Spatial and Spatio-temporal Epidemiology 44 (2023) 100559

Contents lists available at ScienceDirect

Spatial and Spatio-temporal Epidemiology

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

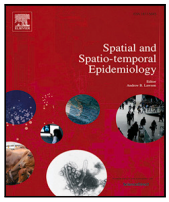

Original research

National lockdowns in England: The same restrictions for all, but do the impacts on COVID-19 mortality risks vary geographically?

Robin Muegge *, Nema Dean, Eilidh Jack, Duncan Lee

School of Mathematics and Statistics, University of Glasgow, United Kingdom

ARTICLE INFO

Keywords:
Bayesian inference
COVID-19 mortality
Lockdowns
Spatio-temporal modelling

ABSTRACT

Quantifying the impact of lockdowns on COVID-19 mortality risks is an important priority in the public health fight against the virus, but almost all of the existing research has only conducted macro country-wide assessments or limited multi-country comparisons. In contrast, the extent of within-country variation in the impacts of a nation-wide lockdown is yet to be thoroughly investigated, which is the gap in the knowledge base that this paper fills. Our study focuses on England, which was subject to 3 national lockdowns between March 2020 and March 2021. We model weekly COVID-19 mortality counts for the 312 Local Authority Districts in mainland England, and our aim is to understand the impact that lockdowns had at both a national and a regional level. Specifically, we aim to quantify how long after the implementation of a lockdown do mortality risks reduce at a national level, the extent to which these impacts vary regionally within a country, and which parts of England exhibit similar impacts. As the spatially aggregated weekly COVID-19 mortality counts are small in size we estimate the spatio-temporal trends in mortality risks with a Poisson log-linear smoothing model that borrows strength in the estimation between neighbouring data points. Inference is based in a Bayesian paradigm, using Markov chain Monte Carlo simulation. Our main findings are that mortality risks typically begin to reduce between 3 and 4 weeks after lockdown, and that there appears to be an urban–rural divide in lockdown impacts.

1. Introduction

The COVID-19 pandemic is the deadliest respiratory disease pandemic since the "Spanish" influenza in 1918, and it is one of at least four newly detected coronaviruses that have emerged since the year 2000 (Morens et al., 2020). The disease is believed to have originated in the Hunan seafood market (Shereen et al., 2020) in Wuhan, China. Without an effective early response strategy, the newly formed coronavirus grew from local chains of infection to a worldwide pandemic, as declared by the World Health Organisation (WHO) on 11th March 2020. As of 9th September 2022, there have been over 6.5 million recorded deaths that were linked to COVID-19, from over 607 million recorded infections worldwide (https://coronavirus.jhu.edu/map.html). The rapid and dramatic development of the pandemic resulted in widespread scientific research. Various statistical analyses modelled and predicted the spread of COVID-19 infections (Dong et al., 2020, Lee et al., 2022), identified the factors that were associated with a higher risk of displaying severe symptoms (Rashedi et al., 2020, Williamson et al., 2020, Wolff et al., 2021), or identified impacts on healthcare (Remuzzi and Remuzzi, 2020). The new insights were particularly vital in the early stages of the pandemic, as they provided governments with the scientific knowledge necessary for developing strategies to contain the virus.

In the first months of the pandemic, there was no effective medicine or vaccine to contain the virus and hence, governments were forced to implement non-pharmaceutical interventions. Mendez-Brito et al. (2021) provided a systematic review of empirical studies comparing the effectiveness of non-pharmaceutical interventions in reducing the number of confirmed COVID-19 cases. Their review, consisting of 34 studies, suggested that the most effective measures included the closing of schools, workplaces, businesses, and venues, as well as banning events. Whilst not ranked amongst the most effective measures, other effective interventions included lockdowns, travel restrictions (national and international), social gathering bans, social distancing, public information campaigns, and mask wearing. Examples of spatio-temporal analyses of non-pharmaceutical interventions include Aravindakshan et al. (2020), Zhang et al. (2021), and Ge et al. (2022).

Enforcing local or national lockdowns was amongst the most common responses of governments who had to react to the increasing

* Correspondence to: School of Mathematics and Statistics, University of Glasgow, Glasgow, G12 8SQ, United Kingdom.
E-mail addresses: Robin.Muegge@glasgow.ac.uk (R. Muegge), Nema.Dean@glasgow.ac.uk (N. Dean), Eilidh.Jack@glasgow.ac.uk (E. Jack), Duncan.Lee@glasgow.ac.uk (D. Lee).

https://doi.org/10.1016/j.sste.2022.100559
Received 22 April 2022; Received in revised form 22 September 2022; Accepted 1 December 2022
Available online 5 December 2022

infection rates and death tolls. While lockdowns might differ by country or region, a general definition of lockdown is "a temporary condition imposed by governmental authorities in which people are required to stay in their homes and refrain from or limit activities outside the home involving public contact" (Merriam-Webster Online Dictionary, 2022). The overriding goal of this study is to investigate spatio-temporal trends in COVID-19 mortality risks following the implementation of three national lockdowns in England to identify geographical differences in the impact of lockdown. While this study adds valuable insights to the observed spatio-temporal changes in mortality risks, it should be noted that the analysis of these trends cannot be used to draw conclusions about causality between lockdowns and mortality risks, as we do not know how the risks might have changed had the lockdowns not been implemented.

Several studies have previously analysed the impact of lockdown on the number of deaths or mortality risks over a particular time and area. Palladino et al. (2021) provided national results on the number of deaths in Italy, France, Spain, and the UK and showed a decreasing trend in the number of deaths upon the implementation of national lockdowns. However, they did not highlight similarities or differences in the temporal trends of COVID-19 death numbers between these countries. Conyon et al. (2020) did provide comparisons between countries, showing a connection between stricter lockdowns and lower numbers of deaths. Gerli et al. (2020) linked the timing of when lockdowns were introduced to the number of deaths, although the exact regulations varied amongst the countries they considered. Coccia (2021) showed an association between the duration of lockdown and fatality rates, but some apparent shortcomings are inherent in the analysis. For example, the fatality rates were computed as the number of deaths divided by the number of infected individuals, suggesting that higher testing capacities will lead to lower fatality rates. While some of these studies showed an impact of lockdown on the number of deaths, they all analysed national data without looking into spatial patterns within the countries under study. Some studies looked at regions within countries rather than nations as a whole, such as Silva et al. (2020) who modelled the numbers of deaths for four state capitals in Brazil, and Siqueira et al. (2020) who estimated mortality risks in autonomous communities in Spain. However, neither of these studies assumed underlying spatio-temporal autocorrelation structures to inform the trends in the number of deaths or mortality risks across the respective areas and time periods.

In England, our study region, there have been three national lockdowns between the start of 2020 and the end of 2021. We aim to explore the spatio-temporal trends in mortality risks after the introduction of these lockdowns, using weekly counts of deaths on a local authority district level in England. Since weekly death counts have been relatively low at the local authority level, they are subject to substantial random variability. Hence, we apply a spatio-temporal smoothing model to obtain more stable risk estimates by borrowing strength across neighbouring data points in space and time.

To our knowledge, our study is the first comprehensive investigation of the spatio-temporal trends in COVID-19 mortality risks following the implementation of national lockdowns in England, and it is the first to consider all three national lockdowns that have occurred thus far. The main questions that we will answer are (i) How long after the implementation of lockdown did mortality risks reduce at a national level, and did this vary by lockdown? (ii) How did the temporal trends in mortality risks differ by region in England in the weeks following the implementation of lockdown? (iii) Which local authorities shared similar temporal trends in mortality risks?

2. Materials and methods

2.1. Data

2.1.1. Study duration and region

The time frame of our study is from 1st February 2020 (the week of the first registered death due to COVID-19 in England) to 14th May 2021 (seven weeks after the third lockdown was lifted), and the study region is mainland England which is partitioned into 312 local authority districts (LADs). The average population size for a single LAD in our study is 179,945, and sizes range from 9721 to 1,141,816 people.

2.1.2. England's national lockdowns

During the observed time frame, there have been three national lockdowns in England, which ranged from 26th March to 12th May 2020 (48 days), 5th November to 2nd December 2020 (28 days), and 5th January to 28th March 2021 (83 days). It should be noted that the weeks spent in lockdown do not align with the weeks in our data, where each week is defined to range from Saturday to Friday. Therefore, each week in our data with at least four days of lockdown is considered as a week during lockdown, while any week with three or fewer days of lockdown is considered as a week outside of lockdown. In our study the lockdowns are therefore defined from 28th March to 15th May 2020 (49 days), 7th November to 4th December 2020 (28 days), and 2nd January to 26th March 2021 (84 days).

2.1.3. Data sources

The Office of National Statistics (ONS) has provided occurrences of deaths as weekly accumulations by local authority districts in England (ONS, 2021b) over the time frame specified in Section 2.1.1. From this data, we extracted the weekly number of deaths with COVID-19 mentioned on the death certificate, which we denote Y_{kt} for LAD k ($k = 1, \dots, 312$) and week t ($t = 1, \dots, 67$). It should be noted that the data include only the deaths of registered people in England, and a person's death is counted towards the LAD in which they were registered, disregarding the actual place where the person had died, should these locations differ.

