A joint hierarchical model for the number of cases and deaths due to COVID-19 across the boroughs of Montreal

Victoire Michal[1a], Leo Vanciu[b], Alexandra M. Schmidt[a]

[a]*McGill University, Department of Epidemiology, Biostatistics and Occupational Health, 1020 Pine Avenue, Montreal, H3A 1A2, QC, Canada*
[b]*Marianopolis College, 4873 Westmount Avenue, Montreal, H3Y 1X9, QC, Canada*

Abstract

Montreal has been the epicentre of the COVID-19 pandemic in Canada with the highest number of deaths in the country. The cumulative numbers of cases and deaths, as of July 25th, 2021, recorded in the 33 areas of Montreal are modelled through a bivariate hierarchical Bayesian model using Poisson distributions. The Poisson means are decomposed in the log scale as the sums of fixed effects and latent effects. The areal median age, the educational level, and the number of beds in long-term care homes are included in the fixed effects. To explore the correlation between cases and deaths inside each area and across areas, three bivariate models are considered for the latent effects, namely an independent one, a conditional autoregressive model, and one that allows for both spatially structured and unstructured sources of variability. As the inclusion of spatial effects change some of the fixed effects, we extend the Spatial+ approach to a Bayesian areal set up to investigate the presence of spatial confounding. Results show that the cumulative totals of cases and deaths are negatively correlated inside and across the boroughs of Montreal. The covariates are not associated in a similar manner for the cases and deaths due to COVID-19. The educational level and the median age seem spatially confounded for the cases and deaths across the boroughs of Montreal.

Keywords: Bayesian inference, Conditional autoregressive distribution, Disease Mapping, Spatial confounding, Spatial+

1. Motivation

As of July 25th, 2021, the coronavirus disease 2019 (COVID-19) pandemic counts a total of 194,248,750 cases and 4,163,599 deaths worldwide (Dong et al., 2020). This disease is caused by an infection with the severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2). The World Health Organization declared COVID-19 to be a pandemic in March 2020

---

[1]Corresponding author.
McGill University, Department of Epidemiology, Biostatistics and Occupational Health, 1020 Pine Avenue, Montreal, H3A 1A2, QC, Canada.
Email address: victoire.michal@mail.mcgill.ca (V. Michal).

(Ghebreyesus, 2020) and since then, countries worldwide have instituted infection control measures (e.g., lockdowns, curfews, mask mandates) to control the spread of the disease. We are interested in studying the COVID-19 spread in the urban agglomeration of Montreal. As of July 25th, 2021, Canada cumulates 1,426,903 cases and 26,547 deaths due to COVID-19 (Government of Canada, 2021) with 9.2% of the cases and 17.9% of the deaths recorded in Montreal ([dataset] Direction régionale de santé publique, 2021). The agglomeration covers approximately 500 km^2 and is the most populous administrative region in the province of Quebec, Canada, with over 1.9 million inhabitants which represents, approximately, 6% of the Canadian population (Statistics Canada, 2017a,b). The agglomeration incorporates 19 boroughs of the City of Montreal and 15 related cities. These boroughs and related cities vary in demography, socio-economic status, and healthcare infrastructure.

This analysis focuses on data on the cumulative number of COVID-19 confirmed cases and deaths in Montreal since the beginning of the pandemic, as of July 25th, 2021. The cases and deaths distributions across the areas of Montreal are shown in Figure 1. This dataset is made publicly available by the regional director of public health ([dataset] Direction régionale de santé publique, 2021). The data consist of 34 boroughs and related cities of Montreal. However, one related city, L'Île-Dorval, did not record cumulative data. We exclude this area from the analysis, yielding a total of 33 areas.

The Ministry of Health and Social Services defines a confirmed case or death due to COVID-19 as follows. Cases are confirmed either by a laboratory or by an epidemiological link. A laboratory confirms a case through the detection of nucleic acids of SARS-CoV-2. To confirm a case by an epidemiological link, the manifestation of clinical symptoms must be compatible with COVID-19 with a high risk of exposure to a case that was confirmed by a laboratory during the period of contagion and no other apparent cause. Deaths may also be confirmed by a laboratory or by an epidemiological link. A laboratory confirms a death due to COVID-19 through the manifestation of clinical symptoms compatible with COVID-19 before death and the detection of nucleic acids of SARS-CoV-2. A death is confirmed to be due to COVID-19 by an epidemiological link if there was a manifestation of clinical symptoms compatible with COVID-19 before death with a high risk of exposure to a case that was confirmed by a laboratory during the period of contagion and no other apparent cause (Ministère de la Santé et des Services sociaux, 2021). The death counts of two areas are censored for privacy issues. Namely, Baie-d'Urfé and Montréal-Ouest recorded between 1 and 4 deaths. We also obtained data on the population size of each borough and related city from the 2016 Canadian Census ([dataset] Ville de Montréal, 2016).

Figure 1: Maps of the COVID-19 cases (left) and deaths (right) across the 33 boroughs of Montreal.

In the literature, many studies identify age as a risk factor of COVID-19 mortality (Zhou et al., 2020; Bonanad et al., 2020; Yanez et al., 2020). Some studies also highlight the relationship between socio-economic status and COVID-19 cases (Hawkins et al., 2020; Gangemi et al., 2020) and COVID-19 deaths (Hawkins et al., 2020; Bermudi et al., 2021). Moreover, apart from COVID-19, there exist important health disparities due to socio-economic inequalities in Montreal. These have been shown to lead to a difference in life expectancy of up to 11 years between areas (Le Blanc et al., 2011). On the other hand, healthcare infrastructure has been positively associated with reactive responses to COVID-19 (Sharma et al., 2021). In Quebec, there exist five classes of senior residences. Three of these fall under the denomination of residential and long-term care centres, which may be public, private and subsidised by the government, or private and non-subsidised. Most senior residences are of the two remaining classes: intermediary resources and private seniors' residences. Although all of these residences can provide similar services, residential and long-term care homes specifically serve seniors with an important loss of autonomy (Gouvernement du Québec, 2021). Hospitals in Quebec, under the direction of the Ministry of Health, were generally better prepared to respond to the COVID-19 pandemic than residential and long-term care centres during the first months of the pandemic (Doucet, 2020). The vast majority of COVID-19 deaths in senior residences in Quebec are linked to residential and long-term care centres. In fact, as of July 25th, 2021, 51% of all COVID-19 deaths in Quebec are linked exclusively to residential and long-term care homes (Institut national de santé publique du Québec, 2021).

This literature review motivates our covariates' selection in order to model the cases and deaths due to COVID-19 in the 33 areas of Montreal. To account for the age profile of each borough, we obtain data on the median age of the population by borough and related city from the 2016 Canadian census ([dataset] Ville de Montréal, 2016). As a proxy for the socio-economic status of each borough, we consider the percentage of the population between 25 and 64 years old with a university diploma. These data are also extracted

from the 2016 Canadian census ([dataset] Ville de Montréal, 2016). Finally, regarding the healthcare infrastructure in Montreal, we include the number of beds in all residential and long-term care homes of each borough (*Centre d'hébergement et de soins de longue durée*, CHSLD). These records are available in public databases ([dataset] Ministère de la Santé et des Services sociaux, 2021a,b).

In this paper, we analyse the *joint* distribution of COVID-19 cases and deaths across the 33 boroughs of Montreal through a Bayesian hierarchical joint model. We are interested in measuring the strength of the associations between the three selected covariates and the cases and deaths due to COVID-19. Also, we aim at investigating how these two outcomes are correlated within and across boroughs. We examine whether the selected covariates impact the risk of being a case or the risk of dying from COVID-19 in a similar fashion. Additionally, we investigate if there is any structure left in the data after accounting for the available covariates. We explore different prior specifications for these local latent effects, ranging from independent to spatially structured ones. It is known that the inclusion of spatially structured latent effects can affect the estimation of the fixed effects (Reich et al., 2006; Khan and Calder, 2020; Dupont et al., 2021); this is known in the literature of Spatial Statistics as spatial confounding. We fit a Bayesian version of the *Spatial+* approach proposed by Dupont et al. (2021) to investigate if there is spatial confounding in the fitted models.

This paper is organised as follows: Section 2 describes the proposed models, the method for spatial confounding adjustment and the inference procedure. In Section 3, the results of the analyses of the COVID-19 cases and deaths are presented. Finally, Section 4 concludes and discusses our findings.

