribution of mobility among residents for a given county and week. Moreover, the five relevant measures of mobility provided by Google (which we average over) may have different intervention effects. We also do not have data on other interventions to reduce the spread of Coronavirus, such as wearing masks, state-mandated closures of different types of businesses, social distancing, washing hands, etc. These are all plausibly correlated with mobility, so it is possible that the intervention effect of mobility incorporates the causal effect of these measures too. The measure of connectivity that we use – that of adjacency between counties – is also somewhat crude. For example, while the first viral hotspot in the United States is thought to have been Seattle, it quickly traveled to New York by air. A more comprehensive measure of connectivity and travel between regions may be helpful to incorporate to the model. The response data also have limitations. Because the availability of Coronavirus tests in the United States has varied over time and space, the number of new positive tests results in a particular county/week might partly reflect the availability of tests. One solution to this would be to use COVID-19-related hospitalizations or deaths as a response. However, these are necessarily less common – providing fewer response data – and they would add to the already substantial lag between the moment of infection and positive test. Finally, as we make the “no unmeasured confounders” assumption, it is possible that there are key confounders that have not been included. Because of this, we include a number of possible confounders, but more can always be done in this regard. Lastly, correlation between variables was not a major issue, as shown in Figure B.1.

As the pandemic continues, many of these data issues are being addressed: surveys are now being done to ascertain levels of mask-usage and other social distancing practices. Publicly available data on the availability of Coronavirus tests by location are emerging. Further, as the number of cases continues to grow, using hospitalizations or deaths as a response becomes more feasible.

Disclaimer

The views expressed in this manuscript are those of the individual authors and do not necessarily reflect the views and policies of the U.S. Environmental Protection Agency. Mention of trade names or commercial products does not constitute endorsement or recommendation for use.

Acknowledgments

The authors thank Sukanya Bhattacharyya, Can Cui, Matt Miller, Parker Trostle, Laura Wendelberger, Stephen Xu and Zun Yin (North Carolina State University), Yawen Guan (University of Nebraska) and Pulong Ma (Duke University) for help compiling the data.

Funding

This work was partially supported by NIH, United States of America grant R01ES031651-01.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.spasta.2022.100711>, which contains a derivation of the conditional expectation of $Y_j(t)$, as well as the correlation among predictors.

References

- Alexeeff, S.E., et al., 2015. Consequences of kriging and land use regression for PM_{2.5} predictions in epidemiologic analyses: insights into spatial variability using high-resolution satellite data. *J. Expo. Sci. Environ. Epidemiol.* 25 (2), 138–144.
- Arino, J., Van den Driessche, P., 2003. A multi-city epidemic model. *Math. Popul. Stud.* 10 (3), 175–193.
- Berres, S., Ruiz-Baier, R., 2011. A fully adaptive numerical approximation for a two-dimensional epidemic model with nonlinear cross-diffusion. *Nonlinear Anal. Real World Appl.* 12 (5), 2888–2903.

