


RESEARCH ARTICLE

A hierarchical modelling approach to assess multi pollutant effects in time-series studies

Marta Blangiardo^{1*}, Monica Pirani¹, Lauren Kanapka², Anna Hansell^{1,3}, Gary Fuller⁴

1 MRC-PHE Centre for Environment and Health, Department of Epidemiology and Biostatistics, Imperial College, St Mary's Campus, London, United Kingdom, **2** Department of Mathematics, Imperial College, South Kensington Campus, London, United Kingdom, **3** Centre for Environmental Health and Sustainability, University of Leicester, Leicester, United Kingdom, **4** MRC-PHE Centre for Environment and Health, Environmental Research Group, King's College, London, United Kingdom

* m.blangiardo@imperial.ac.uk



Abstract

When assessing the short-term effect of air pollution on health outcomes, it is common practice to consider one pollutant at a time, due to their high correlation. Multi pollutant methods have been recently proposed, mainly consisting of collapsing the different pollutants into air quality indexes or clustering the pollutants and then evaluating the effect of each cluster on the health outcome. A major drawback of such approaches is that it is not possible to evaluate the health impact of each pollutant. In this paper we propose the use of the Bayesian hierarchical framework to deal with multi pollutant concentrations in a two-component model: a pollutant model is specified to estimate the 'true' concentration values for each pollutant and then such concentration is linked to the health outcomes in a time-series perspective. Through a simulation study we evaluate the model performance and we apply the modelling framework to investigate the effect of six pollutants on cardiovascular mortality in Greater London in 2011-2012.

OPEN ACCESS

Citation: Blangiardo M, Pirani M, Kanapka L, Hansell A, Fuller G (2019) A hierarchical modelling approach to assess multi pollutant effects in time-series studies. PLoS ONE 14(3): e0212565. <https://doi.org/10.1371/journal.pone.0212565>

Editor: Chon-Lin Lee, National Sun Yat-sen University, TAIWAN

Received: August 24, 2018

Accepted: February 5, 2019

Published: March 4, 2019

Copyright: © 2019 Blangiardo et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: All relevant data are within the paper and its Supporting Information files.

Funding: MB and GF acknowledge partial support through the Natural Environment Research Council/MRC funded project: NE/I00789X/1 "Traffic Pollution and health in London". MP and MB acknowledge partial support through the Medical Research Council funded project: MR/M025195/1 "A general framework to adjust for missing confounders in observational studies". The funders had no role in study design, data collection and

1 Introduction

Short-term air pollution studies aim at evaluating the association between the day-to-day variation in ambient air pollution and the day-to-day variation in a health outcome, such as mortality or hospital admissions. Typically this involves a time-series approach using data from a particular geographical area that contains daily counts of mortality or morbidity, pollution and meteorological measurements. To provide a few recent examples, Xie *et al.* [1] used a data set from Beijing containing counts of daily hospital admissions and mortality from ischaemic heart disease (IHD), measurements of fine particulate matter air pollution and potential confounding meteorological variables such as temperature or relative humidity. Using a Poisson model, they reported a significant association between the pollutant and IHD. Other studies using a time-series approach [2], [3], [4] had mixed results in detecting a short-term association between various air pollutants and a mortality outcome.

It is obvious that the air that we breath contains a number of different pollutants; however, due to the high correlation between these, the typical approach used in this field evaluates the

analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

health effect of one pollutant or at most two pollutants at a time, e.g. [5], [6], [7]. Despite this, recently there have been some attempts to move towards a multi pollutant approach that more realistically depicts the complexity of the air pollutant concentrations. For instance, Pirani *et al.* [8] proposed a Dirichlet process (DP) mixture model to cluster the days by their concentration profiles. The model jointly estimates the covariate patterns and the health effect of each cluster. In addition, the use of the DP allows the number of clusters to be determined by the model and the data, providing extreme flexibility.

In a similar perspective Bobb *et al.* [9] proposed a Bayesian Kernel Machine Regression (BKMR). The idea is to include the pollutants in the model using a smooth function h that is represented through a kernel function. The authors focused on Gaussian kernel as it outperformed linear and ridge regression kernels in simulation studies where h has a complex functional form.

Similarly to the DP approach, the focus of BKMR is to correctly identify the concentration-response relationship rather than to identify the effect of each individual pollutant on the outcome. To partially address this, the authors extended their model to include a framework for variable selection, allowing the inclusion only of pollutants which have an impact on the response. However, such approach does not quantify this impact, to determine exactly how much the pollutants are affecting the response.

An alternative approach to deal with multiple pollutants focuses on composite air quality indexes, that summarises the air pollution concentration; this is commonly carried out by most governments and used to inform people of the health risk posed by air quality. For example, in the UK the Daily Air Quality Index (DAQI) is based on the highest pollutant concentration out of the following regulated ones: sulphur dioxide, ozone, particulate matters, nitrogen dioxide. The concentration is then transformed into pre-determined bands (low, moderate, high, very high). Recently DAQI has been used to evaluate the effect of episodes of high air pollution concentration on respiratory conditions [10].

Finazzi *et al.* [11] proposed a more sophisticated approach by using a hierarchical model based on latent variables, that they called dynamic co-regionalization model. This model aggregates the pollutant data over space and addresses problems such as missing data as well as the presence of an unbalanced network, where not all pollutants are measured at every site. The output of this model can then be used to calculate an index, for instance by taking the maximum, such as the DAQI does. The authors applied their model to Scottish air pollution data and used the maximum to calculate a state-wide index over time.

Very recently Huang *et al.* [12] proposed a two-stage spatio-temporal approach to evaluate the effects of two pollutants on respiratory hospital admissions in Scotland. In the first step they estimate annual concentrations from monitoring stations and output from numerical models and then feed forward the estimates and their uncertainty to the second step to assess the health effects. To include the two pollutants in the second stage, avoiding collinearity, they consider the first pollutant and the residual of the second after accounting for the first through a linear regression, making the approach difficult to be extended to more than two pollutants.