We obtained age–sex specific population data for mid-2020 from ONS (2021c) but used 2019 population data for seven LADs due to a lack of data for 2020. Age–sex specific COVID-19 mortality rates were taken from https://coronavirus.data.gov.uk/details/deaths?areaType=nation&areaName=England on 30th August 2021.

2.1.4. Exploratory analysis using the Standardised Mortality Ratio (SMR)

Risks of COVID-19 mortalities differ by age and sex. For example, many more deaths occur amongst the older population than in younger age groups, and males tend to be at higher risk than females (Biswas et al., 2021). Hence, when comparing the risks of COVID-19 mortality for different areas, the underlying age and sex demographics should be considered, which we do via indirect standardisation by computing the expected weekly number of deaths e_k for LAD k, for which we include a mathematical definition in Section S1 of the supplementary materials. Note that the expected counts do not differ by week as population sizes and mortality rates are not available at that frequency.

The standardised mortality ratio (SMR) is an exploratory measure of disease risk that accounts for the underlying age and sex demographics, and is given by $SMR_{kt} = Y_{kt}/e_k$, for LAD $k = 1, 2, \dots, K(= 312)$ and week $t = 1, 2, \dots, N(= 67)$. For example, $SMR_{kt} = 1.2$ suggests that area k has a 20% elevated risk in week t compared to the national average risk over the study duration.

Fig. 1(a) presents the weekly SMR for all LADs with the weeks of lockdown highlighted in beige, where the black line shows the average SMR over all LADs. The plot shows two distinct waves in SMR values, with peaks in April 2020 and January 2021. The values appear to decline sometime after the introduction of lockdowns 1 and 3. For lockdown 2, the average SMR value did not decrease much, likely because the lockdown was lifted after only four weeks. Fig. 1(b) displays the spatial pattern in the average SMRs across England, measured for each LAD over all weeks in the time frame of our study. The map shows higher average SMRs for urban areas, such as London, Birmingham, Manchester, Liverpool, and Sheffield, while the average SMRs appear lower for more rural areas. The average SMRs appear to change relatively smoothly across the map, suggesting spatial autocorrelation is present in these data.

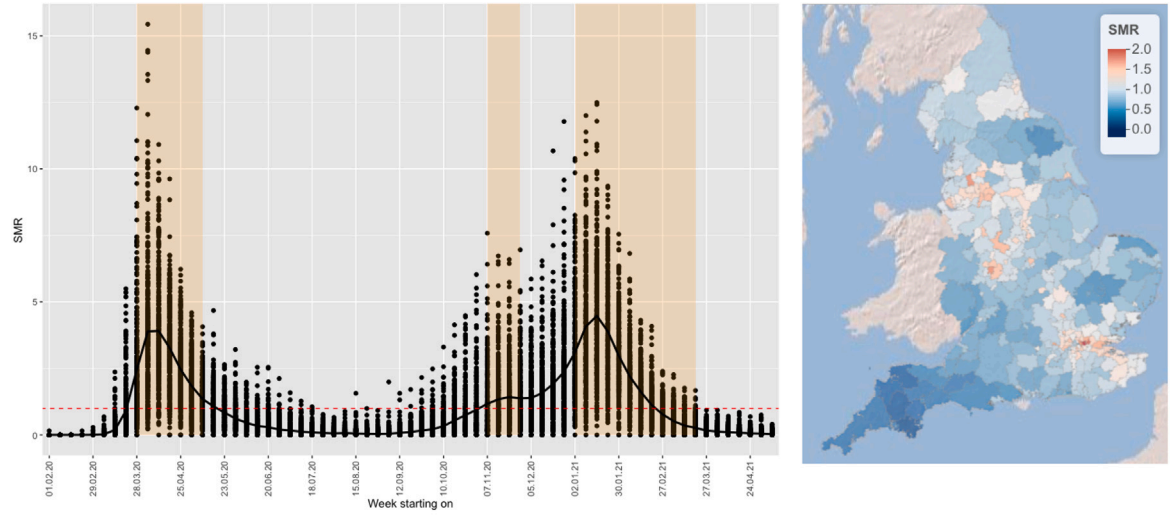

(a) (b)

Fig. 1. (a) SMR by week for all LADs in England; weeks of lockdown are highlighted in beige; the black line shows the average SMR per week; the dashed red line indicates an SMR of 1. (b) Average SMR over the entire study by LAD, for England.

To verify the assumption of spatial autocorrelation in the SMR values, we perform Moran's I tests (Moran, 1950) for the 65 weeks in our study that have seen at least one death. The Moran's I statistics take on values between −0.02142 and 0.66014, with a mean value of 0.2572. A permutation test based on 10,000 random permutations of the data evaluated at a significance threshold of $0.05/65 \approx 0.00077$ (Bonferroni correction, Haynes, 2013) results in rejecting the null hypothesis of no spatial autocorrelation for 48 of the 65 weeks (73.95%). Hence, spatial autocorrelation is present for the majority of weeks in our study. Additional explanations of the methods applied and how to implement them in R can be found in Lee (2020).

We check for temporal autocorrelation in the SMR values by computing temporal autocorrelation coefficients (Chatfield, 2003) for each LAD over the weeks in the study. We apply a Ljung–Box test (Ljung and Box, 1978) to evaluate temporal autocorrelation for lags up to 10 weeks simultaneously. At a significance threshold of $0.05/312 \approx 0.00016$ (Bonferroni correction), the temporal autocorrelation test results in rejecting the null-hypothesis of no temporal autocorrelation for 310 of the 312 LADs (99.36%). Hence, temporal autocorrelation is present in the data.

Having identified spatial and temporal autocorrelation in the SMR values of our data, we now fit a spatio-temporal model to obtain estimates of mortality risks that account for the clear trends and autocorrelations.

2.2. Methods

The observed number of COVID-19 deaths in area k and week t is denoted Y_{kt}, while the expected weekly number of deaths is denoted e_k and does not vary by week.

Since the observed number of COVID-19 deaths are counts of rare events, the natural choice is to fit a Poisson log-linear model to the data. Specifically, we fit the model in a Bayesian setting, where the data likelihood is of the form

$$Y_{kt} \sim \text{Poisson}(e_k \theta_{kt}), \tag{1}$$

$$\ln(\theta_{kt}) = \beta_0 + \phi_{kt}. \tag{2}$$

Here, our goal is to estimate the relative mortality risk θ_{kt} for area k at time t, and the estimated risk can be interpreted as a spatio-temporally smoothed version of the noisy SMR. The natural log of the relative risk $\{\theta_{kt}\}$ is modelled by an intercept term β_0 and a spatio-temporal trend that is modelled by random effects $\{\phi_{kt}\}$. When working with areal-type data, a common method to capture its spatial structure is to incorporate a neighbourhood matrix $\boldsymbol{W}$ in the model. Here, we define $\boldsymbol{W}$ to be a $K \times K$ binary adjacency matrix, with elements $w_{kj} = 1$ if LADs k and j share a border and $w_{kj} = 0$, otherwise.

The spatio-temporal random effects model we implement was developed by Rushworth et al. (2014). This particular model is a good choice for our data since it can capture the spatio-temporal autocorrelations we have identified in Section 2.1.4 without restricting the smoothing over space and time to follow a rigid parametric form or making assumptions about the shape of the temporal trends. It allows the mean of the random effects from one week to depend on the effects from previous weeks, and the (co)variance of the random effects' multivariate Normal distribution depends on the spatial structure of the data captured by the neighbourhood matrix $\boldsymbol{W}$. Note, Rushworth et al. (2014) only considered an autoregressive process of order 1 (AR(1)) but here, we present both AR(1) and AR(2) models. Thus, depending on the version of the model the joint prior distribution of the random effect vectors in $\boldsymbol{\phi} = (\boldsymbol{\phi}_1, \ldots, \boldsymbol{\phi}_N)$ is decomposed as either

$$\begin{aligned} \textbf{AR(1)}: \quad f(\boldsymbol{\phi}_1, \ldots, \boldsymbol{\phi}_N) &= f(\boldsymbol{\phi}_1) \prod_{t=2}^{N} f(\boldsymbol{\phi}_t | \boldsymbol{\phi}_{t-1}) \\ &= N\left(\boldsymbol{0}, \tau^2 \boldsymbol{Q}(\rho, \boldsymbol{W})^{-1}\right) \\ &\quad \times \prod_{t=2}^{N} N\left(\alpha \boldsymbol{\phi}_{t-1}, \tau^2 \boldsymbol{Q}(\rho, \boldsymbol{W})^{-1}\right), \end{aligned} \tag{3}$$

