



Practice of Epidemiology

Hierarchical Regression for Analyses of Multiple Outcomes

David B. Richardson*, Ghassan B. Hamra, Richard F. MacLehose, Stephen R. Cole, and Haitao Chu

* Correspondence to Dr. David B. Richardson, Department of Epidemiology, Gillings School of Global Public Health, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599 (e-mail: david.richardson@unc.edu).

Initially submitted September 2, 2014; accepted for publication February 11, 2015.

In cohort mortality studies, there often is interest in associations between an exposure of primary interest and mortality due to a range of different causes. A standard approach to such analyses involves fitting a separate regression model for each type of outcome. However, the statistical precision of some estimated associations may be poor because of sparse data. In this paper, we describe a hierarchical regression model for estimation of parameters describing outcome-specific relative rate functions and associated credible intervals. The proposed model uses background stratification to provide flexible control for the outcome-specific associations of potential confounders, and it employs a hierarchical “shrinkage” approach to stabilize estimates of an exposure’s associations with mortality due to different causes of death. The approach is illustrated in analyses of cancer mortality in 2 cohorts: a cohort of dioxin-exposed US chemical workers and a cohort of radiation-exposed Japanese atomic bomb survivors. Compared with standard regression estimates of associations, hierarchical regression yielded estimates with improved precision that tended to have less extreme values. The hierarchical regression approach also allowed the fitting of models with effect-measure modification. The proposed hierarchical approach can yield estimates of association that are more precise than conventional estimates when one wishes to estimate associations with multiple outcomes.

cohort studies; epidemiologic methods; models, statistical; Poisson regression; statistics

Abbreviations: HPD; highest posterior density; TCDD, 2,3,7,8-tetrachlorodibenzo-*p*-dioxin.

Occupational and environmental cohort mortality studies often examine associations between an exposure of primary interest and a number of different mortality outcomes. Typically, these exposure-mortality associations are estimated one at a time (1–6). Illustrative examples include analyses of cause-specific mortality among populations exposed to ionizing radiation, dioxin, and benzene (3, 6, 7). Often, the statistical precision of outcome-specific estimates is poor, particularly if there are relatively few exposed events for each type of outcome. In the extreme case, regression models for some outcome types may fail to converge because of sparse data.

A common strategy for dealing with sparse outcomes is to combine several outcome types into a broader category, such as combining specific cancer types into the category “all cancers,” and perform regression analysis on this broader outcome group. However, this strategy does not allow inferences regarding associations between exposure and specific

outcome types, is sensitive to decisions about how to combine outcome types, and imposes the assumption of homogeneity of association across the combined outcome types. An alternative to coalescing several outcomes into a broader category is to employ a hierarchical regression approach. Hierarchical regression can stabilize imprecise estimates of regression model parameters, does not assume homogeneity of association across outcome types, and in some settings allows estimates of cause-specific association to be obtained that were not estimable by fitting a separate model for each outcome type due to sparse data.

While hierarchical models are increasingly being used in epidemiologic analyses to deal with multiple explanatory variables (8, 9), there are fewer examples of their use in regression settings in which there is a primary exposure of interest and multiple outcomes under investigation (10). In the current paper, we describe a hierarchical Poisson regression

model for estimation of parameters describing outcome-specific relative rate functions. We present and illustrate the proposed approach using empirical data.

METHODS

Consider a cohort study in which mortality data for J outcome types have been ascertained over a period of follow-up. We use the term *outcome type* to refer to types or classes of outcomes, as in analyses of cancer deaths where events are classified by type of malignancy. Let Z denote levels of the exposure of primary interest and S denote levels of strata defined by other covariates. A typical analytical data structure generated for the purposes of Poisson regression analyses might consist of cumulative person-time, P , and total number of deaths due to $j = 1 \dots J$ outcome types, D^j , cross-classified by levels of covariates Z and S . A standard approach involves fitting a separate regression model for each type of outcome; for example, analysis of a single outcome type, $j = 1$, under a log-linear rate model may take the form $\lambda^1(\alpha_s^1, \beta^1) = \exp(\alpha_s^1 + \beta^1 Z)$, where $\lambda^1(\cdot)$ is the rate of outcome type $j = 1$, α_s^1 is the baseline rate in stratum $S = s$, and β^1 is the parameter describing the change in the log relative rate of outcome type $j = 1$ as a function of Z . Throughout this paper, the superscript j on parameters denotes their dependence on outcome type.

Alternatively, the J outcome types could be modeled jointly. A general relative rate model for the cause-specific mortality rates can be expressed as

$$\lambda(\alpha_s^j, \beta^j) = \exp(\alpha_s^j) \phi^j(Z; \beta^j), \quad j = 1, 2, \dots, J, \quad (1)$$

where α_s^j are the outcome-type-specific associations with covariates and $\phi^j(Z; \beta^j)$ is the function describing the change in the relative rate for outcome type j as a function of Z . Extensions of expression 1 may allow for effect modification of the relative rate function, $\phi^j(\cdot)$, by covariates.

For each cell of a grouped data tabulation defined by the cross-classification of exposure, Z , and covariates, S , there are D^j events of j outcome types; these are analyzed as though the numbers of observed events are independent Poisson random variables with mean values $P\lambda^j(\cdot)$. The likelihood for such models can be written as the product of the outcome-type-specific terms, each term being identical to the likelihood obtained if that outcome were the only one analyzed, with other outcomes treated as noninformative censoring events; consequently, a simple way to fit this joint model using standard software for maximum likelihood methods applied to Poisson regression is to construct an augmented data set in which outcome type is an additional dimension of the grouped data structure (11). The augmented data describe the cross-classification of person-time and events by levels of exposure, Z , covariates, S , and outcome type, j , where the number of events in each cell depends on outcome type j but the number of person-years at risk does not. The joint likelihood is obtained by fitting a regression model to these augmented data.