- Bradley, J.R., 2020. Joint spatio-temporal analysis of multiple response types using the hierarchical generalized transformation model with application to coronavirus disease 2019 and social distancing. arXiv Preprint (Stat.ME:2002.09983).
- Buckingham-Jeffery, E., Isham, V., House, T., 2018. Gaussian process approximations for fast inference from infectious disease data. *Math. Biosci.* 301, 111–120.
- Burger, R., et al., 2009. Modelling the spatial-temporal evolution of the 2009 A/H1N1 influenza pandemic in Chile. *Math. Biosci. Eng.* 13, 1–17.
- Capasso, V., Di Liddo, A., 1994. Asymptotic behaviour of reaction-diffusion systems in population and epidemic models. *J. Math. Biol.* 32 (5), 453–463.
- Carroll, R., Prentice, C.R., 2021. Using spatial and temporal modeling to visualize the effects of US state issued stay at home orders on COVID-19. *Sci. Rep.* 11 (1), 1–7.
- Centers for Disease Control and Prevention, 2020. Symptoms of Coronavirus. <https://www.cdc.gov/coronavirus/2019-ncov/symptoms-testing/symptoms.html>.
- Chang, S., et al., 2021. Mobility network models of COVID-19 explain inequities and inform reopening. *Nature* (589), 82–87.
- Chen, H., He, J., Song, W., Wang, L., Wang, J., Chen, Y., 2020. Modeling and interpreting the COVID-19 intervention strategy of China: A human mobility view. *PLoS One* 15 (11), e0242761.
- Chinviriyasit, S., Chinviriyasit, W., 2010. Numerical modelling of an SIR epidemic model with diffusion. *Appl. Math. Comput.* 216 (2), 395–409.
- Chudnovsky, A.A., et al., 2012. Prediction of daily fine particulate matter concentrations using aerosol optical depth retrievals from the Geostationary Operational Environmental Satellite (GOES). *J. Air Waste Manage. Assoc.* 62 (9), 1022–1031.
- Cowling, B.J., et al., 2020. Impact assessment of non-pharmaceutical interventions against coronavirus disease 2019 and influenza in Hong Kong: an observational study. *Lancet Public Health* 5 (5), 279–288.
- Dandekar, R., Barbastathis, G., 2020. Quantifying the Effect of Quarantine Control in Covid-19 Infectious Spread Using Machine Learning. Cold Spring Harbor Laboratory Press, <http://dx.doi.org/10.1101/2020.04.03.20052084>, MedRxiv.
- Dehning, J., et al., 2020. Inferring COVID-19 spreading rates and potential change points for case number forecasts. MedRxiv: <http://dx.doi.org/10.1101/2020.04.02.20050922>.
- Dong, E., Du, H., Gardner, L., 2020. An interactive web-based dashboard to track COVID-19 in real time. *Lancet Infect. Dis.* 20 (5), 533–534.
- ESRI, 2020. Definitive Healthcare: USA Hospital Beds. <https://coronavirus-resources.esri.com/datasets/definitivehc::definitive-healthcare-usa-hospital-beds>.
- Giffin, A., 2022. Data and Code for Estimating Intervention Effects on Infectious Disease Control: the Effect of Community Mobility Reduction on Coronavirus Spread. Harvard Dataverse, <http://dx.doi.org/10.7910/DVN/NQNNWE>.
- Giffin, A., et al., 2020. Generalized propensity score approach to causal inference with spatial interference. arXiv Preprint: 2007.00106.
- Google LLC, 2020. Google COVID-19 Community Mobility Reports. <https://www.google.com/covid19/mobility/>. Accessed October 2020.
- Halloran, M.E., Struchiner, C.J., 1991. Study designs for dependent happenings. *Epidemiology* 2 (5), 331–338.
- Hilker, F.M., et al., 2007. A diffusive SI model with Allee effect and application to FIV. *Math. Biosci.* 206 (1), 61–80.
- Hyman, J.M., LaForce, T., 2003. Modeling the spread of influenza among cities. In: *Bioterrorism: Mathematical Modeling Applications in Homeland Security*. SIAM, pp. 211–236.
- Imbens, G.W., 2000. The role of the propensity score in estimating dose-response functions. *Biometrika* 87 (3), 706–710.
- Kermack, W.O., McKendrick, A.G., 1927. A contribution to the mathematical theory of epidemics. *Proc. R. Soc. Lond. Ser. A Contain. Pap. Math. Phys. Charact.* 115 (772), 700–721.
- Kounchev, O., Simeonov, G., Kuncheva, Z., 2020. The TVBG-SEIR spline model for analysis of COVID-19 spread, and a Tool for prediction scenarios. arXiv preprint arXiv:2004.11338.
- Lee, S., Castillo-Chavez, C., 2015. The role of residence times in two-patch dengue transmission dynamics and optimal strategies. *J. Theor. Biol.* 374, 152–164.
- Lee, J., Jung, E., 2015. A spatial-temporal transmission model and early intervention policies of 2009 A/H1N1 influenza in South Korea. *J. Theor. Biol.* 380, 60–73.
- Lee, W.D., Qian, M., Schwanen, T., 2021. The association between socioeconomic status and mobility reductions in the early stage of England's COVID-19 epidemic. *Health Place* 69, 102563.
- Lee, J.M., et al., 2012. The effect of public health interventions on the spread of influenza among cities. *J. Theor. Biol.* 293, 131–142.
- Li, Q., et al., 2020. Early Transmission Dynamics in Wuhan, China, of Novel Coronavirus – Infected Pneumonia. *N. Engl. J. Med.* 382 (13), 1199–1207. <http://dx.doi.org/10.1056/NEJMoa2001316>, PMID: 31995857.
- Livingston, E., Bucher, K., 2020. Coronavirus Disease 2019 (COVID-19) in Italy. *J. Am. Med. Assoc.* 323 (14), 1335. <http://dx.doi.org/10.1001/jama.2020.4344>.
- Lyu, W., Wehby, G.L., 2020. Community use of face masks and COVID-19: Evidence from A natural experiment of state mandates in the US: Study examines impact on COVID-19 growth rates associated with state government mandates requiring face mask use in public. *Health Affairs* 39 (8), 1419–1425.
- Lyu, H., et al., 2020. COVID-19 Time-series Prediction by Joint Dictionary Learning and Online NMF. arXiv preprint arXiv:2004.09112.
- Magdon-Ismail, M., 2020. Machine Learning the Phenomenology of COVID-19 From Early Infection Dynamics. Cold Spring Harbor Laboratory Press, MedRxiv: <http://dx.doi.org/10.1101/2020.03.17.20037309>.

