PHILOSOPHICAL TRANSACTIONS A

rsta.royalsocietypublishing.org

Research

Article submitted to journal

Subject Areas:
Statistics, Bioinformatics, Computational Biology

Keywords:
Principal Component Analysis; COVID-19; multivariate statistics; dimension reduction; spatial epidemiology

Author for correspondence:
Ben Swallow
e-mail: ben.swallow@glasgow.ac.uk

Tracking the national and regional COVID-19 epidemic status in the UK using directed Principal Component Analysis

Ben Swallow[1], Wen Xiang[2], Jasmina Panovska-Griffiths[3,4]

[1]School of Mathematics and Statistics, University of Glasgow, Glasgow, UK
[2]Department of Statistics, London School of Economics and Poltical Science, London, UK
[3]The Big Data Institute, Nuffield Department of Medicine, University of Oxford, Oxford, UK
[4]The Queen's College, University of Oxford, Oxford, UK

One of the difficulties in monitoring an ongoing pandemic is deciding on the metric that best describes its status when multiple highly inter-correlated measurements are available. Having a single measure, such as whether the effective reproduction number R, has been useful in tracking whether the epidemic is on the incline or the decline and for imposing policy interventions to curb the increase. We propose an additional metric for tracking the UK epidemic across all four nations, that can capture the different spatial scales. This paper illustrates how to derive the principal scores from a weighted Principal Component Analysis using publicly available data. We show the detectable impact of interventions on the state of the epidemic and suggest that there is a single dominant trend observable through the principal score, but this is different across nations and waves. For example, the epidemic status can be tracked by cases in Scotland at a countrywide scale, whereas across waves and disjoint nations, hospitalisations are the dominant contributor to principal scores. Thus, our results suggest that hospitalisations may be an additional useful metric for ongoing tracking of the epidemic status across the UK nations alongside R and growth rate.

THE ROYAL SOCIETY PUBLISHING

1. Introduction

There are inherent difficulties in modelling the current state of a pandemic, with challenges associated with censoring, reporting lags, uncertainty in cause of death, asymptomatic cases and testing errors all contributing to possible biases in informed measures of incidence (Shadbolt et al., 2021; Swallow et al., 2021). Several metrics have been used to provide a quantitative evaluation of the current state of an epidemic, with the effective reproduction number, R, being arguably the most commonly used (see Pellis et al. (2021) for further discussion of this). R is a measure of the number of secondary infections stemming from a single infection (UKHSA (2020)) and can reflect the transmissibility of the epidemic. It allows tracking of the status of the epidemic with $R < 1$ suggesting that the epidemic is in the declining phase, whereas $R > 1$ describing increased transmission and a growing epidemic. Since the onset of the pandemic, in the UK a consensus range of R generated from a number of mathematical models (UKHSA (2021)) has been a used to track the epidemic status and to inform and guide policy decision makers in imposing and removing interventions such as lockdowns. However, while R is a useful measure, it is sensitive to, for example, the choice of data being used (e.g., the number of tests being carried out and any delays in reporting of cases) or the method for calculating it (e.g., the combination and type of models used or the length of the time-slice being used for the calculation). Furthermore, once the population is partially or fully vaccinated defining population wide R may not suffice. There are also difficulties associated with the fact that the onset of symptoms often happen a few days after the onset of infection, and hence the current value of R really represents the state of the pandemic at some point in the recent past; in the UK the consensus is that the current value of R is lagged by 2 to 3 weeks (UKHSA (2020)). In this paper we analyse a number of COVID-19 epidemiological time series and explore whether tracking daily cases, hospitalisations, MVBs admissions or deaths related to COVID-19 may be a useful additional to R and growth rate measure for tracking the ongoing national and regional UK epidemic status.

Multivariate projection methods, specifically methods of dimension reduction, aim to find lower-dimensional representations of multiple (correlated) measurements, providing new measurement axes that are weighted linear combinations of the original measurements. Principal Components Analysis (PCA) also has the added representing a geometric rotation of the data into axes that are both orthogonal (or uncorrelated) and maximal in variance retention. In the specific case of measures of epidemic indices, there are inherent correlations between the different trajectories, as there will be significant overlap in patients moving between epidemiological state, although with some lag in time. Larger numbers of infected cases generally lead to higher numbers of those requiring medical intervention or mortality rates. However, the relationship will not be perfectly linear and consistent across measurements.

One of the drawbacks of conventional PCA is that as an unsupervised method, it is often challenging to interpret the often abstract outputs that it produces. Also in many scenarios, known structures in the data can be readily available and are not themselves of interest. One of the interests in these scenarios is finding latent structures in the data that explain hidden correlations between measured quantities. The study of these structures can then inform further on possible important mechanisms within the system under study.

There appears to have been little attempt to account for these inherent between measurements. The formal structures inherent in these models are not always directly applicable to , and can often be computationally complex when conducting inference. Methods that enable practitioners to obtain indicators of the current or recent state of the pandemic in a quick and efficient manner are therefore highly desirable.

We use a similar approach to Xiang and Swallow (2021) in analysing data from the UK COVID-19 dashboard ("Coronavirus (COVID-19) in the UK," 2021) to study simpler representations of the multivariate output of cases, deaths, hospitilisations and intensive care bed occupancy. We use S-Mode and T-mode PCA, with temporal weight matrix calculated from median correlation in residuals following a Generalised Additive Model fitted to a smooth of date index. We conduct

the PCA firstly on deaths, cases, hospitalisations and MV beds, and then a second analysis on just cases and deaths. Each of the two analyses has separate trends for the four nations. We also conduct an additional analysis on all streams, segregated by UK nation and a further analysis looking for differences between the dynamics of the two principal COVID-19 waves currently observed.

2. Methods

(a) Data

Data are extracted from the UK Covid-19 Dashboard ("Coronavirus (COVID-19) in the UK," 2021), consisting of daily measurements of reported cases, deaths, hospitalisations and MV beds occupied between 2020-04-02 to 2021-02-22. Numbers were available at the level of the four individual UK nations, as well as pooled totals across the UK as a whole. Analyses were conducted at both spatial resolutions. Data matrices were imported into R, with rows corresponding to time intervals (days) and columns representing output variable (Covid outcome and/or country). Data were checked for inconsistencies and outliers that may have impacted results, although none were found.

For additional analyses, the data were also subset into nation-specific data matrices and pooled matrix containing only cases and deaths. Finally, matrices corresponding to the two individual waves, corresponding to March to May 2020 and September 2020 and April 2021 respectively (ONS, 2021), were also constructed.

(b) Principal Components Analysis

Multivariate projection and decomposition methods, such as Principal Components Analysis (PCA) enable the extraction of structures in multivariate data through an eigen-decomposition of the correlation or covariance matrix. The eigenvectors form a rotated basis into a new set of uncorrelated axes that are ordered by magnitude of their corresponding eigenvalues, or equivalently the proportion of variation in the original measurements that they explain.

The linear combinations of the variables are a rotation of $\mathbf{X} = (X_1, ..., X_p)^\intercal$, a p dimensional random variable. The objective is to find a linear transformation $\mathbf{Y} = \mathbf{U}^\intercal\mathbf{X}$, where $\mathbf{U}^\intercal = (U_{11}, ..., U_{p1})^\intercal$ is a matrix of constants such that the $\mathbb{V}ar(\mathbf{Y}_1)$ is maximised, subject to the normalising constraint $\mathbf{U}^\intercal\mathbf{U} = \mathbf{I}$. It can therefore be seen both as a variance-maximisation projection of the covariance matrix or a linear transformation into an orthogonal set of bases.

