This is a repository copy of *Adapting non-parametric spline representations of outdoor air pollution health effects associations for use in public health benefits assessment*.

White Rose Research Online URL for this paper:
https://eprints.whiterose.ac.uk/222405/

Version: Published Version

Adapting non-parametric spline representations of outdoor air pollution health effects associations for use in public health benefits assessment

Richard Burnett[1] · Michael Cork[2] · Neal Fann[3] · Hong Chen[1] · Scott Weichenthal[1,4]

Abstract

The magnitude and shape of the association between outdoor air pollution concentrations and health need to be characterized in order to estimate public health benefits from proposed mitigation strategies. Specialized parametric functions have been proposed for this characterization. However, non-parametric spline models offer more flexibility, less bias, and predictive power, in describing these associations and are thus preferred over relatively simple parametric formulations. Unrestricted spline representations are often reported but many are not suitable for benefits analysis due to their erratic concentration-response behavior and are usually not presented in a format consistent with the requirements necessary to conduct a benefits analysis. We propose a method to adapt non-parametric spline representations of concentration-response associations that are suitable for public health benefits analysis by transforming spline predictions and its uncertainty over the study exposure range to a new spline formulation that is both monotonically increasing and restricted to concentration-response patterns suitable for use in health benefits assessment. We selected two examples of the association between long-term exposure to fine particulate matter and mortality in Canada and the USA that displayed spline fits that were neither monotonically increasing nor suitable, we suggest, for benefits analysis. We suggest our model is suitable for benefits analysis and conduct such analyses for both Canada and the USA, comparing benefits estimates to traditional models. Finally, we provide guidance on how to report spline fitting results such they can be used either in benefits analysis directly, or to fit our new model.

Keywords Air pollution health effects · concentration-response models

Introduction

Outdoor air quality has been recognized as a leading cause of disease and death globally (GBD 2019). As such, strategies to improve air quality have been proposed (Johnston et al. 2012; Liang et al. 2018; Gu et al. 2018; Lelieveld et al. 2019). These strategies are often evaluated in terms of their societal monetary costs and extent to which they improve

public health (World Bank 2016; Heo et al. 2017; Muller 2018). Public health improvements are often measured by the number of attributable adverse health events (i.e., incidence of disease or death) associated with a reduction in exposure to outdoor concentrations of pollution over a fixed time period (i.e., year). We can quantitatively represent the attributable number, A N, by the following equation:

$$\mathrm{A}N(z_C, z_F) = P \times \mathrm{BR} \times \mathrm{PAF}_\beta(z_C, z_F)$$

where z_C is the current air pollution concentration, z_F the future concentration predicted under a specific mitigation strategy, P the target population size for which the concentration applies, BR the baseline rate of disease or death, and $\mathrm{PAF}_\beta(z_C, z_F)$, the population attributable fraction defined by the following equation:

$$\mathrm{PAF}_\beta(z_C, z_F) = 1 - \frac{1}{R_\beta(z_C, z_F)}$$

✉ Richard Burnett
  rtburnett1@gmail.com

1 Health Canada, Ottawa, Canada

2 Department of Biostatistics, Harvard T.H. Chan School of Public Health, Boston, MA, USA

3 Office of Air Quality Planning and Standards, U.S. Environmental Protection Agency, Research Triangle Park, Durham, NC, USA

4 McGill University, Montreal, Canada

with $R_\beta(z_C, z_F)$ the relative risk between z_C and z_F, indexed by a vector of parameters β. Here, $R_\beta(z_C, z_F)$ represents the ratio of the probability of an adverse event over a fixed time period for a population which is exposed to z_C to the corresponding probability if that population was contrary to the fact exposed to z_F.

The magnitude, shape, and uncertainty in $R_\beta(z_C, z_F)$ need to be characterized in order to conduct a health benefits analysis. The relative risk function is usually defined with respect to a counterfactual concentration, z_{cf} say, denoted by $R_\beta(z_C|z_{cf})$ with $R_\beta(z_{cf}|z_{cf}) = 1$. We then have

$$R_\beta(z_C, z_F) = \frac{R_\beta\left(z_C\middle|z_{cf}\right)}{R_\beta\left(z_F\middle|z_{cf}\right)}$$

Uncertainty in $R_\beta(z_C|z_{cf})$ is induced by uncertainty in the estimate of β given the algebraic form of how R varies with z. Let $\hat{\beta}$ be the estimate of β with distribution $F_{\hat{\beta}}$. Then the uncertainty distribution of $R_{\hat{\beta}}(z_C, z_F)$ is obtained by generating a (usually) large number, N, of realizations of $\hat{\beta}$ from $F_{\hat{\beta}}$, denoted by $\{\hat{\beta}^{(i)}, i = 1, \ldots, N\}$, thus yielding N estimates

$$R_{\hat{\beta}^{(i)}}(z_C, z_F) = \frac{R_{\hat{\beta}^{(i)}}\left(z_C\middle|z_{cf}\right)}{R_{\hat{\beta}^{(i)}}\left(z_F\middle|z_{cf}\right)}$$

for $i = 1, \ldots, N$. These N estimates of relative risk are used to calculate N estimates of the PAF. Uncertainty in the estimate of attributable deaths can be determined by the joint uncertainty in P, BR, and PAF (GBD 2019).

Previous relative risk models

The most common form of $R_\beta(z_C, z_F)$ is the following equation

$$R_\beta(z_C, z_F) = \frac{R_\beta\left(z_C\middle|z_{cf}\right)}{R_\beta\left(z_F\middle|z_{cf}\right)} = \frac{\exp\{\beta z_C - \beta z_{cf}\}}{\exp\{\beta z_F - \beta z_{cf}\}} = \exp\{\beta(z_C - z_F)\}$$

for scalar parameter β. This characterization is often termed the Log − Linear model (Cohen et al. 2004) since the logarithm of the relative risk function is linear in concentration. We note that under the Log − Linear model, $R_\beta(z_C, z_F)$ is independent of the counterfactual concentration z_{cf} and only depends on the magnitude of the difference in concentrations $z_C - z_F$, but not the magnitude of each concentration.

The estimate of β is usually obtained from the analysis of primary health data with a normal uncertainty distribution, $N(\hat{\beta}, \hat{\sigma}_{\hat{\beta}})$, where $\hat{\sigma}_{\hat{\beta}}$ is the standard error of $\hat{\beta}$. In the case of the Log − Linear model, we have

$$\beta(z_C - z_F) \sim N(\hat{\beta}(z_C - z_F), \hat{\sigma}_{\hat{\beta}}(z_C - z_F))$$

since

$$\sqrt{\text{var}\left(\hat{\beta}z_C - \hat{\beta}z_F\right)} = \sqrt{\text{var}\left(\hat{\beta}z_C\right) + \text{var}\left(\hat{\beta}z_F\right) - 2\,\text{cov}\left(\hat{\beta}z_C, \hat{\beta}z_F\right)} =$$

$$\sqrt{\hat{\sigma}_{\hat{\beta}}^2 z_C^2 + \hat{\sigma}_{\hat{\beta}}^2 z_F^2 - 2\hat{\sigma}_{\hat{\beta}}^2 z_C z_F} = \hat{\sigma}_{\hat{\beta}}(z_C - z_F)$$

The uncertainty distribution in $R_\beta(z_C, z_F)$ is completely characterized by $N(\hat{\beta}, \hat{\sigma}_{\hat{\beta}})$ and the concentrations z_C and z_F.

In recent years, interest has focused on the shape of the association between air pollution exposure and health, specifically assessing evidence of departures from linearity. An early attempt (Cohen et al. 2004) was based on the Log − Log model of the form:

$$\ln R_\beta(z, z_{cf}) = \beta\left(\ln(z+1) - \ln(z_{cf}+1)\right) = \beta\ln\left(\frac{z+1}{z_{cf}+1}\right)$$

where the logarithm of the relative risk varied with the logarithm of concentration. Note that the relative risk function is dependent on the counterfactual concentration, a departure from the Log − Linear model.