2. Methods

Let C_i be the recorded number of COVID-19 cases in borough i, for $i = 1, \ldots, n$, where $n = 33$ is the number of non-overlapping boroughs in Montreal. Let D_i be the number of deaths due to COVID-19 in borough i. We model the deaths and cases through a joint hierarchical Bayesian model in order to accommodate the natural relationship one expects between cases and deaths of a disease. First, we define a marginal distribution for the cases, and then, conditional on the number of cases, we define a distribution for the number of recorded deaths. Sahu and Böhning (2021) modelled the cases and deaths due to COVID-19 in England in a similar fashion, using a 2-stage Bayesian hierarchical model. However, their interest lied in the temporal structure of the data and did not consider joint latent effects. Our approach differs because we do not conduct a spatio-temporal analysis, and we are interested in measuring the correlation between cases and deaths inside and across the boroughs of Montreal. To that end, we include a correlation parameter that helps borrow strength from the cases and deaths inside and across boroughs. More specifically, we assume,

$$C_i \sim \mathrm{Pois}(E_i^C \lambda_i^C) \quad \text{and} \quad D_i \mid C_i \sim \mathrm{Pois}(E_i^D \lambda_i^D), \tag{1}$$

where Pois stands for the Poisson distribution, λ_i^C and λ_i^D denote, respectively, the relative risk of the cases and the deaths in area i, whereas E_i^C and E_i^D are offsets. The offset for the number of cases is computed based on the population size P_i of each borough, while

the offset for the number of deaths given the number of cases is computed based on the number of cases observed in borough i. More specifically, $E_i^C = P_i \left(\sum_{j=1}^n C_j / \sum_{j=1}^n P_j \right)$ and $E_i^D = C_i \left(\sum_{j=1}^n D_j / \sum_{j=1}^n C_j \right)$. To model the number of deaths conditionally on the number of cases, the Poisson distribution is used as an approximation to a Binomial distribution with size C_i and probability p_i. This approximation is reasonable because of high values of C_i's and small p_i's as estimated in an exploratory data analysis (Wakefield, 2013, section 7.3.2). In the next step, we model the log relative risks as follows:

$$\log(\lambda_i^C) = \beta_0^C + \boldsymbol{x}_i \boldsymbol{\beta}^C + b_i^C \quad \text{and} \quad \log(\lambda_i^D) = \beta_0^D + \boldsymbol{x}_i \boldsymbol{\beta}^D + b_i^D, \tag{2}$$

where β_0^C and β_0^D denote, respectively, the overall mean log risks of cases and deaths, $\boldsymbol{x}_i$ is a vector of $K = 3$ covariates used to model both the cases and deaths associated with the K-dimensional vectors of coefficients, $\boldsymbol{\beta}^C$ and $\boldsymbol{\beta}^D$, and b_i^C and b_i^D are latent random effects that accommodate whatever is left after accounting for the available covariates. Further, the inclusion of the latent random effects in both log risks' decompositions allows for the accommodation of overdispersion which is observed in the exploratory data analysis.

We explore bivariate models for the latent random effects, $\boldsymbol{b}_i = [b_i^C, b_i^D]^\top$. The models that we consider are special cases of the general formulation $\boldsymbol{b}_i = \Gamma_u \boldsymbol{u}_i + \Gamma_v \boldsymbol{v}_i$, where the vector of independent random effects, $\boldsymbol{v}_i = [v_i^C, v_i^D]^\top \overset{i.i.d.}{\sim} \mathcal{N}_2(\boldsymbol{0}, \boldsymbol{I}_2)$, $i = 1, \ldots, n$, is independent of $\boldsymbol{u}_i = [u_i^C, u_i^D]^\top$, the vector of spatially structured effects following independent CAR models (Besag, 1974), and where I_d is the d-dimensional identity matrix and

$$\Gamma_u \Gamma_u^\top = \Sigma_u = \begin{bmatrix} \sigma_{C,u}^2 & \rho_u \sigma_{C,u} \sigma_{D,u} \\ \rho_u \sigma_{C,u} \sigma_{D,u} & \sigma_{D,u}^2 \end{bmatrix} \quad \text{and} \quad \Gamma_v \Gamma_v^\top = \Sigma_v = \begin{bmatrix} \sigma_{C,v}^2 & \rho_v \sigma_{C,v} \sigma_{D,v} \\ \rho_v \sigma_{C,v} \sigma_{D,v} & \sigma_{D,v}^2 \end{bmatrix}.$$

Let the CAR distributed effects $u_i^\ell \mid \boldsymbol{u}_{(-i)}^\ell \sim \mathcal{N} \left(\sum_{j=1}^n w_{ij} u_j^\ell / d_i, 1/d_i \right)$, $\ell = C, D$, where $\boldsymbol{u}_{(-i)}^\ell = [u_1^\ell, \ldots, u_{i-1}^\ell, u_{i+i}^\ell, \ldots, u_n^\ell]^\top$, $d_i = \sum_{j=1}^n w_{ij}$ and $\boldsymbol{W} = [w_{ij}]$ is the matrix of weights that defines the neighbourhood structure. Commonly, we define $w_{ij} = 1$ if areas i and j share a border, denoted by $i \sim j$, and $w_{ij} = 0$, otherwise. This neighbourhood structure is used throughout this paper. Let $\boldsymbol{D} = diag(d_i)$ and $\boldsymbol{Q} = \boldsymbol{D} - \boldsymbol{W}$, we may write $\boldsymbol{u}^\ell \sim \mathcal{N}_n(\boldsymbol{0}, \boldsymbol{Q}^-)$, which is not properly defined since $\boldsymbol{Q}$ is not positive definite (Banerjee et al., 2014).

The first special case of the model for $\boldsymbol{b}_i$ that we consider is one with independent random effects across the areas, denoted the IID model: $\boldsymbol{b}_i = \Gamma_v \boldsymbol{v}_i \overset{i.i.d.}{\sim} \mathcal{N}_2(\boldsymbol{0}, \Sigma_v)$. Hence, ρ_v denotes the correlation between the latent effects for cases and deaths. When $\rho_v = 0$, the IID joint model results in independent random effects for the cases and deaths inside each area. Further, if one assumes prior independence between the fixed effects in both models, then it does not make a difference to perform the inference about the two processes independently or jointly. In other words, there is no borrow of strength between the number of cases and deaths within boroughs. If $\rho_v \neq 0$, this bivariate IID model allows for dependence between the deaths and cases within a particular borough.

The IID model, however, does not allow for spatial autocorrelation. Yet, one may expect the cases from neighbouring areas to be more correlated than cases from further apart

boroughs, and similarly for the numbers of deaths. To adjust for this possible spatial autocorrelation, the second special case that we consider for $\boldsymbol{b}_i$ is a multivariate CAR model (Lawson, 2020): $\boldsymbol{b}_i = \Gamma_u \boldsymbol{u_i}$. One may write the joint distribution for the long vector of latent effects, $\boldsymbol{b} = [b_1^C, \ldots, b_n^C, b_1^D, \ldots, b_n^D]$, as $\boldsymbol{b} \sim \mathcal{N}_{2n}(\boldsymbol{0}, \Sigma_b \otimes \boldsymbol{Q}^-)$ (Jin et al., 2007). From this joint formulation, we see that the multivariate CAR model allows for a dependence between the cases and deaths of a particular borough as well as between the counts of neighbouring areas.

