Supplementary Materials for "A single-index model with a surface-link for optimizing individualized dose rules"

Hyung Park*, Eva Petkova, Thaddeus Tarpey
and
R. Todd Ogden†

A Supporting information for Sections 4 and 5 of the main manuscript

A.1 Proof of Proposition 1 in Section 4

Let $Q(g,\beta) = \mathbb{E}\big[Yg(\beta^\top X, A) - b(g(\beta^\top X, A))\big]$ (i.e., the criterion function in (24) of the main manuscript), where the expectation is with respect to the distribution of (Y, A, X), in which their relationships are specified by model (22) of the main manuscript. We have

$$\begin{aligned}
Q(g,\beta) &= \mathbb{E}\Big[\{\mu_0(X) + g_0(\beta_0^\top X, A)\}g(\beta^\top X, A) - b(g(\beta^\top X, A))\Big] \\
&= \mathbb{E}\Big[\mathbb{E}_{A|X}\big[\{\mu_0(X) + g_0(\beta_0^\top X, A)\}g(\beta^\top X, A) - b(g(\beta^\top X, A)) \mid X\big]\Big] \\
&= \mathbb{E}\Big[\mathbb{E}_{A|X}\big[\mu_0(X)g(\beta^\top X, A) + g_0(\beta_0^\top X)g(\beta^\top X, A) - b(g(\beta^\top X, A)) \mid X\big]\Big] \\
&= \mathbb{E}\Big[g_0(\beta_0^\top X, A)g(\beta^\top X, A) - b(g(\beta^\top X, A))\Big] \\
&= \mathbb{E}\Big[\mathbb{E}\big[g_0(\beta_0^\top X, A)g(\beta^\top X, A) - b(g(\beta^\top X, A)) \mid \beta^\top X, A\big]\Big] \\
&= \mathbb{E}\Big[\mathbb{E}\big[g_0(\beta_0^\top X, A)|\beta^\top X, A\big]g(\beta^\top X, A) - b(g(\beta^\top X, A))\Big],
\end{aligned} \tag{S.1}$$

where the second equality is from an application of the iterated expectation rule to condition on X, and the fourth equality is from: $\mathbb{E}_{A|X}[\mu_0(X)g(\beta^\top X, A)|X] = \mu_0(X)\mathbb{E}_{A|X}[g(\beta^\top X, A)|X] = 0$ $(a.s.)$, as a result of the constraint in (24) of the main manuscript that we impose on g. Notice, from the last line of (S.1), that $Q(g,\beta)$ is independent of the "nuisance" term $\mu_0(X)$.

For each fixed β and given $(\beta^\top X, A)$, the maximizer g of the right-hand side of (24) of the main manuscript must satisfy (Ravikumar *et al.*, 2009):

$$\frac{\partial Q(g,\beta)}{\partial g} = \mathbb{E}\big[g_0(\beta_0^\top X, A)|\beta^\top X, A\big] - \dot{b}(g(\beta^\top X, A)) = 0 \quad a.s.\ (X, A), \tag{S.2}$$

*Division of Biostatistics, Department of Population Health, New York University, New York, NY 10016, USA parkh15@nyu.edu
†Department of Biostatistics, Columbia University, New York, NY 10032, USA

that is, the maximizer g must satisfy $\dot{b}(g(\beta^\top X, A)) = \mathbb{E}\big[g_0(\beta_0^\top X, A)|\beta^\top X, A\big]$ $(a.s.)$. Equivalently, the maximizer g must satisfy

$$g(\beta^\top X, A) = \dot{b}^{-1}\big\{\mathbb{E}\big[g_0(\beta_0^\top X, A)|\beta^\top X, A\big]\big\} \quad a.s.\ (X, A), \tag{S.3}$$

based on which we can profile out g in the last line of (S.1). This profilization of $g(\beta^\top X, A)$, for each given $(\beta^\top X, A)$, yields the following profiled objective function (S.1):

$$\mathbb{E}\bigg[\mathbb{E}\big[g_0(\beta_0^\top X, A)|\beta^\top X, A\big]\dot{b}^{-1}\big\{\mathbb{E}\big[g_0(\beta_0^\top X, A)|\beta^\top X, A\big]\big\} - b\big(\dot{b}^{-1}\big\{\mathbb{E}\big[g_0(\beta_0^\top X, A)|\beta^\top X, A\big]\big\}\big)\bigg], \tag{S.4}$$

as a function only of $\beta \in \Theta$. In (S.4), for the notational simplicity, let us reparametrize

$$\theta(\beta^\top X, A) = \dot{b}^{-1}\big\{\mathbb{E}\big[g_0(\beta_0^\top X, A)|\beta^\top X, A\big]\big\}. \tag{S.5}$$

Then, we can re-write (S.4) as:

$$\begin{aligned}\mathbb{E}\bigg[\mathbb{E}\big[g_0(\beta_0^\top X, A)|\beta^\top X, A\big]\theta(\beta^\top X, A) - b(\theta(\beta^\top X, A))\bigg] &= \mathbb{E}\bigg[\mathbb{E}\big[g_0(\beta_0^\top X, A)\theta(\beta^\top X, A) - b(\theta(\beta^\top X, A))|\beta^\top X, A\big]\bigg] \\ &= \mathbb{E}\bigg[g_0(\beta_0^\top X, A)\theta(\beta^\top X, A) - b(\theta(\beta^\top X, A))\bigg].\end{aligned} \tag{S.6}$$

The last line corresponds to the negative of the cross-entropy (or equivalently, the Kullback-Leibler (KL) divergence) between one member of the exponential family with the conditional mean $g_0(\beta_0^\top X, A)$ and another member of the exponential family with the *canonical parameter* $\theta(\beta^\top X, A)$. It is clear that the maximizer, which we denote as β^*, of the negative of the KL divergence occurs at β_0 (i.e., when the two random variables $\beta^{*\top} X$ and $\beta_0^\top X$ agree with each other). Plugging $\beta = \beta^* = \beta_0$ into equation (S.3) implies that the maximizer, which we denote as g^*, of the objective function $Q(g, \beta_0)$ must satisfy: $g^*(\beta_0^\top X, A) = \dot{b}^{-1}\big\{g_0(\beta_0^\top X, A)\big\}$ $(a.s.)$. Since $\dot{b}^{-1}\{\cdot\} = h\{\cdot\}$ (by the property of the exponential family) where $h\{\cdot\}$ is the *canonical link* function, the desired expression, $g^* = h \circ g_0$, of Proposition 1, for the maximizer of the objective function (S.1) follows.

A.2 Reparametrization of the simulation scenarios in Section 5

In this subsection, we illustrate how the underlying models in Scenarios 1–4 can be reparametrized to fit in the SIMSL framework (4) of the main manuscript. Specifically, we can cast those mean models in terms of μ (X main effects) and g (X-by-A interactions) as in the SIMSL (4), where their g terms satisfy the identifiability condition (5) (or, its more general version $\mathbb{E}[g(X, A)|X] = 0$ if their g term cannot be expressed in terms of a 1-dimensional projection $\beta^\top X$).

For Scenario 1, the g term of SIMSL (4) corresponds to $g(\beta^\top X, A) := -25(f_{\text{opt}}(X) - A)^2 - \mathbb{E}[-25(f_{\text{opt}}(X) - A)^2|X] = -25\{A^2 - 2Af_{\text{opt}}(X) + 2f_{\text{opt}}(X) - 4/3\}$, where $f_{\text{opt}}(X) = 1 + \sqrt{0.5}\beta^\top X$, in which $\beta = (0.5, 0.5, 0, 0, \ldots, 0)^\top/\sqrt{0.5}$. The term g satisfies the conditional mean zero (identifiability) constraint (5). On the other hand, the μ term of SIMSL (4) corresponds to $\mu(X) := 8 + 4X_1 - 2X_2 - 2X_3 + \mathbb{E}[-25(f_{\text{opt}}(X) - A)^2|X] = 8 + 4X_1 - 2X_2 - 2X_3 - 25\{f_{\text{opt}}^2(X) - 2f_{\text{opt}}(X) + 4/3\}$. Given X, the component $g(\beta^\top X, A)$ (as a function of A) is maximized at $A = f_{\text{opt}}(X)$, implying that $f_{\text{opt}}(X)$, specified in Scenario 1, indeed corresponds to the optimal individualized dose rule specified in Eq.(2) of the main manuscript.