In expression 1, the parameters, α_s^j , are allowed to depend freely on strata. Because the number of parameters α_s^j may be

large, we employ conditional Poisson regression (12). This is equivalent to background stratification to obtain adjustment for the outcome-specific associations indexed by the parameters α_s^j . Conditioning on parameters α_s^j allows us to focus on estimation of the relative rate function, $\phi^j(\cdot)$.

Hierarchical models for cause-specific rate ratios

One approach to dealing with statistical imprecision in cause-specific estimates of association is to coalesce several sparse outcome types into a broader outcome category. In a cohort mortality analysis in which each individual can have a single outcome (as when decedents are classified according to their underlying cause of death), the practice of combining J outcome types into a broader category is equivalent to fitting a joint model for these outcome types with a constraint that these outcome-specific estimates of association are identical, $\beta^1 = \beta^2 = \dots = \beta^J = \beta$. Such constraints tend to confer the desired statistical stabilization; however, the implications of such constraints may be unappealing and implausible.

An alternative is to use a hierarchical model that allows the β^j parameters to be a function of the overall mean association and residual variation, such as

$$\beta^j \sim N(\delta, \tau^2), \quad \text{for } j = 1 \dots J, \quad (2)$$

where β^j describes the association between exposure and the j th outcome type, δ is the prior mean and is interpreted as the common mean association between exposure and the J outcome types, and τ^2 is the prior variance that allows for deviation of the outcome-specific association from the common mean. The dose-response coefficients β^1, \dots, β^J are shrunk toward the common mean value, δ ; however, the model has sufficient flexibility to automatically allow the outcome-specific estimates to deviate from this mean value if there is substantial evidence in the data to support it. As a consequence, the hierarchical regression approach will, in many situations, tend to result in outcome-specific estimates of association, $\hat{\beta}^1, \dots, \hat{\beta}^J$, that have lower mean squared error than those estimates obtained via a standard (one at a time) exposure-response analysis of each outcome type (9, 13, 14).

Extensions of the regression model in expression 2 may allow that the associations between exposure and the J outcome types are modified by 1 factor, or a set of factors, that define(s) background strata, S . Investigation of modification of associations between exposure and cause-specific mortality tend to exacerbate problems of statistical imprecision and, in many settings, result in problems of nonconvergence when analyses are conducted one outcome type at a time. Hierarchical regression models can yield stabilization of estimated parameters and may permit evaluation of effect-measure modifiers in settings where such estimates would be excessively unstable if the outcome types were modeled one at a time.

The degree to which the outcome-type-specific estimates are shrunk towards the common mean depends upon τ^2 , the variance for the outcome-specific associations (14). When τ^2 is large, there will be little shrinkage, and the outcome-specific estimates will be close to those obtained by fitting outcome types one at a time, as in a standard analysis; as τ^2 approaches 0, the fitted exposure-response associations will approach an

analysis conducted under a constrained model (e.g., an analysis that constrains all outcomes to have an association equal to the common mean). This variance parameter, τ^2 , can be treated as an unknown parameter in the hierarchical regression model so that the data directly inform the estimated value of τ^2 (14, 15). Details regarding implementation of this hierarchical model using the MCMC procedure in the SAS statistical software package (SAS Institute, Inc., Cary, North Carolina) are provided in Appendix 1. Simulation studies of hierarchical regression approaches have been reported previously (13, 16, 17); to illustrate the performance of the approach described in this paper, we present simulations in the Web Appendix (available at <http://aje.oxfordjournals.org/>), which includes Web Table 1.

Example 1: TCDD exposure and solid cancer mortality

To illustrate this approach, we use data from a cohort study that included 3,538 male workers who had been employed between January 1, 1942, and December 31, 1984, at 8 US industrial plants that produced chemicals contaminated with 2,3,7,8-tetrachlorodibenzo-*p*-dioxin (TCDD) (7, 18). Vital status follow-up was conducted through December 31, 1993, and deaths were classified by underlying cause of death, coded according to the revision of the *International Classification of Diseases* in effect at the time of death.

Prior analyses of this cohort have focused on the association between TCDD and all cancer mortality (7, 18). In the current analyses, we examined deaths due to 19 different categories of cancer. Estimated daily TCDD exposure was derived using a job-exposure matrix; the exposure scores reflected a quantitative exposure ranking of workers with respect to relative level of TCDD exposure among workers, rather than assignment of a specific dose of TCDD. For each worker, these scores were accumulated over time to derive a cumulative TCDD exposure score; and, for comparability with previously reported analyses of these data, cumulative exposure scores were log-transformed and lagged 15 years (7). A tabulation of the number of person-years and events was constructed and analyzed using Poisson regression methods with background stratification on attained age (in 5-year groups), birth cohort (in decades), race (nonwhite vs. white), and outcome type. First, we estimated a separate parameter for each outcome-specific association using maximum likelihood methods for Poisson regression, where the relative rate function conformed to the exponential model, $\varphi^j(d; \beta^j) = \exp(\beta^j d)$, and where d denotes the natural log of the cumulative TCDD score, lagged 15 years. Second, we fitted a hierarchical model in which the estimates, β^j , were shrunk to a grand mean, as in expression 2. The hierarchical regression model was fitted using a Markov chain Monte Carlo algorithm; in each model, computations were run for a minimum 100,000 iterations, with the first 10,000 iterations discarded to allow for initial convergence. A diffuse prior was specified for δ , and the variance parameter, τ^2 , was assumed to follow a uniform (0.001, 10) distribution. We chose these prior distributions to allow the data to drive inference as much as possible; however, for settings in which prior information is available, it should be incorporated. From Markov chain Monte Carlo samples, we estimated model coefficients and derived an estimate of the

associated 95% highest posterior density (HPD) credible interval, which is a Bayesian analog to the frequentist confidence interval. We performed sensitivity analyses in which the inverse of the variance parameter, τ^2 , was assumed to follow a gamma (0.01, 0.01) distribution.