The multivariate CAR model assumes *a priori* that the latent effects are necessarily distributed according to a spatial structure. To relax this assumption, we consider the multivariate extension of the BYM model (Besag et al., 1991) that allows for both unstructured and spatially structured sources of variability. This corresponds to the general formulation, $\boldsymbol{b}_i = \Gamma_u \boldsymbol{u}_i + \Gamma_v \boldsymbol{v}_i$. The long vector of latent effects is distributed as $\boldsymbol{b} \sim \mathcal{N}_{2n}(\boldsymbol{0}, \Sigma_b)$, for

$$\Sigma_b = \begin{bmatrix} \sigma_{C,u}^2 \boldsymbol{Q}^- + \sigma_{C,v}^2 \boldsymbol{I}_n & \rho_u \sigma_{C,u} \sigma_{D,u} \boldsymbol{Q}^- + \rho_v \sigma_{C,v} \sigma_{D,v} \boldsymbol{I}_n \\ \rho_u \sigma_{C,u} \sigma_{D,u} \boldsymbol{Q}^- + \rho_v \sigma_{C,v} \sigma_{D,v} \boldsymbol{I}_n & \sigma_{D,u}^2 \boldsymbol{Q}^- + \sigma_{D,v}^2 \boldsymbol{I}_n \end{bmatrix}.$$

Let $\boldsymbol{Y}_i = [C_i, D_i]^\top \equiv [Y_i^C, Y_i^D]^\top$. Although the marginal distribution of $\boldsymbol{Y}_i$ with respect to the latent effects is not available in closed form, it is possible to obtain the marginal moments using the properties of conditional expectations and the law of total covariance. Table 1 shows the *marginal* moments of the cases and deaths obtained from the models discussed above. The detailed computation of these marginal moments is available in Appendix A. From Table 1 it is clear that the IID model is only able to capture correlation between cases and deaths within a borough. The covariance will be negative if $\rho_v < 0$. On the other hand, the CAR and BYM models are able to accommodate correlations both within and among neighbouring boroughs. In the CAR model, if $\rho_u < 0$ the correlation within and among neighbouring boroughs will be negative, and positive if $\rho_u > 0$. In the BYM model, there are two parameters capturing the correlation within a borough, and a negative correlation results if $\rho_u \sigma_{C,u} \sigma_{D,u} [\boldsymbol{Q}^-]_{ii} + \rho_v \sigma_{C,v} \sigma_{D,v} < 0$. On the other hand, for neighbouring boroughs the correlation is captured only by the parameter ρ_u so that if $\rho_u < 0$ a negative correlation results in this case.

2.1. Investigating spatial confounding: a Bayesian alternative to Spatial+

When modelling disease risks through a Poisson model, Clayton et al. (1993) noted that covariates effects may change in the presence of spatially structured latent effects. This issue is termed spatial confounding (Reich et al., 2006). Spatial confounding corresponds to the situation where spatially structured latent effects are correlated with the covariates, resulting in biased estimates of the fixed effects (Reich et al., 2006; Page et al., 2017). Reich et al. (2006) proposed a restricted spatial regression (RSR) model to overcome this issue, by removing the collinearity between the covariates and the spatial effects. However, Khan and Calder (2020) point out that RSR models are challenging to use. Additionally to leading to a loss of computational efficiency, these models cannot be naturally extended to other spatial models than the conditional autoregressive ones. Khan and Calder (2020) also note that their use may entail an increased type-S error for both correctly and incorrectly specified

Model	Moment
IID	$\mathbb{E}(Y_i^\ell) = E_i^\ell \exp\left(\beta_0^\ell + \boldsymbol{x_i}\boldsymbol{\beta}^\ell + \sigma_{\ell,v}^2/2\right) \equiv \mu_i^\ell$ $\mathbb{V}(Y_i^\ell) = \mu_i^\ell\left[1 + \mu_i^\ell\left\{\exp(\sigma_{\ell,v}^2) - 1\right\}\right]$ $\mathrm{Cov}(C_i, D_j) = \begin{cases} \mu_i^C\mu_i^D\left[\exp(\rho_v\sigma_{C,v}\sigma_{D,v}) - 1\right], & \text{if } i = j, \\ 0, & \text{otherwise.} \end{cases}$
CAR	$\mathbb{E}(Y_i^\ell) = E_i^\ell \exp\left(\beta_0^\ell + \boldsymbol{x_i}\boldsymbol{\beta}^\ell + \sigma_{\ell,u}^2[\boldsymbol{Q}^-]_{ii}/2\right) \equiv \mu_i^\ell$ $\mathbb{V}(Y_i^\ell) = \mu_i^\ell\left[1 + \mu_i^\ell\left\{\exp\left(\sigma_{\ell,u}^2[\boldsymbol{Q}^-]_{ii}\right) - 1\right\}\right]$ $\mathrm{Cov}(C_i, D_j) = \begin{cases} \mu_i^C\mu_i^D\left[\exp\left(\rho_u\sigma_{C,u}\sigma_{D,u}[\boldsymbol{Q}^-]_{ii}\right) - 1\right], & \text{if } i = j, \\ \mu_i^C\mu_j^D\left[\exp\left(\rho_u\sigma_{C,u}\sigma_{D,u}[\boldsymbol{Q}^-]_{ij}\right) - 1\right], & \text{if } i \sim j, \\ 0, & \text{otherwise.} \end{cases}$
BYM	$\mathbb{E}(Y_i^\ell) = E_i^\ell \exp\left(\beta_0^\ell + \boldsymbol{x_i}\boldsymbol{\beta}^\ell + \sigma_{\ell,v}^2/2 + \sigma_{\ell,u}^2[\boldsymbol{Q}^-]_{ii}/2\right) \equiv \mu_i^\ell$ $\mathbb{V}(Y_i^\ell) = \mu_i^\ell\left[1 + \mu_i^\ell\left\{\exp(\sigma_{\ell,v}^2) - 1\right\}\left\{\exp\left(\sigma_{\ell,u}^2[\boldsymbol{Q}^-]_{ii}\right) - 1\right\}\right]$ $\mathrm{Cov}(C_i, D_j) = \begin{cases} \mu_i^C\mu_i^D\left[\exp\left(\rho_u\sigma_{C,u}\sigma_{D,u}[\boldsymbol{Q}^-]_{ii} + \rho_v\sigma_{C,v}\sigma_{D,v}\right) - 1\right], & \text{if } i = j, \\ \mu_i^C\mu_j^D\left[\exp\left(\rho_u\sigma_{C,u}\sigma_{D,u}[\boldsymbol{Q}^-]_{ij}\right) - 1\right], & \text{if } i \sim j, \\ 0, & \text{otherwise,} \end{cases}$

Table 1: Marginal moments of $\boldsymbol{Y}_i = [C_i, D_i]^\top = [Y_i^C, Y_i^D]^\top$, $i = 1, \ldots, n$, for each model, where $\ell = C, D$.

models, where a type-S error occurs when a posterior confidence interval for a particular coefficient does not include 0 even though the true value is 0. Recently, Dupont et al. (2021) proposed an approach to accommodate spatial confounding called Spatial+. The proposal is to "regress away" the spatial structure from the covariates and only include the residuals from such regression in the modelling of the variable of interest. The Spatial+ approach of Dupont et al. (2021) is based on generalised additive models (GAM). As pointed out by Dupont et al. (2021), because of the relationship between splines and Markov random fields, next we follow Schmidt (2021) and describe how to adapt the Spatial+ approach when the latent effect follow a CAR prior distribution. Dupont et al. (2021) suggest to use Spatial+ as a tool to investigate for the presence of spatial confounding. If after fitting a model with and without spatial effects the fixed effects estimates are similar, this might be an indication that there is no spatial confounding.

We consider the CAR extension that Dupont et al. (2021) proposed for their method and adopt a fully Bayesian approach. Let $\boldsymbol{x}_{\cdot k} = [x_{1k}, \ldots, x_{nk}]^\top$ be the vector for the kth covariate. For each covariate $k = 1, \ldots, K$, we model the covariates through Gaussian distributions with latent spatial effects that follow a CAR distribution *a priori*. More specifically, let

$$x_{ik} \sim \mathcal{N}(\beta_{0k} + f_{ik}, \sigma_k^2), \tag{3}$$

where, *a priori*, $\boldsymbol{f}_k = [f_{1k}, \ldots, f_{nk}] \sim \mathcal{N}_n(\boldsymbol{0}, \sigma_{fk}^2\boldsymbol{Q}^-)$. After assigning prior distributions to the parameters for the model of each $\boldsymbol{x}_{\cdot k}$ we follow the Bayesian paradigm and obtain a sample from the resultant posterior distribution. Then, we define r_{ik} the posterior mean of the residual $x_{ik} - \widehat{\beta_{0k}} - \widehat{f_{ik}}$, $i = 1, \ldots, N$, $j = 1, \ldots, K$, where $\widehat{\beta_{0k}}$ and $\widehat{f_{ik}}$ are, respectively,

the point estimates (posterior means) of β_{0k} and f_{ik}. Finally, each vector of potentially spatially confounded covariate, $\boldsymbol{x}_{\cdot k}$ is replaced by $\boldsymbol{r}_{\cdot k} = [r_{1k}, \ldots, r_{nk}]^\top, k = 1, \ldots, K$ in the models for the cases and deaths that include spatial effects.