These models were extended to a broader family of shapes by the Shape Constrained Health Impact function (SCHIF) (Nasari et al. 2016) which has the form:

$$\ln R_\beta(z, z_{cf}) = \theta\frac{f(z - z_{cf})}{1 + \exp\left(-(z - z_{cf} - \mu)/\tau\right)}$$

where f is either a linear or logarithmic function of concentration. The SCHIF can take near-linear, sub, and supra-linear forms in addition to sigmodal shapes. Ensemble modelling methods were used to estimate the unknown parameters (θ, μ, τ) and form of f (Nasari et al. 2016). Like the Log − Log model, the SCHIF is also dependent on the counterfactual concentration.

The Log − Linear, Log − Log, and SCHIF were intended to be used on primary health data. However, interest has also focused on combining information from selected studies to form a common relative risk model within a meta-regression framework. Log − Linear models have been used for this purpose (Chen and Hoek 2020) in addition to non-linear models, including the integrated exposure-response (IER) (Burnett et al. 2014) with the form

$$\ln R_\beta(z, z_{cf}) = \ln\left(1 + \alpha\left(1 - \exp\left(-\beta\tilde{z}^\rho\right)\right)\right)$$

the Global Exposure Mortality Model (GEMM) (Burnett et al. 2018), that generalized the SCHIF, with the form

$$\ln R_\beta(z, z_{cf}) = \theta \frac{\ln\left(\tilde{z}/\alpha + 1\right)}{1 + \exp\left(-\left(\tilde{z} - \mu\right)/\tau\right)}$$

and the Fusion model (Burnett et al. 2022)

$$\ln R_\beta(z, z_{cf}) = \gamma \times \left(\min\left(\tilde{z}, \mu\right) + \mathcal{F}\left(\tilde{z}|\rho, \mu, \theta\right) + \rho\theta \ln\left(\max\left(\tilde{z}, \theta\right)/\theta\right)\right)$$

where

$$\mathcal{F}\left(\tilde{z}|\rho, \mu, \theta\right) = \int_\mu^{min(\theta, \tilde{z})} \left(1 + \frac{1-\rho}{\rho}\left(\frac{x-\mu}{\theta-\mu}\right)^{\frac{\theta-\mu}{\theta(1-\rho)}}\right)^{-1} dx$$

with $\tilde{z} = z - z_{cf}$. These models are based on a small number of unknown parameters (three for the IER, and four for both the GEMM and Fusion). All three models report 1000 sets of parameter estimates, in addition to 1000 sets of relative risk predictions over the concentration range, thus enabling calculation of 1000 PAFs based on any two concentrations (Burnett et al. 2014:2018:2022). This PAF uncertainty distribution can be used in benefits assessment.

Most researchers now use non-parametric spline methods to both characterize the shape and magnitude of the concentration-response function, either based on primary data (Brunekreef et al. 2021; Brauer et al. 2022; Dominici et al. 2022), or meta-regression combining information from multiple primary data studies (GBD 2019), in order to gain additional flexibility compared to parametric models with a few parameters, such as the IER, GEMM, and Fusion models. Splines have the general form

$$s_\beta(z) = \sum_{j=1}^J \beta_j b_j(z|\kappa)$$

where the basis functions $b_j(z|k_1, \ldots, k_K)$ are transformations of concentration specified by K knot locations: $\kappa = (k_1, \ldots, k_K)$. The most common forms of $b_j(z|\kappa)$ are cubic polynomials, such as those used to define natural (Bartels et al. 1998) or restricted (Harrel 2015) cubic splines such that the second derivative of the spline is continuous at the knot values. They typically involve a limited number of terms, J, with the number of terms selected based on fitting criteria. These regression splines have been extended to include a penalty term based on the integral of the second derivative of the spline that smooths the estimates of the β_j. Smoothing splines (Eilers and Marx 1996) typically have many basis functions and estimates of the β_j are a function of the penalty parameter, with larger values inducing more smoothing on these estimates.

For spline representations, the quantity of interest in benefits assessment is as follows:

$$\ln R_{\hat{\beta}}(z_C, z_F) = \left(\hat{s}_{\hat{\beta}}(z_C) - \hat{s}_{\hat{\beta}}(z_{cf})\right)$$
$$- \left(\hat{s}_{\hat{\beta}}(z_F) - \hat{s}_{\hat{\beta}}(z_{cf})\right) = \hat{s}_{\hat{\beta}}(z_C)$$
$$- \hat{s}_{\hat{\beta}}(z_F) = \sum_{j=1}^J \hat{\beta}_j\left(b_j(z_C|\kappa) - b_j(z_F|\kappa)\right)$$

independent of the counterfactual concentration z_{cf}. Researchers typically report the results of fitting splines with the exponential of the mean spline prediction over the concentration range and the 95% confidence intervals of those exponentiated predictions (Brunekreef et al. 2021; Brauer et al. 2022; Dominici et al. 2022). The exponential of the mean spline prediction is shifted such that the prediction at the minimum concentration is unity.

The uncertainty in $\hat{s}_{\hat{\beta}}(z_C) - \hat{s}_{\hat{\beta}}(z_F)$ is given by the following:

$$\text{var}\left\{\hat{s}_{\hat{\beta}}(z_C) - \hat{s}_{\hat{\beta}}(z_F)\right\} = \left(b(z_C|\kappa) - b(z_F|\kappa)\right)'\hat{V}\left(b(z_C|\kappa) - b(z_F|\kappa)\right)$$

where $b(z|\kappa) = (b_1(z|\kappa), \ldots, b_J(z|\kappa))'$ and $\hat{V}$ is the estimated covariance matrix of $\hat{\beta} = (\hat{\beta}_1, \ldots, \hat{\beta}_J)'$, the estimates of the spline parameters.

We note that

$$\text{var}\left\{\hat{s}_{\hat{\beta}}(z_C) - \hat{s}_{\hat{\beta}}(z_F)\right\} = \text{var}\left(\hat{s}_{\hat{\beta}}(z_C)\right) + \text{var}\left(\hat{s}_{\hat{\beta}}(z_F)\right)$$
$$- 2\,\text{cov}\left(\hat{s}_{\hat{\beta}}(z_C), \hat{s}_{\hat{\beta}}(z_F)\right)$$

cannot be determined from the mean spline predictions over the concentration range and their associated variance since $\text{cov}\left(\hat{s}_{\hat{\beta}}(z_C), \hat{s}_{\hat{\beta}}(z_F)\right)$ cannot be calculated from these summary results. This is true even for the case where burden assessments are of interest. Here, all future concentrations are set to a counterfactual. However, the uncertainty in the spline prediction at the counterfactual is often positive.

One can form relative risk predictions such that $\text{var}\left(\hat{s}_{\hat{\beta}}(z_{cf})\right) = 0$ and thus $\text{cov}\left(\hat{s}_{\hat{\beta}}(z_C), \hat{s}_{\hat{\beta}}(z_{cf})\right) = 0$. We then have: $\text{var}\left\{\hat{s}_{\hat{\beta}}(z_C) - \hat{s}_{\hat{\beta}}(z_{cf})\right\} = \text{var}\left(\hat{s}_{\hat{\beta}}(z_C)\right)$, that can be determined from the spline prediction confidence intervals given knowledge of the distribution of $\hat{s}_{\hat{\beta}}(z_C)$. Unfortunately, uncertainty in the difference between spline predictions at any two concentrations greater than the counterfactual cannot be determined alone from the uncertainty in spline predictions at the two concentrations even if $\text{var}\left(\hat{s}_{\hat{\beta}}(z_{cf})\right) = 0$.