Scenarios 2, 3 and 4 can also be similarly formulated in terms of μ and g, however, in these scenarios, if we assume SIMSL (4), then the g term of the SIMSL is misspecified (i.e., the term g cannot be expressed in terms of a single-index $\beta^\top X$). Thus, a more general function $g(X, A)$ (over the $p + 1$ dimensional space), rather than $g(\beta^\top X, A)$, needs to be employed for representing the heterogeneous treatment effect (and the more general identifiability condition $\mathbb{E}[g(X, A)|X] = 0$ needs to be considered). In these scenarios, SIMSL (4), with optimization (9), provides the optimal 1-dimensional projection-based approximation to $g(X, A)$ (with respect to the KL divergence, see (24) of the main manuscript).

As an additional note, in Remark 1 of Supplementary Materials C.1, we provide an explicit expression of the conditional expectation $\mathbb{E}[g(\beta^\top X, A)|X]$ for Scenario 4, where the expectation is taken with respect to the distribution of A given X (a truncated normal distribution), in which case A depends on X.

B Application to estimation of a pollutant-season interaction on mortality from air-pollution data

In this section, we consider a data analysis example illustrating an application of the generalized single-index regression method (described in Section 4 of the main manuscript) that models interactions between a set of predictors (X) and a variable (A) in their effects on a non-normal response (counts) variable.

Several time series studies of air pollution and health have provided compelling evidence of a positive association between short-term variation in ambient levels of particulate matter and daily mortality counts (see, e.g., Pope *et al.*, 1995; Dockery and Pope, 1996; Bell *et al.*, 2004). Peng and Dominici (2008) noted that the short-term effects of particulate matter on mortality might exhibit seasonal (time) variation. In particular, the characteristics of the particulate matter mixture can vary seasonally throughout the year. Patterns of human activity also change from season to season, and as a result, an air pollution concentration in one season may lead to a different effect in a different season. Other potential time-varying confounding and modifying factors (such as temperature and influenza epidemics) can also impact effects of air pollution on mortality differently in different seasons (Peng and Dominici, 2008). Therefore, the relationships between mortality and air pollution levels can vary considerably across seasons. In this section, we illustrate the utility of the SIMSL approach to estimating the interactions between seasonal variation (in this case, A) and particulate matter (in this case, X) on their effects on mortality.

The data are from Peng and Welty (2004) and publicly available from the R (R Core Team, 2019) package `gamair` (Wood, 2019). The response of interest is the daily number of deaths in Chicago over 14 years. The outcomes Y_t represent a time series of daily mortality counts (indicating the number of deaths that occurred on day t), with the conditional mean, m_t, that depends on particulate matter levels and season. As typical for modeling time series of counts, we assume a Poisson distribution for the counts Y_t. To model seasonal variation, we introduce the "day of a year" variable that ranges from 0 to 364, defined as $A_t := t \mod 365$, in which $A_t = 0$ corresponds to the first calendar day of a year.

To study the pollutant-season interaction effects, we consider the following SIMSL (see model (21) of the main manuscript):

$$m_t \;=\; \mathbb{E}[Y_t|X_t, X_{t-1}, \ldots, X_{t-K}, A_t] \;=\; \mu_0(X_t, X_{t-1}, \ldots, X_{t-K}) + g_0\left(\sum_{j=0}^{K} \beta_j X_{t-j}, A_t\right) \qquad (t = 1, \ldots, n) \tag{S.7}$$

where the single-index $\beta^\top X$ is defined in terms of distributed lag model:

$$\sum_{j=0}^{K} \beta_j X_{t-j} \quad (t = 1, \ldots, n)$$

where $X_t \in \mathbb{R}$ is the time series of daily (day t) particulate matter levels, and K is the maximum lag. Hence, the model includes multiple lags of pollution, rather than the pollution level only on the day t itself. This K-lagged model is reasonable, as any effects would likely take some time to manifest themselves via the aggravation of existing medical conditions (Wood, 2017). This accumulated particulate matter over K days enters into the function $g_0(\cdot, \cdot)$ in model (S.7). The coefficients $\beta_j \in \mathbb{R}$ associated with this weighted sum need to be estimated from data, and we are particularly interested in summarizing the variability in the lagged particulate matter levels $(X_t, X_{t-1}, \ldots, X_{t-K})$ that is related to the seasonal variation A_t in their effects on the outcome Y_t, rather than the "nuisance" variability unrelated to A_t. The function μ_0 on the right-hand side of (S.7) represents the pollutant "main" effect unrelated to A_t. Our focus is on estimating the interaction effect term term of model (S.7) to study the pollutant-season interactions on mortality.

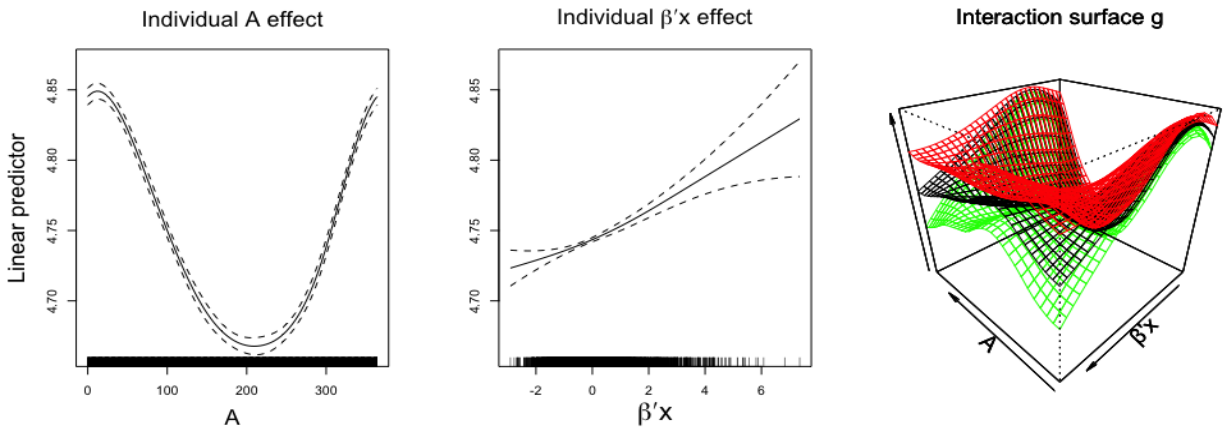

Figure S.1: Estimated individual effect for A (left panel) and that for $\beta^\top X$ (center panel), adjusted for the effect of $\beta^\top X$ and that of A, respectively, conditioning on the estimated single-index $\beta^\top X$. The third panel: Estimated link surface, $g(\cdot, \cdot)$ of the generalized SIMSL (23) of the main manuscript (with the logarithmic link function $h(\cdot) = \log(\cdot)$), for the pollutant ($\beta^\top X$) and season (A) interaction; the red and green surfaces are at plus or minus one standard error from the estimate (the black surface in the middle).

We estimate the pollutant-season interactions based on the working model (23) of the main manuscript, with the logarithmic link $h(\cdot) = \log(\cdot)$:

$$\log(m_t) = g(\beta^\top X_t, A_t) \quad (t = 1, \ldots, n) \tag{S.8}$$

subject to the constraint $\mathbb{E}[g(\beta^\top X, A)|X] = 0$, in which $X := (X_t, X_{t-1}, \ldots, X_{t-K})^\top$ and $\beta := (\beta_0, \ldots, \beta_K)^\top$, where we omit the subscript t in writing X and A, for notational simplicity. In this illustration, we take $K = 4$. We model the seasonal effect to vary smoothly over the course of a year, but we constrain the effect to be periodic across years. We accomplish this by utilizing a set of cyclic cubic spline basis functions $\check{B}$ (whose left and right ends match, up to a second derivative) to define the model matrix $\check{\boldsymbol{B}}$ associated with the variable A (which is again subject to the linear constraint (15) of the main manuscript, for the "orthogonality" condition), and use this cyclic basis matrix $\check{\boldsymbol{B}}$ to construct the tensor product basis matrix $\boldsymbol{D}$ in (14) of the main manuscript for representing g.