$$\begin{aligned} \textbf{AR(2)}: \quad f(\boldsymbol{\phi}_1, \ldots, \boldsymbol{\phi}_N) &= f(\boldsymbol{\phi}_1) f(\boldsymbol{\phi}_2) \prod_{t=3}^{N} f(\boldsymbol{\phi}_t | \boldsymbol{\phi}_{t-1}, \boldsymbol{\phi}_{t-2}) \\ &= N\left(\boldsymbol{0}, \tau^2 \boldsymbol{Q}(\rho, \boldsymbol{W})^{-1}\right) N\left(\boldsymbol{0}, \tau^2 \boldsymbol{Q}(\rho, \boldsymbol{W})^{-1}\right) \\ &\quad \times \prod_{t=3}^{N} N\left(\alpha_1 \boldsymbol{\phi}_{t-1} + \alpha_2 \boldsymbol{\phi}_{t-2}, \tau^2 \boldsymbol{Q}(\rho, \boldsymbol{W})^{-1}\right), \end{aligned} \tag{4}$$

where $\boldsymbol{\phi}_t = (\phi_{1t}, \ldots, \phi_{Kt})$ denotes the vector of random effects for areas $1, \ldots, K$ in week t, and the precision matrix is defined as $\boldsymbol{Q}(\rho, \boldsymbol{W}) = \rho(\text{diag}(\boldsymbol{W}\boldsymbol{1}) - \boldsymbol{W}) + (1-\rho)\boldsymbol{I}$ (Leroux et al., 2000), where $\boldsymbol{W}$ is the binary adjacency matrix from above, $\boldsymbol{I}$ is a $K \times K$ identity matrix, and $\boldsymbol{1}$ is a $K \times 1$ vector of ones. The random effects in $\boldsymbol{\phi}_1$ (AR(1)) or in $\boldsymbol{\phi}_1$ and $\boldsymbol{\phi}_2$ (AR(2)) are assigned solely spatial Leroux CAR prior distributions. The full conditional prior distribution of the random effect ϕ_{kt} for area k and time period t, given the random effects of all other areas in that

time period can then be expressed as

$$\phi_{kt}|\boldsymbol{\phi}_{-kt} \sim N\left(\frac{\rho\sum_{j=1}^{K} w_{jk}\phi_{jt}}{\rho\sum_{j=1}^{K} w_{jk}+1-\rho}, \frac{\tau^2}{\rho\sum_{j=1}^{K} w_{jk}+1-\rho}\right), \quad (5)$$

where $\boldsymbol{\phi}_{-kt}$ is the vector of random effects for all areas except for area k in time period t, the parameter ρ measures the autocorrelation imposed by the spatial structure captured in $\boldsymbol{W}$, and τ^2 is a variance parameter. The mean from the prior distribution of $\phi_{kt}|\boldsymbol{\phi}_{-kt}$ is a weighted average of random effects in time period t from areas adjacent (neighbours) to area k, where the weight is determined by the dependence parameter ρ, and the variance term is smaller when there are more areas adjacent to area k.

For the later time periods ($t = 2, \ldots, N$ in the AR(1) model, or $t = 3, \ldots, N$ in the AR(2) model), the mean of the prior joint distribution of $\boldsymbol{\phi}_t$ depends on the effects from the preceding time period (Eq. (3)) or the last two preceding time periods (Eq. (4)), thus temporally smoothing the risks. Spatial smoothness is induced by the covariance matrix, which is specified by the Leroux CAR prior. In the AR(1) model, α is a temporal dependence parameter that takes on a value in the interval $[-1, 1]$. Thus, $\alpha = 0$ indicates temporal independence, and $\alpha = 1$ indicates strong temporal autocorrelation and makes the distribution a first order random walk. For the AR(2) model, α_1 and α_2 are again temporal dependence parameters that determine the relationship between spatio-temporal random effects that are temporal neighbours of order 1 and 2, respectively. Note that $\alpha_1 = 2$ and $\alpha_2 = -1$ define a second-order random walk.

The mean of the log-transformed mortality risk β_0 from Eq. (2) is assigned a Normal prior distribution with mean zero and large variance, i.e. $\beta_0 \sim N(\mu_0 = 0, \sigma_0^2 = 10,000)$. This weakly informative prior distribution reflects that we have no prior knowledge for the underlying mortality risk of COVID-19 in England in the given time frame. In both the AR(1) and AR(2) models, the spatial autocorrelation parameter ρ is assigned a flat Uniform$(0, 1)$ prior, and the variance parameter τ^2 is assigned an Inverse-Gamma$(a = 1, b = 0.01)$ prior distribution. The temporal autocorrelation parameter α in the AR(1) model is assigned a flat Uniform$(0, 1)$ prior while the temporal dependence parameters α_1 and α_2 in the AR(2) model are assigned a flat improper joint prior distribution, i.e. $f(\alpha_1, \alpha_2) \propto 1$. These are the default prior and hyper-prior distributions specified in the function `ST.CARar` from the package `CARBayesST` (Lee et al., 2018), which we used to fit the Bayesian hierarchical model in R, via MCMC simulation.

3. Results

3.1. Model fitting

We fitted both the AR(1) and AR(2) versions of the model outlined above to the COVID-19 mortality data and obtained estimated risks from the posterior mean of the fitted values divided by the expected counts. The MCMC algorithm produced 2,200,000 samples for each parameter in these models, and we discarded the first 200,000 simulations as the burn-in period. We thinned the simulations by saving only every $1,000$th simulation to reduce the autocorrelation in the Markov chains, which resulted in 2,000 simulated values for each parameter.

Geweke diagnostics (Geweke, 1992) between (−2,2) and an examination of the corresponding trace plots (those for the AR(2) model are included in Section S2 of the supplementary materials) suggested no evidence of a lack of convergence in the algorithm for either the AR(1) or AR(2) model. For example, for the AR(2) model the Geweke diagnostics took on the values 1.4 for β_0, −0.8 for τ^2, −0.5 for ρ, 0.0 for α_1, and −0.2 for α_2. Since both models seem to have converged, we use the Deviance Information Criterion (DIC) (Spiegelhalter et al., 2002) to compare the two models. The AR(1) model has a DIC value of 69,067, while the AR(2) model has a DIC value of 68,772. Hence, the model with second-order temporal autocorrelation fitted the data slightly better, so we use its estimated relative risks in the following analysis. Note that we have confirmed that the AR(2) model fits the data appropriately, via posterior predictive checks, which can be found in Section S3 of the supplementary materials. Additionally, we assessed the sensitivity of our results to prior choice, which is displayed in Section S4 of the supplementary materials.

3.2. How long after the implementation of lockdown did mortality risks reduce?

The first question in this study is the most important from an epidemiological perspective because it quantifies how long lockdowns have to be in place before mortality risks reduce. We answer the question by analysing the panels in Fig. 2, which show boxplots of the distributions of mortality risks across all LADs in the weeks preceding, during and after each of the three lockdowns (left), and plots with 95% credible intervals around the median average estimated risk across England (right). Panels (a), (b), and (c) show these results for the first, second, and third lockdown, respectively, and all the risks presented are relative to the average risk across England for the entire study period. The weeks after each lockdown was lifted are included to see if the lockdowns had a lasting impact on mortality risk. To analyse risk over similar periods of time, we set the panels for the first and third lockdowns to contain 13 weeks from the start of lockdown (weeks $0, \ldots, 12$). However, the panel for the second lockdown ranges over fewer weeks, as week 7 is already the week before the third lockdown was introduced. The plots on the right were included to visualise the variability in the accumulated risks by week across the entire country. Since the 95% credible intervals are very narrow, there does not appear to be much variation in the averages of the simulated estimated risks across all LADs.

The first striking feature apparent from Fig. 2 is that the second lockdown (panel (b)) exhibits a very different trend in risk over time compared with the first and third lockdowns. The second lockdown shows a constant or slightly increasing trend in risk throughout, while the overall risks in the first and third lockdowns show an increasing trend in the first three weeks of lockdown, followed by a decreasing trend thereafter. The second lockdown of only four weeks was much shorter than the first lockdown of seven weeks and the third lockdown of twelve weeks, and the short duration of lockdown 2 may have prevented it from having a sizeable impact on reducing risk. Additionally, the lack of an increasing trend in the first few weeks after the introduction of lockdown 2 may suggest that the pandemic was not yet approaching a severe new wave in terms of mortality. A possible reason for the lack of an increasing trend could be that England had implemented a three-tier system of mobility restrictions that started on 14th October before lockdown 2, with varying degrees of restrictions for the LADs in the three tiers, which (Davies et al., 2021) suggested has had a sizeable effect in reducing the number of COVID-19 deaths. Altogether, since lockdown 2 is markedly different from the other lockdowns in that it did not show any reductions in risks, the remainder of this section will focus on lockdowns 1 and 3.

The second key feature from Fig. 2 is that for lockdowns 1 and 3, the risk of mortality increased for the first three weeks after the introduction of the lockdown, before it started to reduce from the fourth week onward. The reason for this delayed reduction in mortality risk is the lag between a COVID-19 infection and mortality, which the ONS estimate is between 21 and 25 days (3–4 weeks) on average (ONS, 2021a). Thus the high numbers of infections in the last few weeks before the introduction of lockdown would transfer to the high mortality risks observed three to four weeks into lockdown.