2.2. Inference procedure

In the fully Bayesian framework that we consider, regardless of the prior specification of the parameters, none of the Poisson models fitted to the data result in closed form posterior distributions. Hence, we approximate these posterior distributions using computational methods, specifically, Markov Chains Monte Carlo (MCMC) methods. The MCMC procedures are run in R using the `Nimble` package (de Valpine et al., 2017, 2021). The advantage of `Nimble` is that censored outcome values are naturally accounted for, as is necessary for the areas Baie-d'Urfé and Montréal-Ouest.

The CAR prior is not a proper distribution and is invariant to the addition of a constant (Rue and Held, 2005). This implies there is an identifiability issue between the spatially structured effects and the global mean, the intercept. To ensure the posterior distribution is proper, we impose a sum-to-zero constraint for each of the CAR distributed effects. For each of the MCMC iterations, we impose $(1/n) \sum_{i=1}^{n} u_i^\ell = 0$, $\ell = C, D$. Lawson (2020) discusses how to implement the multivariate CAR model in `Nimble`. The `Nimble` codes used for this analysis are all available in Appendix C.

3. Results

For each of the 33 boroughs and related cities of Montreal, we have available the case and death counts due to COVID-19 that were recorded until July 25th, 2021. Let C_i be the cases count in area $i = 1, \ldots, 33$, and D_i, the deaths count. The spatial distributions of the cases and deaths are shown in Figure 1 in Section 1. To model the cases and deaths, we consider as covariates the number of long-term care home beds, the percentage of the population with a university diploma, and the median age in each area. Figure 2 shows the maps of the variables as included in the models. The median age is scaled and the log of the number of beds is computed. Note that the log scale is used to assume a normal approximation of the number of beds. This is necessary, when fitting the Spatial+ models, as we assume that the covariates are normally distributed when adjusting for potential spatial confounding following Dupont et al. (2021). The resulting covariates from the spatial confounding adjustment are shown in Figure B.6 in Appendix B.

Figure 2: Maps of the three covariates included in the models

We are interested in assessing the association of the selected covariates with the number of cases and deaths. We are also interested in measuring the correlation between the cases and deaths inside a particular area or between neighbouring areas. To that end, we allow for a spatial autocorrelation between the cases and deaths, as described in section 2. Hence, we fit six different Poisson models, as described in Section 2. Table 2 presents the models characteristics, namely the form of the latent effects, b_i^ℓ, $i = 1, \ldots, 33$, $\ell = C, D$, and the covariates included. For all these models, the same priors are defined for the parameters involved in the fixed effects: $\beta_0^C, \beta_0^D, \beta_k^C, \beta_k^D \sim \mathcal{N}(0, 10^2)$, $k = 1, 2, 3$. Additionally, the relevant standard deviations that appear in the latent effects' prior distributions are given a half-Cauchy prior, $\sigma_{\ell,v}, \sigma_{\ell,u} \sim \mathcal{HC}(0,1)$, $\ell = C, D$. Finally, when appropriate, the latent effects' correlation parameters are assumed to follow a noninformative uniform prior distribution: $\rho_v, \rho_u \sim \mathcal{U}(-1, 1)$.

	Latent effects		Covariates included in the fixed effects	
	Unstructured	Structured	Original	Adjusted for Spatial Confounding
	$\Gamma_v \boldsymbol{v}_i$	$\Gamma_u \boldsymbol{u}_i$	$\boldsymbol{x}_{\cdot k}$	$\boldsymbol{r}_{\cdot k}$
Simple	-	-	✓	-
IID	✓	-	✓	-
CAR	-	✓	✓	-
BYM	✓	✓	✓	-
CAR+	-	✓	-	✓
BYM+	✓	✓	-	✓

Table 2: List of the fitted models to the number of cases and deaths due to COVID-19 across the boroughs of Montreal. The symbol ✓ denotes which components were included in the respective model.

The models are fitted through the R package Nimble (de Valpine et al., 2021). The

MCMC procedure consists of 2 chains of 800,000 iterations each, with a burn-in period of 400,000 iterations and a thinning factor of 160. An elliptical slice sampler is used for the regression coefficients, as these are normally distributed *a priori*. This sampler was chosen to efficiently sample from these posterior full conditionals. Metropolis-Hastings algorithms are used for the independent random effects in the IID, BYM and BYM+ models. For the spatial effects included in the CAR, BYM, CAR+ and BYM+ models, the default CAR_normal sampler from Nimble is used. The chains have mixed well, as assessed by the trace plots, effective sample sizes and the statistic proposed by Gelman and Rubin (1992). The reason for the need of these long chains seems to be related to the possible collinearity between the variable diploma and the models that include a spatial effect. The IID model and the ones that accommodate spatial confounding showed convergence before reaching the burn in of 400,000 iterations. For consistency, all models were fitted with the same number of iterations.

The posterior summaries for the fixed effects' coefficients are presented in Figure 3 for each model. The solid circles correspond to the estimated posterior means, while the segments are the limits of the 95% posterior credible intervals. The first row corresponds to the fixed effects in the log relative risk of cases and the second row, to the ones in the log relative risk of recorded deaths due to COVID-19. Clearly, the inclusion of random effects in each of the equations change the importance of each of the covariates. The covariate beds is negatively associated with the log relative risk of cases under the Simple model and as a random effect is included in the model, zero falls within the 95% posterior credible interval, suggesting that the number of beds is not associated with the cases. For the log relative risk of deaths, on the other hand, the association with the number of beds is strictly positive. This is expected as the majority of deaths in Quebec are connected to residential and long-term care centres.

The variable diploma is negatively associated with the log relative risk of cases. The range of the 95% posterior credible interval is much wider under the CAR+ and BYM+ models than the ones obtained under the IID and CAR models. Although diploma results in a positive association with the log relative risk of death under the Simple model, once a random effect is included, zero is within the 95% posterior credible of the coefficient associated with diploma. This suggests that the average educational level of the borough is not associated with the risk of death.

For the median age of the borough, the Simple and IID models suggest a negative association with the log relative risk of cases and a positive association with the log relative risk of deaths. For the log relative risk of cases once a spatial structure is included in the model, zero falls within the 95% posterior credible interval. This is also true for the CAR+ and BYM+ models which account for spatial confounding.

Figure 3: Posterior summaries for the model coefficients across all the fitted models. Solid circles: posterior means; Vertical lines: 95% posterior credible intervals; Dashed line: indicates no association.

We model the cases and deaths jointly in order to borrow strength from the different recordings both inside each area (e.g., IID model) and between areas (e.g., CAR model). Figure 4 shows the posterior summaries (posterior mean and posterior 95% credible interval) for the correlation coefficients, ρ_v and ρ_u, respectively included in the latent effects' IID, CAR, CAR+, BYM or BYM+ models. The BYM and BYM+ yield posterior 95% credible intervals that include 0 for ρ_u and ρ_v. On the other hand, the IID, CAR and CAR+ result in negative correlations between the latent effects for the cases and for the deaths ($[-0.73, -0.08]$, $[-0.72, -0.07]$ and $[-0.76, -0.02]$, respectively). This suggests that modelling the cases and deaths jointly is adequate. Note that for the BYM and BYM+ models, if the posterior 95% credible limits for ρ_u and ρ_v include zero, this does not necessarily imply that the marginal correlations will also include zero. From Table 1, it is clear that the correlation between cases and deaths within and across boroughs also depends on the latent effects' standard deviations as well as the spatial structure, when appropriate. Figure B.7 in Appendix B shows the posterior summaries for the intra-borough correlation between cases and deaths for each model. Although ρ_u and ρ_v have posterior credible intervals that include zero, the BYM model yields strictly negative intervals in 19 boroughs.

Figure 4: Posterior summaries for the correlation coefficient of the latent effects. Solid circles: posterior means; Vertical lines: 95% posterior credible intervals. Dashed line: $\rho_u = 0$ and $\rho_v = 0$.