Example 2: ionizing radiation and cancer mortality

To further illustrate this approach, we used data from a recent analysis of associations between radiation dose and incidence of solid tumors among 80,180 people included in the Life Span Study, a study of Japanese atomic bomb survivors. Cohort members had an assigned estimated radiation dose based on the most recent DS02 dosimetry system, were alive and not known to have had cancer before January 1, 1958, and had been in Hiroshima or Nagasaki, Japan, at the time of the atomic bombings in 1945 (19, 20). We used a publicly available cross-tabulation of person-time and counts of solid cancers by city, sex, radiation dose, distance from the hypocenter (within 3 km or 3–10 km), and categories of calendar time, attained age, and age at exposure. We focused on analysis of mortality due to 17 types of solid cancer.

A quantitative exposure score was assigned to estimate the change in relative rate given a 1-unit change in that exposure score. Rather than assign the same score for all outcome types, the assigned score, d^j , varies across outcome types because organ-specific doses have been estimated, based on the DS02 dosimetry system (21), and different organ-specific doses are used as the exposure score for different outcome types (22). Radiation dose–mortality associations were quantified using Poisson regression in which the relative rate function followed the linear excess relative rate model, $\varphi^j(d^j, \beta^j) = 1 + \beta^j d^j$. This linear excess relative rate model has been the preferred model for analyses of these data in recent major publications and National Academy of Sciences reports (23–25). As in prior analyses of these data (19, 20), parameters were included describing effect-measure modification by sex (a binary variable that took a value of –1 for males, else 1), s , attained age (centered at age 70 years), a , and age at exposure (centered at age 30 years), e , leading to a model of the form

$$\varphi^j(s, a, e, d^j) = 1 + \beta^j d^j \exp(\gamma_1^j a + \gamma_2^j e)(1 + \gamma_3^j s).$$

First, we estimated a separate parameter for each outcome-specific association using standard maximum likelihood; second, we fitted a hierarchical model in which the estimates, β^j , were shrunk to a grand mean, as in expression 2. As in prior analyses of these data (20), the parameters describing effect-measure modification by attained age, age at exposure, and sex, γ_1 , γ_2 , and γ_3 , were constrained to be equal to the values obtained for all solid cancers (i.e., these parameters were constrained to equal the values –1.50, –0.21, and 0.24, respectively). Finally, we fitted a model that allowed variation by outcome type in estimates of the parameters for effect-measure modification by attained age and age at exposure, γ_1 and γ_2 , under the model

$$\varphi^j(s, a, e, d^j) = 1 + \beta^j d^j \exp(\gamma_1^j a + \gamma_2^j e)(1 + \gamma_3^j s),$$

where $\beta^j \sim N(\delta, \tau^2)$, for $j = 1 \dots J$, and $\gamma_k^j \sim N(\theta_k, \sigma^2)$, for $k = 1, 2$ and $j = 1 \dots J$. The hierarchical regression models were fitted using a Markov chain Monte Carlo algorithm in SAS. A diffuse prior was specified for δ and θ_k , and τ^2 and σ^2 were assumed to follow a uniform (0.001, 10) distribution (14, 15). From Markov chain Monte Carlo samples, we estimated model coefficients and the 95% HPD interval. We performed sensitivity analyses in which the inverse of the variance parameter, τ^2 , was assumed to follow a gamma (0.01, 0.01) distribution.

RESULTS

TCDD cohort

When considering results obtained by standard Poisson regression, the maximum likelihood estimates of association between cumulative TCDD exposure (lagged 15 years) and cancers of the brain, pancreas, connective tissue, and larynx were among those of largest magnitude. The maximum likelihood estimates for cancers of the esophagus, rectum, peritoneum, and buccal cavity were negative. Figure 1A provides a Q-Q plot of the maximum likelihood estimates of the 19 cause-specific associations. These estimates appear approximately normally distributed; the mean and variance of the outcome-specific estimates of association were 0.06 and 0.03, respectively.

Cause-specific estimates of associations between cumulative TCDD exposure and site-specific cancer mortality obtained by hierarchical regression tended to have substantially tighter credible intervals than estimates obtained by standard Poisson regression (Figure 2A). None of the hierarchically

stabilized central estimates were negative. The estimated association between TCDD and lung cancer obtained using the hierarchical regression, which was relatively precise in the cause-specific analyses conducted by maximum likelihood, was very similar to the central estimate obtained by hierarchical regression; there was a small improvement in precision, as reflected by a slightly narrower 95% credible interval than the 95% confidence interval obtained using maximum likelihood methods. The estimates of association between TCDD and cancers of the brain, pancreas, larynx, and connective tissue (which were among the largest-magnitude associations estimated by standard Poisson regression) were shrunk towards the overall mean and stabilized (as reflected by a narrower confidence interval). The overall mean estimated exposure-response association was 0.10 (95% HPD credible interval: 0.01, 0.19); the variance parameter, τ^2 , was estimated as 0.01 (95% HPD credible interval: 0.00, 0.03). Similar results were obtained in sensitivity analyses specifying a gamma distribution for $1/\tau^2$.