Our paper is set in a similar perspective as [12] and we develop a two-components Bayesian hierarchical model that quantifies the health effect of several pollutants. In the first component we account for process variance in the observed air pollution measurements and for correlation among pollutants; based on this we estimate the corresponding latent 'true' concentration values. However, our paper novelty lays on its fully Bayesian framework, as the two components are jointly estimated so that uncertainty from the concentration estimates can feed forward into the health effect estimates; at the same time information from the outcome can feedback to the air pollution estimates. We do rely on the joint estimation process, on the hierarchical nature of the model and on informative priors on the health effect parameters to

overcome the collinearity among the pollutants; this make the framework extendible to any number of pollutants and capable to disentangle synergic or antagonistic effects of pollutants which would be not detectable in the common single-pollutant modelling framework. The developed approach is used to evaluate the effect of five pollutants (carbon monoxide—CO, nitrogen dioxide—NO₂, ozone—O₃, sulphur dioxide—SO₂ and fine particulate matter, smaller in size than 2.5 μ m—PM_{2.5}) and particle number concentration—PCNT on daily cardiovascular mortality in Greater London for 2011–2012.

The remainder of the paper is structured as follows: section 2 presents the data and the model, section 3 introduces the simulation study, while in section 4 we present the results of our analysis and section 5 covers areas of discussion and concluding remarks.

2 Material and methods

2.1 Data description

Daily measurements of CO, NO₂, O₃, SO₂, PM_{2.5} and PCNT were obtained from a monitoring site in North Kensington, London (UK) over the period 1 January 2011 to 31 December 2012. The London North Kensington site (lat 5131015.78000 N, long 012048.57100 W) is part of both the London Air Quality Network and the national Automatic Urban and Rural Network and is owned and part-funded by the Royal Borough of Kensington and Chelsea. The facility is located within a self-contained cabin on a school ground in a mainly residential area. It has been used in previous time-series studies on air pollution health effects [5, 8, 2]. The same monitoring site has also been used extensively as a background measurement site for source apportionment [13] and also to track the outcome of policies to improve London air pollution [14].

CO, NO₂, O₃ and SO₂ were measured using CEN mandated methods eg EN 14211 for NO₂. Fortnightly calibrations enabled the traceability of measurements to national meteorological standards. PM_{2.5} were measured by TEOM-FDMS (Tapered Element Oscillating Microbalance—Filter Dynamics Measurement System) which is considered equivalent to the EU reference method. Particle number concentration was measured by condensation particle counter (TSI 3022).

We focus on these five pollutants as they are already regulated in ambient air; as a result, they are well monitored, have documented associations to health outcomes [15] and have been showed to need National Ambient Air Quality Standards [16]. In addition several papers have focused on one or more of these: for instance Mills *et al.* [17] presented a systematic review of the effects on NO₂ where particulate matter is also controlled for. Besides the five pollutants, we also investigate the effect of PCNT, as this metric was previously associated to adverse short-term health outcome in London [5].

As a health outcome we consider the daily count of mortality due to cardiovascular diseases (CVD) over the same period obtained from the UK Office of National Statistics and available through the Small Area Health Statistics Unit (SAHSU). These cardiovascular causes were derived from the International Statistical Classification of Diseases, 10th Revision (ICD-10, Chapter I).

To adjust for potential confounding effect of weather variables, we use daily average temperature and relative humidity obtained from a meteorological station close to the North Kensington monitoring site. Table 1 provides descriptive statistics of the variables considered in the analysis. The time-series of health counts and pollutant concentrations used for the analyses are available in Table A in S2 File.

Table 1. Descriptive statistics of the variables included in the study.

	Number of Days	Percentiles					IQR
		10th	25th	50th	75th	90th	
Mortality	731	28	32	37	42	47	10
Meteorological data:							
Temperature (°C)	731	5.1	8.0	11.7	15.5	18.1	7.4
Relative Humidity (%)	731	61.6	69.6	78.0	84.2	88.5	14.5
Pollutants:							
CO (mg/m ³)	715	0.1	0.2	0.2	0.3	0.4	0.1
NO ₂ (μg/m ³)	706	18.2	23.2	33.3	46.9	57.9	23.6
O ₃ (μg/m ³)	695	11.4	24.3	39.1	51.1	64.9	26.8
SO ₂ (μg/m ³)	717	0.0	0.4	1.8	2.6	3.6	2.2
PCNT (p/cm ³)	636	7.8	9.7	12.1	14.9	17.9	5.2
PM _{2.5} (μg/m ³)	730	5.0	6.0	9.0	14.0	25.0	8.0

<https://doi.org/10.1371/journal.pone.0212565.t001>

2.2 Model specification

Our modelling framework consists of two components jointly estimated: a pollutant model and a health model, which we describe in details in this section.

2.2.1 Pollutant model. We start specifying Y_{pt} as the measured concentration level of pollutant p ($p = 1, \dots, P = 6$) on day t ($t = 1, \dots, T = 731$) from the monitoring site. As different pollutants are typically characterised by different scales we recommend standardisation to make them comparable. As Y_{pt} is a continuous variable, and after standardisation it can assume any value in $\mathbb{R}$, it is reasonable to assume the following Normal distribution:

$$Y_{pt} \sim N(\mu_{pt}, \sigma_p^2) \quad (1)$$

where σ_p^2 is the process variance, which is specific for each pollutant. On μ_{pt} a linear model is specified as follows:

$$\mu_{pt} = \gamma_{0p} + \gamma_{1p}X_{\text{temp},t} + \gamma_{2p}X_{\text{temp},t}^2 + \gamma_{3p}X_{\text{rh},t} + \gamma_{4p}X_{\text{rh},t}^2 + \theta_{pt} \quad (2)$$