The single-index coefficient β in model (S.8) is estimated as $(0.77, 0.42, 0.13, 0.44, 0.09)^\top$. The approximate 95% bootstrap confidence intervals (based on 500 bootstrap replications; see Section C.3 for description of the approximate bootstrap confidence interval construction procedure) associated with the lag 0 and the lag 1 (i.e., the first two elements of β), and also the lag 3's particulate matter concentration levels do not include 0, indicating that (not surprisingly) the particulate matter level on the day itself and one day prior are determined to be important for modeling the pollutant-season interaction. In this analysis, three day prior also appears to be important.

The estimated surface-link $g(\beta^\top X, A)$ displayed in the right-most panel in Figure S.1 indicates that for a small value of $\beta^\top X$, the mortality is largely explained by the U-shaped individual effect of the season (i.e., the variable A) (see the left-most panel in Figure S.1), which exhibits the usual pattern of a low mortality rate during the summer and a high mortality rate during the winter. Moreover, the middle panel in Figure S.1 indicates that mortality, not surprisingly, increases monotonically with the accumulated particulate matter ($\beta^\top X$). However, the estimated surface-link on the right-most panel in Figure S.1 visualizes the pollutant-season interactions on mortality, which indicates that the effect the accumulated particulate matter ($\beta^\top X$) has a larger effect in the spring (e.g., around $A = 100$) and smaller effects in the other seasons (e.g., around $A = 300$), showing an interesting pollutant-season interaction pattern. In particular, for a large value of $\beta^\top X$ (say, $\beta^\top X > 2$) the contribution of the accumulated particulate matter ($\beta^\top X$) on mortality in the spring time (around $A = 100$) can be quite substantial.

C Simulation experiments

In this section, we perform additional simulation studies to examine the performance of the proposed approach to optimizing the individualized dose rules.

C.1 Supplementary simulation

We consider a set of simulations that is supplemental to the simulation example in Section 5 of the main manuscript. We generate $X = (X_1, \ldots, X_p)^\top$, with $X_j \sim \text{Unif}(-1, 1)$ and $p = 10$. Given patient characteristics $X \in \mathbb{R}^{10}$, the optimal dose level $f_{\text{opt}}(X)$ is set to be:

$$f_{\text{opt}}(X) = 2\exp\big(-(\beta^\top X)^2\big) \quad (\in [0, 2]), \tag{S.9}$$

i.e., we consider a nonlinear dose rule (a linear dose rule case was considered in Section 5 of the main manuscript), where $\beta = (8, 4, 2, 1, 0, \ldots, 0)^\top/\sqrt{85} \in \Theta$. We set the dose $A \in \mathcal{A} = [0, 2]$ to follow the truncated normal distribution with mean $f_{\text{opt}}(X)$, lower bound 0, upper bound 2, and standard deviation $\sigma = 0.5$, i.e.,

$$A \sim \text{TruncN}\left(f_{\text{opt}}(X), 0, 2, 0.5\right). \tag{S.10}$$

(The distribution of the dose A in (S.10) mimicks an observational study where doctors can guess at the patient-specific optimal dose level $f_{\text{opt}}(X)$ to a certain extent, and prescribe the dose accordingly.) We generate the outcome $Y \in \mathbb{R}$ from the mean model (4) of the main manuscript, i.e.,

$$Y \;=\; \mu(X) + g(\beta^\top X, A) + \epsilon \tag{S.11}$$

with $\epsilon \sim \mathcal{N}(0, 1)$, where

$$\begin{aligned} g(\beta^\top X, A) &= -5(f_{\text{opt}}(X) - A)^2 - \mathbb{E}[-5(f_{\text{opt}}(X) - A)^2|X] \\ \mu(X) &= \delta \sum_{j=1}^{10} \exp(X_j)/2 \;+\; \mathbb{E}[-5(f_{\text{opt}}(X) - A)^2|X], \end{aligned} \tag{S.12}$$

constructed in a way that satisfies the specific identifiability condition $\mathbb{E}[g(\beta^\top X, A)|X] = 0$ in (5) of the main manuscript (to elaborate this, the conditional expectation $\mathbb{E}[-5(f_{\text{opt}}(X)-A)^2|X]$ in (S.12) is subtracted from $-5(f_{\text{opt}}(X) - A)^2$ and then added to $\delta\sum_{j=1}^{10}\exp(X_j)/2$, to make model (S.11) satisfy the "identifiability" constraint $\mathbb{E}[g(\beta^\top X, A)|X] = 0$). In (S.12), the parameter $\delta \in \{1, 2\}$ regulates the contribution of the X main effect on the variance of Y: a larger δ implies a larger variance of the "nuisance" component $\mu(X)$, in comparison to that of the "signal" component $g(\beta^\top X, A)$.

Remark 1 *Under the additive construction (S.11), the term $\mathbb{E}[-5(f_{\text{opt}}(X) - A)^2|X]$ that appears in the components g and μ in (S.12) cancels each other, and thus we do not have to evaluate this term when we generate the response Y. However, to complete our specification of model (S.11) and to calculate the "signal" to "noise" ratio based on the variances of $g(\beta^\top X, A)$ and $\mu(X)$, let us provide a more explicit expression of the term:*

$$-\mathbb{E}[5(f_{\text{opt}}(X)-A)^2|X] = -5\mathbb{E}[f_{\text{opt}}(X)^2-2f_{\text{opt}}(X)A+A^2|X] = -5\big\{f_{\text{opt}}(X)^2-2f_{\text{opt}}(X)\mathbb{E}[A|X]+\mathbb{E}[A^2|X]\big\}, \tag{S.13}$$

in which the expectation given X is evaluated with respect to the distribution of A given in (S.10). By the property of a truncated normal distribution, $\mathbb{E}[A|X]$ and $\mathbb{E}[A^2|X]$ are given by (See, e.g., Johnson et al. *(1994)):*

$$\mathbb{E}[A|X] = f_{\text{opt}}(X) + \frac{\phi(z_1) - \phi(z_2)}{\Phi(z_2) - \Phi(z_1)}\sigma$$

and

$$\mathbb{E}[A^2|X] = \{f_{\text{opt}}(X)\}^2 + \sigma^2 + \sigma^2\frac{\phi'(z_2) - \phi'(z_1)}{\Phi(z_2) - \Phi(z_1)} - f_{\text{opt}}(X)\frac{\phi(z_2) - \phi(z_1)}{\Phi(z_2) - \Phi(z_1)},$$

where $z_1 = \frac{-f_{\text{opt}}(X)}{\sigma}$ *and* $z_2 = \frac{2-f_{\text{opt}}(X)}{\sigma}$ *(where* $\sigma = 0.5$*), the functions* $\phi(\cdot)$ *and* $\Phi(\cdot)$ *are the probability density function (PDF) and the cumulative distribution function (CDF) of the standard normal distribution, respectively, and* ϕ' *represents the first derivative of* ϕ*. These two expressions can be used to evaluate the quantity (S.13) given* X*, and thus compute the terms* $g(\beta^\top X, A)$ *and* $\mu(X)$ *in (S.12).*

Only the g term of model (S.11) is important for estimating f_{opt}, since the optimal dose rule $f_{\text{opt}}(X)$ in (S.9) is (which can be shown to be expressed as):

$$f_{\text{opt}}(X) = \underset{a \in \mathcal{A}}{\text{argmax}}\ g(\beta^\top X, A), \tag{S.14}$$

which does not depend on the μ term. Therefore, under model (S.11), we can define the "signal" to "noise" ratio as: $\text{var}\big\{g(\beta^\top X, A)\big\}/\text{var}\big\{\mu(X) + \epsilon\big\}$. In particular, under (S.12), the "signal" to "noise" ratio is about 1 when $\delta = 1$, and is about 0.3 when $\delta = 2$. When $\delta = 2$, the variance of the μ term dominates that of the "signal" term g; in such a case, estimation of f_{opt} is more difficult compared to the case of $\delta = 1$.