Life Span Study

In analyses of the Life Span Study data, point estimates for all cancer types obtained through maximum likelihood regression were positive, except for cancer of the gallbladder. Figure 1B provides a Q-Q plot of the 17 outcome-specific estimates of association; the mean and variance of the outcome-specific estimates of association were 0.42 and 0.09, respectively. For most cancer outcomes, central estimates obtained using hierarchical regression were similar in magnitude to the maximum likelihood estimates obtained by standard Poisson




Figure 1. A) Q-Q plot of estimated change in the log relative rate of death due to 19 types of solid cancer per 1-unit increase in cumulative dioxin exposure (lagged 15 years) (obtained by maximum likelihood) in a cohort of US male chemical workers exposed to 2,3,7,8-tetrachlorodibenzo-*p*-dioxin, 1942–1993. B) Q-Q plot of the sex-averaged excess relative rate of death due to 19 types of solid cancer per sievert at an attained age of 70 years following exposure to the 1945 atomic bombings in Hiroshima and Nagasaki, Japan, at age 30 years (obtained by maximum likelihood), 1950–2000.

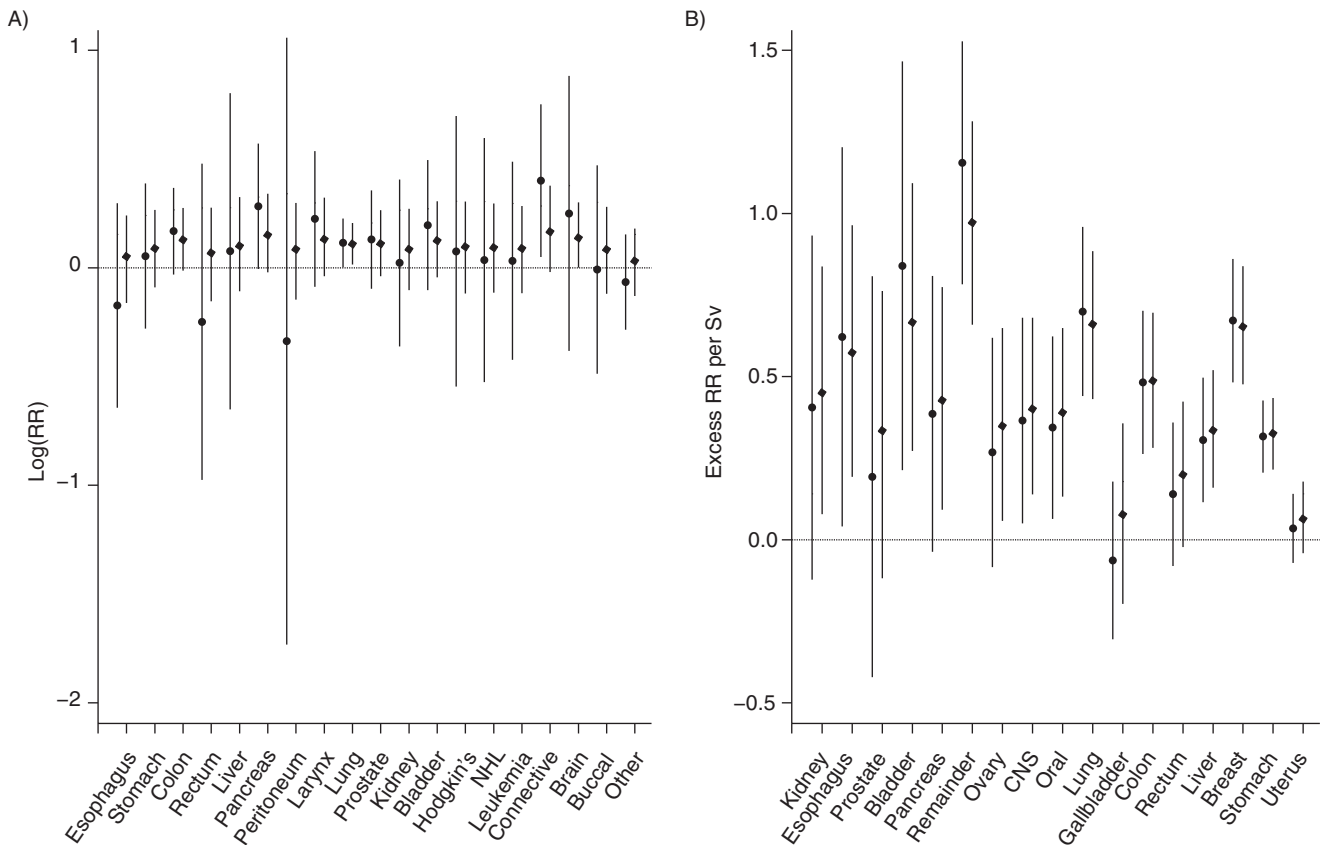


Figure 2. A) Regression model estimates of the change in the log relative rate (RR) of death due to 19 types of solid cancer per 1-unit increase in cumulative dioxin exposure (lagged 15 years) obtained by maximum likelihood (diamonds; 95% confidence intervals indicated by whiskers) and hierarchical regression with “shrinkage” towards the common mean (circles; 95% highest posterior density credible intervals indicated by whiskers) in a cohort of US male chemical workers exposed to 2,3,7,8-tetrachlorodibenzo-*p*-dioxin, 1942–1993. B) Regression model estimates of the change in the excess RR of death due to 17 types of solid cancer per 1-unit increase in radiation dose obtained by maximum likelihood (diamonds; 95% confidence intervals indicated by whiskers) and hierarchical regression with “shrinkage” towards the common mean (circles; 95% highest posterior density credible intervals indicated by whiskers) among survivors of the 1945 atomic bombings in Hiroshima and Nagasaki, Japan, 1950–2000. (“Remainder” represents cancers other than the site-specific cancers listed.) CNS, central nervous system; NHL, non-Hodgkin’s lymphoma.

regression (Figure 2B); however, cause-specific estimates obtained through hierarchical regression tended to have narrower credible intervals than the confidence intervals obtained by standard Poisson regression (Table 1). The estimated overall mean association was an excess relative rate per sievert of 0.43 (95% confidence interval: 0.26, 0.59); the estimate of τ was 0.09 (95% confidence interval: 0.02, 0.19). Similar results were obtained in sensitivity analyses specifying a gamma distribution for $1/\tau^2$.