where γ_{0p} is the pollutant specific intercept, while γ_p are the regression coefficients linking the time-dependent covariates X_t to the pollutant levels; as descriptive plots suggest the presence of a non-linear relationship between the covariates and the pollutant concentration levels (see Figs A-C in [S1 File](#)), we include a linear and quadratic effect of both temperature and relative humidity. In (2) $\{\theta_{1t}, \dots, \theta_{pt}\}$ account for the residual temporal effects (see Figs D-I in [S1 File](#) for the time-series of each pollutant) and for the correlation among pollutants; they are modelled following a multivariate Normal specification with an autoregressive structure, following Shaddick *et al.* [18]:

$$(\theta_{1t}, \dots, \theta_{pt})' \sim \text{MVN}\left((\theta_{1,t-\ell}, \dots, \theta_{p,t-\ell})', \Sigma_p\right) \quad (3)$$

where $t - \ell$ provides the temporal lag of ℓ days for the t -th day. For each pollutant, the concentration at time t depends on the values at time $t - \ell$, while the diagonal of the covariance matrix of the errors Σ_p allows each pollutant to have a different amount of temporal dependence, with smaller values indicating weaker dependence. The off-diagonals represent the temporal dependence between the pollutants, allowing the model to incorporate and maintain the correlation structure in the estimation of the ‘true’ pollutant levels.

Note that this specification has the added benefit of providing a natural way to deal with missing data in the pollutant concentration. As seen in Table 1, there are some days where the concentration is not available for one or more pollutants; the model could impute directly the concentration based on i) the correlation with the other pollutants and ii) the temporal dependency. It seems reasonable to assume that the missingness is completely at random and we jointly model (2) and (3) for measured and missing pollutant concentrations. This means that when a measurement is missing, information from the other pollutants for the same day and from the same and other pollutants in different days are driving the imputation through the correlation structure in θ as well as the relationship with the covariates included in (3).

2.2.2 Health model. The second component of the model links the ‘true’ latent value of the pollutant concentrations μ_{pt} with the counts of the health outcome within a time-series epidemiological framework. Let O_t be the observed number of CVD deaths for the day t , we specify a Poisson distribution as:

$$O_t \sim \text{Poisson}(\lambda_t E_t) \quad (4)$$

where E_t represents the expected number of CVD deaths. Following Pirani *et al.* [8], we take it to be the average mortality over the whole period, hence $E_t = E$. Then λ_t represents the relative risk of CVD death on day t compared to the average. In a previous analysis performed on the same data set, Atkinson *et al.* [2] considered the association between CVD and 1-day lagged pollutant concentrations, thus we coherently adopt the same lag. We therefore specify a regression model on the log link transformed λ_t :

$$\log(\lambda_t) = \beta_0 + \sum_p \beta_p \mu_{p(t-1)} + \sum_i s(X_{ti}, \psi_i) + \delta_{it} + \epsilon_t \quad (5)$$

so that $\exp(\beta_p)$ is the multiplicative change in relative risk of CVD death for a unit increase in the pollution concentration obtained from (2). To be able to interpret the health effects on the correct scale, we back transform the pollutant concentration μ_{pt} estimated from (2) to the original scale. In (5) $s(\cdot, \psi_i)$ denote smooth functions of daily average temperature, relative humidity, as well as of calendar time to account for any residual seasonality and long-term trends. These confounding factors are included in the model through flexible nonparametric penalised spline functions [19]. In particular, we consider a mixed model framework and following Crainiceanu *et al.* [20] we specify a low-rank thin plate spline basis over other options, which tends to show a smaller posterior correlation between parameters. By letting X_{ti} be the i -th confounder on day t , we have the following spline representation:

$$s(X_{ti}, \psi_i) = \alpha_i X_{ti} + \sum_{k=1}^{K_i} b_{ki} |X_{ti} - \kappa_{ki}|^3 \quad (6)$$

where $\psi_i = (\alpha_i, b_{1i}, \dots, b_{K_i i})'$ are the regression coefficients, $X_{ti} - \kappa_{ki}$ are the set of basis functions of the cubic spline and K_i is the number of knots for confounder i , with knot locations $\kappa_{1i} < \kappa_{2i} < \dots < \kappa_{K_i i}$. Based on Atkinson *et al.* [2] we select 3 knots for temperature and relative humidity and 6 for time. Additionally, to account for any holiday effect, we include in the model the linear term I_t which classifies the days according to weekday or weekend/holiday. Finally to account for overdispersion, that is typically present when a Poisson distribution is assumed on the data, we include an additional random effect $\epsilon_t \sim N(0, \sigma_\epsilon^2)$.

2.3 Prior specification

The last step in the model specification consists of is the choice of prior distributions. Minimally informative Normal distributions are specified on all the regression coefficients $\gamma_0, \gamma_p, \beta_0, \delta$ and

α , centered on 0 and with a variance equal to 10^3 . Given the high correlation present among the pollutants, we take advantage of an informative prior on β_p ; we choose a $N(0, 0.1)$ covering a range of values on the relative rate scale from 0.82 to 1.22, which is plausible with what has been seen in the literature on cardio-respiratory diseases.

On the standard deviation for the process variance σ_p and for the random effect σ_e , a Uniform prior is specified ranging between 0 and 100, to ensure minimal information.

The covariance matrix Σ_p is given a P -dimensional inverse Wishart prior, $IW(D, d)$, where D is a symmetric and positive-definite scale matrix and d is the degrees of freedom parameter. We follow the specification presented in Lunn *et al.* [21] and to ensure the weakest information we fix $d = P$; as the prior mean for the inverse Wishart is $d^{-1} D$, D is chosen to be d times the prior estimate of the correlation matrix.

We penalise the random coefficients associated to the basis functions, b_{ki} , shrinking them towards zero to avoid over-fitting. This is achieved by specifying a hierarchical prior structure, so that $b_{ki} \sim N(0, \sigma_{b_i}^{-2})$. This latter parameter controls the amount of smoothness and is supplemented with a $\text{Gamma}(a, b)$ prior distribution, where $a = 1; b = 0.001$.