In this simulation study, as we vary $n \in \{500, 1000, 1500\}$ and $\delta \in \{1, 2\}$ (i.e., the magnitude of the μ term) in (S.12), we will study: 1) the impact of accounting for the X "main" effect (i.e., the μ term) on the estimation of (S.14); and 2) the impact of imposing the constraint $\mathbb{E}[g(\beta^\top X, A)|X] = 0$ (i.e., constraint (5) in the main manuscript) on the estimation of (S.14). We consider the following four approaches.

1. "SIMSL(w.o. X main)", i.e., the SIMSL approach *without* incorporating any X main effect (the μ term) to the estimation of $f_{\text{opt}}(X)$ in (S.14). (This is the primary approach that we considered in Sections 3 and 4 of the main manuscript.)

2. "SIMSL(w. X main)", where we approximate the μ term in (S.11) by a linear model

$$\mu(X) \approx \xi^\top X \tag{S.15}$$

for some unknown $\xi \in \mathbb{R}^p$, and incorporate the approximation (S.15) to the estimation of (S.14). To account for this non-zero approximation (S.15) for the μ term in (S.12) in the estimation of (S.14), we modify the estimation procedure in Section 4 of the main manuscript. We replace the squared error term $\|Y_{n\times 1} - g(\boldsymbol{X}\beta, A_{n\times 1})\|^2 = \|Y_{n\times 1} - \boldsymbol{D}\boldsymbol{\theta}\|^2$ of the penalized least square (PLS) criterion $Q(\boldsymbol{\theta}, \beta)$ in the main manuscript Eq. (16) by $\|Y_{n\times 1} - \boldsymbol{X}\xi - g(\boldsymbol{X}\beta, A_{n\times 1})\|^2 = \|Y_{n\times 1} - \boldsymbol{X}\xi - \boldsymbol{D}\boldsymbol{\theta}\|^2 = \|Y_{n\times 1} - \tilde{\boldsymbol{D}}\tilde{\boldsymbol{\theta}}\|^2$, where $\tilde{\boldsymbol{D}} = [\boldsymbol{X}; \boldsymbol{D}]$ (the enlarged model matrix) and $\tilde{\boldsymbol{\theta}} = (\xi^\top, \boldsymbol{\theta}^\top)^\top$ (the enlarged coefficient vector). The constraint (5) is still imposed on g, based on the reparametrization of the design matrix $\boldsymbol{D}$ in (14) that incorporates the constraint (15). We then optimize the corresponding PLS criterion over $(\tilde{\boldsymbol{\theta}}, \beta)$, using the iterative procedure in Section 3.2 of the main manuscript.

Remark 2 *This (misspecified) linear model approximation (S.15) reflects a practical situation where we do not know the true functional form of* μ *(which is specified in (S.12)). Throughout the main manuscript (except in Section 6), we have simply treated* μ *as a zero function, i.e., we use approximation* $\mu(X) \approx 0$*.*

Remark 3 *For a general exponential family response* Y*, we can incorporate the approximation (S.15) for the* μ *term via the enlarged working model:* $h\big(m(X, A)\big) = \xi^\top X + g(\beta^\top X, A)$ *(again, subject to constraint (5) on* g*) replacing the working model (23) of the main manuscript. As Step 1 of the estimation procedure, for fixed* β*, we optimize* ξ *and* g *under the generalized additive model (GAM) framework, where we represent the "linear predictor" of GAM by the enlarged linear predictor* $\tilde{\boldsymbol{D}}\tilde{\boldsymbol{\theta}}$*, and optimize* $\tilde{\boldsymbol{\theta}}$ *via penalized likelihood maximization (in which the smoothing parameters are estimated via optimization of REML or GCV). As Step 2 of the estimation procedure, for fixed* $\tilde{\boldsymbol{\theta}}$*, we estimate* β*, via weighted least squares (as in Section 4 of the main manuscript) based on the working residuals and weights from the GAM fit obtained from Step 1; these two steps are iterated until convergence of the estimates for* β*.*

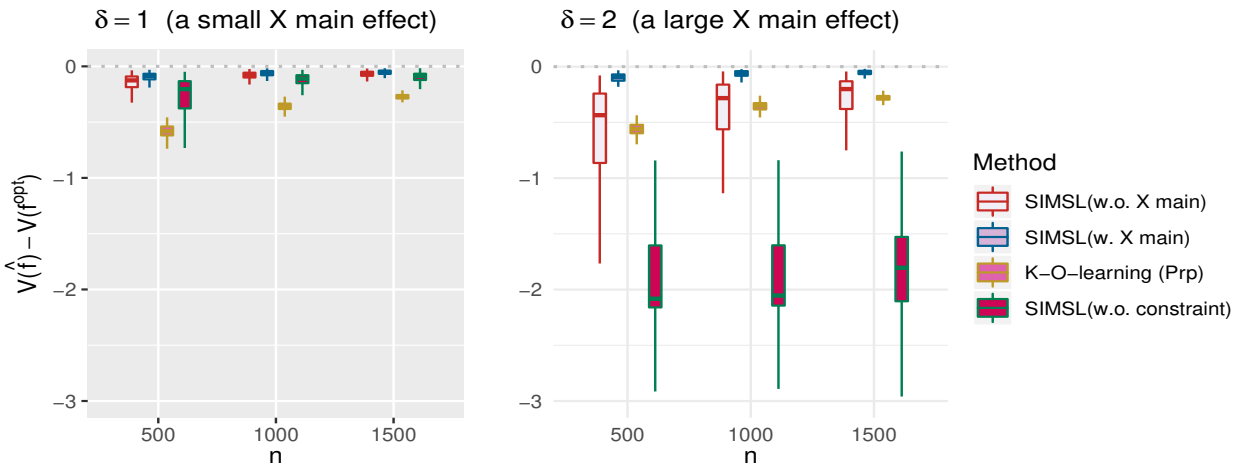

Figure S.2: Boxplots obtained from 200 Monte Carlo simulations comparing 4 approaches to estimating $\mathcal{D}^{opt}$, given each scenario indexed by $n \in \{500, 1000, 1500\}$ and $\delta \in \{1, 2\}$. A larger value is preferred.

3. "K-O-learning (Prp)", i.e., the propensity score-adjusted Gaussian kernel-based outcome-weighted learning, which was also considered in Sections 5 and 6, as a comparison method.

4. "SIMSL(w.o. constraint)", i.e., the unconstrained SIMSL *without* imposing constraint (5) (i.e., without constraint $\mathbb{E}[g(\beta^\top X, A)|X] = 0$ on g) on the estimation of (S.14). For the tensor product representation of g that we employed in this paper, the model without constraint (5) corresponds to the representation of g without applying the linear constraint (15) on the marginal basis matrix $\breve{\boldsymbol{B}}$ in (14) of the main manuscript; all the other estimation procedure given in Section 3.2 of the main manuscript remains unchanged for this approach.