We also fitted a model that allowed variation in estimates of the parameters γ_1 and γ_2 by outcome type (Table 1). For most cancer outcomes, central estimates obtained using this hierarchical regression were similar to the maximum likelihood estimates obtained by Poisson regression; and cause-specific estimates obtained by means of this hierarchical regression model also tended to have narrower credible intervals than the confidence intervals obtained by standard Poisson regression. The mean values for the parameters for modification by age at exposure and attained age were -0.17 and -1.63 , respectively.

DISCUSSION

Increasingly, hierarchical models are used to deal with multiple explanatory variables in an analysis, particularly when data are sparse or exposures are correlated (8, 26, 27). Less common is use of hierarchical models in a setting in which there is a single exposure variable of primary interest but multiple outcome types. However, often in occupational and environmental cohort studies there is a primary exposure of interest and a range of diseases or causes of death under investigation. This is also true in other substantive areas of epidemiology where investigators examine associations between a primary exposure variable of interest and multiple outcome variables.

Our proposed approach involves joint modeling of associations between the exposure and multiple outcome types. Some authors have advocated for joint modeling of occupational or environmental exposure-disease associations as an approach to reduce concerns about interpretation of *P* values and confidence intervals in settings of multiple comparisons

Table 1. Sex-Averaged Excess Relative Rate of Death Due to 19 Types of Solid Cancer per Sievert at an Attained Age of 70 Years Following Exposure to the 1945 Atomic Bombings in Hiroshima and Nagasaki, Japan, at Age 30 Years, Obtained Using 2 Different Statistical Methods, 1950–2000^a

Cancer Site or Type	Statistical Method					
	Maximum Likelihood (Model 1) ^b		Hierarchical Regression			
	ERR	95% CI	Model 2 ^c		Model 3 ^d	
ERR			95% HPD CrI	ERR	95% HPD CrI	
Kidney	0.40	−0.12, 0.93	0.46	0.09, 0.85	0.43	0.06, 0.80
Esophagus	0.62	0.04, 1.20	0.56	0.17, 0.97	0.52	0.14, 0.91
Prostate	0.19	−0.42, 0.81	0.33	−0.11, 0.76	0.33	−0.09, 0.76
Bladder	0.84	0.21, 1.47	0.60	0.22, 1.02	0.64	0.25, 1.06
Pancreas	0.39	−0.04, 0.81	0.42	0.09, 0.78	0.40	0.09, 0.74
Remainder ^e	1.15	0.78, 1.53	0.95	0.65, 1.25	0.87	0.53, 1.24
Ovary	0.27	−0.08, 0.62	0.35	0.05, 0.68	0.37	0.06, 0.72
CNS	0.37	0.05, 0.68	0.40	0.11, 0.68	0.37	0.10, 0.70
Oral	0.34	0.06, 0.62	0.39	0.14, 0.65	0.37	0.12, 0.64
Lung	0.70	0.44, 0.96	0.67	0.44, 0.92	0.67	0.46, 0.92
Gallbladder	−0.06	−0.31, 0.18	0.09	−0.17, 0.39	0.09	−0.18, 0.38
Colon	0.48	0.26, 0.70	0.49	0.28, 0.69	0.45	0.24, 0.69
Rectum	0.14	−0.08, 0.36	0.20	−0.01, 0.43	0.21	−0.03, 0.45
Liver	0.31	0.12, 0.50	0.33	0.16, 0.52	0.34	0.14, 0.56
Breast	0.67	0.48, 0.86	0.66	0.49, 0.85	0.69	0.45, 0.96
Stomach	0.32	0.21, 0.43	0.33	0.22, 0.44	0.33	0.20, 0.45
Uterus	0.03	−0.07, 0.14	0.07	−0.04, 0.19	0.06	−0.06, 0.18

Abbreviations: CI, confidence interval; CrI, credible interval; CNS, central nervous system; ERR, excess relative rate; HPD, highest posterior density.

^a Estimates were obtained by maximum likelihood (with 95% CI) and hierarchical regression with “shrinkage” towards the common mean (with 95% HPD CrI).

^b $\phi^j(s, a, e, d^j) = 1 + \beta^j d^j \exp(\gamma_1^j a + \gamma_2^j e)(1 + \gamma_3^j s)$, where $\gamma_1 = -1.50$, $\gamma_2 = -0.21$, and $\gamma_3 = 0.24$.

^c $\phi^j(s, a, e, d^j) = 1 + \beta^j d^j \exp(\gamma_1^j a + \gamma_2^j e)(1 + \gamma_3^j s)$, where $\gamma_1 = -1.50$, $\gamma_2 = -0.21$, $\gamma_3 = 0.24$, and $\beta^j \sim N(\delta, \tau^2)$.