Another goal of this analysis, other than investigate the association of the log relative risks of cases and deaths with the available covariates, is to estimate the correlation between the cases and deaths within and across boroughs. The posterior summaries for the intra-borough correlations between cases and deaths obtained for each model are available in Appendix B, in Figure B.7 and the posterior means for the correlation between areas estimated under the CAR, BYM, CAR+ and BYM+ models are available in Appendix B, in Figure B.8, Figure B.9, Figure B.10 and Figure B.11, respectively. We note that all the estimated correlations are negative. Figure 5 maps the posterior means of the correlations between cases and deaths inside each borough for each model that include latent effects. The IID model does not allow for a correlation across boroughs but does accommodate an intra-borough correlation (see Table 1). All the models estimate negative correlations between the cases and deaths on average *a posteriori*. For example, the correlation's posterior mean in Senneville is -0.2 for the CAR model and -0.1 for the BYM model. In Côte-des-Neiges-Notre-Dame-de-Grâce, the posterior mean is -0.4 for the CAR+ model and -0.3 for the BYM. For the IID, CAR and CAR+ models, the correlations' posterior 95% credible intervals (see Figure B.7 in Appendix B) are always negative. In Figure 5, there does not seem to be a spatial structure in the negative posterior means of the intra-borough correlations between the cases and deaths due to COVID-19. Additionally, Figure B.12 in Appendix B maps the relative risks of cases and deaths as estimated by each model. The negative correlations between cases and deaths are illustrated in this Figure B.12. For example, the south west region of Montreal has low estimated risks of cases but high estimated risks of deaths for each model.

Figure 5: Maps of the posterior means of the intra-borough correlation between cases and deaths for each model. These correlations were estimated based on the equations for the covariance under each model shown in Table 1.

Finally, we compare the models' performances using the WAIC (Watanabe and Opper, 2010) in Table 3. In terms of WAIC, for which smaller values are preferred, the Simple model is the least adequate among the fitted ones, with a value of 3,320 and 265 effective number of parameters. This suggests that there is a need for latent effects in the modelling of cases and deaths due to COVID-19 in Montreal. However, there does not seem to be a need for a spatial structure in these random effects as the CAR, BYM, CAR+ and BYM+

all perform worse than the IID model. The IID model yields a WAIC of 205 whereas the others yield a WAIC of 212, 225, 243 and 227, respectively.

	Simple	IID	CAR	BYM	CAR+	BYM+
WAIC	3320.1	*205.0*	212.2	224.9	242.7	227.1
p_W	265.2	*57.2*	60.7	67.2	75.3	68.0

Table 3: Values of WAIC and effective number of parameters for each fitted model. The smallest value indicates the best model among fitted ones (in italics).

4. Discussion

In this paper, we analyse the case and death counts due to COVID-19 across the 33 boroughs of Montreal, as of July 25th, 2021. We fit six different models with conditional Poisson distributions for the cases and for the deaths in each borough. The areal means are decomposed in the log scale as the sum of an offset, fixed effects, and latent random effects. One model is fitted without latent effects. We then allow for a correlation between the cases' latent effects and the deaths' latent effects within an area. In one model, we assume that the latent effects are independent across the areas, whereas in the other four models, we accommodate a potential spatial autocorrelation between the latent effects of neighbouring areas. In two of these models, we set multivariate CAR and BYM priors for the latent effects. In the final two models, we first adjust the covariates for potential spatial confounding following Dupont et al. (2021) and then consider again multivariate CAR and BYM priors for the latent effects.

Interestingly, the posterior summaries for the fixed effects parameters are reversed between the cases and deaths sections of the models. Regarding the log number of beds, the coefficient has posterior 95% credible intervals that include zero for the cases (through the IID, CAR, BYM, CAR+ and BYM+ models) but the posterior credible intervals are strictly positive for the deaths. We used the number of beds in CHSLDs as a proxy for the healthcare infrastructure in Montreal. This result suggests that this covariate has no link with the recorded number of cases but that, given the number of cases, an increased number of beds corresponds to a higher risk of death due to COVID-19. This result is consistent with the COVID-19 situation in Montreal and in the province of Quebec. In particular, long-term care homes were the epicentre of the COVID-19 crisis in Quebec (Hsu et al., 2020; Institut national de santé publique du Québec, 2021). The associations between the percentage of the population with a university diploma and the cases and deaths due to COVID-19 in Montreal are also different for each outcome considered. There does not seem to be a relationship between the percentage of people with diploma and the risk of death due to COVID-19, given the cases, for the IID, CAR, BYM, CAR+ and BYM+ models. On the other hand, each model results in negative posterior credible intervals for the effects of the diploma on the number of cases. The diploma variable is used as a proxy to the socio-economic situation of each borough in Montreal. This result seems to be aligned with other studies that also found that level of education is negatively associated with the log

risk of cases (see e.g., Hawkins et al. (2020)). Finally, the risk of COVID-19 cases does not seem to be associated with age for the CAR, BYM, CAR+ and BYM+ models whereas older populations seem correlated with higher risks of death due to COVID-19. Once again, this result agrees with the current knowledge on COVID-19: all may contract COVID-19, but younger people are less likely to die from this disease (Williamson et al., 2020).

As the importance of the covariates changed when a latent spatially structured random effect was considered, we also investigated if there is spatial confounding by fitting a Bayesian version of the method proposed by Dupont et al. (2021). We find that the log number of beds does not seem to be spatially confounded, whereas the diploma and median age do. As suggested by Dupont et al. (2021), we believe there is an interest in adjusting for spatial confounding even when there is no intuition of spatial confounding, as their method can be used as a test for the presence or absence of such confounding. However, in terms of WAIC, the model with only IID latent effects seems to perform the best among the fitted ones. This suggests that there is no need to account for spatially structured latent effects. Note that the adjustment for spatial confounding is unnecessary in this IID model.

We believe that there is an interest to the data analysis conducted in this paper as we are able to provide an estimate of the correlations between the cases and deaths due to COVID-19 within, and across, boroughs. Fitting separate models for cases and deaths, or even joint models that assume parameters to be independent *a priori*, such as in the Simple model, impose independence between cases and deaths. Our proposed approach allows the data to drive the inference procedure, and if there is correlation between the two outcomes, this can be estimated through the marginal correlations computed from the results shown in Table 1. Note that in terms of WAIC, we find that the Simple model performs worse than the five other models that include latent effects which allow for marginal correlations between cases and deaths both within and across boroughs. Finally, we find that the correlation between the cases and deaths due to COVID-19 within and between boroughs is always negative. This may be due to the aggregation of cases and deaths over time. We only have available aggregated data, which may yield different correlations than what would be observed at a certain time or through a temporal analysis. Looking at aggregated provincial time series of cases and deaths due to COVID-19, the number of cases and deaths were positively correlated at the beginning of the pandemic. We hypothesise that the estimated negative correlation between cases and deaths are related to different policies that were instituted throughout the pandemic and, moreover, due to the fact that people started being vaccinated. For instance, in the province of Quebec, the number of cases increased from 573 to 1683 between March 15th, 2021, and April 15th, 2021, whereas the vaccination coverage for the population aged 80 and older increased from 57.9% to 91.1% in the same period (Institut national de santé publique du Québec, 2021). It can be observed in the aggregated time series that in this period, the daily number of deaths somewhat stabilised at low numbers.

In this paper, we used the median age of the boroughs as a proxy to the age population profile. If data for different age groups are available, one could consider models for case and death counts for each age stratum, allowing for different effects of the available variables and correlation parameters across the different age groups. In this framework, one would need to carefully consider the inclusion of the number of long-term care homes' beds in the

various strata.

Funding

This work was supported by the Natural Sciences and Engineering Research Council (NSERC) of Canada [Discovery Grant RGPIN-2017-04794] and the IVADO Fundamental Research Project [PRF-2019-6839748021].

Declaration of Competing Interests

The authors report no conflict of interest or financial interests.

References

Banerjee, S., Carlin, B.P., Gelfand, A.E., 2014. Hierarchical modeling and analysis for spatial data. CRC press. doi:10.1201/b17115.

Bermudi, P.M.M., Lorenz, C., de Aguiar, B.S., Failla, M.A., Barrozo, L.V., Chiaravalloti-Neto, F., 2021. Spatiotemporal ecological study of COVID-19 mortality in the city of São Paulo, Brazil: shifting of the high mortality risk from areas with the best to those with the worst socio-economic conditions. Travel Medicine and Infectious Disease 39, 101945. doi:10.1016/j.tmaid.2020.101945.