The final important finding from Fig. 2 is the times it took for the risks to reduce to baseline levels. For a baseline level of a risk of 1, it took nine weeks after the introduction of both lockdowns 1 and 3 for the median risk across England to reduce to that level. As an alternative comparison, the figure also reveals that it took ten (lockdown 1) and

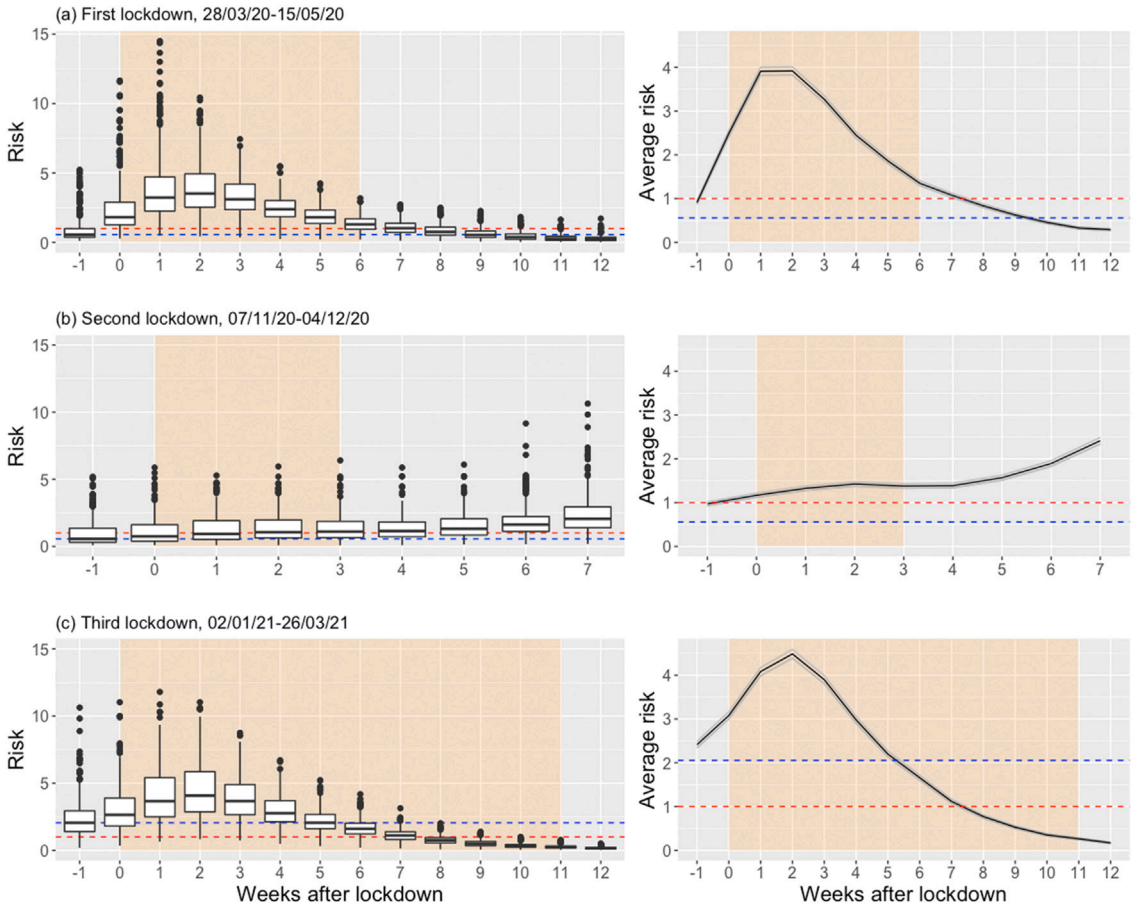

Fig. 2. Left: Boxplots of the posterior median estimated risks across LADs, by week after lockdown. Right: Median line with 95% credible intervals for the average estimated risk across LADs, by week after lockdown. In each panel, week −1 is the week preceding the lockdown, week 0 is the onset week of lockdown, week 1 is the first week after the onset of lockdown, etc. The weeks coloured in beige comprise the lockdowns. The y-axes measuring estimated risk and average estimated risk are on the same scale for the three panels to allow comparison across the lockdowns. The blue dashed line shows the median estimated risk across England from the last week before lockdown, the red dashed line represents a risk of 1. Note that week 7 in panel (b) is the same as week −1 in panel (c).

six (lockdown 3) weeks after the introduction of the lockdowns for the median risks across England to reduce back to their respective pre-lockdown levels. However, the median risks across England were 0.56 and 2.05 in the week preceding lockdowns 1 and 3, respectively, so it might have taken longer for the median risk in lockdown 1 to reduce to its pre-lockdown level simply because this initial level was much lower than that preceding lockdown 3.

In conclusion for question (i), lockdowns appear to reduce mortality risk after approximately four to five weeks. However, based on the observations for lockdown 2, the timing and duration of lockdown are likely to be decisive factors for lockdowns to impact mortality risks.

3.3. How did the temporal trends in mortality risks differ by region in England?

Differences in the temporal trends of mortality risks can be explored at a larger geographical scale by analysing regional data. The 312 LADs from the English mainland are nested exactly within the following nine regions: East of England, East Midlands, London, North East, North West, South East, South West, West Midlands, Yorkshire and The Humber. Note that a map of the regions of England with outlines of the LADs can be found in Section S5 of the supplementary materials. We compute weighted averages of risks for region r and week t as $\hat{\theta}_{rt} = \frac{1}{P_r} \sum_{k \in r} P_k \times \hat{\theta}_{kt}$, where $\hat{\theta}_{kt}$ denotes the estimated risk from area k and week t, and $k \in r$ indicates that LAD k falls into region r, while P_k and P_r denote the population sizes of LAD k and region r, respectively. The reason for taking the weighted averages with regard to population size is that the risk in a more populous LAD will have a more substantial impact on the region's risk than the risk in a less populous LAD. We have created plots showing 95% credible intervals on the population-based weighted averages of estimated risks by region, which we present in Section S6 of the supplementary material. These credible intervals are again quite narrow, suggesting that there is little variability in the regional average estimated risks.

As an alternative measure, we divide the weekly risk from each LAD by its risk from the week before the respective lockdown was implemented to obtain scaled risks. Thus, each area starts at a scaled risk of 1 in the week preceding the introduction of lockdown and the scaled risks in the succeeding weeks are relative to the initial risk levels, allowing for a more fair comparison of temporal trends in mortality risks in different areas. The scaled estimated risk is computed as $\hat{s}_{kl} = \hat{\theta}_{kl} / \hat{\theta}_{k(-1)}$, for $l = -1, 0, 1, 2, \ldots$, where $\hat{\theta}_{kl}$ denotes the estimated risk from LAD k and week l after the introduction of lockdown. So $l = -1$ denotes the week before the introduction of lockdown, $l = 0$ denotes the week during which the lockdown was introduced, $l = 1$ denotes the first week after the introduction of lockdown, and so forth. Thus, the scaled risk should be interpreted as the risk from a specific week during or after lockdown relative to the risk in the week preceding lockdown. For example, if the scaled risk of a particular week is 2.5, then this suggests that the risk was 2.5 times as large in that week relative to the week preceding lockdown. From the scaled estimated risks, we compute weighted averages of scaled risk for region r and week l of lockdown as $\hat{s}_{rl} = \frac{1}{P_r} \sum_{k \in r} P_k \times \hat{s}_{kl}$, with notation analogous to that of the weighted average risks above.