Extensions to standard PCA relax the *a priori* assumption of unknown structure and allow users to account for existing spatial and/or temporal structures inherent in the data through the use of spatial and/or temporal weighting matrices. Accounting for existent temporal structures in the data allow the extraction of important residual joint structures that can be more readily interpreted than if these known structures are not accounted for. The S-Mode PCA aims to find dominant temporal trends across the four data streams. Conversely, T-Mode PCA aims to find different patterns in the data and the associated time points at which they occur.

(c) Flow directed PCA

(i) S-Mode and T-Mode Principal Components Analysis

Next we aim to discover important global spatial and temporal trends in cases and deaths from COVID-19. To extract the important trends, PCA and similar dimension reduction techniques are an obvious choice. PCA conducts an eigendecomposition of the covariance (or correlation) of a data matrix, with eigenvalues ordered by magnitude to reduce a set of p correlated variables to a smaller set of $k < p$ orthogonal variables. Versions of PCA for spatio-temporal data were referred to by (Richman, 1986), S-mode and T-mode PCA, the particular mode depending on whether the columns of are time points (T-mode) or countries (S-mode).

S-Mode PCA aims to find dominant temporal trends across the spatial locations, highlighting a small number of dominant temporal trends across all countries. Conversely, T-Mode PCA aims to find different spatial patterns in the data and the associated time points at which they occur. In general, however, PCA finds unsupervised structures in the data by conducting an eigen decomposition of the correlation of covariance matrix. Whilst this can often be useful in visualising data in lower dimensions, it is not possible to guide the structure of the new axes using prior information or independent data. In order to account for known spatio-temporal correlations inherent in the data, we use spatio-temporally weighted S-mode and T-mode PCA, which aim to find dominant temporal and spatial patterns respectively. Gallacher et al. (2017) extended these approaches to account for known spatio-temporal structures in river flow systems through the use of spatial and/or temporal weight matrices to inform spatio-temporal structure.

The modelling approach can be summarised as follows. Assuming the data matrix is a $n \times p$ matrix, a $p \times p$ column weight matrix $\boldsymbol{\Omega}$ and $n \times n$ row weight matrix $\boldsymbol{\Phi}$ can be constructed so that PCA is applied to a transformed matrix $\tilde{\boldsymbol{X}} = \boldsymbol{\Phi X \Omega}$. Residual spatial or temporal structure may then be associated with the impact of any intervention, such as lockdowns or vaccination programmes, applied at those times. The temporal weight matrix is applied to the original data matrix to scale the columns appropriately, then dimension reduction is conducted on the new scaled matrix. In our analyses here, we only consider the temporal column matrix, $\boldsymbol{\Omega}$, however we describe the full process for completeness.

Let $\tilde{\boldsymbol{X}}$ denote the scaled data matrix, weighted by spatial and/or temporal weight matrices as follows.

$$\tilde{\boldsymbol{X}} = \boldsymbol{\Phi X \Omega} = \tilde{\boldsymbol{U}} \tilde{\boldsymbol{D}} \tilde{\boldsymbol{V}}^{\intercal},$$

where the decomposed matrices generated on the scaled variables are denoted with a tilde. The PCs of the new weighted variables therefore become $\boldsymbol{X \Omega \tilde{V}}$, and the loadings are $\boldsymbol{\Omega}^{-\intercal} \tilde{\boldsymbol{V}}$.

Analyses were conducted using the stpca package in R.

(ii) Temporal weight matrix

The temporal weight matrix is constructed similarly to Gallacher et al. (2017) using independent Generalised Additive Models (GAMs) (Wood, 2017). These are fitted by restricted maximum likelihood (REML) in the mgcv package (Wood, 2012) to each of the time series with an intercept and an univariate smooth function of date as predictor. As in Xiang and Swallow (2021), this aims to remove trends specific to each stream, with residual variation used to determine the principal scores. Remaining correlation in the model residuals between time $[1, \ldots, (n-1)]$ and $[2, \ldots, n]$ is calculated for each stream and then the median value is used as an estimate of ρ, the average global temporal correlation. The median residual temporal correlations ρ are estimated for each of the output time series (these will vary in each analysis depending on the spatial resolution/measurements used). The ith row and jth column element of the temporal weight matrix $T_{i,j}$ is then specified as $\rho^{|i-j|}$ for all time indices in the original data matrix. The temporal weight matrix is applied to the data matrix and dimension reduction is then conducted to create a new uncorrelated set of bases. The method is identifiable up to a change in sign, so in some cases similar trends are apparent, only inverted.

3. Results

(a) Pooled UK analysis

The results at the UK-wide scale suggest that a single combined index of the four different measurements is able to explain approximately 72% of variation (Figure 1), whilst when removing the smoothed temporal trend, the remaining variation accounts for around 40% of the variation.

Our results suggest that there is a general trend of decreasing state of prevalence until August 2020, a short period of consistently low prevalence until early September 2020, which then peaks in November 2020 and is followed by a much larger peak in January 2021 (Figure 1). Finally, the

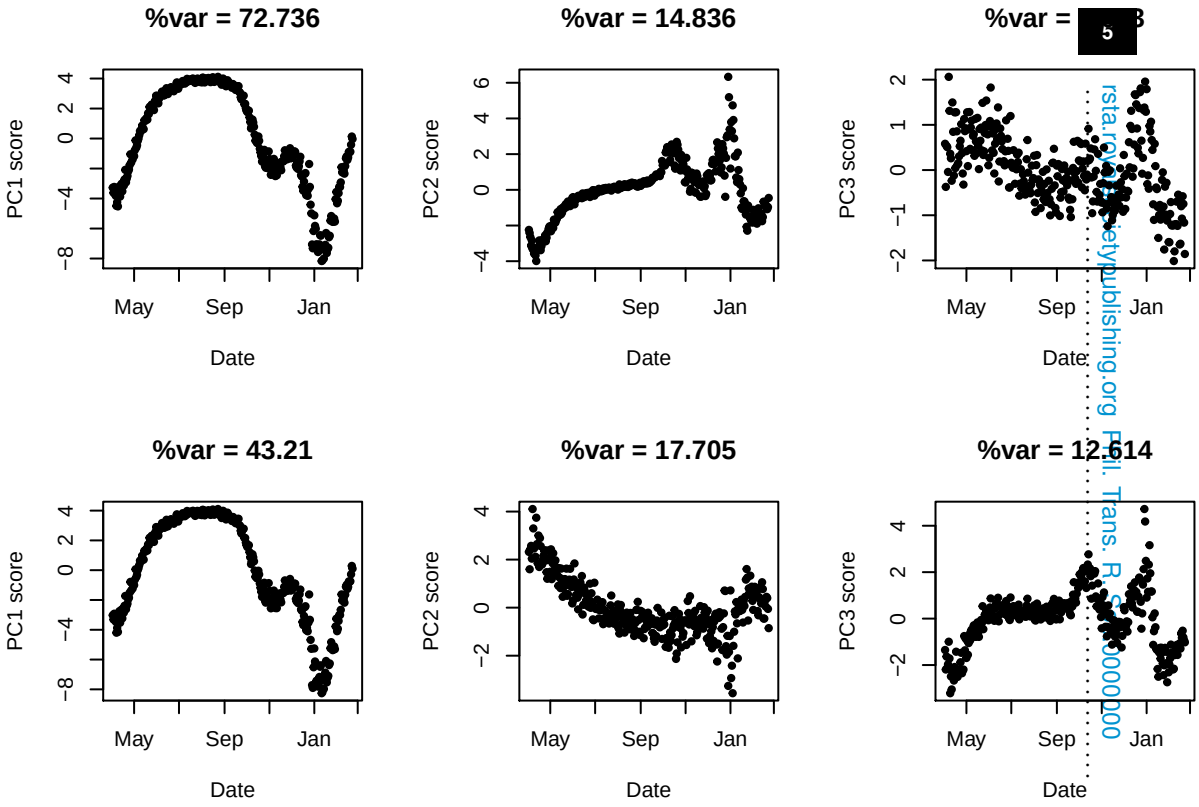

Figure 1. First 3 PC scores for unweighted S-Mode PCA (top 3 panels) and temporally weighted S-Mode PCA (bottom 3 panels).