2.4 Implementation and sensitivity analysis

The model is run using a Markov Chain Monte Carlo simulative method in R, through the R2OpenBUGS and Coda packages; we discarded the first 50,000 iterations of the MCMC and retained the following 10,000 to estimate the posterior distribution of the parameters. We considered two chains and checked for convergence of the parameters visually (see Figs J-L in [S1 File](#)) and analytically (evaluating the MC error below 5% of the standard deviation of the posterior estimates as well as the Gelman-Rubin diagnostic tool).

It is important to stress that this is the first paper to consider jointly the pollutant and the health components; this results in uncertainty on μ_{pt} affecting the estimates of the relative risks $\beta_1, \dots, \beta_p$, while at the same time the information from the outcome is fed backwards into the latent concentration values. This is a crucial point as in this way the correlation between the pollutants is naturally accounted for through the hierarchical structure and through the input from the outcome.

To evaluate the robustness of our modelling framework we changed the prior specification of all the parameters where a non informative prior was assumed. In particular, on the regression coefficients (except the β that estimate the health effects of the pollutants on which we keep an informative prior to deal with co-linearity) we specified a Normal distribution centered on zero and with a variance equal to 10^6 , while the process error and random effect variance were set to Inverse Gammas with parameters 1 and 0.001; finally we put a Gamma with parameters $a = b = 0.001$ on the precisions of the random coefficients associated to the basis functions.

A key aspect in air pollution time-series health studies is represented by the inclusion of the smoothing functions for the time-varying confounding factors. Here, we need to adequately control for their potential non-linear confounding effect while retaining sufficient information for estimating the concentration effects. To perform model checking on the knots we also ran the model with 14 knots on time (7 for each year, [22]), and 9 on temperature and humidity, which we think is large enough to account for a high degree of non linearity, while at the same time not leading to oversmoothing. As model selection tool we used the Deviance Information Criterion (DIC, Spiegelhalter *et al.* [23]), one of the suggested methods [24] to choose the degree of smoothness for time-series studies of air pollution and mortality.

3 Simulation study

We carried out a simulation study to evaluate if the proposed modelling framework is able to estimate the relative risk of highly correlated pollutants on a health outcomes.

3.1 Simulation set-up

We simulated mortality and air pollution concentration for 2000 days and considered 6 pollutants. For the sake of simplicity we did not include any confounder factor (e.g. meteorology) in the pollutant or health components of the model. We fixed the correlation among the pollutants to be equivalent to that observed on the time-series data from Greater London:

$$P_p = \begin{bmatrix} 1 & 0.737 & -0.535 & 0.442 & 0.515 & 0.630 \\ & 1 & -0.606 & 0.510 & 0.730 & 0.659 \\ & & 1 & -0.260 & -0.394 & -0.396 \\ & & & 1 & 0.390 & 0.490 \\ & & & & 1 & 0.420 \\ & & & & & 1 \end{bmatrix}. \quad (7)$$

where the order of the pollutants is CO, NO₂, O₃, SO₂, PCNT, PM_{2.5}. The following steps were used to simulate the data on concentration and outcome:

1. Using the above correlation matrix we generated the true pollutant levels assuming an autoregressive structure of order 1, as specified in (3), $\mu_{pt} \sim N(\mu_{p(t-1)}, P_p)$. This represents the gold standard concentration.
2. At the same time we also simulated the measured concentration for the six pollutants, which we assumed centered on the true latent concentration, but with a process variance equal to 0.1 ($Y_{pt} \sim N(\mu_{pt}, 0.1)$).
3. We then simulated the daily number of events for a health outcome using a Poisson distribution, where the mean is given by $E_t \lambda_t$, with E_t fixed to the average of the mortality time-series presented in Section 2.1 and

$$\log(\lambda_t) = 1 + 0.2\mu_{t1} + 0.2\mu_{t2} - 0.2\mu_{t3} + 0\mu_{t4} + 0\mu_{t5} + 0\mu_{t6}$$

so that we are able to assess if the model can capture true effects as well as the lack thereof.


4. We repeated the process 100 times.

We ran our modelling framework, hereafter called “hierarchical two-component model” (*H2Mjoint*), and compared it with a standard Poisson model, here named as “measurement error model” (*ME*), where the true concentration is replaced by the measured one:


$$\begin{aligned} O_t &\sim \text{Poisson}(\lambda_t E_t) \\ \log(\lambda_t) &= \beta_0 + \sum_p \beta_p Y_{p(t-1)} + \sum_i s(X_{it}, \psi_i) + \delta_{it} + \epsilon_t \end{aligned} \quad (8)$$

We used this as benchmark, given that it is the model commonly specified in epidemiological studies to study short-term health effects of air pollution. The graphical representation of ME is presented in Fig 1(b) and shows a direct link between Y_{pt} and λ_t .

The H2Mjoint model jointly estimates the pollutant concentrations and their health effects through a fully specified Bayesian framework. As an additional comparison we specified an alternative model where the two components are fitted separately, called *H2M*: for this model



(a) Hierarchical Two-Component Model



(b) Measurement Error Model

Fig 1. Graphical representation of the modelling frameworks. (a) shows the proposed two component model: the left hand side represents the pollutant component, while the right hand side the health component. The latent concentration for each pollutant and day, μ_{pt} , obtained from the pollutant component enters the health model as predictor. The specification of the link between μ_{pt} and λ_t makes the difference between H2M and H2Mjoint. In the former the uncertainty from μ_{pt} goes forward into the health model, but there is no feedback from O_t ; in the latter the uncertainty goes forward, while at the same time information from the mortality count O_t can influence back μ_{pt} . (b) shows the ME model: the pollutant component is not there and the measured pollutant concentration Y_{pt} is now directly linked to λ_t . For both (a) and (b) the circles denote latent random variables, while the rectangles are observed quantities; single rectangles are random variables, while the double rectangles enter the model as data and are not characterised by a probability distribution i.e. the Y_t in (b).