For each simulation run, we obtain an estimate $\hat{f}$ for f_{opt} based on a training set (of size $n \in \{500, 1000, 1500\}$), and we evaluate its value (1) (i.e., $\mathcal{V}(\hat{f}) = \mathbb{E}[\mathbb{E}[Y|A = \hat{f}(X), X]]$) based on a testing set of size 10^4, using a Monte Carlo approximation. Since we know the true data-generating model in simulation studies, the optimal rule f_{opt} can be determined for each simulation run and accordingly, the optimal value $\mathcal{V}(f_{\text{opt}})$ can be computed. Given each estimate $\hat{f}$ for f_{opt}, we report $\mathcal{V}(\hat{f}) - \mathcal{V}(f_{\text{opt}})$ as the performance measure of $\hat{f}$. (We subtract the optimal value $\mathcal{V}(f_{\text{opt}})$ from the value of each estimate $\mathcal{V}(\hat{f})$ to facilitate the comparison across different simulation settings specified by different $\delta \in \{1, 2\}$.) A larger value of the measure indicates better performance.

The results in Figure S.2 indicate that, when the magnitude of the X main effect is small (i.e., when $\delta = 1$), all three SIMSL approaches perform at a similar level. However, when the magnitude of the X main effect is large (i.e., when $\delta = 2$), we can identify a substantial difference. Generally, if the variance of the X main effect dominates that of the A-by-X interaction effect, then the X main effect tends to heavily influence the estimation of the A-by-X interaction effect. In such a case (when $\delta = 2$), the SIMSL without the identifiability constraint (5), i.e., "SIMSL(w.o. constraint)," suffers a lot from the misspecification of the X main effect. This deterioration of the performance is due to the lack of the identifiability condition and the lack of the orthogonality between the X main effect and the A-by-X interaction effect, which lead to a serious confounding between the two effects in the estimation. On the other hand, despite the fact that "SIMSL(w.o. X main)" does not also model the X main effect, this approach suffers relatively little from the increased magnitude of the X main effect, in comparison to "SIMSL(w.o. constraint)." This indicates that the SIMSL with the proposed model identifiability constraint makes the approach robust to misspecification of the μ term.

As discussed in Section 6 of the main manuscript, the performance of SIMSL can be generally improved if the X main effect (i.e., the μ term) is incorporated to the estimation. By modeling the X main effect, we

can reduce the noise level in the model, and thereby improve the efficiency of the estimator for the g term. In Figure S.2, the SIMSL approach "SIMSL(w. X main)," where the X main effect is incorporated to the estimation, outperforms both "SIMSL(w.o. X main)" and "K-O-learning (Prp)," especially when a large X main effect is present in the underlying model (i.e., when $\delta = 2$).

C.2 Simulation illustration: Cumulative logit single-index models for optimizing dose rules, when the outcome is an ordinal categorical variable

In this section, we consider an extension of SIMSL when the treatment response Y is an ordinal categorical variable. When the treatment response's value exists on an arbitrary scale where only the relative ordering between different values is significant, assigning a numerical meaning to the expected treatment response (1) of the main manuscript (i.e., the so-called "value" function) and optimizing it over a space of f may only yield a suboptimal rule.

On the other hand, as discussed in Remark 1 of the main manuscript, the proposed SIMSL can be extended and appropriately modified to account for the ordinal and categorical nature of a response Y when optimizing individualized dose rules. In this section, we demonstrate the performance of the cumulative logit SIMSL with a simulation illustration.

We will use the same data generation models for the covariates $X \in \mathbb{R}^{10}$ and the dose A as in Section C.1, i.e., $f_{\text{opt}}(X) = 2\exp\big(-(\beta^\top X)^2\big)$ and $\beta = (8, 4, 2, 1, 0, \ldots, 0)^\top/\sqrt{85} \in \Theta$, which is the same optimal dose function as in Section C.1, and the same component functions g and μ used in (S.12).

The only difference is in the ordered categorical response $Y \in \{1, 2, \ldots, K\}$ generation. To generate Y, we first consider a latent response model

$$\tilde{Y} = -\big[\mu(X) + g(\beta^\top X, A)\big] + \tilde{\epsilon}, \tag{S.16}$$

and based on this, we generate a pre-categorized response $\tilde{Y} \in \mathbb{R}$. We then divide $\tilde{Y}$ into K categories, using the rule:

$$\alpha_{k-1} < \tilde{Y} \le \alpha_k \iff Y = k \quad (k = 2, \ldots, K-1) \tag{S.17}$$

(and $Y = 1$ for $\tilde{Y} \le \alpha_1$; $Y = K$ for $\tilde{Y} > \alpha_{K-1}$), and obtain an ordered categorical response $Y \in \{1, \ldots, K\}$. In (S.17), $\alpha_1, \alpha_2, \ldots, \alpha_{K-1}$ are $K-1$ cut-points, satisfying $\alpha_1 < \alpha_2 < \ldots < \alpha_{K-1}$, that determine the categorization of $\tilde{Y}$ in (S.16).

In model (S.16), let us assume that $\tilde{\epsilon}$ follows the standard logistic distribution (which is slightly more heavy-tailed than the standard normal distribution but is similar in shape). This formulation makes the latent model (S.16) with categorization (S.17) equivalent to the cumulative logit SIMSL (27) discussed in the main manuscript. Remark 4 below describes the equivalence.

Remark 4 *Model (S.16) implies that, for any $u \in \mathbb{R}$, we have*

$$P\big(\tilde{Y} \le u | X, A\big) = P\big(\tilde{\epsilon} \le u + \mu(X) + g(\beta^\top X, A) | X, A\big) = F\big(u + \mu(X) + g(\beta^\top X, A)\big), \tag{S.18}$$

where $F(u) = 1/(1+\exp(-u))$ is the cumulative distribution function (CDF) of the standard logistic distribution. Furthermore, (S.17) implies that $P\big(\alpha_{k-1} < \tilde{Y} \le \alpha_k | X, A\big) = P\big(Y = k | X, A\big)$. Then, by (S.18),

$$P(Y \le k | X, A) = P(\tilde{Y} \le \alpha_k | X, A) = F(\alpha_k + \mu(X) + g(\beta^\top X, A)) \quad (k = 1, \ldots, K-1), \tag{S.19}$$

where $F(\cdot)$ satisfies $F^{-1}(u) = \log(u/(1-u)) = \mathit{logit}(u)$, and thus we can write:

$$\mathit{logit}\big\{P(Y \le k | X, A)\big\} = \alpha_k + \mu(X) + g(\beta^\top X, A) \quad (k = 1, \ldots, K-1), \tag{S.20}$$

which is the cumulative logit SIMSL presented in the main manuscript.

To categorize $\tilde{Y}$ in (S.16) to $K = 5$ categories, we (arbitrarily) set $\{\alpha_1, \alpha_2, \alpha_3, \alpha_4\} = \{-2, -0.5, 0.3, 3\}$ and obtain Y that takes a value in $\{1, 2, 3, 4, 5\}$, under the rule (S.17). For this simulation study, as in

Section C.1, we vary $n \in \{500, 1000, 1500\}$ and $\delta \in \{1, 2\}$ (i.e., the magnitude of the X "main" effect $\mu(X)$ in (S.20) controlled by $\delta \in \{1, 2\}$ in (S.12)). For model (S.20), without loss of generality, let us assume that a larger value of Y is desired. The cumulative logit SIMSL (S.20) (and thus model (S.16)) indicates that a *smaller* value of $\{\mu(X) + g(\beta^\top X, A)\}$ implies a *greater* probability of $\{Y > k\}$ (for any $k = 1, \ldots, K-1$), and thus the cumulative logit SIMSL-based optimal decision rule is:

$$f_{\text{opt}}(X) \;=\; \underset{a \in \mathcal{A}}{\text{argmin}}\, \{\alpha_k + \mu(X) + g(\beta^\top X, a)\} \;=\; \underset{a \in \mathcal{A}}{\text{argmin}}\, \{g(\beta^\top X, a)\}.$$

In this simulation illustration, we will consider 4 approaches. 1) "SIMSL (Ordinal cat.)": the SIMSL approach based on the cumulative logit model (S.20) that properly accounts for the ordinal categorical structure of Y, optimized based on the procedure described in Remark 1 of the main manuscript; 2) "SIMSL (Continuous)": the SIMSL approach that simply treats the ordinal categorical response Y as a continuous variable, optimize based on the procedure described in Section 3.2 of the main manuscript; 3) "K-O-learning (Prp)": the Gaussian kernel-based outcome-weighted learning (propensity score-adjusted); 4) "L-O-learning (Prp)": the linear kernel-based outcome-weighted learning (propensity score-adjusted). In both "SIMSL (Ordinal cat.)" and "SIMSL (Continuous)", we incorporate the X main effect through a linear model approximation (S.15) for μ.