principal score shows a reduction to the end of the study period. The second principal score still accounts for a reasonable amount of variability and suggests that there is a relatively smooth increasing trend up to January 2021 overall.

(b) Deaths and cases only for the UK

The results for deaths and cases only were relatively consistent with those from all four measurements. This is not too surprising given the the results of the former were predominantly driven by the number of cases, a stream that is also present in this reduced analysis.

(c) Nation-specific analysis

For the nation-specific analyses, the first PC was able to explain between 70% and 95% of the total variation in the data, suggesting that there is only really a single independent trajectory occurring (Figures 5 to 8). The corresponding loadings show that this axis consists predominantly of the number of hospitalised patients. This was consistent across the individual nations and both unweighted and temporally weighted analyses.

All analyses showed peaks in April 2020 and January 2021, with a smaller peak in November 2021 that showed initial signs of reduction before increasing again to the major peak in January 2021. The lockdown that was introduced in November 2020 seems to have been particularly beneficial in Scotland, which showed a reduction down to similar low levels as over the summer months in 2020. The sequential significant growth up to January 2021 follows temporary Christmas reductions in distancing measures and may be a result of this. A combined impact of strict intervention measures after the Christmas holidays and in the New Year combined with

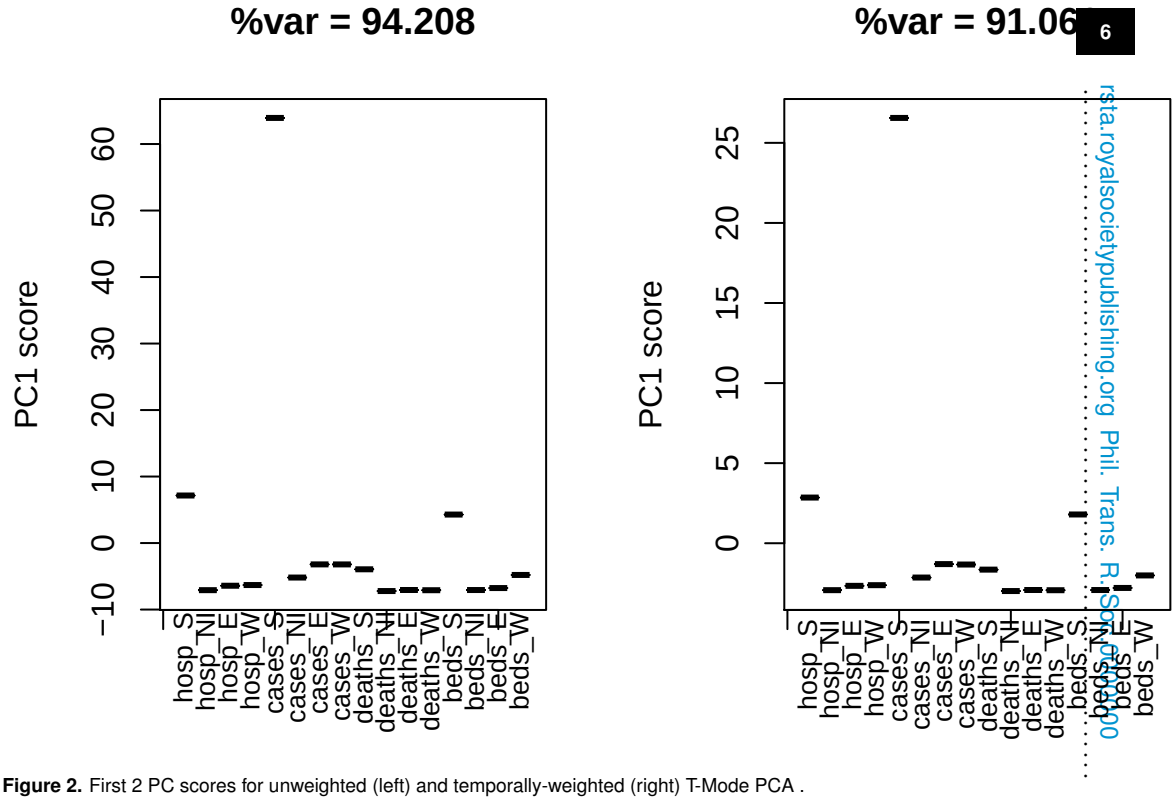

Figure 2. First 2 PC scores for unweighted (left) and temporally-weighted (right) T-Mode PCA .

the rolling out of the vaccination programme lead to a reduction across all trends and this is highlighted in both the first and second PC scores.

Our results suggest that hospitalisations are consistently important as representative of the state of the pandemic at any given time point. This matches the fact that these are likely to be relatively unbiased compared to cases (Vekaria et al., 2021) and less age- or health-specific than MVB occupation. MVBs are also expected to be limited by capacity unlike the overall hospital admissions, the former of which policy has changed significantly over the course of the pandemic (NICE (2021)).

(d) Comparison of waves

We also rerun the analyses to compare dynamics of trends separately for the two principal waves of Covid-19 over the studied period. We denote wave 1 to consist of data prior to 31st May 2020 and wave 2 to correspond to data from 1st September 2020 to the end of March 2021 (ONS, 2021). The results were consistent with the full analysis in terms of which variables dominated the first principal component, however there were clearly different general temporal trends in each of these periods. Particularly noteworthy is the principal linear downward trend seen in PC1 for wave 1 dynamics. A similar linear reduction is seen at the end of wave 2. Both of these time periods correspond to strict lockdown periods in the UK.

The loadings between the two analyses, which represent the contribution of an individual measurement to each linear transformation, showed some interesting trends. Specifically, in wave 2, the loadings decreased for hospitalisations in all nations; deaths in England; and cases in Wales and Northern Ireland. All other measurements were more dominant in wave 1.