We conclude from these observations that the spline parameter estimates and their covariance matrix should be reported in addition to knot locations and type of spline used. Alternatively, many spline curves could be reported based on random draws of the spline parameters.

Spline representations often describe the association between concentration and response in a complex manner, including steep changes or waviness in the association over narrow concentration intervals. These types of associations may be of questionable biological plausibility and thus are

not desirable features in a benefits assessment. An extension of the SCHIF (eSCHIF) was proposed to transform spline predictions into a form thought to be more suitable for benefits assessment (Brauer et al. 2022). Here

$$\text{lneSCHIF}\left(\tilde{z}\right) = \theta_0 \ln\left(\frac{\tilde{z}}{\alpha_0} + 1\right) + \theta_1 \frac{\ln\left(\frac{\tilde{z}}{\alpha_1} + 1\right)}{1 + \exp\left(-\frac{\tilde{z} - \mu}{\tau}\right)}$$

This algebraic function was fit to each of the 1000 series of spline predictions over the concentration range 0 to 20 μg/m³ (Brauer et al. 2022). Restrictions on the parameters $(\alpha_0, \alpha_1, \tau)$ were imposed in order to smooth each series of spline predictions. However, no restrictions were made on (θ_0, θ_1), resulting in fitted curves that were not necessarily monotonically increasing.

In this paper, we propose a new approach to using spline predictions of concentration-response associations that are monotonically increasing in the mean predictions over concentration, which we suggest is an additional desirable feature for benefits assessment. We allow the uncertainty in these predictions to possibly include the null or negative associations at any concentration, a property like the case if the original study data were in fact fit with a monotonically increasing spline (Pya and Wood 2015).

Monotonically increasing shape constrained health impact functions (mSCHIF)

Our approach is to generalize the eSCHIF to be able to characterize a wider variety of shapes and constrain the fit to be monotonically increasing. Under such constraints, the logarithmic function may not be flexible enough to model all the possible desirable shapes predicted by splines. In addition, using two terms may also not be flexible enough to model all possible spline predictions and their associated uncertainty under the additional constraint of monotonicity. We thus propose a new spline formulation defined by the following:

$$\text{mSCHIF}\left(\tilde{z}\right) = \exp\left\{ \sum_{l=1}^{L} \theta_l \mathcal{M}\left(\tilde{z}|\alpha_l, \lambda_l, \mu_l, \tau_l\right) \right\}$$

w i t h t e r m s :
$\mathcal{M}\left(\tilde{z}|\alpha_l, \lambda_l, \mu_l, \tau_l\right) = \mathcal{F}\left(\tilde{z}|\alpha_l, \lambda_l\right)\mathcal{L}\left(\tilde{z}|\mu_l, \tau_l\right), l = 1, \ldots, L$, composed of the product of two functions

$$\mathcal{F}\left(\tilde{z}|\alpha, \lambda\right) = \int_0^{\tilde{z}} \left(1 + \left(\frac{x}{\alpha}\right)^\lambda\right)^{-1} \partial x$$

and

$$\mathcal{L}\left(\tilde{z}|\mu, \tau\right) = \left(1 + \exp\left(-\frac{\tilde{z} - \mu}{\tau}\right)\right)^{-1}$$

with $\theta_l \geq 0$, $0 < \alpha_l, \mu_l, \tau_l \leq r$, and $0 \leq \lambda_l \leq r$, $(l = 1, \ldots, L)$, where r is the range in concentration.

$\mathcal{F}$ is a monotonically increasing supra-linear function while $\mathcal{L}$ is a two-parameter logistic function taking either sub-linear or sigmodal shapes over the concentration range depending on the values of μ and τ. When $\lambda = 0$, $\mathcal{F}$ is linear in concentration, when $\lambda = 1$, $\mathcal{F}$ is log-linear in concentration, and when $\lambda = 2$, $\mathcal{F}$ is an arctangent function. For all other values of λ, an explicit form of $\mathcal{F}$ is not available; thus, numerical integration is required.

$\mathcal{F}$ has the property that it is nearly linear when $\tilde{z} < \alpha$, with this approximation improving as λ increases. The decline in the derivative of $\mathcal{F}$ increases as λ increases for $\tilde{z} > \alpha$, resulting in a relatively flat functional form for $\mathcal{F}$ over this concentration range. The smaller value of τ, the less increase in $\mathcal{L}$ for $\tilde{z} < \mu$. These properties of $\mathcal{F}$ and $\mathcal{L}$ facilitate adding mSCHIF terms together under the constraint $(\theta_l \geq 0, l = 0, \ldots, L)$. We view the mSCHIF as a spline-type formulation with $\mathcal{F}\left(\tilde{z}|\alpha_l, \lambda_l\right)\mathcal{L}\left(\tilde{z}|\mu_l, \tau_l\right)$ as their basis functions. They can take on a wide variety of shapes depending on the values of their associated parameters, a property not shared by traditional spline models.

mSCHIF parameter estimation and inference

We estimate the mSCHIF parameters by non-linear regression (nlxb routine in R package: nlsr 2020) where the response is the mean spline predictions, $\hat{s}_{\hat{\beta}}(z_i) - \hat{s}_{\hat{\beta}}(z_{cf})$, $i = 1, \ldots, I$, over the concentration range, with z_{cf} representing the minimum concentration of interest. We estimate the $L \times L$ covariance matrix, Σ, of $(\hat{\theta}_l, l = 1, \ldots, L)$ by equating the covariance of the spline predictions across concentrations to the covariance of the mSCHIF predictions. That is:

$$X'\hat{V}X = T'\Sigma T,$$

where X is a $J \times I$ matrix consisting of the J spline bases functions evaluated at the I , and T is a $L \times I$ matrix consisting of the L mSCHIF bases functions evaluated at the I concentrations. We then have the following:

$$\hat{\Sigma} = \left(TT'\right)^{-1}TX'\hat{V}XT'\left(TT'\right)^{-1}$$

The mSCHIF predictions will in most cases not fit the spline predictions perfectly. We incorporate this uncertainty by adding the squared difference between the two model predictions at each concentration to the covariance of the spline predictions. We then have the following:

$$\hat{\Sigma} = \left(TT'\right)^{-1} T \left(X'\hat{V}X + \hat{D}\right) T' \left(TT'\right)^{-1}$$

where $\hat{D}$ is a $I \times I$ diagonal matrix of the squared differences in predictions among the two models.

We assign all the uncertainty in the mSCHIF predictions to L parameters, $\hat{\theta} = \left(\hat{\theta}_l, l = 1, \ldots, L\right)$, assuming that $\left(\hat{\alpha}_l, \hat{\lambda}_{l1}, \hat{\mu}_l, \hat{\tau}_l, l = 1, \ldots, L\right)$ are known without error. However, we are capturing the total amount of uncertainty characterized by the spline by equating the spline prediction uncertainty to that of the mSCHIF. We assume the difference between lnmSCHIF predictions at z_C and z_F is normally distributed with mean

$$\hat{m}\left(\tilde{z}_C, \tilde{z}_F\right) = \left(T(\tilde{z}_C) - T(\tilde{z}_F)\right)' \hat{\theta}$$

and standard deviation

$$\hat{\sigma}\left(\tilde{z}_C, \tilde{z}_F\right) = \sqrt{\left(T(\tilde{z}_C) - T(\tilde{z}_F)\right)' \hat{\Sigma} \left(T(\tilde{z}_C) - T(\tilde{z}_F)\right)},$$

with $\tilde{z}_C = z_C - z_{cf}$ and $\tilde{z}_F = z_F - z_{cf}$.