Besag, J., 1974. Spatial interaction and the statistical analysis of lattice systems. Journal of the Royal Statistical Society: Series B (Methodological) 36, 192–225. doi:10.1111/j.2517-6161.1974.tb00999.x.

Besag, J., York, J., Mollié, A., 1991. Bayesian image restoration, with two applications in spatial statistics. Annals of the Institute of Statistical Mathematics 43, 1–20. doi:10.1007/BF00116466.

Bonanad, C., García-Blas, S., Tarazona-Santabalbina, F., Sanchis, J., Bertomeu-González, V., Fácila, L., Ariza, A., Núñez, J., Cordero, A., 2020. The effect of age on mortality in patients with COVID-19: a meta-analysis with 611,583 subjects. Journal of the American Medical Directors Association 21, 915–918. doi:10.1016/j.jamda.2020.05.045.

Clayton, D.G., Bernardinelli, L., Montomoli, C., 1993. Spatial correlation in ecological analysis. International Journal of Epidemiology 22, 1193–1202. doi:10.1093/ije/22.6.1193.

de Valpine, P., Paciorek, C., Turek, D., Michaud, N., Anderson-Bergman, C., Obermeyer, F., Wehrhahn Cortes, C., Rodríguez, A., Temple Lang, D., Paganin, S., 2021. NIMBLE: MCMC, particle filtering, and programmable hierarchical modeling. doi:10.5281/zenodo.1211190.

de Valpine, P., Turek, D., Paciorek, C., Anderson-Bergman, C., Temple Lang, D., Bodik, R., 2017. Programming with models: writing statistical algorithms for general model structures with NIMBLE. Journal of Computational and Graphical Statistics 26, 403–417. doi:10.1080/10618600.2016.1172487.

[dataset] Direction régionale de santé publique, 2021. Situation du coronavirus (COVID-19) à Montréal. URL: https://santemontreal.qc.ca/population/coronavirus-covid-19/situation-du-coronavirus-covid-19-a-montreal/. Accessed July 25, 2021.

Dong, E., Du, H., Gardner, L., 2020. An interactive web-based dashboard to track COVID-19 in real time. The Lancet Infectious Diseases 20, 533–534. doi:10.1016/S1473-3099(20)30120-1. Accessed October 4, 2021.

Doucet, H., 2020. "Responding to vulnerability": Ethics and CHSLD in the age of COVID-19. Éthique & Santé 17, 142–146. doi:10.1016/j.etiqe.2020.07.003.

Dupont, E., Wood, S.N., Augustin, N., 2021. Spatial+: a novel approach to spatial confounding (with discussion). To appear in Biometrics .

Gangemi, S., Billeci, L., Tonacci, A., 2020. Rich at risk: socio-economic drivers of COVID-19 pandemic spread. Clinical and Molecular Allergy 18, 1–3. doi:10.1186/s12948-020-00127-4.

Gelman, A., Rubin, D.B., 1992. Inference from iterative simulation using multiple sequences. Statistical Science 7, 457–472. doi:10.1214/ss/1177011136.

Ghebreyesus, T.A., 2020. WHO Director-General's opening remarks at the media briefing on COVID-19 - 11 March 2020. URL: `https://www.who.int/director-general/speeches/detail/who-director-general-s-opening-remarks-at-the-media-briefing-on-covid-19---11-march-2020`. Accessed September 30, 2021.

Gouvernement du Québec, 2021. Act respecting health services and social services, RSQ, chapter s-4.2. URL: `http://legisquebec.gouv.qc.ca/en/ShowDoc/cs/S-4.2`. Accessed September 15, 2021.

Government of Canada, 2021. COVID-19 daily epidemiology update. URL: `https://health-infobase.canada.ca/covid-19/epidemiological-summary-covid-19-cases.html?stat=num&measure=total&map=pt#a2`. Accessed October 5, 2021.

Hawkins, R.B., Charles, E., Mehaffey, J., 2020. Socio-economic status and COVID-19-related cases and fatalities. Public Health 189, 129–134. doi:`10.1016/j.puhe.2020.09.016`.

Hsu, A., Lane, N., Sinha, S., Dunning, J., Dhuper, M., Kahiel, Z., Sveistrup, H., 2020. Understanding the impact of COVID-19 on residents of canada's long-term care homes—ongoing challenges and policy responses. International Long Term Care Policy Network , 15.

Institut national de santé publique du Québec, 2021. Données COVID-19 par région sociosanitaire. URL: `https://www.inspq.qc.ca/covid-19/donnees/par-region`. Accessed October 1, 2021.

Jin, X., Banerjee, S., Carlin, B.P., 2007. Order-free co-regionalized areal data models with application to multiple-disease mapping. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 69, 817–838. doi:`10.1111/j.1467-9868.2007.00612.x`.

Khan, K., Calder, C.A., 2020. Restricted spatial regression methods: Implications for inference. Journal of the American Statistical Association , 1–13doi:`10.1080/01621459.2020.1788949`.

Lawson, A.B., 2020. NIMBLE for Bayesian disease mapping. Spatial and Spatio-temporal Epidemiology 33, 100323. doi:`10.1016/j.sste.2020.100323`.

Le Blanc, M.F., Raynault, M.F., Lessard, R., 2011. 2011 Report of the Director of Public Health. Social Inequalities in Health in Montreal. Progress to Date. Synthesis Report. Direction de santé publique: Agence de la santé et des services sociaux de Montréal.

Ministère de la Santé et des Services sociaux, 2021. Définition de cas de COVID-19 - Québec. URL: `https://msss.gouv.qc.ca/professionnels/documents/coronavirus-2019-ncov/definition-nosologique-COVID-19_2021-07-05.pdf`. Accessed September 15, 2021.

[dataset] Ministère de la Santé et des Services sociaux, 2021a. Fichier cartographique des installations - M02. URL: `https://publications.msss.gouv.qc.ca/msss/document-001644/`. Accessed July 25, 2021.

[dataset] Ministère de la Santé et des Services sociaux, 2021b. Installations : Lieux physiques. URL: `https://m02.pub.msss.rtss.qc.ca/M02ListeInstall.asp`. Accessed July 25, 2021.

Page, G.L., Liu, Y., He, Z., Sun, D., 2017. Estimation and prediction in the presence of spatial confounding for spatial linear models. Scandinavian Journal of Statistics 44, 780–797. doi:`10.1111/sjos.12275`.

Reich, B.J., Hodges, J.S., Zadnik, V., 2006. Effects of residual smoothing on the posterior of the fixed effects in disease-mapping models. Biometrics 62, 1197–1206. doi:`10.1111/j.1541-0420.2006.00617.x`.

Rue, H., Held, L., 2005. Gaussian Markov random fields: theory and applications. CRC press. doi:`10.1201/9780203492024`.

Sahu, S.K., Böhning, D., 2021. Bayesian spatio-temporal joint disease mapping of Covid-19 cases and deaths in local authorities of England. Spatial Statistics , 100519doi:`10.1016/j.spasta.2021.100519`.

Schmidt, A.M., 2021. Discussion on "Spatial+: a novel approach to spatial confounding" by Emiko Dupont, Simon N. Wood and Nicole A. Augustin. To appear in Biometrics .

Sharma, A., Borah, S.B., Moses, A.C., 2021. Responses to COVID-19: The role of governance, healthcare infrastructure, and learning from past pandemics. Journal of Business Research 122, 597–607. doi:`10.1016/j.jbusres.2020.09.011`.

Statistics Canada, 2017a. Canada [country] and Canada [country] (table). 2016 Census. Statistics Canada Catalogue no. 98-316-X2016001. URL: `https://www12.statcan.gc.ca/census-recensement/2016/dp-pd/prof/index.cfm?Lang=E`. Accessed September 15, 2021.

Statistics Canada, 2017b. Montréal [economic region], Quebec and Quebec [province] (table). Census profile. 2016 Census. Statistics Canada Catalogue no. 98-316-X2016001. URL: `https://www12.statcan.gc.ca/`

census-recensement/2016/dp-pd/prof/index.cfm?Lang=E. Accessed September 15, 2021.

[dataset] Ville de Montréal, 2016. Profils sociodémographiques 2016. URL: http://ville.montreal.qc.ca/portal/page?_pageid=6897,68149755&_dad=portal&_schema=PORTAL. Accessed September 15, 2021.