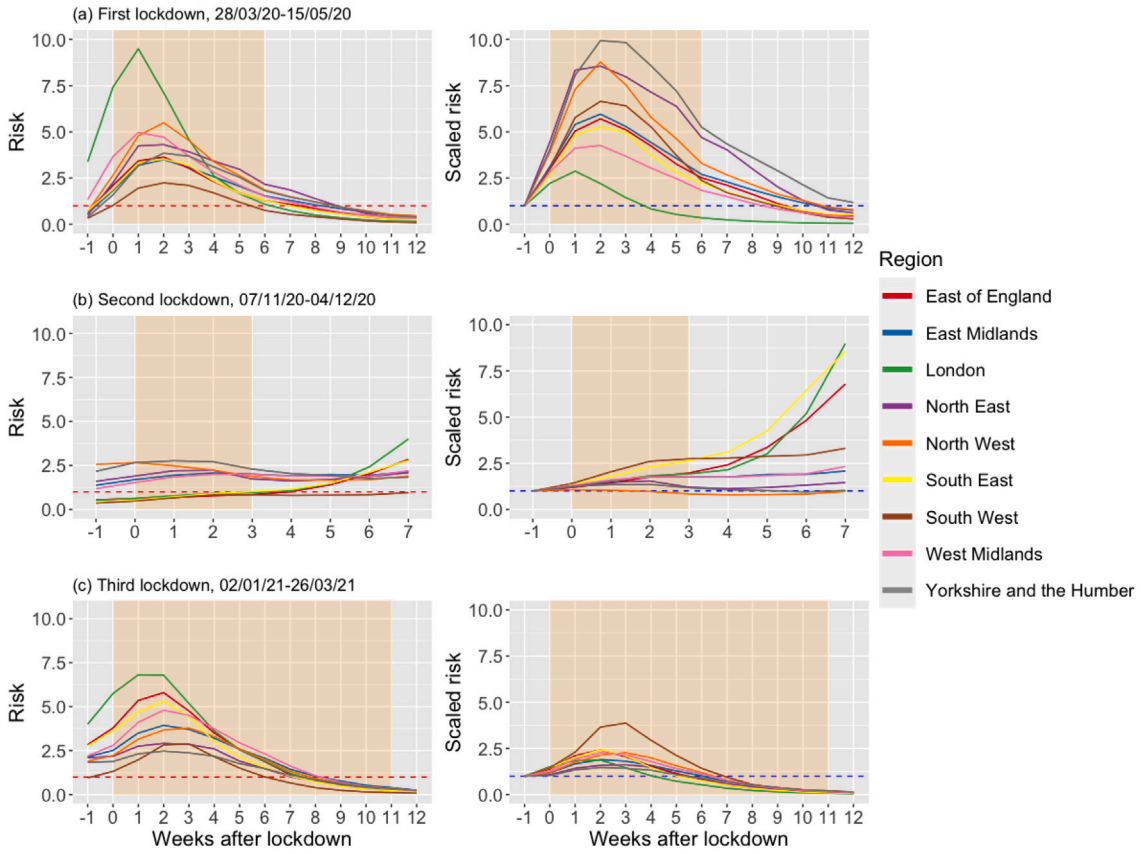

Fig. 3. Average estimated risk (left) and average scaled estimated risk (right) in weeks following the implementation of national lockdown, by region. In each panel, week −1 is the week preceding the lockdown, week 0 is the onset week of lockdown, week 1 is the first week after the onset of lockdown, etc. The weeks coloured in beige comprise the lockdowns. The red dashed line represents a risk of 1, and the blue dashed line represents a scaled risk of 1. Note that week 7 in panel (b) is the same as week −1 in panel (c).

Fig. 3 shows line plots for the population-based weighted averages of estimated risk and scaled estimated risk by week and region. For lockdown 2 the average mortality risks stayed relatively constant for all regions when the lockdown was in place. However, the risks increased rapidly in London, South East, and East of England after the lockdown was lifted. Note that the Alpha variant of COVID-19 was detected in the Kent area (South East) in September 2020, and this variant was estimated to be 1.5 times more transmissible, while it featured mortality risks that were 1.6 times higher than that of the earlier variants of COVID-19 (Page and McNamara, 2021). Hence, it is likely that the Alpha variant drove the distinct temporal trends in mortality risks in southeast England after lockdown 2.

The most striking feature of lockdowns 1 and 3 is that the average risk was highest in London during their first three weeks of lockdown. The average risk in London reduced quickly after the implementation of each lockdown. In the average risk plot (left) for lockdown 1, this is particularly noticeable as London had an exceptionally high peak in the second week of lockdown but was the second region to reach an average risk of 1 in the seventh week of lockdown. According to Batty et al. (2021), the proportion of essential workers in the UK is approximately 23.6%, while that of London is only 16%. Hence, the smaller proportion of essential workers might explain why mortality risks reduced so quickly in London after the first weeks of lockdowns 1 and 3, while they reduced more slowly for some of the other regions.

The scaled risk plots provide some additional insights into the temporal trends of risks. In lockdown 1, they reveal that the average risk in London did not increase as drastically when put in relation to the high level where it started, and that the average risk in London was quickest to return to its initial level within five weeks. Relative to their lower initial risk levels, the regions Yorkshire and The Humber, North East, and North West saw the most drastic increases in average risks. It also took the longest time for the average risks to return to their initial levels in these regions. In lockdown 3 the only region that stands out on the scaled risk plot is the South West, but this can be explained by its particularly low initial risk level. The scaled risk plot emphasises that the temporal trends in risks were very similar during the third lockdown for the remaining eight regions.

In conclusion to question (ii), our analysis suggests regional differences in the temporal trends in mortality risks for lockdowns 1 and 2 and no substantial differences for lockdown 3. In lockdown 1, it took longer for risks to reduce in the northern regions of England (North East, North West, Yorkshire and The Humber) compared to the rest of the country. For lockdown 2, we observed sharp increases in average risks in southeast England (London, East of England, South East) after the lockdown was lifted, while the average risks remained relatively constant for the other regions. Since none of the regions showed a substantial reduction in mortality risks in the weeks during or after lockdown 2, the remainder of this study will consider only lockdowns 1 and 3.

3.4. Which local authorities shared similar temporal trends in mortality risks?

Question (iii) requires us to check if there are groups of LADs with similar temporal risk trends in the weeks after lockdowns 1 and 3 (note that a sequence of maps showing the estimated risks by LAD in those weeks can be found in Section S7 of the supplementary materials). Clustering methods can be used to identify these groups. Here, we apply k-means clustering (Hartigan and Wong, 1979) for lockdowns 1 and 3, assigning the LADs to between one and ten clusters according to similarities in their estimated risks. We consider the same number of weeks after the introduction of each lockdown to allow a comparison

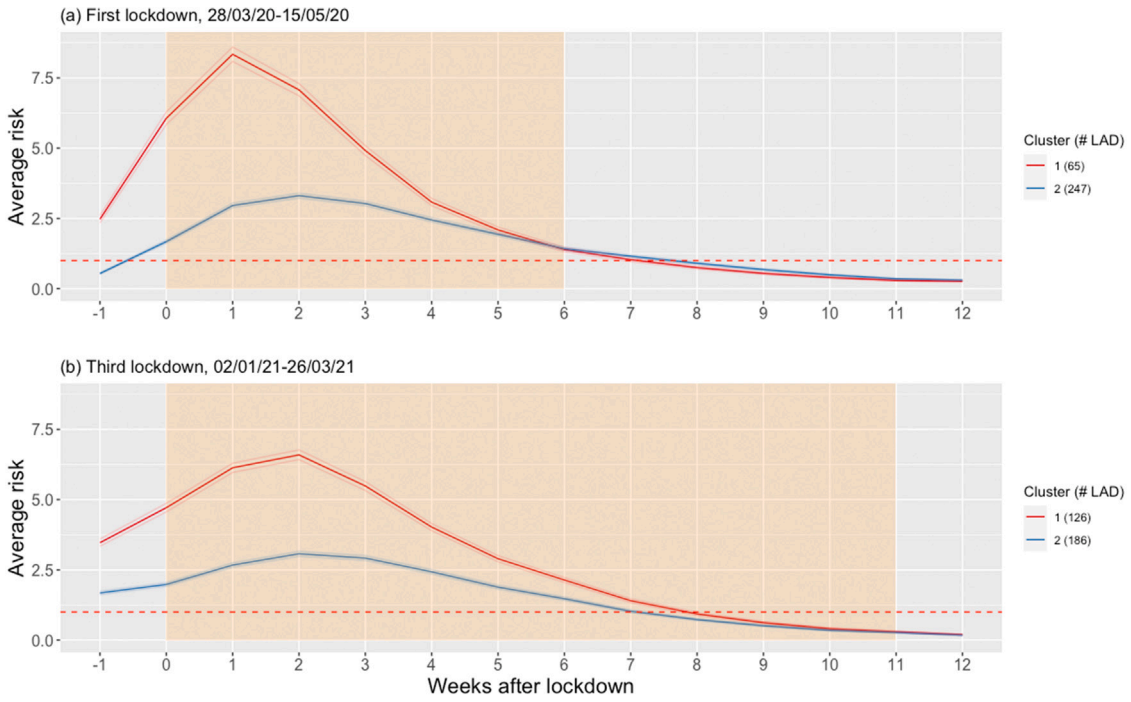

Fig. 4. Median average estimated risks with 95% credible intervals in the weeks following the implementation of national lockdown, by cluster. The red dashed line indicates an estimated risk of 1.

of their temporal trends. The weeks we consider range from the week before lockdown to the twelfth week after the introduction of lockdown to include all weeks from the longer lockdown 3. Note that it is common to standardise the observations from each variable before applying k-means clustering. However, since we are interested in distinguishing LADs both by peak risk and time period of when the peak risk occurred, we do not standardise the estimated risks for each week so that the k-means algorithm can pick up these differences.

We performed a sensitivity check of the k-means method for our data by running the algorithm on reduced datasets that excluded potential outliers identified with agglomerative hierarchical clustering using the single linkage method (Gower and Ross, 1969). The resulting clusterings were very similar for the full and reduced datasets, suggesting that the clustering we obtained for the whole dataset is not affected by extreme values. Thus, the following analysis uses the results for the whole dataset.

We analysed the within-cluster sum of squares and the average silhouette width (Rousseeuw, 1987) for the clusters obtained from the k-means algorithm, for different numbers of clusters k (the corresponding plots can be found in Section S8 of the supplementary materials). The within-cluster sum of squares plots show a substantial drop when moving from one cluster to two clusters, which is consistent for both lockdowns. We choose the optimal number of clusters by checking the average silhouette plots, which show that the average silhouette width is the largest for 2 clusters in lockdowns 1 and 3.

For the week before each lockdown, the week when the lockdown was introduced, and the 12 weeks thereafter, we compute weighted averages of estimated risk by cluster c and week t as $\hat{\theta}_{ct} = \frac{1}{P_c} \sum_{k \in c} P_k \times \hat{\theta}_{kt}$, where P_c and P_k are the population sizes for cluster c and LAD k respectively, and $k \in c$ denotes that LAD k is in cluster c. We accumulate the risks by cluster since it is not practical to look at the risks of all 312 LADs simultaneously. We take population-size based weighted averages since the risks from the LADs with larger populations contribute stronger to the risk of the entire population in that cluster than the risks from LADs with smaller populations.