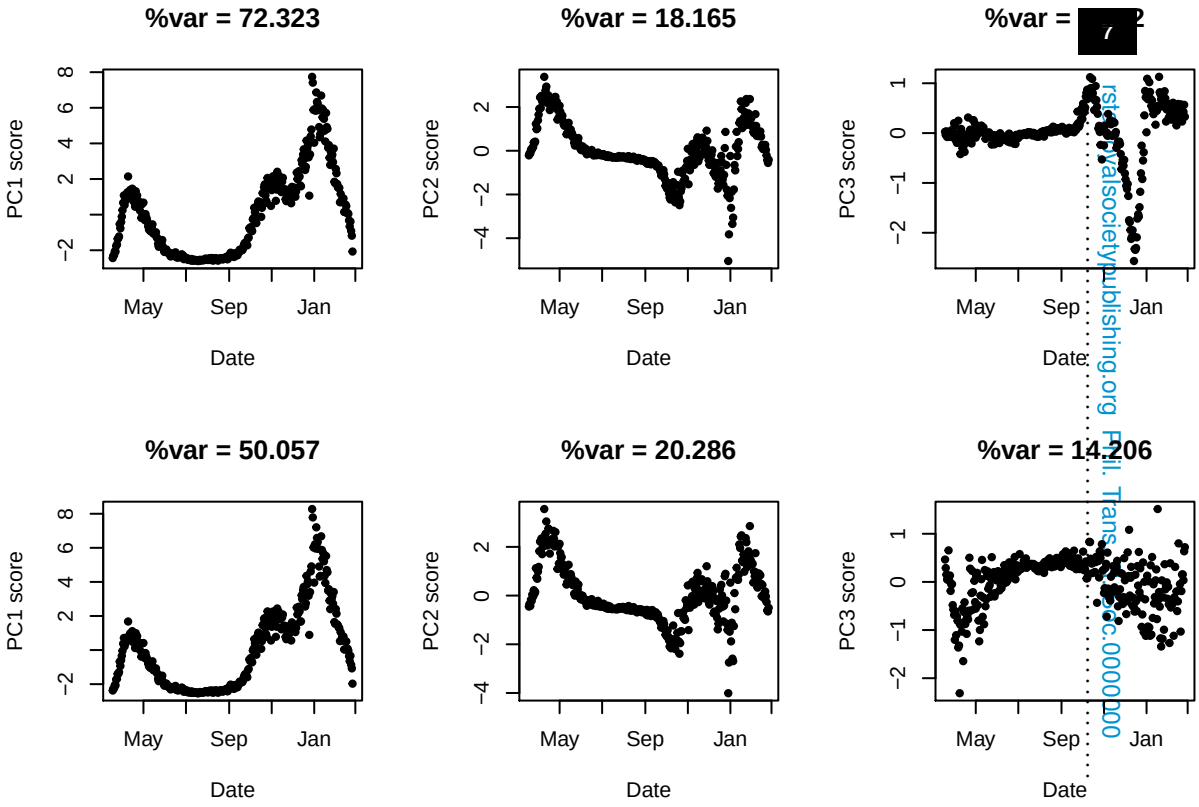

Figure 3. First 3 PC scores for unweighted S-Mode PCA (top 3 panels) and temporally weighted S-Mode PCA (bottom 3 panels) applied to just cases and deaths.

4. Conclusions

In conclusion, representation of the pandemic is shown to vary depending on the spatial resolution at which the analysis is conducted. Looking at the UK as a whole, the trends are largely driven by the number of reported cases (specifically in Scotland). This could possibly be because cases have been better managed and therefore the stream may be an overall better representation of the course of the pandemic as a whole. However, reported numbers of cases can be highly variable and subject to possible under-reporting and strongly dependent on testing, contact-tracing and isolation behaviour and policies. For the nation-specific analyses, numbers of hospitalised patients were highlighted as the dominant trend in the data. Of the four indices used in this analysis, this matches previous considerations over the validity of the measures as representative of the quantity of interest. Death figures have also been challenging during the pandemic due to delays in reporting and uncertainty over exact cause of death. There have been several different figures adopted, some which relate to mortality within 28 days of a positive test and others where COVID-19 was mentioned on the death certificate. Both these aspects of the data can cause biases in reported deaths related to COVID-19. The number of MVBs, whilst perhaps slightly less biased than the reported numbers of cases and deaths, is likely to be susceptible to changes in policy and severely limited by a carrying capacity at each hospital and movement of patients in and out of these wards depending on need. Hospitalisation figures, however, are less susceptible to either of these drawbacks. In summary, our results suggest that overall hospitalisations may be a good additional metric to adopt for tracking the epidemic on regional or local spatial scales alongside tracking R and growth rate.

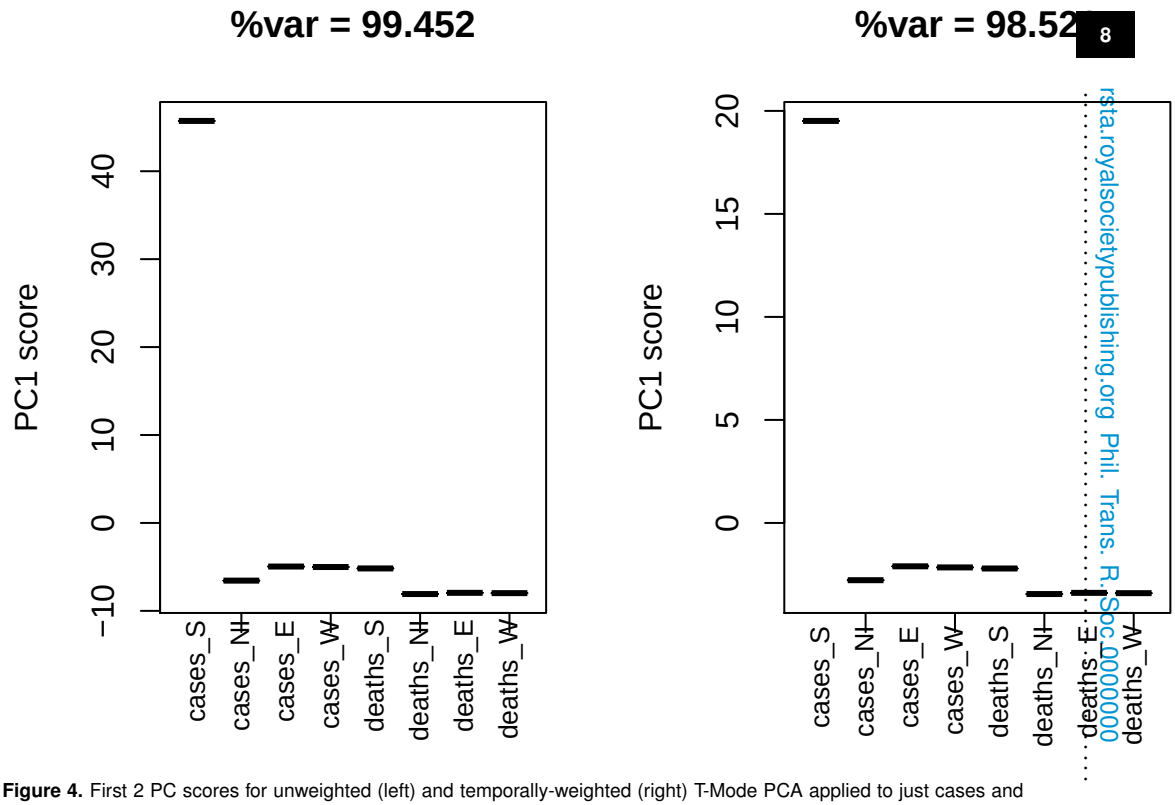

Figure 4. First 2 PC scores for unweighted (left) and temporally-weighted (right) T-Mode PCA applied to just cases and deaths.

Our results also highlight the care that must be taken when choosing the level of aggregation of data for statistical analysis (Jeffery et al., 2014, 2009). For example, our results highlight that aggregating the data to UK level changes the dynamics of the data streams and masks important local dynamics.