We will evaluate each estimator $\hat{f}$ of f_{opt} by evaluating its "value", $\mathcal{V}(\hat{f}) = \mathbb{E}[\mathbb{E}[\tilde{Y}|A = \hat{f}(X), X]]$, based on the latent $\tilde{Y}$ from the underlying model (S.16) (rather than based on the categorical response Y whose expectation does not have a valid numerical meaning). Again, we assume a larger value of $\tilde{Y}$ in (S.16) (and hence a larger value of Y) is desired, and thus a larger "value" $\mathcal{V}(\hat{f})$ is desired. Since we know the true data-generating model (S.16), the optimal rule f_{opt} can be determined for each simulation run and $\mathcal{V}(f_{\text{opt}})$ can be computed accordingly. Given each estimate $\hat{f}$ for f_{opt}, we report $\mathcal{V}(\hat{f}) - \mathcal{V}(f_{\text{opt}})$ as the performance measure of $\hat{f}$.

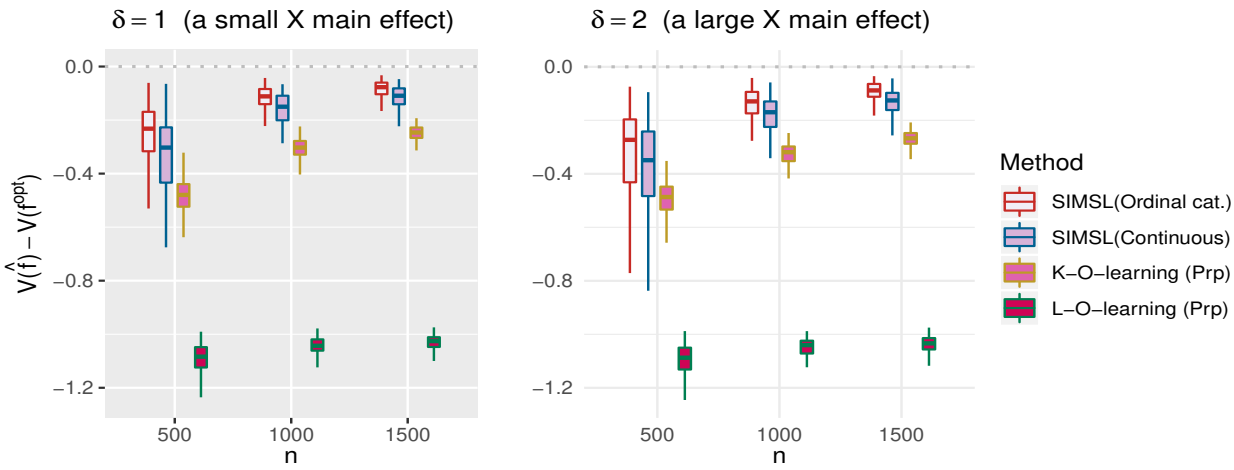

Figure S.3: Boxplots obtained from 200 Monte Carlo simulations comparing 4 approaches to estimating $\mathcal{D}^{opt}$, given each scenario indexed by $n \in \{500, 1000, 1500\}$ and $\delta \in \{1, 2\}$. A larger value is preferred.

The results in Figure S.3 indicate that "SIMSL (Ordinal cat.)" consistently outperforms "SIMSL (Continuous)" and "K-O-learning (Prp)," illustrating the advantage of properly accounting for the ordered categorical nature of Y in the estimation of f_{opt}. Given the nonlinear optimal dose rule (S.9), the linear approach "L-O-learning (Prp)" appears to be ineffective compared to the other more flexible methods. Furthermore, under model (S.20), the exponentiated difference between any two logits for two different values of dose A, say, $\exp\{g(\beta^\top X, a_1) - g(\beta^\top X, a_2)\}$, provides an intuitive cumulative odds ratio interpretation for comparing two dose levels (which can be visualized in terms of a single-index $\beta^\top X$), that can be useful in clinical practice.

C.3 Bootstrap confidence intervals

In Section 6 of the main manuscript, we compute a normal-approximation based 95% bootstrap confidence interval β_j, $\big(\hat{\beta}_j - 1.96\sqrt{\hat{\text{var}}(\hat{\beta}_j)},\ \hat{\beta}_j + 1.96\sqrt{\hat{\text{var}}(\hat{\beta}_j)}\big)$, where $\hat{\text{var}}(\hat{\beta}_j)$ denotes the sampling variance estimate of $\hat{\beta}_j$ $(j = 1, \ldots, p)$, obtained based on 500 bootstrap replicates of $\hat{\beta}$.

For each (the bth, say) bootstrap re-sample, we compute the corresponding estimate $\hat{\beta}^{(b)}$, where its first component $(\hat{\beta}_1^{(b)})$ is restricted to be positive for the model identifiability. This positivity constraint imposed on the estimate of β_1 will force the signs of the other single-index coefficient estimates of $\beta_2, \ldots, \beta_p$ to switch if the sign of β_1 is initially estimated to be negative. If $\beta_1 \approx 0$, then the variability of $\hat{\beta}_j^{(b)}$ $(j = 2, \ldots, p)$ across the bootstrap re-samples due to their sign switches can be rather dramatic. This may inflate the bootstrap-based estimate of $\text{var}(\hat{\beta})$. To prevent this inflation of the variance estimate, given the estimate $\hat{\beta}$ from the original sample and $\hat{\beta}^{(b)}$ from each bootstrap re-sample, we first evaluate two inner products, $\langle \hat{\beta}, \hat{\beta}^{(b)} \rangle$ and $\langle \hat{\beta}, -\hat{\beta}^{(b)} \rangle$, and then take either $\hat{\beta}^{(b)}$ or $-\hat{\beta}^{(b)}$ that gives the positive inner product, as the final bootstrap estimate for that re-sample (i.e., we match the "overall sign" of the bootstrap estimate $\hat{\beta}^{(b)}$ with that of the original estimate $\hat{\beta}$). Based on this, we build a 95% normal-approximated bootstrap confidence interval for β_j, $\big(\hat{\beta}_j - 1.96\sqrt{\hat{\text{var}}(\hat{\beta}_j)},\ \hat{\beta}_j + 1.96\sqrt{\hat{\text{var}}(\hat{\beta}_j)}\big)$, where $\hat{\text{var}}(\hat{\beta}_j)$ is estimated from such 500 bootstrap estimates $\hat{\beta}^{(b)}$ $(b = 1, \ldots, 500)$.

To evaluate the coverage probability of these bootstrap confidence intervals, we conduct a set of simulation experiments. We will use the same response model as in Scenario 1 of Section 5 of the main manuscript, i.e., using the mean model $\mathbb{E}[Y|X, A] = 8 + 4X_1 - 2X_2 - 2X_3 - 25(f_{\text{opt}}(X) - A)^2$. We set $X = (X_1, \ldots, X_5)^\top$ with each $X_j \sim \text{Uniform}[-1, 1]$, $A \sim \text{Uniform}[0, 2]$ and the additive noise $\epsilon \sim \mathcal{N}(0, 1)$, generated independently with each other. We consider two different optimal individualized dose rules, f_{opt}: 1) a *linear* rule, $f_{\text{opt}}(X) = 1 + \beta^\top X$; and 2) a *nonlinear* rule, $f_{\text{opt}}(X) = \cos\big(1 + \beta^\top X\big)$, where, in both cases, we set $\beta = (0.1, 0.2, -0.2, 0.5, -0.5)^\top$, standardized to have a unit L^2 norm. We vary the sample size $n \in \{100, 200, 400, 800, 1600\}$. For each simulation run, we fit the g term of the proposed SIMSL, $\mathbb{E}[Y|X, A] = \mu(X) + g(\beta^\top X, A)$ (subject to constraint (5)), by the optimization framework (9) of the main manuscript. In this simulation set, in order to illustrate the performance of the proposed method when $\mu(X)$ is misspecified (to be a zero function), we do not incorporate the X main effect, i.e., we simply treat μ as a zero function, and obtain an estimate of β of SIMSL without accounting for the X main effect. We compute a 95% normal-approximation bootstrap confidence interval for β using the procedure described above.