^d $\phi^j(s, a, e, d^j) = 1 + \beta^j d^j \exp(\gamma_1^j a + \gamma_2^j e)(1 + \gamma_3^j s)$, where $\gamma_k^j \sim N(\theta_k, \sigma^2)$ for $k=1, 2, \gamma_3 = 0.24$, and $\beta^j \sim N(\delta, \tau^2)$.

^e Cancers other than the site-specific cancers listed.

(28). Others have advocated for joint modeling of exposure-outcome associations as an approach to facilitate modeling of, and formal statistical evaluation of heterogeneity in, the magnitudes of association between exposure and different types of outcomes (11). In contrast, in the current paper we focus on joint modeling of exposure-outcome associations in the context of a hierarchical model used to stabilize estimation of a set of cause-specific exposure-disease associations.

By describing a method that may help with estimation of associations between exposure and specific outcome types, we are not suggesting that there is anything inherently wrong with analyses of broader outcome groups constructed by aggregating several outcome types. Broad groupings of outcomes—for example, analysis of a category such as all cancer deaths—may be of interest from a public health perspective, or because the investigator believes that the sensitivity and specificity of classification of outcomes are acceptable for broad categories (such as all cancers) but not for narrower categories. Nonetheless, it is reasonable that concerns about statistical imprecision lead investigators to aggregate outcomes despite interest in outcome-specific estimates of association. In the latter

setting, a hierarchical regression approach may be useful for reducing mean squared error in the cause-specific estimates of association, and in some instances may permit estimation of associations that are not estimable when fitting regression models one outcome type at a time (29, 30). In the proposed hierarchical regression, the aggregation of outcome types into a broad group serves to help stabilize statistical estimates of outcome-specific associations, with the grouping treated as a type of prior information incorporated into the regression model.

In our first empirical example, there was substantial instability in most of the outcome-specific estimates of association derived using standard maximum likelihood methods, which precluded a strong conclusion regarding the direction and magnitude of most of the exposure-response functions for these outcome types (Figure 1). A hierarchical regression model was fitted that assumed a second-stage parametric function in which outcome-specific associations were random variables distributed around a common mean. This resulted in outcome-specific estimates of association that were stabilized relative to estimates derived via a standard analysis; and

for those outcome types that were relatively precisely estimated, the estimated exposure-response association was similar in magnitude to the estimate obtained via a standard analysis. In contrast, prior analyses of associations between TCDD and cancer mortality focused on the broad outcome group of all cancer mortality; such an analysis implicitly assumes a common association between exposure and all types of cancer. The hierarchical Poisson regression model results suggest that a number of cancer mortality outcomes, including deaths due to cancer of the colon, liver, lung, larynx, and brain, exhibited relatively precise positive associations with cumulative TCDD exposure in the hierarchical regression analysis. This illustrates how a hierarchical regression model provides a balance between these modeling approaches and may yield results that might tend to reduce mean squared error in resultant estimates.

In our second empirical example, involving radiation-cancer mortality associations among Japanese atomic bomb survivors, there was modest evidence of “shrinkage” or stabilization of estimated associations between exposure and different outcome types. Moreover, given the strong evidence of modification of the effects of ionizing radiation by factors such as age at exposure and attained age, we illustrated how the proposed hierarchical modeling approach allows for hierarchical modeling of effect-measure modifiers.

This hierarchical regression approach has similarities to empirical Bayes methods applied to a set of cause-specific estimates of association (10, 20). Under the empirical Bayes approaches employed in some recent papers, the investigator commences by deriving a collection of maximum likelihood estimates of the outcome-specific associations of interest and, in a second-stage analysis, employs an empirical Bayes method to stabilize these estimates by inverse variance weighting (20, 31). Such a 2-stage approach works if cause-specific estimates were reliably and validly estimated in the first stage. In contrast, joint modeling of the outcome in a hierarchical framework can overcome problems that arise when some outcome-specific estimates cannot be obtained because of poor model convergence. Furthermore, as illustrated in our second example, the hierarchical modeling approaches described here readily extend to more flexible modeling of the exposure-outcome associations.

This hierarchical approach is useful when a group of parameters to be modeled as following a normal distribution is carefully chosen, such that exposure-outcome associations within the group are similar. Under a hierarchical regression approach, the model represents the analyst’s belief about the pattern of variation in outcome-specific associations; however, the hierarchical modeling approach allows adjustment of some aspects of the parametric model for the exposure-response association to better conform to the data. In some settings, the exposure of primary interest may be protective for some outcomes and increase risk of other outcomes. In principle, this is not a problem, as long as the group of parameters can be appropriately modeled as exchangeable and as following the specified distribution. More generally, the assumption that a group of parameters can be modeled as following a specified distribution represents prior knowledge incorporated into the analysis; a consequence of assuming this hierarchical structure is a tendency to obtain more precise

credible intervals than would be obtained in the absence of such an assumption. Of course, if a critic disagrees with a choice of prior, this may help to focus attention on a specific issue of disagreement and suggests a way in which sensitivity analysis can be used to assess how different beliefs regarding the prior alter results.

Our illustrative models assumed a single normal distribution of outcome-specific associations. In some cases, the investigator may have a strong prior belief about the categorization of outcomes into specific groups. For example, β^1, \dots, β^k are parameters describing associations between exposure and 1 class of outcome and belong to 1 group, while $\beta^{k+1}, \dots, \beta^J$ are parameters describing associations between exposure and a second class of outcome and belong to a second group; this could be readily handled by extending expression 2 to a model for 2 common distributions, such as $\beta^j \sim N(\delta_1, \tau_1^2)$, for $j = 1 \dots k$, and $\beta^j \sim N(\delta_2, \tau_2^2)$, for $j = k + 1 \dots J$. Alternatively, we could extend expression 2 by applying a mixture prior that allows shrinkage in a more flexible fashion and allows the data to determine the classes to which the coefficients should belong.