There were limited differences in general trends between the first and second wave of COVID-19 across the UK, with both waves showing linear reductions post-lockdown interventions. MVBs generally became less dominant while hospitalisations emerged as more dominant in the second wave. Deaths in England were also more dominant in wave 1. This could relate to a change in policy only intervening in the most serious cases, or an improvement in outcome due to improved understanding of the disease and the impact from the large scale vaccination programme from December 2020.

The fact that most analyses of spatial trends show two major principal components, whilst temporal analyses tend to show only one dominant trend in space, is worth highlighting. The principal spatial trend will generally show the dominant trend over larger regions, whilst the second spatial trend could relate to an increase in testing capacity over the time period, showing an overall increase in numbers detected over the whole, whereas early in the emergence of SARS-CoV-2, only those with severe symptoms would likely be detected. It could also highlight second-order structure in the data, corresponding to higher variability in times of policy changes (e.g., lockdowns, reopening of schools, reopening of public services) or local variations in trends not picked up by the main principal component.

In Xiang and Swallow (2021), the authors analysed data from global trends relating to reported cases of, and deaths resulting from, SARS-CoV-2. Their analyses found a single temporal trend dominated the global spread of the disease. This suggests that there are multi-scale processes occurring, namely a general spread across countries that happens smoothly and consistently, and

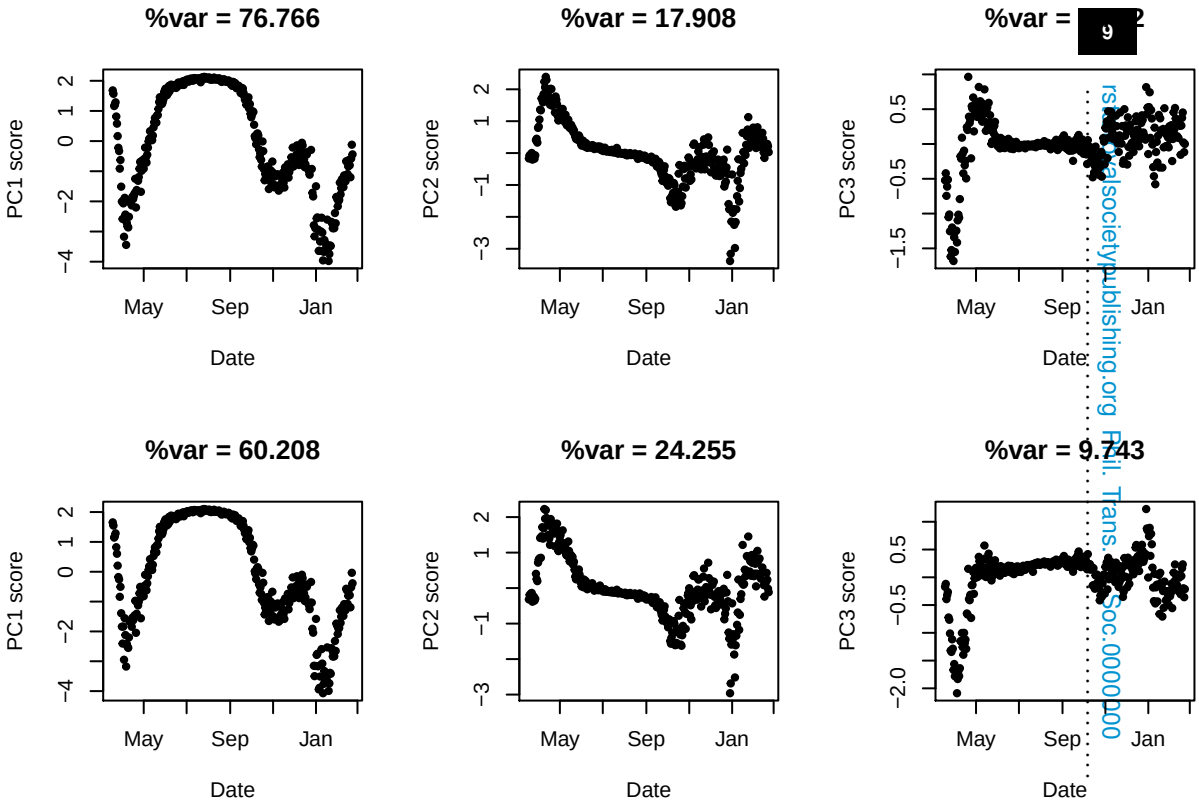

Figure 5. First 3 PC scores for unweighted S-Mode PCA (top 3 panels) and temporally weighted S-Mode PCA (bottom 3 panels) for Scotland.

more local dynamics within a country that occur at a finer resolution. This again highlights the importance of considering appropriate scales when looking at dynamics of infectious diseases (e.g. Garabed et al., 2020). Aggregation of data may be beneficial computationally or in terms of reducing the impact of biases and errors in the data, however it will inevitably change results.

Removing common temporal correlation from the data through a temporal weight matrix reduced the amount of variation explained in all analyses, but had little impact on the overall conclusions. This supports the idea that there is a great deal of similarity across nations within the UK in terms of seasonal dynamics and principal trends, however it is not sufficient to treat the streams as either independent or entirely analogous.

There appear to have been few attempts to look at the structures inherent in multiple measurements of the pandemic and how to reduce these down to single measures, other than the R_0 value widely applied. Rahman et al. (2020) construct a structural equation model, in which dimension reduction is included, and a variable relating to 'pandemic severity' is generated as one of their dimensions of interest. However, this is not the main focus of their research and they only include cases and deaths in their analysis. Their analyses support the idea that treating the measurements as independent is clearly erroneous, as they clearly have large amounts of correlation. Also, errors and biases in each stream will be compounded if independent models are fitted. Using dimension reduction techniques, such as directed PCA, allows for the correlation to be incorporated into the methodology and reduces the impact of any single stream bias by integrating over them all.