Wakefield, J., 2013. Bayesian and frequentist regression methods. Springer Science & Business Media. doi:10.1007/978-1-4419-0925-1_7.

Watanabe, S., Opper, M., 2010. Asymptotic equivalence of bayes cross validation and widely applicable information criterion in singular learning theory. Journal of Machine Learning Research 11.

Williamson, E.J., Walker, A.J., Bhaskaran, K., Bacon, S., Bates, C., Morton, C.E., Curtis, H.J., Mehrkar, A., Evans, D., Inglesby, P., et al., 2020. Factors associated with COVID-19-related death using openSAFELY. Nature 584, 430–436. doi:10.1038/s41586-020-2521-4.

Yanez, N.D., Weiss, N.S., Romand, J.A., Treggiari, M.M., 2020. COVID-19 mortality risk for older men and women. BMC Public Health 20, 1–7. doi:10.1186/s12889-020-09826-8.

Zhou, F., Yu, T., Du, R., Fan, G., Liu, Y., Liu, Z., Xiang, J., Wang, Y., Song, B., Gu, X., et al., 2020. Clinical course and risk factors for mortality of adult inpatients with COVID-19 in Wuhan, China: a retrospective cohort study. The Lancet 395, 1054–1062. doi:10.1016/S0140-6736(20)30566-3.

Appendix A. Detailed computation of the covariance between cases and deaths

In this section, we develop the general expression for the correlation between the cases and deaths inside each area and between areas. Recall the general formulation of the model described in section 2 for C_i, the number of cases in area i, $i = 1, \ldots, n$, and D_i, the deaths count:

$$C_i \sim \text{Pois}(E_i^C \lambda_i^C) \quad \text{and} \quad D_i \mid C_i \sim \text{Pois}(E_i^D \lambda_i^D),$$

where Pois stands for the Poisson distribution and where E_i^C and E_i^D are offsets. We further decompose the relative risks, λ_i^C and λ_i^D, in the log scale as follows:

$$\log(\lambda_i^C) = \beta_0^C + \boldsymbol{x}_i \boldsymbol{\beta}^C + b_i^C \quad \text{and} \quad \log(\lambda_i^D) = \beta_0^D + \boldsymbol{x}_i \boldsymbol{\beta}^D + b_i^D,$$

for β_0^C and β_0^D the mean log risks and for $\boldsymbol{x}_i$ the vector of K covariates associated to the coefficients $\boldsymbol{\beta}^C$ and $\boldsymbol{\beta}^D$. The latent random effects b_i^C and b_i^D may be defined in the general form:

$$\boldsymbol{b}_i \equiv [b_i^C, \ b_i^D]^\top = \Gamma_u \boldsymbol{u}_i + \Gamma_v \boldsymbol{v}_i,$$

where $\boldsymbol{v}_i = [v_i^C, \ v_i^D]^\top \overset{i.i.d.}{\sim} \mathcal{N}_2(\boldsymbol{0}, \boldsymbol{I}_2)$ is independent from $\boldsymbol{u}_i = [u_i^C, \ u_i^D]^\top$, for $\boldsymbol{u}^\ell \sim \mathcal{N}_n(\boldsymbol{0}, \boldsymbol{Q}^-)$, $\ell = C, D$, with $\boldsymbol{I}_d$, the $d \times d$ identity matrix, $\boldsymbol{Q}$, the matrix defining the spatial structure, and with

$$\Gamma_u \Gamma_u^\top = \Sigma_u = \begin{bmatrix} \sigma_{C,u}^2 & \rho_u \sigma_{C,u} \sigma_{D,u} \\ \rho_u \sigma_{C,u} \sigma_{D,u} & \sigma_{D,u}^2 \end{bmatrix} \quad \text{and} \quad \Gamma_v \Gamma_v^\top = \Sigma_v = \begin{bmatrix} \sigma_{C,v}^2 & \rho_v \sigma_{C,v} \sigma_{D,v} \\ \rho_v \sigma_{C,v} \sigma_{D,v} & \sigma_{D,v}^2 \end{bmatrix}.$$

Hence, the distribution for the long vector $\boldsymbol{b} = [b_1^C, \ldots, b_n^C, b_1^D, \ldots, b_n^D]^\top$ is $\boldsymbol{b} \sim \mathcal{N}_{2n}(\boldsymbol{0}, \Sigma_b)$, for

$$\Sigma_b = \begin{bmatrix} \sigma_{C,u}^2 \boldsymbol{Q}^- + \sigma_{C,v}^2 \boldsymbol{I}_n & \rho_u \sigma_{C,u} \sigma_{D,u} \boldsymbol{Q}^- + \rho_v \sigma_{C,v} \sigma_{D,v} \boldsymbol{I}_n \\ \rho_u \sigma_{C,u} \sigma_{D,u} \boldsymbol{Q}^- + \rho_v \sigma_{C,v} \sigma_{D,v} \boldsymbol{I}_n & \sigma_{D,u}^2 \boldsymbol{Q}^- + \sigma_{D,v}^2 \boldsymbol{I}_n \end{bmatrix}.$$

Given the fixed effects, we can compute the *marginal* covariances between cases and deaths, marginalising over the random effects:

$$\begin{aligned} \text{Cov}(C_i, D_j) &= \text{Cov}_{b_i^C, b_j^D} \left[\mathbb{E}(C_i \mid b_i^C), \mathbb{E}(D_j \mid b_j^D) \right] + \mathbb{E}_{b_i^C, b_j^D} \left[\text{Cov}(C_i, D_j \mid b_i^C, b_j^D) \right] \\ &= E_i^C \exp(\beta_0^C + \boldsymbol{x}_i \boldsymbol{\beta}^C) E_j^D \exp(\beta_0^D + \boldsymbol{x}_j \boldsymbol{\beta}^D) \text{Cov}_{b_i^C, b_j^D} \left[\exp(b_i^C), \exp(b_j^D) \right] \\ &= E_i^C E_j^D \exp(\beta_0^C + \boldsymbol{x}_i \boldsymbol{\beta}^C) \exp(\beta_0^D + \boldsymbol{x}_j \boldsymbol{\beta}^D) \\ &\quad \times \exp\left(\frac{1}{2} [\Sigma_{b,ii} + \Sigma_{b,(n+j)(n+j)}] \right) \left[\exp(\Sigma_{b,i(n+j)}) - 1 \right] \\ &= \mathbb{E}(C_i) \mathbb{E}(D_j) \left[\exp(\Sigma_{b,i(n+j)}) - 1 \right], \end{aligned}$$

with marginal expectation for the cases $\mathbb{E}(C_i) = E_i^C \exp(\beta_0^C + \boldsymbol{x}_i \boldsymbol{\beta}^C + 1/2[\sigma_{C,u}^2 [\boldsymbol{Q}^-]_{ii} + \sigma_{C,v}^2])$ and for the deaths $\mathbb{E}(D_j) = E_j^D \exp(\beta_0^D + \boldsymbol{x}_j \boldsymbol{\beta}^D + 1/2[\sigma_{D,u}^2 [\boldsymbol{Q}^-]_{jj} + \sigma_{D,v}^2])$. Hence, inside a particular area i, for $i = j$, we have

$$\text{Cov}(C_i, D_i) = \mathbb{E}(C_i) \mathbb{E}(D_i) \left[\exp(\rho_u \sigma_{C,u} \sigma_{D,u} [\boldsymbol{Q}^-]_{ii} + \rho_v \sigma_{C,v} \sigma_{D,v}) - 1 \right].$$

Across areas, if i and j are not neighbouring areas, then $\text{Cov}(C_i, D_j) = 0$ and if $i \sim j$, we obtain

$$\text{Cov}(C_i, D_i) = \mathbb{E}(C_i) \mathbb{E}(D_i) \left[\exp(\rho_u \sigma_{C,u} \sigma_{D,u} [\boldsymbol{Q}^-]_{ij}) - 1 \right].$$

Appendix B. Additional results from the analysis of cases and deaths due to COVID-19 in Montreal

The maps of the covariates after adjusting for the spatial confounding are shown in Figure B.6. Figure B.7 shows the posterior summaries for the intra-borough correlation between the cases and deaths, as estimated by the IID, CAR, BYM, CAR+ and BYM+ models. Figures B.8, B.9, B.10 and B.11 plot the posterior means for the correlations between boroughs. The posterior means for the relative risks from each model are mapped in Figure B.12.