In some cases, reporting a large number, N, of spline predictions over the concentration range is more convenient than reporting a potentially very large number of bases functions and parameter estimates, as would be the case for some smoothing splines. We first subtract the spline prediction at the counterfactual for each of the N series of predictions, resulting in a prediction of zero for each of the N series. We do this since the lnmSCHIF prediction at the counterfactual is zero with zero uncertainty. We then replace $X'\hat{V}X$ by the empirical covariance among the N predictions for all pairs of concentrations.

Illustrative examples

We provide two examples to illustrate some of the properties of our method. Both examples examine the association between mortality and exposure to outdoor concentrations of fine particulate matter ($PM_{2.5}$) based on primary data from the Medicare cohort (Dominici et al. 2022) and the 2006 Canadian Census Health and Environment Cohort (CanCHEC) (Chen et al. 2022). In both examples, annual $PM_{2.5}$ estimates were based on the same geostatistical model that incorporated satellite retrievals, chemical transport model estimates, and ground data (Meng et al. 2019).

Medicare cohort

The Medicare cohort consists of subjects enrolled in the Medicare program in the USA over the age of 65 years (Dominici et al. 2022). For this example, the cohort includes over 74 million enrollees entered from 2000 to 2016. Subject-level information included age, sex, race/ethnicity, and Medicaid eligibility (measure of lower income). Zip code area (ZIP) measures include median household income, median house value, proportion of residents in poverty, proportion of residents that own their own house, and proportion of residents with a high school diploma. The databases include the fact and date of death from any cause and the time-varying ZIP code of the enrollees mailing address. Subjects were followed until their date of death or until December 31, 2016, the end of the study period. A Poisson regression model was used with the response defined by the count of all-cause deaths during each follow-up year, calendar year, ZIP code, age (5-year groupings), sex, and race; offset by the corresponding total person-time under observation. Annual $PM_{2.5}$ concentrations were assigned to the same year of follow-up. A smoothing spline was used to characterize the magnitude and shape of the $PM_{2.5}$-mortality association with the smoothing parameter estimated from the data controlling for the five ZIP code area variables as linear terms in addition to the four category Census Region variable.

A smoothing spline was fit within the Generalized Additive Model framework using the R package "gam" and the $s(PM_{2.5}, bs = "tp")$ specification with the smoothing parameter estimated from the data by cross-validation. We extracted 1000 draws of the smoothing spline parameters and generated 1000 sets of spline predictions over the concentration range. We then calculated the empirical covariance matrix based on the spline predictions at all possible pairs of concentrations among the 1000 sets of predictions and used this matrix to estimate the covariance matrix Σ.

The mean spline predictions over the 1000 sets of predictions from the minimum concentration (0 µg/m³) to the 99th percentile (17.4 µg/m³) (solid red line, Fig. 1, panel A) suggest a complex association between $PM_{2.5}$ and non-accidental mortality. A near linear increase over the $0-3$ µg/m³ range, flattening of the association over $3-6$ µg/m³ range, sharply increasing over the $6-9$ µg/m³ range, and finally a less steep increase above 9 µg/m³. The smoothing spline mean predictions display several waves over the concentration range that are not present in the mSCHIF mean predictions (solid blue line – panel A). The lack of monotonically increasing spline predictions and the presence of waves makes it less suitable for benefits analysis. The mSCHIF 95% confidence intervals (gray shaded area) are of similar magnitude to the spline confidence intervals (dashed red lines).

The two mSCHIF terms, $\exp\left(\hat{\theta}_0 \mathcal{F}\left(\tilde{z}|\hat{\alpha}_1, \hat{\lambda}_1\right) \mathcal{L}\left(\tilde{z}|\hat{\mu}_1, \hat{\tau}_1\right)\right)$ (solid red line) and $\exp\left(\hat{\theta}_1 \mathcal{F}\left(\tilde{z}|\hat{\alpha}_2, \hat{\lambda}_2\right) \mathcal{L}\left(\tilde{z}|\hat{\mu}_2, \hat{\tau}_2\right)\right)$ (solid black line) are plotted against $PM_{2.5}$ in Fig. 1 (panel B). A linear increase is observed for $\exp\left(\hat{\theta}_1 \mathcal{F}\left(\tilde{z}|\hat{\alpha}_1, \hat{\lambda}_1\right) \mathcal{L}\left(\tilde{z}|\hat{\mu}_1, \hat{\tau}_1\right)\right)$ from 0 to 2 µg/m³ with another linear increase with smaller slope for concentrations

Fig. 1 Smoothing spline mean $PM_{2.5}$ predictions over concentration range 0 to 17.4 $\mu g/m^3$ (solid red line; panel **A**) for the Medicare cohort with corresponding 95% confidence intervals (dashed red lines) in addition to mSCHIF mean predictions (solid blue line) and 95% confidence intervals (gray shaded area), mSCHIF terms $(\exp\left(\hat{\theta}_1 \mathcal{M}\left(\tilde{z}|\hat{\alpha}_1, \hat{\lambda}_1, \hat{\mu}_1, \hat{\tau}_1\right)\right)$ —solid red line; $\exp\left(\hat{\theta}_2 \mathcal{M}\left(\tilde{z}|\hat{\alpha}_2, \hat{\lambda}_2, \hat{\mu}_2, \hat{\tau}_2\right)\right)$ —solid black line; $\exp\left(\hat{\theta}_1 \mathcal{M}\left(\tilde{z}|\hat{\alpha}_1, \hat{\lambda}_1, \hat{\mu}_1, \hat{\tau}_1\right) + \hat{\theta}_2 \mathcal{M}\left(\tilde{z}|\hat{\alpha}_2, \hat{\lambda}_2, \hat{\mu}_2, \hat{\tau}_2\right)\right)$ —dashed blue line) are presented in panel **B**
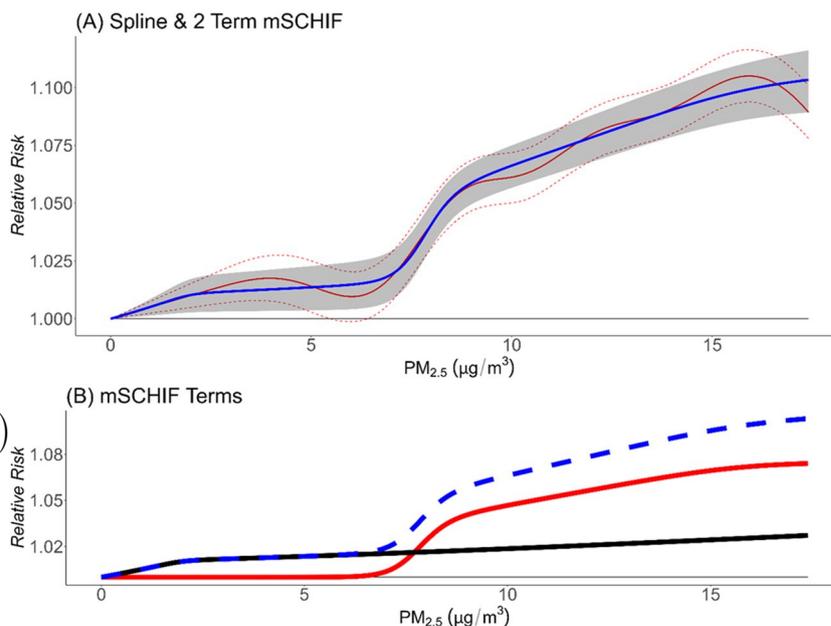
greater than 2 $\mu g/m^3$. The second mSCHIF term, $\exp\left(\hat{\theta}_2 \mathcal{F}\left(\tilde{z}|\hat{\alpha}_2, \hat{\lambda}_2\right)\mathcal{L}\left(\tilde{z}|\hat{\mu}_2, \hat{\tau}_2\right)\right)$, displays a sigmodal shape with inflection point $\hat{\mu}_2 = 7.76$ $\mu g/m^3$.