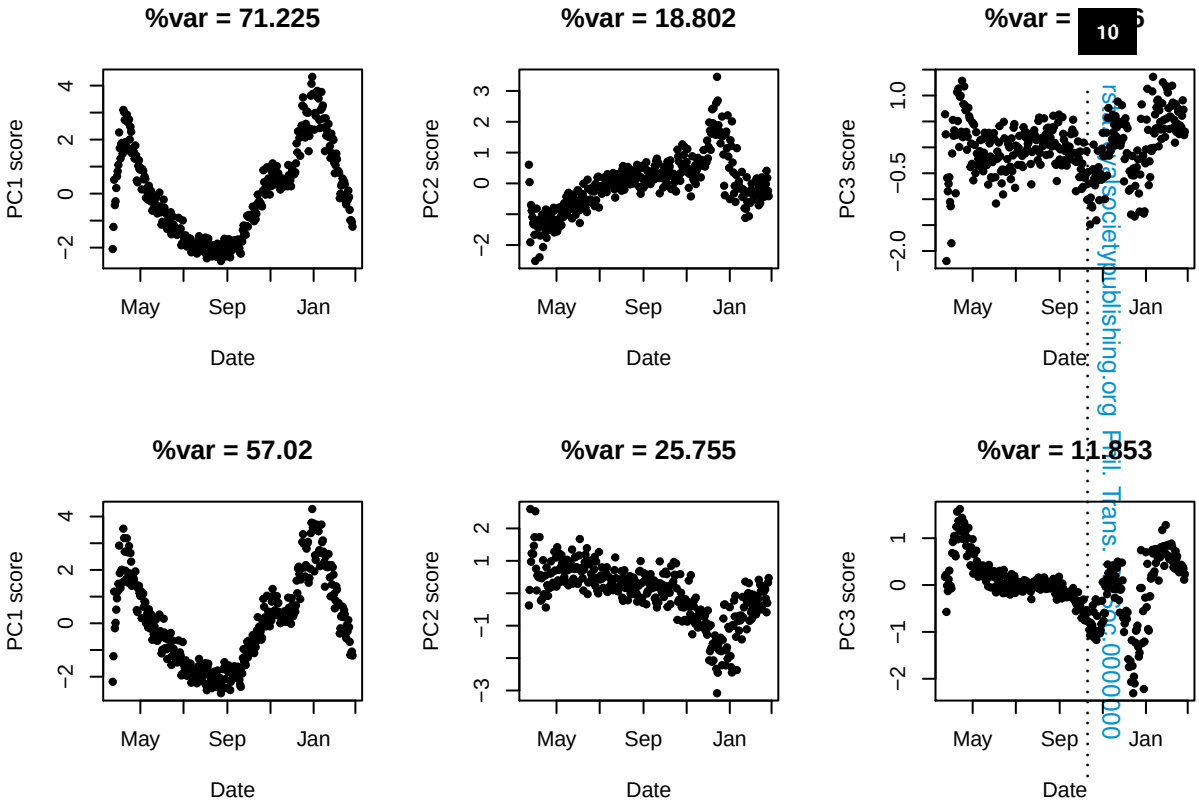

Figure 6. First 3 PC scores for unweighted S-Mode PCA (top 3 panels) and temporally weighted S-Mode PCA (bottom 3 panels) for Wales.

5. References

Coronavirus (COVID-19) in the UK [WWW Document], 2021. [WWW Document]. URL https://coronavirus.data.gov.uk/details/download

Gallacher, K., Miller, C., Scott, E.M., Willows, R., Pope, L., Douglass, J., 2017. Flow-directed PCA for monitoring networks. Environmetrics 28, e2434. https://doi.org/10.1002/env.2434

Garabed, R.B., Jolles, A., Garira, W., Lanzas, C., Gutierrez, J., Rempala, G., 2020. Multi-scale dynamics of infectious diseases. Interface Focus 10, 20190118. https://doi.org/10.1098/rsfs.2019.0118

Jeffery, C., Ozonoff, A., Pagano, M., 2014. The effect of spatial aggregation on performance when mapping a risk of disease. International Journal of Health Geographics 13, 9. https://doi.org/10.1186/1476-072X-13-9

Jeffery, C., Ozonoff, A., White, L.F., Nuño, M., Pagano, M., 2009. Power to detect spatial disturbances under different levels of geographic aggregation. Journal of the American Medical Informatics Association 16, 847–854. https://doi.org/10.1197/jamia.M2788

NICE, 2021. COVID-19 rapid guideline: managing COVID-19. https://www.nice.org.uk/guidance/ng191

ONS, 2021. 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/coronaviruscovid19infectionsurveytechnicalarticlewavesandlagsofcovid19inenglandjune2021

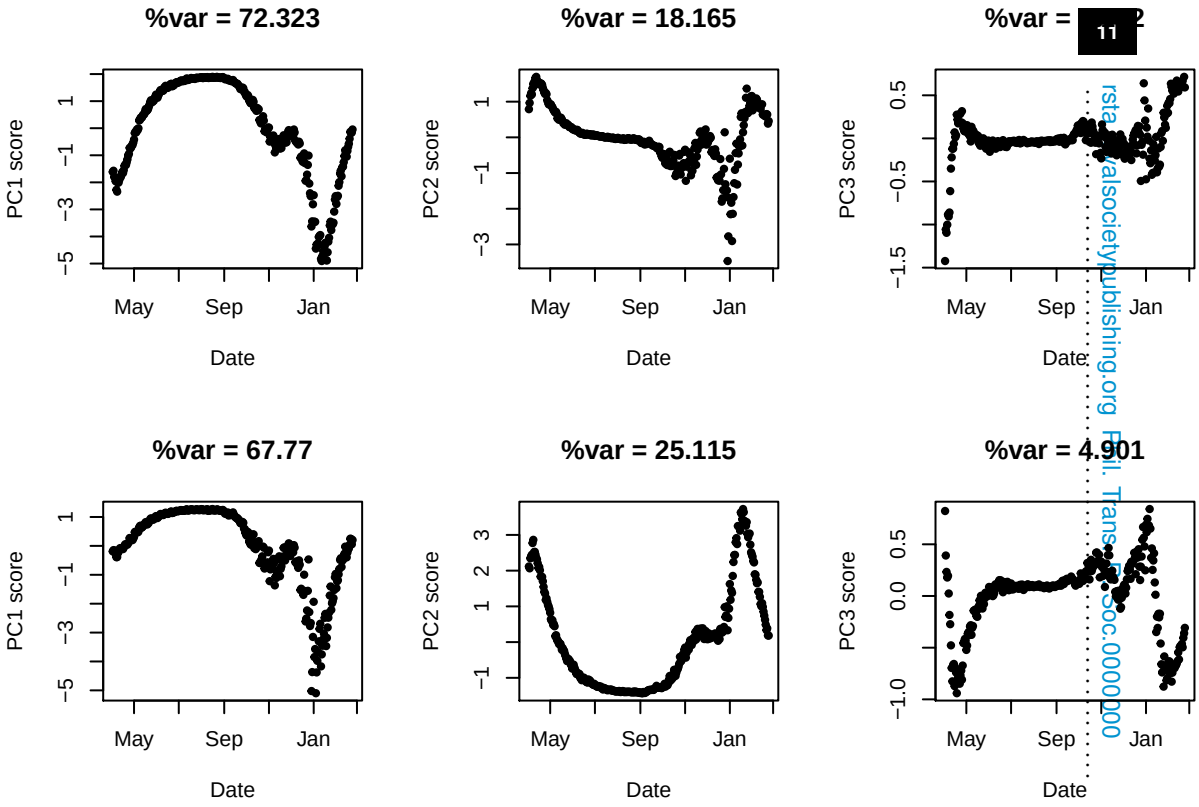

Figure 7. First 3 PC scores for unweighted S-Mode PCA (top 3 panels) and temporally weighted S-Mode PCA (bottom 3 panels) for England.

Pellis, L., Birrell, P.J., Blake, J., Overton, C.E., Scarabel, F., Stage, H.B., Brooks-Pollock, E., Danon, L., Hall, I., House, T.A., Keeling, M., Read, J.M., JUNIPER consortium, 2021. Estimation of reproduction numbers in real time: Conceptual and statistical challenges.

Rahman, M.M., Thill, J.-C., Paul, K.C., 2020. COVID-19 pandemic severity, lockdown regimes, and people's mobility: Early evidence from 88 countries. Sustainability 12. https://doi.org/10.3390/su12219101

Richman, M.B., 1986. Rotation of principal components. Journal of Climatology 6, 293–335. https://doi.org/https://doi.org/10.1002/joc.3370060305

Shadbolt, N., Brett, A., Chen, M., Mario, G., McKendrick, I.J., Panovska-Griffiths, J., Pellis, L., Reeve, R., Swallow, B., 2021. The challenges of data in future pandemics.