Fig. 4 presents plots that show the median average estimated risks with 95% credible intervals by cluster and week, over the 2,000 simulations obtained from the fitted model. The 95% credible intervals are barely visible, suggesting that there is very little variation in the simulated average risks by cluster. In the first three to four weeks of each lockdown, the LADs in cluster 1 were on average at a substantially higher risk than those in cluster 2. The peak average estimated risk for cluster 1 of lockdown 3 was not as high (approximately 6.59) as for cluster 1 of lockdown 1 (approximately 8.34). The peak average estimated risk for cluster 2 was similar for lockdowns 1 (approximately 3.31) and 3 (approximately 3.07). Lastly, the average estimated risks from the two clusters were very similar towards the last weeks of each lockdown.

Further investigation of the clustering assignments shows that 45 of the 65 LADs in cluster 1 of the first lockdown were also in cluster 1 of the third lockdown, while 81 LADs switched from cluster 2 of lockdown 1 to cluster 1 of lockdown 3. The much larger number of LADs in the higher risk cluster for lockdown 3 is likely to have caused a relatively low Rand index (Rand, 1971) of 0.56, which suggests that the two cluster assignments are not very similar.

Our ultimate goal for this part of the study is to see any geographical patterns for LADs with similar temporal risk trends. Therefore, Fig. 5 shows maps for lockdowns 1 and 3 that display the cluster memberships of the LADs in the two lockdowns.

In lockdown 1, most of the separation into the two clusters appears to be explained by an urban/rural divide. The peak estimated risk was higher in LADs close to London (9.3 million people), Manchester (2.73 million people), Birmingham (2.6 million people), Liverpool (902,000 people), and Sunderland (341,000 people), while lower in most rural areas. In lockdown 3 some urban/rural divide might still be possible, but the most prominent feature is that most LADs at higher risk levels are in the southeast of England in London (82% of LADs in cluster 1), South East (62% of LADs in cluster 1), and East of England (71% of LADs in cluster 1) which might be explained by the more accelerated spreading of the Alpha variant of COVID-19 which originated in Kent in the southeast. Additionally, Grint et al. (2021) suggested that many of the early Alpha variant cases of COVID-19 were observed in the North West, which could explain the higher risks in the adjacent LADs of Eden, Allerdale, and Carlisle.

Note that COVID-19 mortality risks had previously been linked to deprivation, such that more deprived areas tend to be exposed

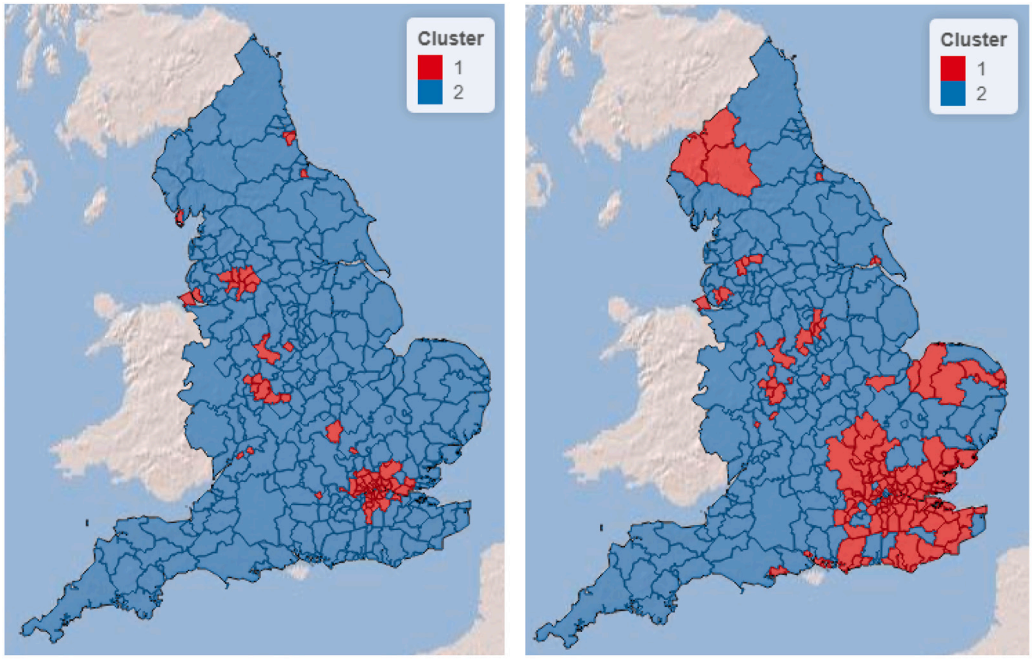

(a) First lockdown (b) Third lockdown

Fig. 5. Maps showing clusters that were formed according to weekly estimated risk, by lockdown. The LADs in cluster 1 had a higher peak estimated risk, on average.

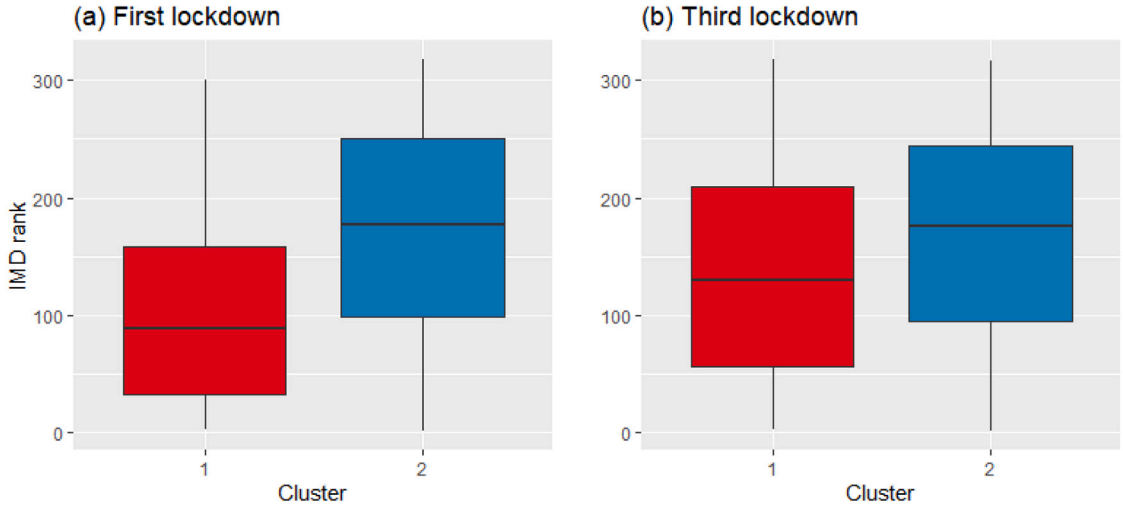

Fig. 6. Boxplots for the IMD rankings of LADs in the two clusters of lockdowns 1 and 3, where a lower IMD score implies higher deprivation.

to greater mortality risks (Tinson, 2021). The government website (Ministry of Housing, 2019a) provides a 2019 ranking of lower layer super output areas (LSOAs) with regard to their index of multiple deprivation (IMD), which combines deprivation scores from several domains (Ministry of Housing, 2019c provides an infographic for the IMD). They also provide population-based weighted averages of LSOA ranks for each LAD (Ministry of Housing, 2019b), from which we created a deprivation ranking for the LADs in our data, from 1 (most deprived) to 312 (least deprived).

Fig. 6 shows the distribution of IMD rankings in each cluster for lockdowns 1 and 3. Cluster 1 of lockdown 1 has a median IMD rank of 89, while that of cluster 2 is 176. In comparison to the median ranks from lockdown 1, the median IMD rank in cluster 1 of lockdown 3 is substantially higher at 130, while that of cluster 2 is the same at 176. Hence, there might be an association between higher deprivation and increased risks during lockdown 1, while in lockdown 3 the difference is less pronounced.

4. Discussion

Our motivation for this study was to explore the temporal trends in COVID-19 mortality risks after the introduction of national lockdowns in England, with the main quantity of interest being how long it took for them to be effective at reducing mortality risks below the levels from before the lockdowns were introduced.

At a national level (Section 3.2), it took around three weeks after the introduction of lockdowns 1 and 3 for mortality risks to stop increasing, likely due to a 21 to 25 day lag time between disease onset and mortality (ONS, 2021a). The risks reduced to pre-lockdown levels after ten (lockdown 1) and six weeks (lockdown 3) of lockdown, respectively. Lockdown 2 did not lead to a meaningful reduction in mortality risk. However, in the first weeks of lockdown 2 the risks did not increase like they did in lockdowns 1 and 3. Hence, it is not possible to quantify the efficacy of lockdown 2 based on the observed changes

in mortality risks. Additionally, it should be noted that lockdown 2 was lifted after only four weeks. Al-Zoughool et al. (2022) have studied the association between lockdown duration and timing with hospital cases and infection rates rather than mortality risks. Nevertheless, their suggestion that a lockdown of fewer than 30 days might not lead to substantial reductions in hospital cases and infection rates was not conflicted by our results.