<https://doi.org/10.1371/journal.pone.0212565.g001>

first (1)–(3) are run and the posterior distribution for the pollution concentration is estimated. The distribution is then fed forward into the health component ((4)–(6)), so that similarly to the H2Mjoint the health effects account for the uncertainty which derives from the estimates of the air pollution concentrations (through γ and most of all θ and Σ_p which model the

Table 2. Results of the simulation study: The table shows the bias, root mean square error (RMSE), 95% credible intervals (CI) width and coverage for the ME, H2M and H2Mjoint. The bias and RMSE are substantially reduced for all the 6 pollutant coefficients using H2Mjoint / H2M. Coverage improves and at the same time width of the 95% credible interval increases, suggesting that the uncertainty is larger for the hierarchical two-component modelling framework, as expected, given that this comes also from the pollutant component. The comparison of H2Mjoint with H2M shows how the influence of the outcome helps reduce the bias, while at the same time the uncertainty does not increase.

	Bias			RMSE		
	ME	H2M	H2Mjoint	ME	H2M	H2Mjoint
β_1	-0.021	-0.007	-0.002	0.003	0.002	0.002
β_2	-0.036	-0.006	0.002	0.004	0.005	0.004
β_3	0.013	0.004	0.000	0.003	0.004	0.002
β_4	0.008	0.003	0.002	0.001	0.002	0.001
β_5	0.021	0.002	-0.001	0.002	0.002	0.002
β_6	0.022	0.002	-0.001	0.002	0.002	0.002
	95% CI width			95% CI coverage		
	ME	H2M	H2Mjoint	ME	H2M	H2Mjoint
β_1	0.16	0.20	0.20	65	92	93
β_2	0.20	0.30	0.30	53	97	97
β_3	0.16	0.17	0.16	71	92	97
β_4	0.13	0.16	0.16	77	98	99
β_5	0.16	0.19	0.22	61	95	94
β_6	0.15	0.18	0.19	65	97	99

<https://doi.org/10.1371/journal.pone.0212565.t002>

dependence in time and across pollutants). At the same time in H2M, the feedback from the health outcome is not allowed to influence the air pollutant concentration estimates. H2Mjoint and H2M are both represented in Fig 1(a): the former has two links between μ_{pt} and λ_t , going each in one direction (uncertainty feeding forward and backwards), while for the latter there is only one link and the arrow only points from μ_{pt} to λ_t as no feedback is allowed.

The model comparison is carried out in terms of bias, root mean square error (RMSE), 95% credible interval (CI) coverage and 95% CI width.

3.2 Simulation results

Table 2 presents the results of the simulation study in terms of the indexes above. It is clear that the hierarchical two-component model framework (H2Mjoint / H2M) outperforms the model that uses measured air pollution concentration (ME, as in (8)); across the six pollutants the bias is reduced by 3 to 11 fold and the coverage of the CI95% is always above 90% (compared to 53-77% for the ME model). In terms of precision, the RMSE is generally smaller for the Bayesian model, indicating better accuracy in the estimates, while the width of the confidence interval is larger, which can be explained by the additional uncertainty included in the concentration estimates and that feeds forward into the health component. This can also be seen in the 95% CI plot for the β coefficients (Fig 2). Comparing H2Mjoint with H2M shows that there is an advantage in allowing for the joint specification of the two model components: in H2Mjoint the bias is smaller, which is clearer when the true effects are different from 0 (β_1 - β_3); at the same time there is no increase in the estimate uncertainty, as the widths of the 95% credible intervals do not change substantially.

4 Real application results

We focused on H2Mjoint to evaluate the effects of the six pollutants (CO, NO₂, O₃, SO₂, PM_{2.5} and PCNT) on daily CVD mortality in Greater London for 2011-2012, as the simulation




Fig 2. 95% posterior credible intervals for β under ME, H2M and H2Mjoint. The H2Ms show smaller levels of uncertainty, as this influence the coefficients from the pollutant estimates as well as from the health model itself. At the same time the ME model shows a larger bias in the estimates, due to the measurement error, while H2Mjoint model shows a median estimate virtually equal to the true values, suggesting how the feedback from the outcome can play a role in reducing the corresponding bias.

<https://doi.org/10.1371/journal.pone.0212565.g002>

showed that allowing for the feedback from the outcome leads to an improvement in the estimates in terms of bias. In addition to the multi pollutant model (H2Mjoint), we ran single pollutant models as a comparison, given that this is the typical approach in the field. The results are robust to the changes in prior and in the number of knots on the spline specification, but as the DIC was smaller for the model with 6 knots on time and 3 on temperature and humidity

Table 3. Posterior mean and 95% credible interval of the percent increase in mortality for an IQR change in pollutant concentration: (left) multi pollutant H2Mjoint model; (centre) single pollutant H2Mjoint model; (right) single pollutant frequentist model (Atkinson *et al.*, 2016). Note that all the pollutants are measured in $\mu\text{g}/\text{m}^3$ except for PCNT which is measured in p/cm^3 and CO which is measured in mg/m^3 .

Pollutant	IQR	Multi Pollutant (H2Mjoint) % Increase (95% CI)		Single Pollutant (H2Mjoint) % Increase (95% CI)		Single Pollutant (Atkinson <i>et al.</i> , 2016) % Increase (95% CI)	
CO	0.10	-1.67	(-4.72, 1.65)	-1.59	(-3.89, 0.84)	-1.47	(-2.94, 0.01)
NO ₂	23.65	9.40	(3.06, 16.03)	-0.25	(-2.90, 2.43)	-1.69	(-3.97, 0.64)
O ₃	26.85	3.46	(0.18, 6.71)	2.61	(0.02, 5.32)	3.31	(0.83, 5.84)
SO ₂	2.20	-1.94	(-6.59, 2.80)	-1.13	(-4.96, 3.15)	-2.33	(-4.18, -0.45)
PCNT	5.18	-2.89	(-6.36, 1.05)	-0.31	(-3.56, 3.35)	*	*
PM _{2.5}	8.00	-1.24	(-3.45, 0.92)	-0.79	(-2.06, 0.47)	-0.9	(-2.09, 0.25)

* CO and PCNT were not analysed in Atkinson *et al.*, 2016.