Table S.1: **Coverage proportion of bootstrap confidence intervals for single-index coefficients, when X main effect is not incorporated into the model**: The proportion of time (out of 200 simulation runs) that the 95% normal-approximation bootstrap contains the true value of β_j $(j = 1, \ldots, 5)$, with varying $n \in \{100, 200, 400, 800, 1600\}$.

	Coverage proportion (when X main effect is not incorporated)									
	Linear rule					Nonlinear rule				
n	β_1	β_2	β_3	β_4	β_5	β_1	β_2	β_3	β_4	β_5
100	0.905	0.930	0.925	0.920	0.930	0.970	0.685	0.705	0.715	0.720
200	0.945	0.955	0.950	0.945	0.950	0.925	0.805	0.825	0.820	0.820
400	0.945	0.955	0.955	0.940	0.955	0.915	0.890	0.845	0.900	0.890
800	0.925	0.960	0.960	0.935	0.955	0.890	0.885	0.875	0.880	0.870
1600	0.980	0.975	0.940	0.955	0.960	0.935	0.930	0.950	0.945	0.945

The results in Table S.1 indicate that, for the linear rule case, the "actual" coverage probabilities are

reasonably close (although there are some deviations) to the "nominal" coverage probability of 0.95, suggesting that the normal-approximation bootstrap confidence intervals are reasonably accurate for a simple linear dose rule. For the nonlinear rule case in which the single-index is related to the optimal rule in a nonlinear way, it tends to require a larger sample size in comparison to the linear rule case for the coverage probabilities to get close to the nominal level of 0.95. However, in both scenarios, the proposed bootstrap procedure appears to provide a reasonably accurate confidence interval with an increasing sample size.

To study the case where the X main effect is incorporated in the estimation, we perform an additional simulation experiment. We again utilize a linear model approximation (S.15) for μ to account for the X main effect, and estimate the "main effect-augmented" SIMSL, using the procedure described in Section C.1. The results in Table S.2 indicate that the coverage property of the bootstrap confidence intervals, especially for the nonlinear rule cases (which are generally more difficult to estimate), tends to improve when the X main effect $\mu(X)$ is accounted for, in comparison to the results in Table S.1. By explaining the variance associated with the X main effect, the main effect augmentation has the effect of reducing the underlying noise variance, and appears to facilitate an improved normal approximation for the distribution of the estimator for β.

Table S.2: **Coverage proportion of bootstrap confidence intervals for single-index coefficients, when X main effect is incorporated into the model**: The proportion of time (out of 200 simulation runs) that the 95% normal-approximation bootstrap contains the true value of β_j ($j = 1, \ldots, 5$), with varying $n \in \{100, 200, 400, 800, 1600\}$.

	Coverage proportion (when X main effect is incorporated)									
	Linear rule					Nonlinear rule				
n	β_1	β_2	β_3	β_4	β_5	β_1	β_2	β_3	β_4	β_5
100	0.915	0.965	0.950	0.955	0.960	0.915	0.915	0.935	0.940	0.940
200	0.940	0.930	0.925	0.940	0.925	0.975	0.940	0.955	0.960	0.940
400	0.920	0.925	0.940	0.945	0.925	0.935	0.925	0.925	0.970	0.970
800	0.905	0.945	0.945	0.940	0.935	0.940	0.935	0.975	0.950	0.940
1600	0.965	0.965	0.935	0.955	0.965	0.980	0.945	0.975	0.960	0.970

C.4 Potential shape constraints on the link surface g

Depending on context, for scientific interpretability of the model, it may be sometimes desirable to consider shape constraints on the surface-link, such as convexity of g as a function of A given X. For univariate smoothing, shape constraints such as monotonicity restrictions and convexity/concavity can be incorporated into smoothing within generalized additive models, via penalized smoothers based on appropriately re-parameterized B-splines and penalties (Pya and Wood, 2015). Literature on shape-constrained splines includes Ramsey (1988); Bollaerts *et al.* (2006); Meyer (2012). Specifically, Pya and Wood (2015) proposed shape-constrained P-splines, based on a novel nonlinear extension of the P-splines of Eilers and Marx (1996) with discrete penalties on the basis coefficients. Furthermore, using the concept of tensor product spline bases, Pya and Wood (2015) considered smooth functions of multiple covariates under the monotonicity constraint, where monotonicity may be assumed on either all or a selection of the covariates. For the concavity (convexity) constraint, Mazumdera *et al.* (2019) considered fitting a nonparametric smooth multivariate concave (convex) function, using smooth convex approximations to non-smooth convex least squares estimators. Other literature on the concavity (convexity) constrained multivariate smoothing includes Seijo and Sen (2011); Lim and Glynn (2012); Hannah and Dunson (2014).

Although the proposed SIMSL is developed on the basis of the established generalized additive regression framework, the extra nonlinearity induced by the use of shape-constrained P-splines (or in other constrained nonparametric smoothing) does not allow the unconstrained generalized additive modeling methods to be

reused or simply modified (Pya and Wood, 2015); substantially modified algorithms are required, especially for implementation of multidimensional smoothing. For this reason, in this paper we do not pursue the development of the concavity/convexity (or possibly monotonicity) shape-constrained SIMSL framework and leave that as future work. Instead, for an illustration of the potential applicability of shape-constrained SIMSL, we consider in this section a special case where each 1-dimensional slice of the surface-link $g(u, A)$ for a fixed value of $u = \beta^\top X$ is restricted to be a quadratic function (rather than a concave/convex function) of A.

We can implement this quadratic restriction of the function over the domain A, by using a quadratic polynomial marginal basis for the 2-dimensional tensor basis construction (12) of the main manuscript, instead of using the originally employed cubic B-splines basis ($\check{B}$) for the variable A. Under tensor product representation (12) for surface g, if we fix the single-index u and restrict $\check{B}$ to a quadratic polynomial basis, then $g(u, A)$ is a quadratic function of A.

An alternative approach to obtaining such a quadratic conditional fit, which we employ in our demonstration below, is to utilize the property of P-splines: the limit of a P-splines fit with strong smoothing is a polynomial (Eilers and Marx, 1996). Specifically, for large values of smoothing parameter and a penalty of order k, the fitted P-splines series will approach a polynomial of degree $k-1$ (if the degree of the B-splines is equal to or higher than k). Thus, given a cubic (i.e., $k = 3$) B-splines marginal basis $\check{B}$ for A, we will apply the difference penalty of order 3 (rather than that of order 2, which we have employed throughout the paper), and impose a very large smoothing parameter $\check{\lambda}$ (say $\check{\lambda} = 10^6$) associated with the 3rd order penalty matrix; this essentially restricts the g function over the domain A to be quadratic. On the other hand, we keep the surface along the single-index (u) direction to be flexible by leaving the associated smoothing parameter λ to be chosen based on optimizing RMLE from observed data.