We compare the mSCHIF (solid blue line) and eSCHIF (solid brown line) predictions for this example as both models consist of two terms (Figure S1). The eSCHIF predictions are not monotonically increasing with concentration. The mSCHIF predictions more accurately reflected the mean spline predictions than the eSCHIF predictions in this example. This is due to the added flexibility of the $\mathcal{F}$ function incorporated into the mSCHIF compared to the logarithmic function used by the eSCHIF.

2006 Canadian Census Environment and Health Cohort (*CanCHEC*)

The 2006 CanCHEC (Chen et al. 2022) consists of 2,663,645 non-institutionalized subjects aged 30–79 years who completed the 2006 Canadian census long-form. Subject-level information included age, sex, education, income, occupation, marital status, and visible minority status. Area level contextual information included multiple measures of social and economic status, in addition to community size and region of Canada. Home addresses were annually identified by linkage to income tax records. Each subject was followed over time from time from January 1, 2007 to December 31, 2016 to determine their vital status by linkage to mortality records, resulting in 25,730,790 person-years of observation. Cox proportional hazards models were used to relate time-varying air pollution exposures, based on a 3-year moving average lagged one prior to follow-up year, to survival from

all non-accidental causes of death (213,882), adjusting for subject and area level mortality risk factors. A restricted cubic spline (Harrel 2015) (RCS) was used to characterize the magnitude and shape of the $PM_{2.5}$-mortality association based on minimizing the Akaike information criterion (Akaike 1974), resulting in the selection of 10 knots.

We illustrate a modelling feature of the mSCHIF by selectively characterizing specific patterns in the mean spline predictions and associated uncertainty. The RCS mean predictions from the minimum concentration (2.5 $\mu g/m^3$) to the 99th percentile (11.9 $\mu g/m^3$) (solid red line, Fig. 2, panel A1) suggest a complex association between $PM_{2.5}$ and non-accidental mortality with a sub-linear increase over the $2.5 - 5$ $\mu g/m^3$ range, a wavey pattern in the association with a slight increase in magnitude over the $5 - 10$ $\mu g/m^3$ range, followed by a much more rapid increase above 10 $\mu g/m^3$. We would like to capture the sub-linear increase between 2.5 and 5 $\mu g/m^3$, replace the wavey pattern between 5 and 10 $\mu g/m^3$ with a linear increase, and finally closely model the rapid increase in risk above 10 $\mu g/m^3$. We can model these features by limiting the allowable range in the μ_l mSCHIF parameters.

We start with two mSCHIF terms with the following parameter limitations: $0 < \mu_1 < 2.5$ and $2.5 < \mu_2 < r$. Note that the μ_l values are defined with respect to $PM_{2.5} - 2.5$ $\mu g/m^3$. We select 2.5 $\mu g/m^3$ for the upper bound on μ_1 since the logit function $\mathcal{L}\left(\tilde{z}|\mu_1, \tau_1\right)$ is sub-linear for concentrations below μ_1 and supra-linear above. The resulting mean mSCHIF predictions are displayed in Fig. 2 (panel A1 – solid blue line), yielding a concentration-mortality pattern near the one we are designing for. However, the two-term mSCHIF
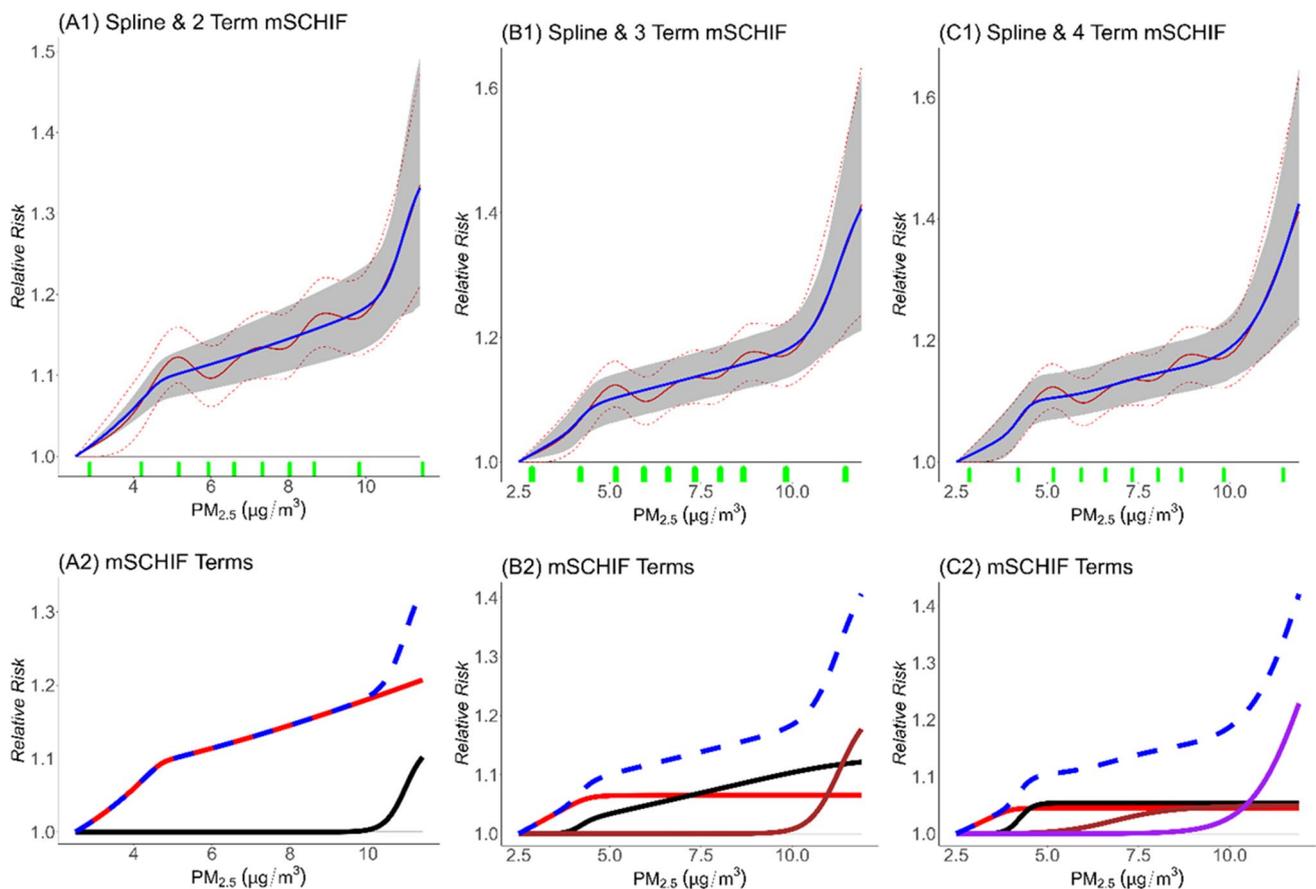
Fig. 2 Mean restricted cubic spline PM$_{2.5}$ predictions over concentration range 2.5 to 11.9 μg/m^3 (solid red line) for the 2006 CanCHEC cohort with corresponding 95% confidence intervals (dashed red lines) in addition to mSCHIF mean predictions (solid blue line) and 95% confidence intervals (gray shaded area) for 2-, 3-, and 4-term mSCHIF in panels **A**1, **B**1, and **C**1 respectively. Spline knot locations represented by green tick marks on *x*-axis. Corresponding mSCHIF terms presented in panels **A**2, **B**2, and **C**2

uncertainty (gray shaded area) is clearly underestimated compared to that of the spline (dashed red lines) when PM$_{2.5} < 5 \ \mu g/m^3$, but gives a reasonable uncertainty estimate when PM$_{2.5} > 5 \ \mu g/m^3$.