<https://doi.org/10.1371/journal.pone.0212565.t003>

(9257), thus we present the results of this model specification (which was also the one used in Atkinson *et al.* [2]). For the models with increased number of knots and different prior the results are presented in Tables A and B in S1 File, together with their DIC.

Given the standardisation of the pollutant concentrations, we found that the specification of a Normal distribution was reasonable. Nevertheless, we evaluated the suitability of the logarithm using Hinkley's normality statistics (Table C in S1 File) and found that for O₃ and SO₂ it was detrimental. In addition, we produced residuals plots to check the fit of the model (Fig M in S1 File), and as they are scattered around 0 for all the six pollutants we concluded that modelling on the original scale was reasonable.

As the pollutants are on different scales, to make their effects comparable, in Table 3 we present the results in terms of percent increase for a interquartile range (IQR) change in air pollution concentrations, defined as:

$$\% \text{increase} = (e^{\beta_p \times \text{IQR}} - 1) \times 100\%$$

Out of the six metrics NO₂ and O₃ shows an increased risk of CVD mortality, with credible intervals entirely above 0, suggesting strong evidence of an effect. For the remaining pollutants the point estimates of percent change are slightly below 0, but there is a high degree of uncertainty on the results and the credible intervals include 0. There is high correlation between measured and latent pollutant concentration and, as expected, the latter is slightly less extreme due to shrinkage intrinsic in the modelling framework (see Fig N in S1 File for a plot comparing the posterior mean of μ_{pt} with the measured concentration Y_{pt} for the six pollution metrics). In addition, the process variance on the standardised metrics is presented on Table 4 and shows the lowest values for NO₂ and O₃, both around 0.04, while it is 0.08 for PM_{2.5} and it increases between 0.14 to 0.45 for the remaining pollutants. This suggests that the model is able to account almost entirely for the variability of NO₂, O₃ and PM_{2.5}, while it would potentially point towards some residual confounding for SO₂, CO and PCNT.

By contrast the time-series single pollutant model shows a positive posterior mean with a credible interval above zero only for O₃. For the remaining pollutants the point estimates are negative and their intervals cross 0, pointing towards lack of substantive evidence of an effect. Process variances are larger, with a posterior mean spanning from 0.11 to 0.66, suggesting that the multi pollutant model borrow strength across pollutants to improve the accuracy of the concentration estimates.

Table 4. Posterior mean and 95% credible interval for the process variance σ^2 for the H2Mjoint framework: Multi pollutant model (left) and single pollutant model (right). Note that all the pollutants are measured in $\mu\text{g}/\text{m}^3$ except for PCNT which is measured in p/cm^3 and CO which is measured in mg/m^3 .

Pollutant	Multi Pollutant Model Posterior Mean (95% CI)		Single Pollutant Model Posterior Mean (95% CI)	
CO	0.18	(0.15, 0.22)	0.44	(0.36, 0.52)
NO ₂	0.03	(0.01, 0.05)	0.20	(0.17, 0.27)
O ₃	0.04	(0.02, 0.06)	0.16	(0.12, 0.21)
SO ₂	0.45	(0.29, 0.51)	0.66	(0.55, 0.79)
PCNT	0.14	(0.10, 0.17)	0.59	(0.49, 0.68)
PM _{2.5}	0.08	(0.04, 0.12)	0.11	(0.04, 0.18)

<https://doi.org/10.1371/journal.pone.0212565.t004>

The single pollutant results are in line with Atkinson *et al.* [2], who analysed the same period for Greater London and are reported on the right hand side of Table 3. There is slightly more uncertainty in the H2Mjoint framework, as expected, as the pollutant component contributes to it. This translates into point estimates which are generally closer to zero and wider credible intervals; it is particularly interesting to note how accounting for uncertainty shift SO₂ estimates towards zero, so that the protective effect seen in Atkinson *et al.* [2] disappears.

5 Discussion

In this paper we proposed a fully Bayesian hierarchical model to assess the health effect of multi pollutant concentrations in a time-series perspective, allowing for the integration of uncertainty on the concentration and health components. We deal with the common issue of multi-collinearity among pollutants as the joint hierarchical specification of (2)–(6) allows to i) directly estimate and incorporate the correlation when modelling the ‘true’ latent concentration from the measured ones at the monitoring site; ii) incorporate such correlation in the link between concentration and health. It is worth noting that we are focusing on concentration as opposed to exposure which is a characteristic of individuals; in order to estimate the latter we would need data at the individual level (e.g. personal monitors).

The use of a hierarchical model has been shown to provide stable health effect estimates [25] and it allows to specify an informative prior, which acts as a constraint on the parameter estimates, helping deal with the potential colinearity among the pollutants. In addition, a Bayesian approach naturally accounts for missing data in the estimation process; this means that we were able to use the entire 731 days of the time-series, while other analyses [2] were based on less data points as the days with missing pollutant concentrations were removed. In our case, as we are considering six pollutants at the same time, this would mean removing 125 days as one or more pollutants did not have concentration value recorded. We assumed a completely at random mechanism of missingness, meaning that we do not expect the complete case analysis to be biased. Nevertheless the ability to impute the missing data through the hierarchical framework leads to higher power in detecting the health effects. However, the single pollutants within our Bayesian framework, where the imputation has been performed, but there is no dependency across pollutants, show very similar results to [2]. This seems to suggest that it is the multi pollutant framework to improve the estimation of the health effects, rather than the solely ability to use more data through the imputation.

Note that the meteorological covariates (temperature and humidity) are included in the concentration component as well as in the health component of the model. This was done according to the approach proposed by Cefalu *et al.* [26] who discussed how the covariates

included in the exposure prediction model need to be included as confounders in the epidemiological model to avoid biased results.