Figure B.6: Maps of the three covariates included in the models after adjusting for spatial confounding

Figure B.7: Posterior summaries for the intra-borough correlation between cases and deaths for each of the IID, CAR, BYM, CAR+ and BYM+ models

Figure B.8: Posterior means of the correlation between deaths and cases inside and between the borough for the CAR model

Figure B.9: Posterior means of the correlation between deaths and cases inside and between the borough for the BYM model

Figure B.10: Posterior means of the correlation between deaths and cases inside and between the borough for the CAR+ model

Figure B.11: Posterior means of the correlation between deaths and cases inside and between the borough for the BYM+ model

Figure B.12: Maps of the posterior means for the relative risks estimated from the six models for the cases and deaths. Rows 1 & 2: Simple, CAR and IID models; Rows 3 & 4: BYM, CAR+ and BYM+ models.

Appendix C. Nimble codes

The Nimble codes used to perform the analyses presented in this paper are listed below.

Listing 1: Nimble code for the joint model with independent latent effects (IID)

```
1  IID_Code <- nimbleCode({
2    for(i in 1:N){
3      # Likelihood
4      cases[i] ~ dpois(exp(b[i,1]))
5      lambda_cases[i] <- exp(b[i,1])
6      censored[i] ~ dinterval(deaths[i], c[i,])
7      deaths[i] ~ dpois(exp(b[i,2]))
8      lambda_deaths[i] <- exp(b[i,2])
9
10     # Priors
11     mean_b[i,1] <- log(E_cases[i]) + beta0_cases + inprod(X[i,1:p], beta_cases[1:p])
12     mean_b[i,2] <- log(E_deaths[i]) + beta0_deaths + inprod(X[i,1:p], beta_deaths[1:p])
13
14     b[i,1:2] ~ dmnorm(mean=mean_b[i,1:2], cholesky=Chol_Sigma_b[1:2,1:2], prec_param=0)
15     b_cases[i] <- b[i,1] - mean_b[i,1]
16     b_deaths[i] <- b[i,2] - mean_b[i,2]
17   }
18
19   Sigma_b[1,1] <- sigma_b_cases^2
20   Sigma_b[1,2] <- rho*sigma_b_cases*sigma_b_deaths
21   Sigma_b[2,1] <- rho*sigma_b_cases*sigma_b_deaths
22   Sigma_b[2,2] <- sigma_b_deaths^2
23   Chol_Sigma_b[1:2,1:2] <- chol(Sigma_b[1:2,1:2])
24
25   beta0_cases ~ dnorm(0, sd=10)
26   beta0_deaths ~ dnorm(0, sd=10)
27   for(k in 1:p){
28     beta_cases[k] ~ dnorm(0, sd=10)
29     beta_deaths[k] ~ dnorm(0, sd=10)
30   }
31   sigma_b_cases ~ T(dt(0,1,1),0,)
32   sigma_b_deaths ~ T(dt(0,1,1),0,)
33   rho ~ dunif(-1,1)
34 })
```

Listing 2: Nimble code for the joint model with spatial latent effects (CAR)

```
1  CAR_Code <- nimbleCode({
2    for(i in 1:N){
```

```
3      # Likelihood
4      cases[i] ~ dpois(lambda_cases[i])
5      lambda_cases[i] <- exp(log(E_cases[i]) + beta0_cases + inprod(X[i,1:p], beta_cases[1:p])
6                               + s[i,1])
7      censored[i] ~ dinterval(deaths[i], c[i,])
8      deaths[i] ~ dpois(lambda_deaths[i])
9      lambda_deaths[i] <- exp(log(E_deaths[i]) + beta0_deaths
10                               + inprod(X[i,1:p], beta_deaths[1:p]) + s[i,2])
11
12     s[i,1:2] <- Achol[1:2,1:2] %*% u[i, 1:2]
13     s_cases[i] <- s[i,1]
14     s_deaths[i] <- s[i,2]
15   }
16
17   # Priors
18   beta0_cases ~ dnorm(0, sd=10)
19   beta0_deaths ~ dnorm(0, sd=10)
20
21   u[1:N,1] ~ dcar_normal(adj[1:L], weights[1:L], num[1:N], 1, zero_mean=1)
22   u[1:N,2] ~ dcar_normal(adj[1:L], weights[1:L], num[1:N], 1, zero_mean=1)
23
24   Sigma_b[1,1] <- sigma_b_cases^2
25   Sigma_b[1,2] <- rho*sigma_b_cases*sigma_b_deaths
26   Sigma_b[2,1] <- rho*sigma_b_cases*sigma_b_deaths
27   Sigma_b[2,2] <- sigma_b_deaths^2
28   Achol[1:2,1:2] <- t(chol(Sigma_b[1:2,1:2]))
29
30   sigma_b_cases ~ T(dt(0,1,1),0,)
31   sigma_b_deaths ~ T(dt(0,1,1),0,)
32   rho ~ dunif(-1,1)
33
34   for(k in 1:p){
35     beta_cases[k] ~ dnorm(0, sd=10)
36     beta_deaths[k] ~ dnorm(0, sd=10)
37   }
38 })
```

Listing 3: Nimble code for the joint model with spatial and independent latent effects (BYM)

```
1 BYM_Code <- nimbleCode({
2   for(i in 1:N){
3      # Likelihood
4      cases[i] ~ dpois(exp(theta[i,1] + s[i,1]))
```

```
5     lambda_cases[i] <- exp(theta[i,1] + s[i,1])
6     censored[i] ~ dinterval(deaths[i], c[i,])
7     deaths[i] ~ dpois(exp(theta[i,2] + s[i,2]))
8     lambda_deaths[i] <- exp(theta[i,2] + s[i,2])
9
10    s[i,1:2] <- Achol[1:2,1:2] %*% u[i, 1:2]
11    s_cases[i] <- s[i,1]
12    s_deaths[i] <- s[i,2]
13
14    mean_theta[i,1] <- log(E_cases[i]) + beta0_cases + inprod(X[i,1:p], beta_cases[1:p])
15    mean_theta[i,2] <- log(E_deaths[i]) + beta0_deaths + inprod(X[i,1:p], beta_deaths[1:p])
16    theta[i,1:2] ~ dmnorm(mean_theta[i,1:2], cholesky=Chol_Sigma_theta[1:2,1:2])
17    theta_cases[i] <- theta[i,1] - mean_theta[i,1]
18    theta_deaths[i] <- theta[i,2] - mean_theta[i,2]
19  }
20
21  # Priors
22  beta0_cases ~ dnorm(0, sd=10)
23  beta0_deaths ~ dnorm(0, sd=10)
24
25  u[1:N,1] ~ dcar_normal(adj[1:L], weights[1:L], num[1:N], 1, zero_mean=1)
26  u[1:N,2] ~ dcar_normal(adj[1:L], weights[1:L], num[1:N], 1, zero_mean=1)
27
28  Sigma_s[1,1] <- sigma_s_cases^2
29  Sigma_s[1,2] <- rho_s*sigma_s_cases*sigma_s_deaths
30  Sigma_s[2,1] <- rho_s*sigma_s_cases*sigma_s_deaths
31  Sigma_s[2,2] <- sigma_s_deaths^2
32  Achol[1:2,1:2] <- t(chol(Sigma_s[1:2,1:2]))
33
34  sigma_s_cases ~ T(dt(0,1,1),0,)
35  sigma_s_deaths ~ T(dt(0,1,1),0,)
36  rho_s ~ dunif(-1,1)
37
38  Sigma_theta[1,1] <- sigma_theta_cases^2
39  Sigma_theta[1,2] <- rho_theta*sigma_theta_cases*sigma_theta_deaths
40  Sigma_theta[2,1] <- rho_theta*sigma_theta_cases*sigma_theta_deaths
41  Sigma_theta[2,2] <- sigma_s_deaths^2
42  Chol_Sigma_theta[1:2,1:2] <- chol(Sigma_theta[1:2,1:2])
43
44  sigma_theta_cases ~ T(dt(0,1,1),0,)
45  sigma_theta_deaths ~ T(dt(0,1,1),0,)
46  rho_theta ~ dunif(-1,1)
47
```

```
48    for(k in 1:p){
49      beta_cases[k] ~ dnorm(0, sd=10)
50      beta_deaths[k] ~ dnorm(0, sd=10)
51    }
52  })
```