To improve our estimate of uncertainty when PM$_{2.5} < 5 \ \mu g/m^3$, we add more mSCHIF terms such that the μ upper bounds are less than 2.5 μg/m^3. The three-term mSCHIF fit is displayed in Fig. 2, panel B1, with an improvement in uncertainty estimation. The four-term mSCHIF (Fig. 2, panel C1) gives an even better approximation to the spline uncertainty. The corresponding 2-,3-, and 4-term mSCHIF terms are displayed in Fig. 2, panels A2, B2, and C2 respectively.

Comparison with the *Log-Linear* model

The Log − Linear is the most frequently used concentration-response model to estimate public health benefits due to exposure to outdoor air pollution. We estimated

the slope parameter of the log-linear model, β, for each cohort using the exact same dataset as was used to fit the splines. Here we compare estimates of attributable deaths due to selected percentage reductions in outdoor concentrations to PM$_{2.5}$ based on the mSCHIF and Log − Linear model representations of the relative risk function (Fig. 3).

We extracted PM$_{2.5}$ estimated outdoor concentrations for the 3-year average (2015–2017) from the Air Quality Benefits Assessment Tool (AQBAT) (Judek et al. 2023) used by Health Canada for each of the 293 census divisions in Canada in addition to the corresponding number of non-accidental deaths in Canadians over the age of 25 years. We also extracted PM$_{2.5}$ estimated outdoor concentrations for 2018 at 45,011 1km grid cells for the continental USA from Environmental Benefits Mapping and Analysis Program – Community Edition (BenMAP – CE) (U.S. EPA 2023) with corresponding number of all causes of death over the age of 65 years for 2020. These death categories correspond to those used to estimate the relative risk models of the two cohorts.

Fig. 3 CanCHEC mSCHIF mean predictions (solid blue line) and 95% confidence limits (gray shaded area) over Canadian population based $PM_{2.5}$ distribution (orange vertical bars) overlaid with $Log-Linear$ model mean predictions (solid black line) and 95% confidence intervals (dashed black lines) in panel **A**1. Panel **A**2 presents mean marginal changes in relative risk per $0.1ug/m^3$ (mSHCIF—solid blue line; $Log-Linear$—solid black line) and 95% confidence intervals (mSHCIF—gray shaded area; $Log-Linear$—dashed black lines) (panel **A**2) over Canadian population based $PM_{2.5}$ distribution. Attributable deaths due to selected percentage reductions in $PM_{2.5}$ presented in panel **A**3. Corresponding representation for Medicare cohort presented in panels **B**1–**B**3
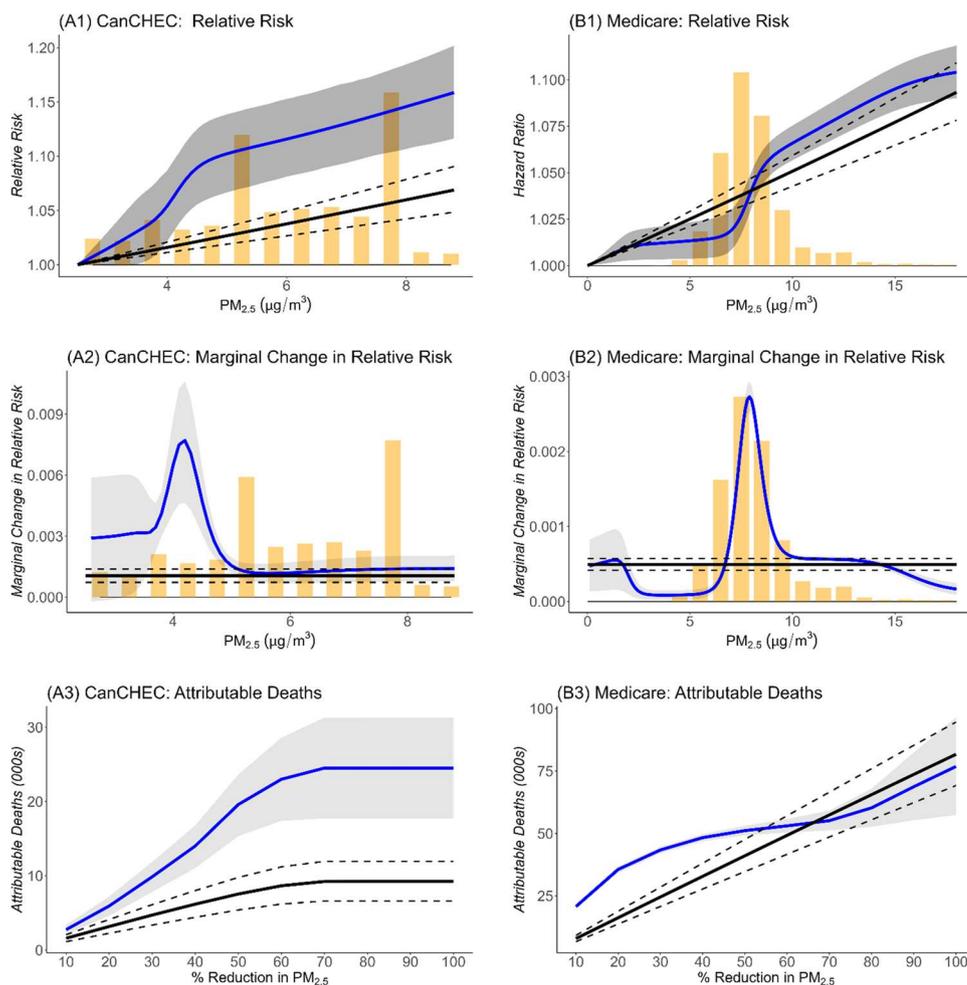
In CanCHEC, any $PM_{2.5}$ person-year exposure value less than $2.5 \mu g/m^3$ was set to $2.5 \mu g/m^3$, representing 0.5% of all person-years of follow-up (Brauer et al. 2022). We made the same assignment for the census division concentrations, affecting 1% of the total deaths. To represent the population exposure to $PM_{2.5}$, we constructed a death-weighted $PM_{2.5}$ distribution by either census division or grid cell (orange vertical bars, Fig. 3). The Canadian mean death-weighted exposure was $6.1 \mu g/m^3$ (sd=1.60) while the USA mean was $8.1 \mu g/m^3$ (sd=1.69).

The CanCHEC mSCHIF and $Log-Linear : \hat{\beta} = 0.01056 \left(\hat{\sigma}_{\hat{\beta}} = 0.00166\right)$ relative risk functions are plotted against $PM_{2.5}$ concentrations in Fig. 3 (panel A1). The mSCHIF mean predictions (solid blue line) are much larger than the corresponding $Log-Linear$ throughout the concentration range due to the steep increase in risk prediction over the 2.5 to $5 \mu g/m^3$ range.

This example illustrates that large differences in magnitude between non-parametric spline and parametric model predictions can occur due to how these models use data. The parametric model uses all the data together to estimate the unknown parameters while the spline model primarily uses different segments of the data to estimate parameters associated with basis functions whose support corresponds to those segments. The spline model predicts a relative risk of 1.10 between 5 and $2.5 \mu g/m^3$ while the $Log-Linear$ predicts a 1.026 relative risk. This difference in relative risk is maintained throughout the concentration range. The 2.5 to $5 \mu g/m^3$ range represents 20% of the person-years of follow-up in the cohort; however, this rapid increase in relative risks appears to have minimal influence on the estimate of the $Log-Linear$ parameter β.