We apply this quadratic shape-restricted SIMSL to the warfarin data considered in Section 6 of the main manuscript. The estimated surface g is given in Figure S.4, and the single-index coefficient β is estimated as $(0.08, -0.15, 0.05, -0.57, -0.19, 0.18, -0.22, 0.48, -0.38, -0.28, -0.27, 0.01, -0.02)^\top$. In this example, although the surface g along the variable A (for each slice $u = \beta^\top X$) is restricted to be quadratic, the overall estimated surface is quite similar to that in Figure 1 of the main manuscript (where no quadratic restriction along the direction A is imposed), although the detailed surface shape is different. However, the estimated dose rule from the fitted model was outperformed by the other SIMSL rules considered in Section 6 of the main manuscript: the mean (and sd) of the estimated Value for the dose rule, evaluated based on the 100 randomly split testing sets considered in Section 6 of the main manuscript is -0.246 (0.03), compared to -0.233(0.03) for the SIMSL with X main effect (see Figure 2 of the main manuscript for comparison). Generally, flexible smoothing of the link surface g of SIMSL helps us to identify an appropriate representation of the A-by-X interaction surface (which is optimized in conjunction with β), which can possibly inspire a more parsimonious parametrization of the link surface based on, for example, a low-degree polynomial.

D Supplementary information for Section 6 of the main manuscript

The fitted single-index coefficients β (and their 95% bootstrap confidence intervals) obtained with or without incorporating the X main effects for the application example discussed in Section 6 of the main manuscript are provided below.

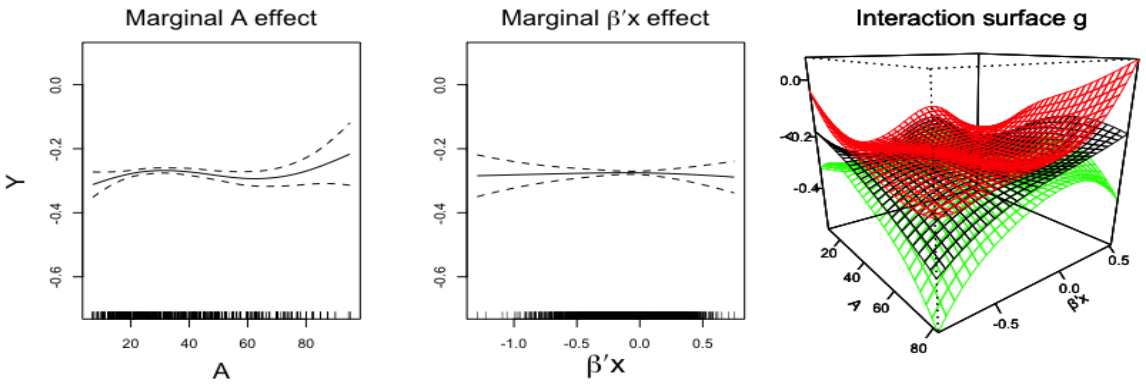

Figure S.4: The first two panels: the marginal effect of dose A (left panel) and that of the estimated single-index (middle panel) with 95% confidence bands (dashed curves) given the estimated $\beta^\top X$. The third panel: the estimated link surface (g) for the dose (A) and index ($\beta^\top X$) interaction; the red and green surfaces are at ± 2 standard error from the estimated surface (the black) in the middle, conditioning on the estimated single-index.

Table S.3: **Patient characteristics and the estimated single-index coefficients**: The patient covariates "Weight", "Height" and "Age" are standardized to mean 0 and unit variance when the single-index coefficient β (Coef.) is computed, and all the other patient covariates are indicator (0/1) variables (for indicator variables, mean refers to proportion). Associated with the estimated single-index coefficients are 95% bootstrap confidence intervals indicated by their lower bounds (LB) and their upper bounds (UB).

Patient	Mean	X main effect incorporated				X main effect not incorporated			
characteristics	(SD)	Coef.	LB	UB	*(signf.)	Coef.	LB	UB	*(signf.)
Weight (kg)	82.21 (22.93)	0.103	0.003	0.202	*	0.055	-0.005	0.116	
Height (cm)	170.21 (10.47)	-0.178	-0.341	-0.014	*	-0.149	-0.254	-0.044	*
Age (decade)	5.79 (1.50)	0.053	-0.038	0.144		0.033	-0.014	0.081	
Use of cytochrome P450 inducers (0/1)	0.02 (0.14)	-0.599	-1.065	-0.133	*	-0.528	-0.909	-0.147	*
Use of amiodarone (0/1)	0.09 (0.28)	-0.205	-0.531	0.120		-0.056	-0.309	0.197	
Sex (0: female; 1: male)	0.59 (0.49)	0.209	-0.012	0.429		0.146	0.015	0.278	*
African or black race (0/1)	0.19 (0.39)	-0.211	-0.454	0.031		-0.019	-0.207	0.170	
Asian race (0/1)	0.22 (0.41)	0.454	0.116	0.792	*	0.694	0.339	1.049	*
VKORC1 A/G genotype (0/1)	0.35 (0.47)	-0.319	-0.546	-0.091	*	0.021	-0.166	0.208	
VKORC1 A/A genotype (0/1)	0.27 (0.44)	-0.189	-0.472	0.094		0.321	0.073	0.570	*
CYP2C9 *1/*2 genotype (0/1)	0.13 (0.34)	-0.340	-0.658	-0.021	*	-0.257	-0.563	0.049	
CYP2C9 *1/*3 genotype (0/1)	0.09 (0.28)	-0.012	-0.262	0.238		-0.006	-0.233	0.220	
Other CYP2C9 genotypes† (0/1)	0.03 (0.15)	-0.088	-0.541	0.366		0.135	-0.244	0.513	

†The other CYP2C9 genotypes do not include the CYP2C9 *1/*1 genotype, which is taken as the baseline genotype.

References

Bell, M., Samet, J., and Dominici, F. (2004). Time-series studies of particulate matter. *Annual Review of Public Health* **25**:247–280.

Bollaerts, K., Eilers, P., and van Mechelen, I. (2006). Simple and multiple p-splines regression with shape constraints. *British Journal of Mathematical and Statistical Psychology* **59**:451–469.

Dockery, D. W. and Pope, C. A. (1996). *Epidemiology of acute health effects: Summary of time-series studies*, chapter Particles in Our Air, pp. 123–147. Harvard University Press.

Eilers, P. and Marx, B. (1996). Flexible smoothing with B-splines and penalties. *Statistical Science* **11**:89–121.

Hannah, L. A. and Dunson, D. B. (2014). Multivariate convex regression with adaptive partitioning. *Journal of Machine Learning Research* **14**:3261–3294.

Johnson, N. L., Kotz, S., and Balakrishnan, N. (1994). *Continuous univariate distributions*. New York: Wiley, 6th edition.

Lim, E. and Glynn, P. W. (2012). Consistency of multidimensional convex regression. *Operations Research* **60**:196–208.

Mazumdera, R., Choudhuryb, A., Iyengarc, G., and Sen, B. (2019). A computational framework for multivariate convex regression and its variants. *Journal of the American Statistical Association* **114**:318–331.

Meyer, M. (2012). Constrained penalized splines. *The Canadian Journal of Statistics* **40**:190–206.

Peng, R. and Dominici, F. (2008). *Statistical Methods for Environmental Epidemiology with R*. Springer.

Peng, R. and Welty, L. (2004). The NMMAPS data package. *R News* **4**.

Pope, C. A., Dockery, D. W., and Schwartz, J. (1995). Review of epidemiological evidence of health effects of particulate air pollution. *Inhalation Toxicology* **7**:1–18.

Pya, N. and Wood, S. (2015). Shape constrained additive models. *Statistics and Computing* **25**:543–559.

R Core Team (2019). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

Ramsey, J. (1988). Monotone regression splines in action (with discussion). *Statistical Science* **3**:425–461.

Ravikumar, P., Lafferty, J., Liu, H., and Wasserman, L. (2009). Sparse additive models. *Journal of Royal Statistical Society: Series B* **71**:1009–1030.

Seijo, E. and Sen, B. (2011). Nonparametric least squares estimation of a multivariate convex regression function. *Annals of Statistics* **39**:1633–1657.

Wood, S. N. (2017). *Generalized Additive Models: An Introduction with R*. Chapman & Hall/CRC, second edition.

Wood, S. N. (2019). gamair: Data for GAMs: An introduction with R. *R package version 1.0.2* .