- Mbuvha, R., Marwala, T., 2020. Bayesian Inference of COVID-19 Spreading Rates in South Africa. Cold Spring Harbor Laboratory Press, MedRxiv: <http://dx.doi.org/10.1101/2020.04.28.20083873>.
- Milner, F.A., Zhao, R., 2008. SIR model with directed spatial diffusion. *Math. Popul. Stud.* 15 (3), 160–181.
- Nouvellet, P., Bhatia, S., Cori, A., Ainslie, K.E., Baguelin, M., Bhatt, S., Boonyasiri, A., Brazeau, N.F., Cattarino, L., Cooper, L.V., et al., 2021. Reduction in mobility and COVID-19 transmission. *Nature Commun.* 12 (1), 1–9.
- Nychka, D., Furrer, R., Sain, S., 2014. *Fields: Tools for Spatial Data*. R Package Version 7.1. Vol. 10. Accessed Online.
- Paeng, S.-H., Lee, J., 2017. Continuous and discrete SIR-models with spatial distributions. *J. Math. Biol.* 74 (7), 1709–1727.
- Prem, K., et al., 2020. The effect of control strategies to reduce social mixing on outcomes of the COVID-19 epidemic in Wuhan, China: a modelling study. *Lancet Public Health* 5 (5), 261–270.
- Punn, N.S., Sonbhadra, S.K., Agarwal, S., 2020. COVID-19 Epidemic Analysis Using Machine Learning and Deep Learning Algorithms. Cold Spring Harbor Laboratory Press, MedRxiv: <http://dx.doi.org/10.1101/2020.04.08.20057679>.
- Rashed, E.A., Hirata, A., 2021. One-year lesson: Machine learning prediction of COVID-19 positive cases with meteorological data and mobility estimate in Japan. *Int. J. Environ. Res. Public Health* 18 (11), 5736.
- Reich, B.J., et al., 2020. A review of spatial causal inference methods for environmental and epidemiological applications. *arXiv Preprint: 2007.02714*.
- Reluga, T., 2004. A two-phase epidemic driven by diffusion. *J. Theor. Biol.* 229 (2), 249–261.
- Robinson, M., Stilianakis, N.I., Drossinos, Y., 2012. Spatial dynamics of airborne infectious diseases. *J. Theor. Biol.* 297, 116–126.
- Rubin, D.B., 1974. Estimating causal effects of treatments in randomized and nonrandomized studies. *J. Educ. Psychol.* 66 (5), 688.
- Sattenspiel, L., Dietz, K., et al., 1995. A structured epidemic model incorporating geographic mobility among regions. *Math. Biosci.* 128 (1), 71–92.
- Sattenspiel, L., Herring, D.A., 2003. Simulating the effect of quarantine on the spread of the 1918–19 flu in central Canada. *Bull. Math. Biol.* 65 (1), 1–26.
- Sobel, M.E., 2006. What do randomized studies of housing mobility demonstrate? Causal inference in the face of interference. *J. Amer. Statist. Assoc.* 101 (476), 1398–1407.
- Sparks, A., Hengl, T., Nelson, A., 2017. GSODR: Global Summary Daily Weather Data in R. *J. Open Source Softw.* 2 (10), 177.
- Stein, M.L., 2005. Space-time covariance functions. *J. Amer. Statist. Assoc.* 100 (469), 310–321.
- US Bureau of Labor Statistics, 2019. Health Care and Social Assistance: NAICS 62. <https://www.bls.gov/iag/tgs/iag62.htm>.
- US Census Bureau, 2016. American Community Survey Public Use Microdata Samples. <https://www.census.gov/newsroom/press-releases/2016/cb16-tps144.html>.
- Wang, Y., Wang, J., Zhang, L., 2010. Cross diffusion-induced pattern in an SI model. *Appl. Math. Comput.* 217 (5), 1965–1970.
- Wu, X., Nethery, R.C., Sabath, M., Braun, D., Dominici, F., 2020a. Air pollution and COVID-19 mortality in the United States: Strengths and limitations of an ecological regression analysis. *Sci. Adv.* 6 (45), eabd4049.
- Wu, X., et al., 2020b. Exposure to Air Pollution and COVID-19 Mortality in the United States: A Nationwide Cross-Sectional Study. Cold Spring Harbor Laboratory Press, MedRxiv: <http://dx.doi.org/10.1101/2020.04.05.20054502>. [arXiv:https://www.medrxiv.org/content/early/2020/04/27/2020.04.05.20054502](https://www.medrxiv.org/content/early/2020/04/27/2020.04.05.20054502.full.pdf), URL <https://www.medrxiv.org/content/early/2020/04/27/2020.04.05.20054502>.
- Yilmazkuday, H., 2020. Stay-at-home works to fight against COVID-19: international evidence from google mobility data. *J. Hum. Behav. Soc. Environ.* <http://dx.doi.org/10.1080/10911359.2020.1845903>.