We used vague, weakly informative priors in all analyses to more clearly show the performance of the modeling techniques. While in some situations researchers may have little prior information to include in analyses, this will not always be the case. Model performance can be further improved in many situations by formally incorporating any prior information that exists. These models were readily implemented using the SAS statistical package (Appendix 2). While we focused on Poisson regression analysis of tabulations of person-time and events, a similar approach could be developed for a hierarchical proportional hazards regression, following a similar approach of data augmentation to allow joint modeling of the cause-specific hazards. Further, our approach assumed that the outcomes of interest were independent, and we did not consider the possibility of competing risks. Extensions of this model to incorporate competing risks could make the method more broadly applicable.

In summary, the proposed hierarchical regression approach provides a useful complement to standard approaches for assessing outcome-specific exposure-response functions in epidemiologic cohort settings, and the results of such analyses may help inform conclusions regarding which outcome-specific estimates of association warrant further investigation. This method can be applied in a range of substantive investigations in which an investigator is interested in joint modeling of associations between an exposure of primary interest and a set of categories of cause of death.

ACKNOWLEDGMENTS

Author affiliations: Department of Epidemiology, Gillings School of Global Public Health, University of North Carolina at Chapel Hill, Chapel Hill, North Carolina (David B. Richardson, Stephen R. Cole); Department of Environmental and Occupational Health, School of Public Health, Drexel University, Philadelphia, Pennsylvania (Ghassan B. Hamra); Division of Epidemiology and Community Health, School of

Public Health, University of Minnesota, Minneapolis, Minnesota (Richard F. MacLehose); and Division of Biostatistics, School of Public Health, University of Minnesota, Minneapolis, Minnesota (Haitao Chu).

This project was supported in part by grant R03-OH-010056 from the National Institute for Occupational Safety and Health, Centers for Disease Control and Prevention. This project was also supported in part by an appointment to the Research Participation Program of the Environmental Protection Agency's Office of Research and Development, administered by the Oak Ridge Institute for Science and Education through an interagency agreement between the US Department of Energy and the Environmental Protection Agency.

Conflict of interest: none declared.

REFERENCES

- Pierce DA, Shimizu Y, Preston DL, et al. Studies of the mortality of atomic bomb survivors. Report 12, part I. Cancer: 1950–1990. *Radiat Res.* 1996;146(1):1–27.
- Wing S, Shy CM, Wood JL, et al. Mortality among workers at Oak Ridge National Laboratory. Evidence of radiation effects in follow-up through 1984. *JAMA.* 1991;265(11):1397–1402.
- Cardis E, Vrijheid M, Blettner M, et al. The 15-Country Collaborative Study of Cancer Risk among Radiation Workers in the Nuclear Industry: estimates of radiation-related cancer risks. *Radiat Res.* 2007;167(4):396–416.
- Lubin JH, Pottern LM, Stone BJ, et al. Respiratory cancer in a cohort of copper smelter workers: results from more than 50 years of follow-up. *Am J Epidemiol.* 2000;151(6):554–565.
- Hein MJ, Stayner LT, Lehman E, et al. Follow-up study of chrysotile textile workers: cohort mortality and exposure-response. *Occup Environ Med.* 2007;64(9):616–625.
- Rinsky RA, Hornung RW, Silver SR, et al. Benzene exposure and hematopoietic mortality: a long-term epidemiologic risk assessment. *Am J Ind Med.* 2002;42(6):474–480.
- Steenland K, Piacitelli L, Deddens J, et al. Cancer, heart disease, and diabetes in workers exposed to 2,3,7,8-tetrachlorodibenzo-*p*-dioxin. *J Natl Cancer Inst.* 1999;91(9):779–786.
- Greenland S. Hierarchical regression for epidemiologic analyses of multiple exposures. *Environ Health Perspect.* 1994;102(suppl 8):33–39.
- MacLehose RF, Dunson DB, Herring AH, et al. Bayesian methods for highly correlated exposure data. *Epidemiology.* 2007;18(2):199–207.
- Greenland S. A semi-Bayes approach to the analysis of correlated multiple associations, with an application to an occupational cancer-mortality study. *Stat Med.* 1992;11(2):219–230.
- Pierce DA, Preston DL. Joint analysis of site-specific cancer risks for the atomic bomb survivors. *Radiat Res.* 1993;134(2):134–142.
- Richardson DB, Langholz B. Background stratified Poisson regression analysis of cohort data. *Radiat Environ Biophys.* 2012;51(1):15–22.
- Witte JS, Greenland S. Simulation study of hierarchical regression. *Stat Med.* 1996;15(11):1161–1170.
- Greenland S. Principles of multilevel modelling. *Int J Epidemiol.* 2000;29(1):158–167.
- Gelman A. Prior distributions for variance parameters in hierarchical models. *Bayesian Anal.* 2006;1(3):515–533.
- Thurston SW, Ruppert D, Davidson PW. Bayesian models for multiple outcomes nested in domains. *Biometrics.* 2009;65(4):1078–1086.
- Xiao L, Thurston SW, Ruppert D, et al. Bayesian models for multiple outcomes in domains with application to the Seychelles Child Development Study. *J Am Stat Assoc.* 2014;109(505):1–10.
- Steenland K, Deddens J, Piacitelli L. Risk assessment for 2,3,7,8-tetrachlorodibenzo-*p*-dioxin (TCDD) based on an epidemiologic study. *Am J Epidemiol.* 2001;154(5):451–458.
- Preston DL, Ron E, Tokuoka S, et al. Solid cancer incidence in atomic bomb survivors: 1958–1998. *Radiat Res.* 2007;168(1):1–64.
- Pawel D, Preston D, Pierce D, et al. Improved estimates of cancer site-specific risks for A-bomb survivors. *Radiat Res.* 2008;169(1):87–98.
- Young RW, Kerr GD. *Reassessment of the Atomic-Bomb Radiation Dosimetry for Hiroshima and Nagasaki—DS02.* Hiroshima, Japan: Joint US-Japan Dosimetry Working Group, Radiation Effects Research Foundation; 2005.
- Shimizu Y, Kato H, Schull WJ. Studies of the mortality of A-bomb survivors. 9. Mortality, 1950–1985: part 2. Cancer mortality based on the recently revised doses (DS86). *Radiat Res.* 1990;121(2):120–141.
- Committee on the Biological Effects of Ionizing Radiation (BEIR V), National Research Council. *Health Effects of Exposure to Low Levels of Ionizing Radiation (BEIR V).* Washington, DC: National Academy Press; 1990.
- Committee to Assess Health Risks From Exposure to Low Levels of Ionizing Radiation, National Research Council. *Health Risks From Exposure to Low Levels of Ionizing Radiation: BEIR VII Phase 2.* Washington, DC: The National Academies Press; 2006.
- Preston DL, Shimizu Y, Pierce DA, et al. Studies of mortality of atomic bomb survivors. Report 13: solid cancer and noncancer disease mortality: 1950–1997. *Radiat Res.* 2003;160(4):381–407.
- Witte JS, Greenland S, Haile RW, et al. Hierarchical regression analysis applied to a study of multiple dietary exposures and breast cancer. *Epidemiology.* 1994;5(6):612–621.
- Witte JS, Greenland S, Kim LL. Software for hierarchical modeling of epidemiologic data. *Epidemiology.* 1998;9(5):563–566.
- Rothman KJ, Greenland S, Lash TL. *Modern Epidemiology.* 3rd ed. Philadelphia, PA: Lippincott Williams & Wilkins; 2007.
- Gelman A, Jakulin A, Pittau MG, et al. A weakly informative default prior distribution for logistic and other regression models. *Ann Appl Stat.* 2008;2(4):1360–1383.
- Hoerl AE, Kennard RW. Ridge regression: biased estimation for nonorthogonal problems. *Technometrics.* 1970;12(1):55–67.
- Preston DL, Krestinina LY, Sokolnikov ME, et al. How much can we say about site-specific cancer radiation risks? *Radiat Res.* 2010;174(6):816–824.