The cluster analysis of temporal trends at a local authority level (Section 3.4) revealed an urban/rural divide in risks in lockdown 1, with urban areas tending to have a higher peak risk. Hamidi et al. (2020) partially corroborate our findings by showing higher risks for larger metropolitan areas in the USA but lower risks for more densely populated counties. Hence, the urban/rural divide might be driven by multiple factors such as population density, accessibility to medical aid, hospital occupancy rates, or mobility within and between districts. Further, our analysis of lockdown 1 suggests a possible association between higher deprivation and increased mortality risks, which has previously been shown in several studies, including Williamson et al. (2020) or Sartorius et al. (2021).

For lockdown 3, the cluster analysis indicated a higher value at the peak of the risk in LADs from southeast England, which might be explained by the Alpha variant detected in the Kent area (South East) in September 2020. Grint et al. (2021) adds weight to this hypothesis by finding both an increased early spreading of the Alpha variant in southeast England and higher mortality risks associated with the Alpha variant when compared to earlier versions of COVID-19.

Lastly, although mortality risks were at similar levels in weeks 7 through 12 after the implementation of lockdowns 1 and 3, the first lockdown was lifted in its seventh week, while the third lockdown was only lifted in its twelfth week. The third lockdown continued because the number of patients in hospitals was high (Prime Minister's Office, 2021). Hence, the relationship between mortality risks and the number of patients in hospitals might have changed from lockdown 1 to lockdown 3, and this assumption is supported by Gray et al. (2021).

When communicating these results, it is important to state that this is an exploratory study and thus it could not assess an 'effect' of lockdown on mortality risk, as the counterfactual event of what would have happened without lockdown in place could not be observed. Hence, the study comments on trends alone. Although only lockdowns 1 and 3 showed a reduction in mortality risks, the temporal trends at a national level were quite similar for these lockdowns, which suggests that lockdowns might have a real effect on reducing mortality risks.

Section 2.1.3 mentioned that in the ONS dataset, a person's death was counted towards the LAD where they were registered, rather than the actual place where they died, should these locations differ. A limitation of this study is that we do not know to which extent the results might be affected by such mismatches. For example, especially for the most vulnerable, such as older or severely ill people, it might be possible that they are treated in a specialised hospital in a LAD other than the one where they were registered, for some disease unrelated to COVID-19, before getting infected and dying of COVID-19. Analysing the impact of such mismatches could form an extension to this study if data on the place of death rather than registration were available. Another exciting extension to this study could be to perform a multivariate analysis that also considers the temporal trends in the number of hospitalisations when looking at mortality risks in the weeks following lockdown. Explanatory variables such as vaccination status or the particular variant of COVID-19 that led to death could be included in the analysis, depending on data availability.

Funding

The first author was funded by a University of Glasgow Maclaurin PhD Scholarship.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.sste.2022.100559.

References

Al-Zoughool, M., Oraby, T., Vainio, H., Gasana, J., Longenecker, J., Ali, W.A., AlSeaidan, M., Elsaadany, S., Tyshenko, M.G., 2022. Using a stochastic continuous-time Markov chain model to examine alternative timing and duration of the COVID-19 lockdown in Kuwait: What can be done now? Arch. Public Health 80, 22. http://dx.doi.org/10.1186/s13690-021-00778-y.

Aravindakshan, A., Boehnke, J., Gholami, E., Nayak, A., 2020. Preparing for a future COVID-19 wave: Insights and limitations from a data-driven evaluation of non-pharmaceutical interventions in Germany. Sci. Rep. 10, 20084. http://dx.doi.org/10.1038/s41598-020-76244-6.

Batty, M., Murcio, R., Iacopini, I., Vanhoof, M., Milton, R., 2021. London in lockdown: Mobility in the pandemic city. In: COVID-19 Pandemic, Geospatial Information, and Community Resilience. CRC Press, pp. 229–244.

Biswas, M., Rahaman, S., Biswas, T.K., Haque, Z., Ibrahim, B., 2021. Association of sex, age, and comorbidities with mortality in COVID-19 patients: A systematic review and meta-analysis. Intervirology 64, 36–47. http://dx.doi.org/10.1159/000512592.

Chatfield, C., 2003. The Analysis of Time Series: An Introduction, sixth ed. In: Chapman & Hall/CRC Texts in Statistical Science, Routledge, 2000 Corporate Blvd N.W. Boca Raton, FL 33431, USA, ISBN 978-1-5848-8317-3.

Coccia, M., 2021. The relation between length of lockdown, numbers of infected people and deaths of COVID-19, and economic growth of countries: Lessons learned to cope with future pandemics similar to COVID-19 and to constrain the deterioration of economic system. Sci. Total Environ. 775, 145801. http://dx.doi.org/10.1016/j.scitotenv.2021.145801.

Conyon, M.J., He, L., Thomsen, S., 2020. Lockdowns and COVID-19 deaths in Scandinavia. SSRN Electron. J. http://dx.doi.org/10.2139/ssrn.3616969.

Davies, N.G., Barnard, R.C., Jarvis, C.I., Russell, T.W., Semple, M.G., Jit, M., Edmunds, W.J., 2021. Association of tiered restrictions and a second lockdown with COVID-19 deaths and hospital admissions in England: A modelling study. Lancet Infect. Dis. 21, 482–492. http://dx.doi.org/10.1016/S1473-3099(20)30984-1.

Dong, E., Du, H., Gardner, L., 2020. An interactive web-based dashboard to track COVID-19 in real time. Lancet Infect. Dis. 20, 533–534. http://dx.doi.org/10.1016/S1473-3099(20)30120-1.

Ge, Y., Zhang, W.-B., Liu, H., Ruktanonchai, C.W., Hu, M., Wu, X., Song, Y., Ruktanonchai, N.W., Yan, W., Cleary, E., Feng, L., Li, Z., Yang, W., Liu, M., Tatem, A.J., Wang, J.-F., Lai, S., 2022. Impacts of worldwide individual non-pharmaceutical interventions on COVID-19 transmission across waves and space. Int. J. Appl. Earth Obs. Geoinf. 106, 102649. http://dx.doi.org/10.1016/j.jag.2021.102649.

Gerli, A.G., Centanni, S., Miozzo, M.R., Virchow, J.C., Sotgiu, G., Canonica, G.W., Soriano, J.B., 2020. COVID-19 mortality rates in the European Union, Switzerland, and the UK: Effect of timeliness, lockdown rigidity, and population density. Minerva Med. 111, http://dx.doi.org/10.23736/S0026-4806.20.06702-6.

Geweke, J., 1992. Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments. In: Bernardo, J.M., Berger, J.O., Dawid, A.P., Smith, A.F.M. (Eds.), Bayesian Statistics, Vol. 4. Clarendon Press, Oxford, pp. 169–193.

Gower, J.C., Ross, G.J.S., 1969. Minimum spanning trees and single linkage cluster analysis. Appl. Stat. 18, 54. http://dx.doi.org/10.2307/2346439.

Gray, W.K., Navaratnam, A.V., Day, J., Wendon, J., Briggs, T.W.R., 2021. COVID-19 hospital activity and in-hospital mortality during the first and second waves of the pandemic in England: An observational study. Thorax http://dx.doi.org/10.1136/thoraxjnl-2021-218025.

Grint, D.J., Wing, K., Houlihan, C., Gibbs, H.P., Evans, S.J.W., Williamson, E., McDonald, H.I., Bhaskaran, K., Evans, D., Walker, A.J., Hickman, G., Nightingale, E., Schultze, A., Rentsch, C.T., Bates, C., Cockburn, J., Curtis, H.J., Morton, C.E., Bacon, S., Davy, S., Wong, A.Y.S., Mehrkar, A., Tomlinson, L., Douglas, I.J., Mathur, R., MacKenna, B., Ingelsby, P., Croker, R., Parry, J., Hester, F., Harper, S., DeVito, N.J., Hulme, W., Tazare, J., Smeeth, L., Goldacre, B., Eggo, R.M., 2021. Severity of severe acute respiratory system coronavirus 2 (SARS-CoV-2) alpha variant (b.1.1.7) in England. Clin. Infect. Dis. http://dx.doi.org/10.1093/cid/ciab754.

Hamidi, S., Sabouri, S., Ewing, R., 2020. Does density aggravate the COVID-19 pandemic? J. Am. Plan. Assoc. 86, 495–509. http://dx.doi.org/10.1080/01944363.2020.1777891.

Hartigan, J.A., Wong, M.A., 1979. Algorithm AS 136: A K-means clustering algorithm. Appl. Stat. 28, 100. http://dx.doi.org/10.2307/2346830.

Haynes, W., 2013. Bonferroni correction. In: Encyclopedia of Systems Biology. Springer New York, New York, NY, p. 154. http://dx.doi.org/10.1007/978-1-4419-9863-7_1213.

Ministry of Housing, C..L.G., 2019a. National statistics: English indices of deprivation 2019 (file 1). https://www.gov.uk/government/statistics/english-indices-of-deprivation-2019. (Accessed 16 March 2022).