A key characteristic of our modelling framework is that it involves a joint specification of the two components (pollution model and health model) in (2)–(6). This ensures that the dependence in the pollution component is maintained in the health component and that all the uncertainty is accounted for in the estimation process, differently from classical two-stage models, where the pollution concentration estimates are considered without the associated uncertainty to evaluate their health effects. At the same time, the joint framework leads to feedback from the health outcome to the pollution estimates and their health effects. This is depicted in the graphical representation, which shows the difference between our joint approach and the standard single pollutant model. In the simulation study, we found that the feedback from the outcome provides additional information to estimate the health effects when these are truly different from 0, while at the same time is not introducing bias in any direction when the true health effects are null. Running our model on simulated data we found that the proposed framework caters well for the process error intrinsic in the observed concentrations and is able to estimate the health effects more accurately than the model which considers the observed concentrations as concentration in the health model.

At the same time, on the real data application, we showed that our multi pollutant model is able to capture the short-term harmful effect of a possible synergic mechanism between NO_2 and O_3 . After adjustment for airborne particles and other regulated gases (i.e. CO and SO_2), we found a positive association between a mixture of these two oxidant gases and cardiovascular mortality. The result is in line with Williams *et al.* [27], which considered a two-pollutant model of NO_2 and O_3 , but for daily counts of all cause mortality for Greater London in 2000–2005, hence characterised by more power due to the longer period and larger numbers. The plausibility of the reported associations is also consistent with atmospheric chemistry findings [28] and toxicological results [29]. Williams *et al.* [27] have looked into combining the two pollutants, due to their high correlation and their complex chemistry; for instance NO_2 is a precursor of O_3 , but also scavenging it, which explains why in the centre of cities the level of NO_2 is higher and the level of O_3 are lower. However a clear drawback would be not to be able to disentangle the effects of the two pollutants, which might act on health through different mechanisms. For this reason we think that our approach is beneficial here, as it has shown the ability to identify and quantify the magnitude of the short-term health effect of the simultaneous concentration of multiple air pollutants, that it is not detectable using a traditional single pollutant model [30]. Therefore, from an air quality management perspective we believe that a multi pollutant approach, such as the one proposed in this study, has a potential for suggesting effective control strategies to reduce adverse effects on human health, since it is able to provide insights on the complex trade-offs between different ambient pollutants.

In this paper we showed how the Bayesian hierarchical modelling framework is advantageous for dealing with multi pollutant concentrations in a purely time-series perspective. A natural extension will consist of increasing the number of measurement sites, moving to a spatio-temporal model. This would allow to account for natural spatial variation which can be particularly strong for some of the pollutants (e.g. NO_2) and should help increase the accuracy of the pollution estimates, hence reducing the process variance. We could adapt the approach proposed by Goldman *et al.* [31] who specified a simulative framework to account for different sources of measurement error and evaluated their impact on health effect estimates. In addition, a spatio-temporal extension would allow to evaluate also chronic effects, which are generally dominated by spatial heterogeneity. The proposed framework is also transferable to evaluate the effects of extreme values (e.g. pollutant concentrations above the WHO

thresholds), which would entail modification of the health component specification to include a threshold and could be modelled assuming a linear or non linear concentration-response.

Supporting information

S1 File. Supplementary information on the paper. It includes a series of figures and tables related to descriptive statistics, convergence of the MCMC simulation and sensitivity analysis results.
(PDF)

S2 File. Table for the data used in the analyses.
(XLSX)

Acknowledgments

We thank Margaret Douglass for the health data extraction. The Small Area Health Statistics Unit has ethics and governance approvals to hold and use mortality data from the National Research Ethics Service—reference 7/LO/0846 and from the Health Research Authority Confidentially Advisory Group, reference HRA—24/CAG/1039. We would like to thank colleagues within the Environmental Research Group at King's College London and at the National Physical Laboratory for providing the pollutant measurements. We would also like to thank the anonymous reviewers whose comments have helped us improve the paper. All authors declare no conflict of interests.

Author Contributions

Conceptualization: Marta Blangiardo.

Data curation: Anna Hansell, Gary Fuller.

Formal analysis: Lauren Kanapka.

Investigation: Gary Fuller.

Methodology: Marta Blangiardo, Monica Pirani, Lauren Kanapka.

Resources: Anna Hansell.

Writing – original draft: Marta Blangiardo.

Writing – review & editing: Marta Blangiardo, Monica Pirani, Anna Hansell, Gary Fuller.

References

1. Xie W, Li G, Zao D, Xie X, Wei Z, Wang W, et al. Relationship between fine particulate air pollution and ischaemic heart disease morbidity and mortality. *Heart*. 2014; 101:257–263. <https://doi.org/10.1136/heartjnl-2014-306165> PMID: 25341536
2. Atkinson R, Analitis A, Samoli E, Fuller G, Green D, Mudway I, et al. Short-term exposure to traffic-related air pollution and daily mortality in London, UK. *J Expo Sci Environ Epidemiol*. 2016; 26:125–132. <https://doi.org/10.1038/jes.2015.65> PMID: 26464095
3. Samoli E, Atkinson R, Analitis A, Fuller G, Green D, Mudway I, et al. Associations of short-term exposure to traffic-related air pollution with cardiovascular and respiratory hospital admissions in London, UK. *Occup Environ Med*. 2016; 75:300–307. <https://doi.org/10.1136/oemed-2015-103136>
4. Dai L, Zanobetti A, Koutrakis P, Schwartz J. Associations of Fine Particulate Matter Species with Mortality in the United States: A Multicity Time-Series Analysis. *Environ Health Perspect*. 2014; 122:837–842. <https://doi.org/10.1289/ehp.1307568>