(Appendix follows)

APPENDIX 1

Hierarchical Poisson Regression

Assume we have an analytical data set, *one*, generated for the purposes of Poisson regression analyses consisting of counts of person-time, and counts of each type of event of interest, $d_{j=1}, d_{j=2}, \dots, d_{j=j}$, cross-classified by levels of covariates S and a primary exposure of interest, Z .

We can create an augmented table, *two*, that includes strata defined by S , Z , and j ; the analytical data structure consists of counts of person-time, pyr , and the number of deaths, $dths$, in each stratum defined by the cross-classification of levels of explanatory variables S , Z , and outcome type j . The number of deaths depends upon the outcome type, j , while the number of person-years is the same for all outcome types j within levels of the explanatory variable, Z and S .

In the proposed hierarchical model, background stratification is used to obtain adjustment for associations with covariates; this is implemented by conditional Poisson regression. To facilitate estimation of the conditional Poisson regression model, we can create a transformed analytical data structure, *three*, that includes 1 observation per stratum defined by cross-classification of covariates S and outcome type j , as has been described previously (12). We also create variables $_ncovals$ and $_totcases$ that denote the total number of exposure values and the total number of cases, respectively, in each stratum.

APPENDIX 2

SAS Code for Regression Models

A hierarchical regression model that shrinks the parameters for associations between exposure and outcome-specific mortality to a common mean (with background stratification on the covariates describing cause-specific covariate associations) can be readily fitted using the PROC MCMC procedure in the SAS statistical software package. Illustrative SAS code is provided below.

The arrays $_cases$, $_pt$, and z index the values for the counts of events, person-time, and levels of the exposure variable(s) of interest in each stratum of the analytical data structure. The length of the arrays will depend upon the analytical data structure. In the illustrative code below, the variables $_cases1$ - $_cases14$ describe the numbers of events observed in distinct levels of exposure within strata, $_pt1$ - $_pt14$ describe the associated numbers of person-years at risk within strata, and $z1$ - $z14$ describe the stratum-specific exposure scores. The “parms” statement defines the parameter(s) to be estimated. In this illustrative example, there are 15 outcome types; the parameters $b1$ - $b15$ estimate the outcome-type-specific associations of interest. These parameters are stabilized by specifying that they arise from a single normal distribution with mean δ and variance τ . The rate ratio function conforms to a standard log-linear model.

```
proc mcmc data=three nbi=5000 nmc=150000 ;
array c{*} _cases1-_cases14;
array p{*} _pt1-_pt14;
array z{1,14} z1-z14;
array b{15};
parms b: 1 ; parms delta 1; parms tau .1;
prior b: ~normal(delta,var=tau);
prior delta ~normal(0,var=100);
prior tau ~uniform(0.001,10);
num=1; den=0; nc= _ncovals ;
do i=1 to _ncovals ;
  phi= exp( b{k}*z{1,i} ) ;
  num=num*(phi**c{i}); den=den+phi*p{i};
end;
logl=log(num) - (_totcases*log(den)) ;
model _totcases~general(logl); run;
```