Ministry of Housing, C..L.G., 2019b. National statistics: English indices of deprivation 2019 (file 10). https://www.gov.uk/government/statistics/english-indices-of-deprivation-2019. (Accessed 10 February 2022).

Ministry of Housing, C..L.G., 2019c. The english Indices of Deprivation 2019 (IoD2019). https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/833959/IoD2019_Infographic.pdf. (Accessed 16 March 2022).

Lee, D., 2020. A tutorial on spatio-temporal disease risk modelling in R using Markov chain Monte Carlo simulation and the CARBayesST package. Spatial Spatio-Temporal Epidemiol. 34, 100353. http://dx.doi.org/10.1016/j.sste.2020.100353.

Lee, D., Robertson, C., McRae, C., Baker, J., 2022. Quantifying the impact of air pollution on COVID-19 hospitalisation and death rates in Scotland. Spatial Spatio-Temporal Epidemiol. 42, 100523. http://dx.doi.org/10.1016/j.sste.2022.100523.

Lee, D., Rushworth, A., Napier, G., 2018. Spatio-Temporal Areal Unit Modeling in R with Conditional Autoregressive Priors Using the CARBayesST Package. J. Stat. Softw. 84, http://dx.doi.org/10.18637/jss.v084.i09.

Leroux, B.G., Lei, X., Breslow, N., 2000. Estimation of disease rates in small areas: A new mixed model for spatial dependence. In: Halloran, M., Berry, D. (Eds.), Statistical Models in Epidemiology, the Environment, and Clinical Trials. In: The IMA Volumes in Mathematics and Its Applications, vol. 116, Springer, New York, NY, pp. 179–191. http://dx.doi.org/10.1007/978-1-4612-1284-3_4.

Ljung, G.M., Box, G.E.P., 1978. On a measure of lack of fit in time series models. Biometrika 65, 297–303. http://dx.doi.org/10.1093/biomet/65.2.297.

Mendez-Brito, A., Bcheraoui, C.E., Pozo-Martin, F., 2021. Systematic review of empirical studies comparing the effectiveness of non-pharmaceutical interventions against COVID-19. J. Infection 83, 281–293. http://dx.doi.org/10.1016/j.jinf.2021.06.018.

Merriam-Webster Online Dictionary, 2022. Lockdown definition and meaning. https://www.merriam-webster.com/dictionary/lockdown. (Accessed 14 April 2022).

Moran, P., 1950. Notes on continuous stochastic phenomena. Biometrika 37, 17–23. http://dx.doi.org/10.1093/biomet/37.1-2.17.

Morens, D.M., Breman, J.G., Calisher, C.H., Doherty, P.C., Hahn, B.H., Keusch, G.T., Kramer, L.D., LeDuc, J.W., Monath, T.P., Taubenberger, J.K., 2020. The origin of COVID-19 and why it matters. Am. J. Trop. Med. Hyg. 103, 955–959. http://dx.doi.org/10.4269/ajtmh.20-0849.

ONS, 2021a. Coronavirus (COVID-19) Infection Survey technical article: Waves and lags of COVID-19 in England, June 2021. https://www.ons.gov.uk/peoplepopulationandcommunity/healthandsocialcare/conditionsanddiseases/articles/coronavirusCOVID19infectionsurveytechnicalarticle/wavesandlagsofCOVID19inenglandjune2021. (Accessed 11 March 2022).

ONS, 2021b. Dataset: Death registrations and occurrences by local authority and health board. https://www.ons.gov.uk/peoplepopulationandcommunity/healthandsocialcare/causesofdeath/datasets/deathregistrationsandoccurrencesbylocalauthorityandhealthboard. (Accessed 02 June 2021).

ONS, 2021c. Dataset: Estimates of the population for the UK, England and Wales, Scotland and Northern Ireland. https://www.ons.gov.uk/peoplepopulationandcommunity/populationandmigration/populationestimates/datasets/populationestimatesforukenglandandwalesscotlandandnorthernireland. (Accessed 01 December 2021).

Page, M.L., McNamara, A., 2021. Alpha COVID-19 variant (B.1.1.7). https://www.newscientist.com/definition/uk-COVID-19-variant-b-1-1-7/. (Accessed 11 March 2022).

Palladino, R., Bollon, J., Ragazzoni, L., Barone-Adesi, F., 2021. Effect of implementation of the lockdown on the number of COVID-19 deaths in four European countries. Disaster Med. Public Health Preparedness 15, e40–e42. http://dx.doi.org/10.1017/dmp.2020.433.

Prime Minister's Office, 2021. Oral statement - PM statement to the House of Commons on roadmap for easing lockdown restrictions in England: 22 February 2021. https://www.gov.uk/government/speeches/pm-statement-to-the-house-of-commons-on-roadmap-for-easing-lockdown-restrictions-in-england-22-february-2021. (Accessed 10 March 2022).

Rand, W.M., 1971. Objective criteria for the evaluation of clustering methods. J. Amer. Statist. Assoc. 66, 846. http://dx.doi.org/10.2307/2284239.

Rashedi, J., Poor, B.M., Asgharzadeh, V., Pourostadi, M., Kafil, H.S., Vegari, A., Tayebi-Khosroshahi, H., Asgharzadeh, M., 2020. Risk factors for COVID-19. Le Infezioni in Med. 28, 469–474.

Remuzzi, A., Remuzzi, G., 2020. COVID-19 and Italy: What next? Lancet 395, 1225–1228. http://dx.doi.org/10.1016/S0140-6736(20)30627-9.

Rousseeuw, P.J., 1987. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. J. Comput. Appl. Math. 20, 53–65. http://dx.doi.org/10.1016/0377-0427(87)90125-7.

Rushworth, A., Lee, D., Mitchell, R., 2014. A spatio-temporal model for estimating the long-term effects of air pollution on respiratory hospital admissions in Greater London. Spatial Spatio-Temp. Epidemiol. 10, 29–38. http://dx.doi.org/10.1016/j.sste.2014.05.001.

Sartorius, B., Lawson, A.B., Pullan, R.L., 2021. Modelling and predicting the spatio-temporal spread of COVID-19, associated deaths and impact of key risk factors in England. Sci. Rep. 11, 5378. http://dx.doi.org/10.1038/s41598-021-83780-2.

Shereen, M.A., Khan, S., Kazmi, A., Bashir, N., Siddique, R., 2020. COVID-19 infection: Emergence, transmission, and characteristics of human coronaviruses. J. Adv. Res. 24, 91–98. http://dx.doi.org/10.1016/j.jare.2020.03.005.

Silva, L., Filho, D.F., Fernandes, A., 2020. The effect of lockdown on the COVID-19 epidemic in Brazil: Evidence from an interrupted time series design. Cadernos De Saúde Pública 36, http://dx.doi.org/10.1590/0102-311x00213920.

Siqueira, C.A.d., de Freitas, Y.N.L., de Camargo Cancela, M., Carvalho, M., Oliveras-Fabregas, A., de Souza, D.L.B., 2020. The effect of lockdown on the outcomes of COVID-19 in Spain: An ecological study. PLoS One 15, e0236779. http://dx.doi.org/10.1371/journal.pone.0236779.

Spiegelhalter, D.J., Best, N.G., Carlin, B.P., van der Linde, A., 2002. Bayesian measures of model complexity and fit. J. R. Stat. Soc. Ser. B Stat. Methodol. 64, 583–639. http://dx.doi.org/10.1111/1467-9868.00353.

Tinson, A., 2021. What geographic inequalities in COVID-19 mortality rates and health can tell us about levelling up. https://www.health.org.uk/news-and-comment/charts-and-infographics/what-geographic-inequalities-in-COVID-19-mortality-rates-can-tell-us-about-levelling-up. (Accessed 10 March 2022).

Williamson, E.J., Walker, A.J., Bhaskaran, K., Bacon, S., Bates, C., Morton, C.E., Curtis, H.J., Mehrkar, A., Evans, D., Inglesby, P., Cockburn, J., McDonald, H.I., MacKenna, B., Tomlinson, L., Douglas, I.J., Rentsch, C.T., Mathur, R., Wong, A.Y.S., Grieve, R., Harrison, D., Forbes, H., Schultze, A., Croker, R., Parry, J., Hester, F., Harper, S., Perera, R., Evans, S.J.W., Smeeth, L., Goldacre, B., 2020. Factors associated with COVID-19-related death using openSAFELY. Nature 584, 430–436. http://dx.doi.org/10.1038/s41586-020-2521-4.

Wolff, D., Nee, S., Hickey, N.S., Marschollek, M., 2021. Risk factors for COVID-19 severity and fatality: A structured literature review. Infection 49, 15–28. http://dx.doi.org/10.1007/s15010-020-01509-1.

Zhang, W., Ge, Y., Liu, M., Atkinson, P.M., Wang, J., Zhang, X., Tian, Z., 2021. Risk assessment of the step-by-step return-to-work policy in Beijing following the COVID-19 epidemic peak. Stoch. Environ. Res. Risk Assess. 35, 481–498. http://dx.doi.org/10.1007/s00477-020-01929-3.