5. Atkinson R, Fuller G, Anderson R, Harrison R, Armstrong B. Urban ambient particle metrics and health: a time-series analysis. *Epidemiology*. 2010; 21:501–511. <https://doi.org/10.1097/EDE.0b013e3181debc88>
6. Dong H, Yu Y, Yao S, Lu Y, Chen Z, Li G, et al. Acute effects of air pollution on ischaemic stroke onset and deaths: a time-series study in Changzhou, China. *BMJ Open* 2018; 7(8).
7. Luo L, Zhang Y, Jiang J, Luan H, Yu C, Nan P, et al. Short-Term effects of ambient air pollution on hospitalization for respiratory disease in Taiyuan, China: A time-series analysis. *Int J Environ Res Public Health* 2018; 15(10). <https://doi.org/10.3390/ijerph15102160>
8. Pirani M, Best N, Blangiardo M, Liverani S, Atkinson R, Fuller G. Analysing the health effects of simultaneous exposure to physical and chemical properties of airborne particles. *Environ Int*. 2015; 79:56–64. <https://doi.org/10.1016/j.envint.2015.02.010>
9. Bobb J, Valeri L, Henn B, Christiani D, Wright R, Mazumdar M, et al. Bayesian kernel machine regression for estimating the health effects of multi-pollutant mixtures. *Biostatistics*. 2015; 16:493–508. <https://doi.org/10.1093/biostatistics/kxu058> PMID: 25532525
10. Smith G, Bawa Z, Macklin Y, Morbey R, Dobney A, Vordoulakis S, et al. Using real-time syndromic surveillance systems to help explore the acute impact of the air pollution incident of March/April 2014 in England. *Environ Res*. 2015; 136:500–504. <https://doi.org/10.1016/j.envres.2014.09.028>
11. Finazzi F, Scott M, Fassó A. A model-based framework for air quality indices and population risk evaluation, with an application to the analysis of Scottish air quality data. *J R Stat Soc Ser C*. 2013; 62:287–308. <https://doi.org/10.1111/rssc.12001>
12. Huang G, Scott M, Lee D. Multivariate space-time modelling of multiple air pollutants and their health effects accounting for exposure uncertainty. *Stat Med*. 2017; p. 1–15.
13. Beddows D, Harrison R, Green D, Fuller G. Receptor modelling of both particle composition and size distribution from a background site in London, UK. *Atmospheric Chem Phys*. 2015; 15:10107–10125. <https://doi.org/10.5194/acp-15-10107-2015>
14. Font A, Fuller G. Did policies to abate atmospheric emissions from traffic have a positive effect in London? *Environ Pollut*. 2016; 218:463–474.
15. WHO Europe. Review of evidence on health aspects of air pollution. Copenhagen, Denmark: WHO Regional Office for Europe; 2013.
16. Owens E, Patel M, Kirrane E, Long T, Brown J, Cote I, et al. Framework for assessing causality of air pollution-related health effects for reviews of the National Ambient Air Quality Standards. *Regul Toxicol Pharmacol*. 2017; 88:332–337. <https://doi.org/10.1016/j.yrtph.2017.05.014>
17. Mills I, Atkinson R, Anderson H, Maynard R, Strachan D. Distinguishing the associations between daily mortality and hospital admissions and nitrogen dioxide from those of particulate matter: a systematic review and meta-analysis. *BMJ Open*. 2016; 6(7). <https://doi.org/10.1136/bmjopen-2015-010751>
18. Shaddick G, Wakefield J. Modelling daily multivariate pollutant data at multiple sites. *J R Stat Soc Ser C*. 2002; 51:351–372. <https://doi.org/10.1111/1467-9876.00273>
19. Ruppert D, Wand M, Carroll R. Semiparametric Regression. Cambridge University Press; 2003.
20. Crainiceanu C, Rupert D, Wand M. Bayesian Analysis for Penalized Spline Regression Using WinBUGS. *J Stat Softw*. 2005; 14:1–24. <https://doi.org/10.18637/jss.v014.i14>
21. Lunn D, Jackson C, Best N, Thomas A, Spiegelhalter D. The BUGS Book. Chapman & Hall/CRC; 2012.
22. Dominici F, Samet J, Zeger S. Combining evidence on air pollution and daily mortality from the 20 largest US cities: a hierarchical modelling strategy. *J R Stat Soc Ser A*. 2000; 163:263–284. <https://doi.org/10.1111/1467-985X.00170>
23. Spiegelhalter D, Best N, Carlin B, Van Der Linde A. Bayesian measures of model complexity and fit. *J R Stat Soc Ser B*. 2002; 64(4):583–639. <https://doi.org/10.1111/1467-9868.00353>
24. Peng R, Dominici F, Louis T. Model choice in time series studies of air pollution and mortality. *J R Stat Soc Ser A*. 2006; 169(2):179–203. <https://doi.org/10.1111/j.1467-985X.2006.00410.x>
25. MacLehose R, Dunson D, Herring A, Hoppin J. Bayesian Methods for Highly Correlated Exposure Data. *Epidemiology*. 2007; 18:199–207. <https://doi.org/10.1097/01.ede.0000256320.30737.c0>
26. Cefalu M, Dominici F. Does exposure prediction bias health effect estimation? The relationship between confounding adjustment and exposure prediction. *Epidemiology*. 2014; 25:583–590.
27. Williams M, Atkinson R, Anderson R, Kelly F. Associations between daily mortality in London and combined oxidant capacity, ozone and nitrogen dioxide. *Air Qual Atmos Health*. 2014; 7:407–414. <https://doi.org/10.1007/s11869-014-0249-8>
28. Gamon L, White J, Wille U. Oxidative damage of aromatic dipeptides by the environmental oxidants NO₂^{*} and O₃. *Org Biomol Chem*. 2014; 12:8280–8287. <https://doi.org/10.1039/C4OB01577K>

29. Drechsler-Parks D. Cardiac output effects of O₃ and NO₂ exposure in healthy older adults. *Toxicol Ind Health*. 1995; 11:99–109. <https://doi.org/10.1177/074823379501100109>
30. Mauderly JL, Samet JM. Is there evidence for synergy among air pollutants in causing health effects? *Environ Health Perspect*. 2009; 117(1):1. <https://doi.org/10.1289/ehp.11654> PMID: 19165380
31. Goldman GT, Mulholland AM, Russell AG, Strickland MJ, Klein M, Waller LA, et al. Impact of exposure measurement error in air pollution epidemiology: effect of error type in time-series studies. *Environ Health* 2011; 10(61).