Swallow, B., Birrell, P., Blake, J., Burgman, M., Challenor, P., Coffeng, L.E., Dawid, P., Angelis, D.D., Goldstein, M., Hemming, V., Marion, G., McKinley, T., Overton, C., Panovska-Griffiths, J., Pellis, L., Prober, W., Shea, K., Villela, D., Vernon, I., 2021. Challenges in estimation, uncertainty quantification and elicitation for pandemic modelling.

UKHSA, 2020. The R value and growth rate. https://www.gov.uk/guidance/the-r-value-and-growth-rate

UKHSA, 2020. Reproduction number (R) and growth rate: methodology. https://www.gov.uk/government/publications/reproduction-number-r-and-growth-rate-methodology

Vekaria, B., Overton, C., Wiśniowski, A., Ahmad, S., Aparicio-Castro, A., Curran-Sebastian, J., Eddleston, J., Hanley, N.A., House, T., Kim, J., Olsen, W., Pampaka, M., Pellis, L., Ruiz, D.P., Schofield, J., Shryane, N., Elliot, M.J., 2021. Hospital length of stay for covid-19 patients: Data-driven methods for forward planning. BMC Infectious Diseases 21, 700. https://doi.org/10.1186/s12879-021-06371-6

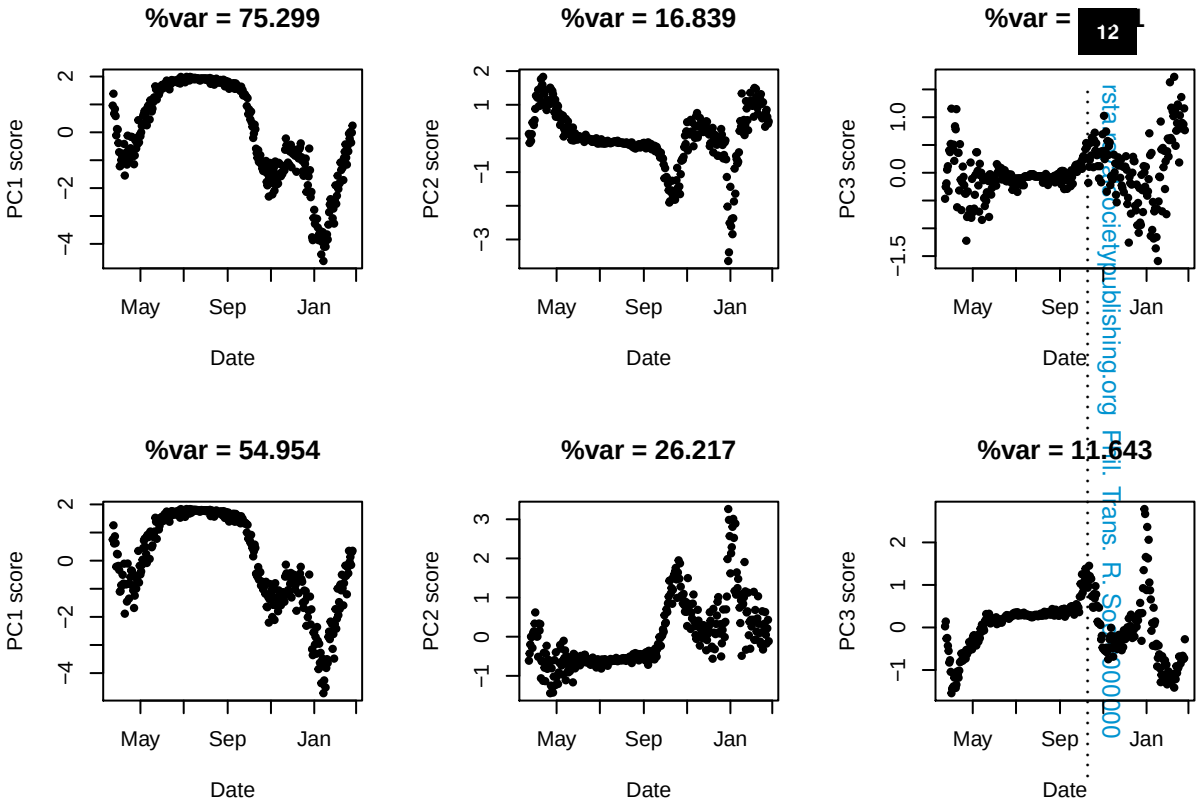

Figure 8. First 3 PC scores for unweighted S-Mode PCA (top 3 panels) and temporally weighted S-Mode PCA (bottom 3 panels) for Northern Ireland.

Wood, S.N., 2017. Generalized additive models: An introduction with R, second edition, Chapman & hall/crc texts in statistical science. CRC Press.

Wood, S.N., 2012. Mgcv: Mixed GAM computation vehicle with GCV/AIC/REML smoothness estimation.

Xiang, W., Swallow, B., 2021. Multivariate spatio-temporal analysis of the global covid-19 pandemic. medRxiv. https://doi.org/10.1101/2021.02.08.21251339

6. Acknowledgements

BS is a member of the Scottish COVID-19 Response Consortium, which was undertaken in part as a contribution to the Rapid Assistance in Modelling the Pandemic (RAMP) initiative, coordinated by the Royal Society.

JPG's work was supported by funding from the UK Health Security Agency and the UK Department of Health and Social Care. This funder had no role in the study design, data analysis, data interpretation, or writing of the report. The views expressed in this article are those of the authors and not necessarily those of the UK Health Security Agency or the UK Department of Health and Social Care.

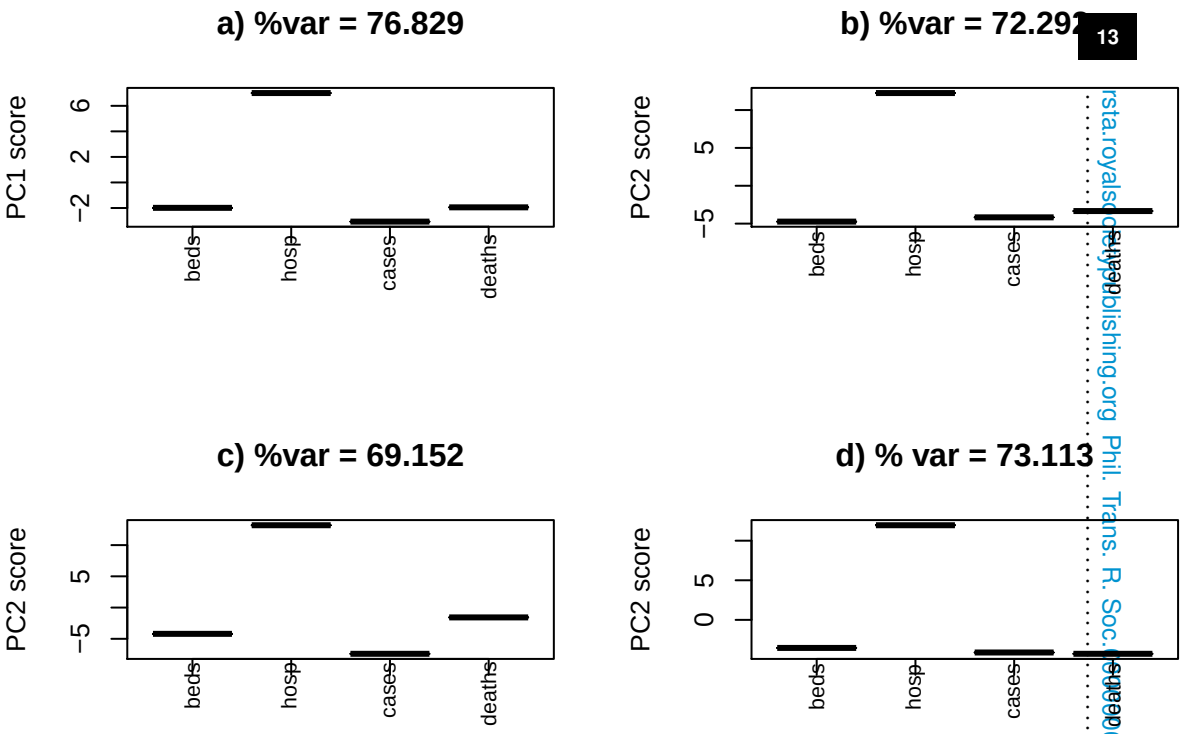

Figure 9. First PC score for temporally weighted T-Mode PCA - a) England; b) Scotland; c) Wales; d) Northern Ireland.