To further examine differences between the two model predictions, we plot the change in the logarithm of the relative risk by $0.1 \mu g/m^3$ increments. These marginal changes in risk indicate a complex pattern between the models over concentration with the mSCHIF displaying a larger change between 2.5 and $5 \mu g/m^3$ compared to the constant change predicted by the $Log-Linear$, but a similar change for concentrations greater than $5 \mu g/m^3$. These differences between the models are reflected in estimates of attributable deaths (panel A3) for selected percent decreases in concentration (10% to 100% by 10% increments) with larger attributable death estimates for the mSCHIF model compared to the

Log − Linear model, and differences increasing with increasing percentage reductions in concentration. Note that the estimates of attributable deaths for 80%, 90%, and 100% are the same since these reductions change the highest concentration of 8.8 μg/m^3 to values less than the counterfactual of 2.5 μg/m^3.

We conducted a similar analysis for the Medicare cohort (panels B1–B3 of Fig. 3). In this cohort, the increase in relative risk over the 0 to 2 μg/m^3 range was similar for both models (panel B1). The Log − Linear : $\hat{\beta} = 0.00494$ $(\hat{\sigma}_{\hat{\beta}} = 0.00040)$ model predictions were larger than the mSCHIF over the 2 to 8 μg/m^3 range, but smaller when PM$_{2.5} > 8$ μg/m^3 (panel B1). The marginal changes in risk were also similar for the two models in the 0 to 2 μg/m^3 range, but the Log − Linear model predicted larger marginal changes in the 2 to 6 μg/m^3 range, smaller changes in the 6 to 10 μg/m^3 range, and similar changes when PM$_{2.5} > 10$ μg/m^3 compared to the mSCHIF predictions (panel B2). Attributable death estimates were lower for the Log − Linear model for the 10–70% reductions compared to the mSCHIF estimate, but larger for the higher percent reductions.

Discussion

In order to estimate the public health burden of exposure to outdoor concentrations of air pollution, the magnitude and shape of the association between outdoor concentrations and a health outcome must be characterized in addition to its uncertainty. Several attempts have been made using parametric functions indexed by a few parameters (Cohen et al. 2004; Chen and Hoek 2020; Burnett et al. 2014, 2018, 2022). However, recent interest has focused on non-parametric approaches using either regression (Brauer et al. 2022; Brunekreef et al. 2021) or smoothing (GBD 2019; Dominici et al. 2022) splines. Typically, no restrictions on the shape of the splines are made, resulting in non-monotonic predictions with multiple changes in direction that make these functions less desirable for public health benefits analysis. Limiting the number of terms in a regression spline can smooth the curve (Brunekreef et al. 2021), thus making it more useful for benefits analysis. However, it is not clear how good this approach is at characterizing the concentration-response shape. Even smoothing splines, if not restricted, can result in non-monotonicity with waviness as evidenced by our Medicare example.

In these cohort studies, an estimate of the outdoor concentration near a subject's home is related to their survival. There exists a complex association, that varies by subject, between their estimate of outdoor concentration and their personal exposure (Hammond et al. 2014), with an equally likely complex association between their personal exposure

and probability of surviving any fixed time period. In addition, it is likely that components/sources of particulate matter have a different, and potentially non-linear association, with survival (Henneman et al. 2023), resulting in a potentially complex concentration-response pattern based on total mass. Even if every component/source of particulate matter had a linear concentration-response, total mass may have a non-linear association due to variations in the atmospheric mixture at different total mass concentrations.

Although simple approximations have been suggested, such as linear or supra-linear, we suggest that we in fact do not know what the true concentration-response pattern should look like and can only rely on data to guide us. In this paper, we suggest a method that first identifies the best fitting unrestricted spline to the data, then transform the spline fit and its uncertainty to a monotonic function suitable for benefits assessment. We make few restrictions on the algebraic form of the fit, other than monotonicity.

We also suggest that it is important to conduct cohort-specific concentration-response analysis to examine the consistency, or lack of, in the shape among cohorts. Pooling these shapes provides a means to constructing a common concentration-response model as was done for the Global Exposure Mortality Model (GEMM) (Burnett et al. 2018).

However, there is interest to examine the concentration-response shape of cohorts representing the population of specific judications; as demographics, disease distribution, and how total mass represents the toxicity of the atmospheric mixture can vary by region. This raises concerns that a common concentration-response based on global studies may be a poor approximation to that of any specific region or country.

The Global Burden of Disease program (GBD 2019) fit a smoothing spline with their Integrated Exposure-Response (IER) meta-regression framework to predict risk from PM$_{2.5}$ exposure to both outdoor and indoor pollution covering a very large range (0 to 1000 μg/m^3) by combining information from multiple studies. They restricted the spline to be monotonically increasing and concave resulting in a supra-linear curve. However, such restrictions for individual studies may be too limiting as evidenced by our two examples. Restricting spline fits to be monotonically increasing may still not lead to a curve useful for benefits analysis as it may contain several concentration-response patterns like step-functions over narrow concentration intervals.

We address these limitations by proposing a new spline formulation, denoted as "monotonically increasing Shape Constrained Health Impact Function" or mSCHIF. The mSCHIF consists of a series of basis functions, not defined by knot values, but by four parameters, that can model a wide variety of shapes including near-linear, sub-linear, supra-linear, and sigmodal. A linear combination of these basis functions, multiplied by a parameter restricted to be

positive, creates a monotonically increasing wide variety of shapes that we suggest can model virtually any monotonically increasing concentration-response pattern of interest. Since the basis functions can take a variety of shapes, far fewer are needed to make acceptable predictions compared traditional spline formulations. We also capture the totality of uncertainty in the spline predictions by equating the covariance matrix of the spline predictions over the concentration range to that of the mSCHIF predictions.

We give guidance to investigators in how to present results from spline model predictions. The typical method of reporting a spline fit is to plot mean predictions and their 95% confidence intervals, transformed in a manner such that the spline prediction at a pre-specified counterfactual concentration, often the minimum, is set to zero. However, if the uncertainty at the counterfactual is also not set to zero, the curve cannot be used for burden (all concentrations reduced to the counterfactual) analysis. Even when the uncertainty at the counterfactual is set to zero, the spline fit cannot be used for benefits analysis (difference in risk between any two specific concentrations) since the uncertainty in risk associated with such changes in concentration cannot be determined from this summary information. The spline parameter estimates, their covariance matrix, knot locations, and type of spline need to be reported. Alternatively, reporting 1000 sets of spline predictions over the concentration range is sufficient for benefits analysis using the spline fit directly. This information can also be used to estimate the mSCHIF. Benefits estimates and its uncertainty can be calculated from the mSCHIF parameter estimates and covariance matrix $\hat{\Sigma}$.

Our comparison between the mSCHIF and Log−Linear model relative risk, marginal change in risk, and attributable deaths estimates, suggest there can be considerable departures in the concentration-response from linearity. These differences could potentially lead to differences in policy development.

Our two examples involved fitting either a regression or smoothing spline to primary data from cohorts. However, there is interest in characterizing concentration-response by combining information from multiple cohorts (GBD 2019; Burnett et al. 2014, 2018, 2022). Our mSCHIF is another option to create a common curve if an unrestricted spline fit is available to the meta-data.

Although we accessed mortality counts and $PM_{2.5}$ concentration data from AQBAT and BenMAP-CE, we preformed the benefits calculations externally to these computer programs. We did this due to how both programs accept input on the relationship between concentration and mortality. They were constructed to use the information provided by the Log−Linear model, namely a single parameter estimate and its standard error, assuming a linear association between the logarithm of the hazard function and concentration. Since the mSCHIF is constructed as the sum of parameter estimates multiplied by transformations in concentration, with uncertainty described in terms of a multivariate normal distribution, it cannot be directly used by these computer programs.