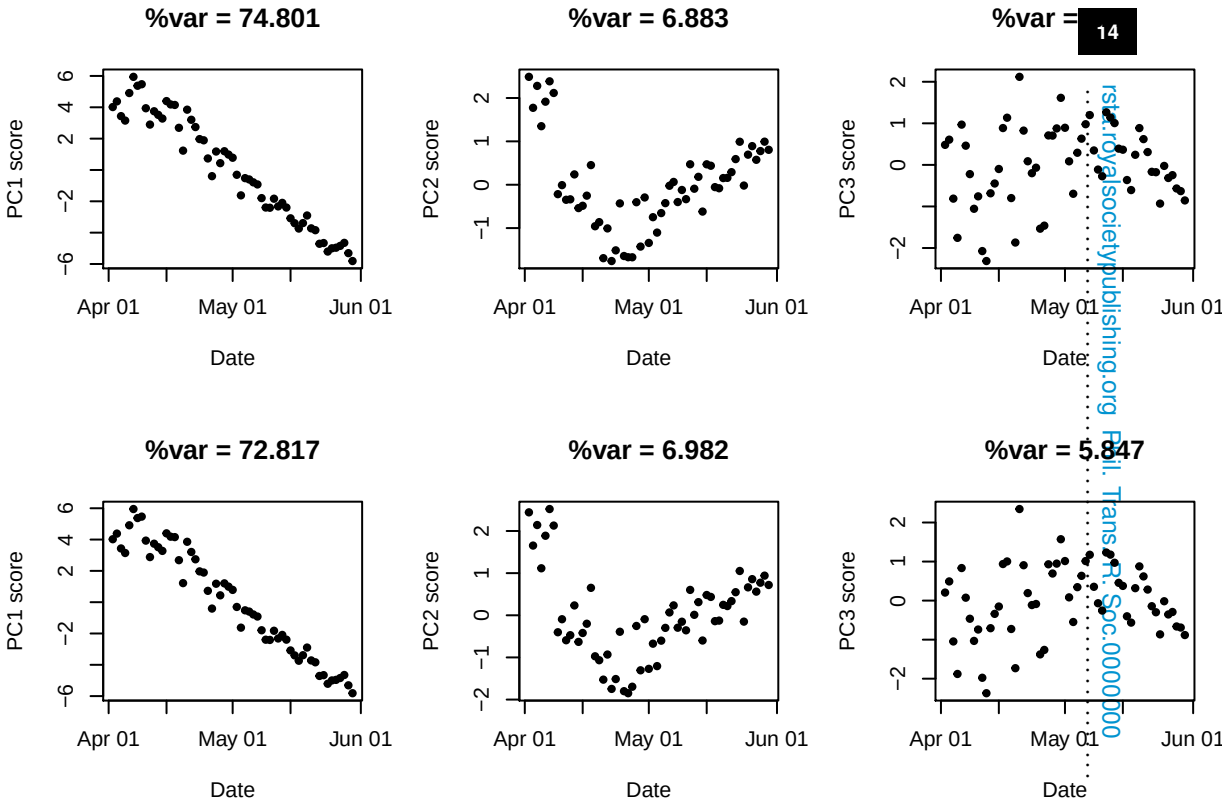

Figure 10. First 2 PC scores for temporally weighted T-Mode PCA covering dates in Covid-19 Wave 1.

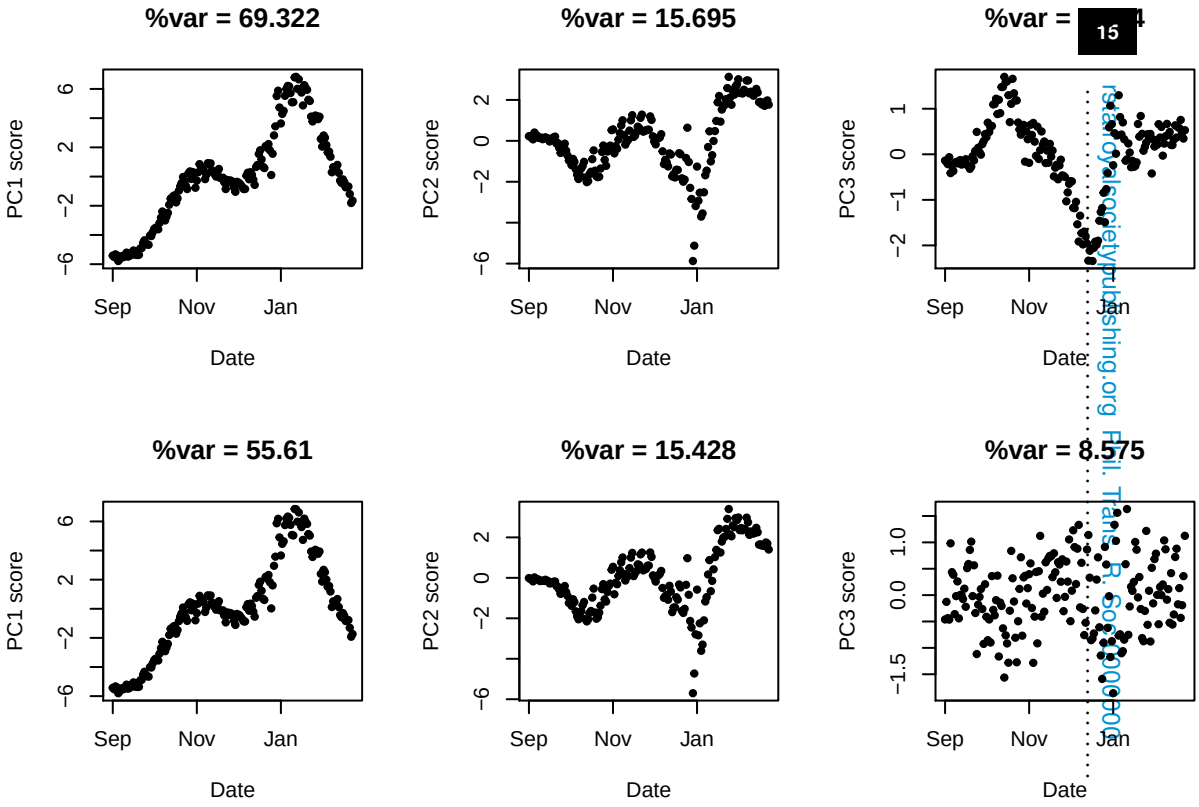

Figure 11. First 2 PC scores for temporally weighted T-Mode PCA covering dates in Covid-19 Wave 2.