We specifically selected spline model fits that met several criteria in order to illustrate properties of our new model. The spline fits had to display complex associations between $PM_{2.5}$ concentrations and mortality in such a manner was it would likely not be suitable for benefits analysis; such as having several changes in direction and thus not being monotonically increasing with concentration. The mSCHIF can accept two formats of information from the spline fit: spline parameter estimates, their covariance matrix, knot locations, and type of spline or multiple sets of spline predictions over the concentration range. We selected a regression spline for the CanCHEC example where we accessed the first type of information and a smoothing spline for the Medicare example where we accessed the second type. We therefore did not select these fits as the most appropriate representation of the $PM_{2.5}$-mortality association in either Canada or the USA. We did, however, further conduct benefits analysis to highlight potential differences in attributable death estimates between the mSCHIF and the most used model for such purposes. We thus highlighted the fact that characterizing the exposure-response relationship in a potentially non-linear manner may influence policy decisions compared to using the traditional benefits analysis Log−Linear model.

References

Akaike H (1974) A new look at the statistical model identification. IEEE Transact Automat Contr 19(6):716–723. https://doi.org/10.1109/TAC.1974.1100705

Bartels RH, Beatty JC, Barsky BA (1998) Hermite and cubic spline interpolation. Ch. 3. In: An introduction to splines for use in

computer graphics and geometric modelling. Morgan Kaufmann, San Francisco, CA, pp 9–17

Brauer M, Brook JR, Christidis T, Chu Y, Crouse DL, Erickson A et al (2022) Mortality-air pollution associations in low exposure environments (MAPLE): phase 2. Research Report 212. Health Effects Institute

Brunekreef B, Strak M, Chen J, Andersen ZJ, Atkinson R, Bauwelinck M et al (2021) Mortality and morbidity effects of long-term exposure to low-level PM2.5, black carbon, NO2 and O3: an analysis of European cohorts. Research Report (Health Effects Institute), p 208

Burnett R, Chen H, Szyszkowicz M et al (2018) Global estimates of mortality associated with long-term exposure to outdoor fine particulate matter. P Natl Acad Sci 115:9592

Burnett RT, Pope CA III, Ezzati M, Olives C, Lim SS, Mehta S, Shin HH, Singh G, Hubbell B, Brauer M, Anderson HR, Smith KR, Balmes J, Bruce N, Kan H, Laden F, Prüss-Ustün A, Turner MC, Gapstur SM et al (2014) An integrated risk function for estimating the global burden of disease attributable to ambient fine particulate matter exposure. Environ Health Perspect. https://doi.org/10.1289/ehp.1307049

Burnett RT, Spadaro JV, Garcia GR, Pope CA (2022) Designing health impact functions to assess marginal changes in outdoor fine particulate matter. Environ Res 204(Part C):112245

Chen H, Quick M, Kaufman JS, Burnett RT (2022) Impact of lowering fine particulate matter from major emission sources on mortality in Canada: a nationwide causal analysis. PNAS. https://doi.org/10.1073/pnas.2209490119

Chen J, Hoek G (2020) Long-term exposure to PM and all-cause and cause-specific mortality: a systematic review and meta-analysis. Environment International. https://doi.org/10.1016/j.envint.2020.105974

Cohen AJ, Anderson HR, Ostro B, Pandey KD, Krzyzanowski M, Kuenzli N, Gutschmidt K, Pope CA, Romieu I, Samet JM, Smith KR (2004) Mortality impacts of urban air pollution. In: Ezzati M, Lopez AD, Rodgers A, Murray CJL (eds) Comparative quantification of health risks: global and regional burden of disease due to selected major risk factors, vol 2. World Health Organization, Geneva, Switzerland

Dominici F, Zanobetti A, Schwartz J, Braun D, Sabath B, Wu X (2022) Assessing adverse health effects of long-term exposure to low levels of ambient air pollution: implementation of causal inference methods. Health Effects Institute. Research Report, p 211

Eilers PHC, Marx B (1996) Flexible smoothing with B-splines and penalties. Stat Sci 11(2):89–121

GBD (2019) Risk Factor Collaborators (2020) Global burden of 87 risk factors in 204 countries and territories, 1990–2019: a systematic analysis for the Global Burden of Disease Study 2019. Lancet 396:1223–1249

Gu Y, Wong TW, Law CK, Dong GH, Ho KF, Yang Y, Yim SHL (2018) Impacts of sectoral emissions in China and the implications: air quality, public health, crop production, and economic costs. Environ Res Lett 13:084008

Hammond D, Croghan G, Shin H, Burnett R, Bard R, Brook RD, Williams R (2014) Cardiovascular impacts and micro-environmental exposure factors associated with continuous personal PM2.5 monitoring. Journal of Exposure Science and Environmental. Epidemiology 24:337–345

Harrel FE (2015) Regression modeling strategies with applications to linear models, logistic and ordinal regression, and survival analysis. Springer

Henneman L, Choirats C, Dedoussia I, Dominici F, Roberts J, Zigler C (2023) Mortality risk from United States coal electricity generation. Science 382:941–946

Heo J, Adams PJ, Gao HO (2017) Public health costs accounting of inorganic PM2.5 pollution in metropolitan areas of the United States using a risk-based source-receptor model. Environ Int 106:119–126

Johnston FH, Henderson SB, Chen Y, Randerson JT, Marlier M, DeFries RS, Kinney P, Bowman DMJS, Brauer M (2012) Estimated global mortality attributable to smoke from landscape fires. Environ Health Perspect 120:695–701

Judek S, Stieb D, Jovic B, Edwards B (2023) Air quality benefits assessment tool user guide. Health Canada, Ottawa

Lelieveld J, Klingmüller K, Pozzer A, Burnett RT, Haines A, Ramanathan V (2019) Effects of fossil fuel and total anthropogenic emission removal on public health and climate. P Natl Acad Sci 116:7192

Liang CK, West JJ, Silva RA, Bian H, Chin M, Davila Y, Dentener FJ, Emmons L, Flemming J, Folberth G, Henze D, Im U, Jonson JE, Keating TJ, Kucsera T, Lenzen A, Lin M, Lund MT, Pan A et al (2018) HTAP2 multi-model estimates of premature human mortality due to intercontinental transport of air pollution and emission sectors. Atmos Chem Phys 18:10497–10520

Meng J, Li C, Martin RV, van Donkelaar A, Hystad P, Brauer M (2019) Estimated long-term (1981-2016) concentrations of ambient fine particulate matter across North America from chemical transport (2019) modeling, satellite remote sensing, and ground-based measurements. Environ Sci Technol. 53(9):5071–5079

Muller NZ (2018) Environmental benefit-cost analysis and the national accounts. J Benefit Cost Anal 9(1):27–66. https://doi.org/10.1017/bca.2017.15

Nasari MN, Szyszkowicz M, Chen H, Crouse DL, Turner MC, Jerrett M, Pope CA III, Hubbell B, Fann N, Cohen A, Gapstur SM, Diver WR, Stieb D, Forouzanfar MH, Kim S-Y, Olives C, Krewski D, Burnett RT (2016) A class of non-linear exposure-response models suitable for health impact assessment applicable to large cohort studies of ambient air pollution. Air Qual Atmos Health 9(8):961–972

nlsr: Functions for Nonlinear Least Squares Solutions - Updated 2022 (r-project.org) 2020

Pya N, Wood SN (2015) Shape constrained additive models. Stat Comput 25:543–559

U.S. EPA (2023). Environmental Benefits Mapping Analysis Program Community Edition (BenMAP-CE). http://www2.epa.gov/benmap. Accessed 18 Jan 2023

World Bank; Institute for Health Metrics and Evaluation (2016) The cost of air pollution: strengthening the economic case for action. World Bank, Washington, DC. © World Bank. https://openknowledge.worldbank.org/handle/10986/25013